ABSTRACT

Title of dissertation: Search for Pair Production of Top Squarks in Proton-Proton Collisions at $\sqrt{s} = 8\,\text{TeV}$

Brian Michael Calvert, Doctor of Philosophy, 2015

Dissertation directed by: Professor Nicholas J. Hadley
Department of Physics

Supersymmetric extensions to the standard model can solve a number of current, unresolved issues in particle physics. In most of these models, the top-squark, the supersymmetric partner to the top quark, plays an integral role in fixing some of these issues. Although the existence of many supersymmetric particles have been strongly constrained by experiments, currently the existence of the top-squark remains largely unconstrained. This dissertation presents several searches for top-squark pair-production in R-parity conserving supersymmetry where the lightest neutralino is assumed to be stable. The data utilized in this search corresponds to 19.66 fb^{-1} of proton-proton collision data collected by the CMS experiment at $\sqrt{s} = 8$ TeV during the 2012 LHC run. The main focus of the dissertation is a search in the dileptonic final state, where the experimental final state is two leptons, two bottom quarks, and missing transverse momentum. Using a cut-based approach, no excess of events above the nominal background expectations is observed. This result is combined with a top-squark search in the semi-leptonic final state to exclude top-squark pair-production at the 95% confidence level for top-squark masses up to 700 GeV, when the lightest neutralino's mass is below 260 GeV. This dissertation also presents a powerful new approach to the dileptonic top-squark search.

Shape-based comparisons, using three complementary discriminating variables, between the observed data and the nominal background expectations achieve much better statistical sensitivity to top-squark pair-production in comparison with the cut-based search. Notably, the shape analysis excludes the existence of top-squarks that are nearly mass-degenerate with the top quark. Currently, no other direct top-squark search can achieve this exclusion. As well, there are a number of observed excesses in the shape analysis. The statistical significances of these excesses are tested against top-squark pair-production models. The subset of models where top-squark decays to a top quark and the lightest neutralino and the mass-splitting between the top-squark and the lightest neutralino is (150$\pm$12.5) GeV are found to fit with a statistical significance of $\sim$ 3.5–4σ. The global significance of these excesses is quantified by correcting for the look-elsewhere effect; the highest post-correction significances are found to be $\sim$ 2.5–3σ.

Search for Pair Production of Top Squarks in Proton-Proton Collisions at $\sqrt{s} = 8\,\mathrm{TeV}$

by

Brian Michael Calvert

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:
Professor Nicholas J. Hadley, Chair/Advisor
Professor Sarah Eno
Professor Hassan Jawahery
Professor Zackaria Chacko
Professor Alice Mignerey, Dean's Representative

Dedication

For my parents, who have been a constant source of love and support. To my father, who has epitomized discipline, dedication, and strength in his efforts to facilitate my mother's recovery. To my mother, who continues to fight her way forward along a long and difficult path.

Acknowledgments

This dissertation represents the culmination of what has been a challenging, scary, exciting, and ultimately integral five years of my life. The work that I poured into this dissertation would not have been even close to feasible without the generous support and guidance of others. This support became especially crucial halfway through my graduate career when, on May 16th, 2013, a horrible accident forced me to contend with the near-death of my mother. Her long, ongoing road to recovery has often required active attention on my part, and the entire experience has certainly colored the second half of my graduate career.

To that end, I must extend my sincere gratitude to my adviser, Professor Nick Hadley. Nick has consistently demonstrated compassion, understanding, and trust in me with regards to how I have managed the interplay of my research obligations and my mother's condition. He has allowed me the academic freedom to delve into my own novel research ideas, while still providing a wise, supervising voice of reason. Moreover, his advice and guidance on varying topics including statistics, particle physics, and intra-collaboration diplomacy have proven to be invaluable in my development as a physicist.

I must also acknowledge two of my other dedicated research mentors, Ted Kolberg and Matthieu Marionneau. They have maintained patience throughout my growth as a research scientist, and I have learned a great deal from our many discussions on physics and other topics.

I also acknowledge my analysis partners outside of UMD, including Chayanit, Ferdinando, Javier, and many others. Special mention should be paid to Pedrame

Bargassa, for his dedicated leadership on the lepton combination, and Geng-Yuan Jeng, for his invaluable knowledge and experience with the (often nuanced) statistical calculations that comprise my main research results.

Beyond those specific mentions, I also acknowledge the thousands of scientists across the LHC team and CMS collaboration, whose work has enabled the many interesting physics results from CMS. Special mention goes to the CMS MET working group, with whom I significantly improved my understanding of the CMS detector; and the CMS SUSY group, who have provided valuable commentary on the top-squark searches presented in this dissertation. I extend my gratitude to the members of my dissertation committee, who have helped me improve the quality of the writing and research presented in this dissertation.

The UMD CMS group has been my professional home for the last five years and I owe its members, both current and former, a great deal. Special mention goes to the professors, especially Sarah Eno and Alberto Belloni, for the many interesting discussions and guidance on research; Marguerite Tonjes and my other fellow cluster administrators for their disciplined dedication in keeping our Tier 3 cluster running (and their patience with me when I was far too busy to give proper help); and my fellow current graduate students, Young Ho Shin, Chris Anelli, as well as former graduate student, Kevin Pedro, for their friendship and mutual support.

Speaking of friends, I have had the good fortune while at UMD to meet and bond with many wonderful people. This includes my physics brothers at Seven Seas: Josh, Thomas, Jeff, Paul, Matt H., and Matt K.; my good friends Wendy and Sitara; my fellow Cubists; and too many others to name the rest individually.

Ali H. and Chara H. deserve special mention for being especially supportive during various points of the past two and a half years. I also especially want to thank Abraham Guevara and Estivins Campos, who, in my frequent trips back home to help my mother, have been family in all but name.

Finally, I acknowledge my family: my parents Greg and Juliette, and my sisters Misha and Stephanie. Although our lives were invariably changed two and a half years ago, my love for you has never been stronger.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

APD	Avalanche Photodiode
ASIC	Application-Specific Integrated Circuits
ATLAS	A Toroidal LHC ApparatuS
BDT	Boosted Decision Tree
BSM	Beyond Standard Model
CERN	European Organization for Nuclear Research
CH	Charged Hadron
CL	Confidence Level
CMS	Compact Muon Solenoid
CP	Charge-Parity
CPU	Central Processing Unit
CSC	Cathode Strip Chamber
CTEQ	Coordinated Theoretical-Experimental Project on QCD
CTF	Combinatorial Track Finder
DT	Drift Tube
DY	Drell-Yan
EB	ECAL Barrel
ECAL	Electromagnetic Calorimeter
EE	ECAL Endcap
EM	Electromagnetic
ES	ECAL Preshower
EWSB	Electroweak Symmetry Breaking
FCNC	Flavor-Changing Neutral Current
FPGA	Field-Programmable Gate Array
FSR	Final-State Radiation
GSF	Gaussian Sum Filter
GUT	Grand Unified Theory
HB	HCAL Barrel
HCAL	Hadron Calorimeter
HE	HCAL Endcap
HF	HCAL Forward
HO	HCAL Outer
HPD	Hybrid Photodiode
HLT	High-Level Trigger
IP	Interaction Point
ISR	Initial-State Radiation
L1	Level 1
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty
LO	Leading Order
LSP	Lightest Supersymmetric Particle
LUT	Lookup Table
MC	Monte Carlo

MET	Missing Transverse Energy
MIP	Minimum Ionizing Particle
MSSM	Minimal Supersymmetric Model
MSTW	Martin-Stirling-Thorne-Watt
MVA	Multivariate
NH	Neutral Hadron
NLO	Next-to-Leading Order
NNLO	Next-to-Next-to-Leading Order
NNLL	Next-to-Next-to-Leading Logarithmic
NPE	Non-parameter Functional Estimator
PC	Priestley-Chao
PDF	Parton Distribution Function
PE	Parametric Estimator
PF	Particle Flow
PMT	Photomultiplier Tube
PS	Proton Synchrotron
PSB	Proton Synchrotron Booster
PU	Pileup
QED	Quantum Electrodynamics
QCD	Quantum Chromodynamics
RF	Radio Frequency
RMS	Root Mean Square
RPC	Resistive Plate Chamber
SM	Standard Model
SPS	Super Proton Synchrotron
SS/OS	Same-Sign/Opposite-Sign
SUSY	Supersymmetry
TC	Track Counting
TEC	Tracker End Cap
TIB	Tracker Inner Barrel
TID	Tracker Inner Disks
TOB	Tracker Outer Barrel
TTC	Timing, Trigger and Control
T2bw	The $\widetilde{\mathrm{t}} \to \mathrm{b}\widetilde{\chi}_1^+ \to \mathrm{bW}^+\widetilde{\chi}_1^0$ top-squark decay-mode
T2tt	The $\widetilde{\mathrm{t}} \to \mathrm{t}\widetilde{\chi}_1^0 \to \mathrm{bW}^+\widetilde{\chi}_1^0$ top-squark decay-mode
VEV	Vacuum Expectation Value
VPT	Vacuum Phototriode
WIMP	Weakly Interacting Massive Particle
WLS	Wavelength-Shifting

Chapter 1: Introduction

For centuries, scientists have sought to understand the formation of the universe.

From our current understanding, the early universe was a very different place. Less than 10^{-10} seconds after the Big Bang, the universe was an expanding, hot – $\mathcal{O}\left(10^{16}\right)$ °C – soup of energy. In this ultra-hot energy bath, pairs of quantum particles popped in and out of empty space freely, traveling short distances before annihilating with one another. The majority of these quantum particles are very heavy and thus no longer trivially observable today, as they rapidly decay to lighter particles when they're at rest.

The photons within this energetic soup constantly interacted with the charged plasma surrounding them. Through these interactions, the photons integrated information about the underlying large-scale distributions of the matter in the early universe, but these photons were completely bound within the charged plasma, and thus this information was effectively inaccessible.

The universe continued to expand and subsequently cool down. Eventually, about 3.8×10^5 years after the Big Bang, it became energetically favorable for neutral atoms to form. At this point, the photons, previously locked in constant

interactions with the charged plasma, could propagate outward into the empty space of the universe, carrying with them the signatures of the underlying structure of the universe at this time.

We see these photons today: they comprise the Cosmic Microwave Background (CMB), a ubiquitous, near isotropic (at the level of 1 part in 10^{-5}) radiation permeating the entire observable universe. These CMB photons contain a copious amount of information about the structure of the early universe, and its evolution since then, information that accelerated the development of observational cosmology, starting in the 1960s and culminating (presently) in the standard model of Big Bang cosmology.

Studying the CMB is a powerful tool for understanding the early universe, but it is not the only one. In particular, through direct collisions of high-energy particles, we can recreate the high-energy conditions that existed less than a second after the Big Bang. In accordance with Einstein's famous equation, $E = mc^2$, these high-energy conditions can directly translate to the production of new, heavy quantum particles. These heavy particles decay to lighter daughter particles. Studying these daughter particles provides information about the heavy particles as well as the nature of their quantum interactions with one another. This tactic has been used to great success throughout the 20th century, notably resulting in the invention, in the 1960s, of the standard model (SM) of particle physics.

In the SM, the physical forces we see, like magnetic attraction or radioactive decay, are macroscopic manifestations of more microscopic quantum interactions that are governed by fundamental mathematical symmetries. Nominally, all of the fundamental particles in the SM would be massless. However, as far as we know, all

the fundamental particles, with the exception of the photon, the gluon, and possibly neutrinos, interact with a special particle, the Higgs boson. These interactions generate effective masses for the individual particles. The SM provides a remarkably accurate description (up to 1 part in 10^{12} [1]) of the quantum world around us.

For all its successes, however, the SM is not without its issues. One of the most important of these issues is the observed relative lightness of the Higgs boson. Just as the Higgs boson gives mass to other fundamental particles, these other fundamental particles can give mass to the Higgs boson through quantum corrections. In the current formulation of the SM, there is, at first glance, a relative imbalance in these quantum corrections. Without additional care, this relative imbalance leads to a predicted Higgs boson mass that is much larger than the experimentally observed mass. In the basic SM, restoring this balance, so that the predicted Higgs boson mass matches the observation, requires excessive fine-tuning of the parameters of the theory.

One proposed solution for this fine-tuning issue is supersymmetry. The crux of supersymmetry is that every fundamental SM particle is actually a member of a supersymmetric pair of particles. The interactions of each member of these pairs with the Higgs boson serve to balance out the quantum corrections to the Higgs boson mass, removing the need for any fine-tuning of parameters in the theory.

In order to test the validity of supersymmetry and other extensions to the SM, as well as to further probe the accuracy of the SM itself, tens of thousands of scientists and engineers have built the Large Hadron Collider (LHC), one of the current technological marvels of the world. In a ring 27 km in diameter, at depths up

to 100 m below ground, thousands of large, superconducting magnets, maintained at temperatures only a few degrees above absolute zero, steer tightly collimated beams of protons moving at 99.999997% of the speed of light. The stored energy in these magnets corresponds to 60 kg of TNT, enough energy to melt 550 kg of copper.

These magnets steer the beams of ultra relativistic protons to collide together in a space less than a meter in length. These collisions produce many particles that we already know about, such as the electron, photon, tau or top quark. As well, however, if SM-extending theories like supersymmetry are valid, these collisions should also produce new particles such as the predicted supersymmetric partners.

This has motivated an extensive data analysis program at the LHC. Analysts regularly search through the available data, rigorously combing through it to find any and all experimental signatures of the production of these proposed new particles.

In order to produce the data for these analyses, two of the major LHC experiments, ATLAS and CMS, use large, comprehensive detectors to record and analyze the proton collisions provided by the LHC. Individual sub-detectors measure all of the various kinds of particles that get produced in these collisions, and statistical claims are made about the likely existence (or non-existence) of new, non-SM particles by comparing our observed data against the expected signatures from the existing SM particles.

This dissertation presents a search for supersymmetry using data from the LHC. Data collected by the CMS experiment are analyzed to look for signatures of the top-squark, a proposed supersymmetric partner to the top quark. The coupling between the Higgs boson and the top quark is the strongest among all the funda-

mental particles in the SM. Thus the top-squark, as the partner of the top quark, plays one of the largest roles in balancing the aforementioned quantum corrections to the Higgs boson mass.

This search is difficult because the expected signature of the top-squark decay is a topology that looks almost exactly identical to the signature of a top quark decay. In the decays of both the top-squark and the top quark, for a wide range of top-squark masses, the same visible particles are produced with approximately the same energies. The main difference between the two particle decays is the presence of additional particles, invisible to normal detection means, in the decay of the top-squark.

Although these additional particles are invisible, their presence in any given collision event can be inferred through carefully measuring the overall momentum imbalance of the visible particles in the event. Special variables, specifically dedicated to exploiting this momentum imbalance and its correlation with the topology of the visible particles, facilitate a strong discrimination between the top quark background and the top-squark signal, ultimately improving the current experimental sensitivity to the existence, or lack thereof, of the top-squark.

An overview of the rest of this dissertation is as follows,

- Chapter 2 provides an overview of the SM.
- Chapter 3 quickly reviews some of the current, notable issues with the SM. As well, this chapter briefly describes supersymmetry, one of the proposed solutions to these issues.

- Chapter 4 gives an overview of the LHC. As well, this chapter contains a detailed description of the CMS experiment.

- Chapter 5 describes the object and event reconstruction algorithms used by CMS to turn its raw collision data into calibrated, precise pictures of each collision event.

- Chapter 6 gives a comprehensive description of a cut-based search for top-squark pair-production in the dileptonic final state. The search capitalizes on a powerful discriminating variable, the dileptonic stransverse mass, $M_{\mathrm{T2}}(\ell\ell)$ [2, 3], that exploits the correlations between the two leptons and the visible momentum imbalance to discriminate between the top quark background and the top-squark signal. This search serves as a baseline foundation for other top-squark searches that are presented in Chapter 7.

- Chapter 7 presents a number of extensions to the basic dileptonic top-squark search from Chapter 6. These extensions include a statistical combination of the cut-based dilepton top-squark search with a search for top-squark pair-production in the single lepton final state. This statistical combination improves the overall experimental sensitivity to top-squark pair-production in interesting areas of the supersymmetry parameter space.

  These extensions also include a upgrade to the cut-based dilepton top-squark search. This direct extension combines correlated, but complementary information from two additional M_{T2} variants into a 3D shape. Shape-based comparisons are then performed between the observed collision data and the

expected background.

These shape comparisons strongly improve the results of the dileptonic top-squark search; In particular, the 3D shape extension restricts top-squark pair-production in an experimentally challenging region of the SUSY parameter space, where the top-squark is nearly mass-degenerate with the top-quark. With this result, the extended dilepton top-squark search represents the *only* search for direct top-squark pair-production, across all possible top-squark final states, from the first run of the LHC that can exclude the existence of top-squark pair-production in this region.

Moreover, excesses are observed in certain regions of the 3D M_{T2} shape. Likelihood-based fits of top-squark signal models are used to quantify the overall statistical significance of these excesses. The strongest observed local significances are found to be around 3–4σ, and are consistent with relatively low-mass top-squarks, $m_{\tilde{\mathrm{t}}_1} \lesssim 300$ GeV, where the mass splitting between the top-squark and the lightest neutralino, another SUSY particle, is around 150 GeV. Two separate methods are used to calculate the global significances of these excesses. Both methods are found to yield similar results, and estimate the global significances to be ~ 2.5–3σ.

Chapter 2: Standard Model

The Standard Model (SM) of particle physics describes the currently known elementary particles and their interactions with one another through all of the fundamental forces except gravity.

2.1 General Overview of the SM

The matter that makes up our universe is comprised of elementary spin-1/2 fermions. The observed macroscopic interactions between these fermions – the electromagnetic, weak nuclear, and strong nuclear forces – are mediated at a quantum level by associated spin-1 vector bosons, where each force is associated with a specific set of vector bosons. The strengths of these interactions are governed by coupling strength parameters for each of the individual forces as well as individual charges associated to each particle.

A spin-0 (scalar) Higgs boson interacts with all of the fermions and some of the vector bosons. These interactions manifest as effective masses for these particles.

All of the particles in the SM have associated anti-particles, denoted with a bar over the particle symbol, à la $t \leftrightarrow \bar{t}$; these anti-particles have the same masses, couplings, and so forth as the associated particles, but their charges are opposite in

magnitude.

The fermions are further separated into two types, quarks and leptons, each of which has three generations, respectively comprised of an electroweak doublet of particles.

Quarks

The first generation of quarks consists of the up and down (u and d) quarks; these quarks, which have masses around a few MeV, are the primary constituents of the nucleons (protons, neutrons) that comprise atomic nuclei. The second generation of quarks are the charm and strange (c and s) quarks, with masses of $\sim$ 1.3 GeV and $\sim$ 0.1 GeV respectively. Finally, the third generation quarks are the top and bottom (t and b), with masses around $\sim$ 173 GeV and $\sim$ 4 GeV. The members of each generation possess electric charges of +2/3 and -1/3 respectively; these electric charges are associated with electromagnetic interactions. Similarly, each quark also possesses a color charge, which is associated with strong nuclear interactions and takes one of three values (red, green, or blue). These strong nuclear interactions drive a process known as hadronization, which results in the formation of hadrons, color-neutral bound states comprised of either three quarks (baryons) or a quark and antiquark (mesons).

Individual quarks can decay spontaneously to lighter quarks, so as long as the baryon number, $B = \frac{1}{3}(n_{\mathrm{q}} - n_{\overline{\mathrm{q}}})$, is preserved in the decay. These decays proceed through the weak force, which is the only interaction that does not preserve quark

flavor (e.g. a top quark can decay to a bottom quark and a bottom quark can decay to a charm or up quark).

Leptons

The three electrically charged leptons, with charge values of -1, are the electron (e), muon (μ), and tau (τ), with masses of 0.5 MeV, 0.1 GeV, and 1.8 GeV respectively. Each of these charged leptons has an associated, electrically neutral, left-handed neutrino; these neutrinos interact only through the weak force and thus are effectively impossible to detect except through specific, dedicated experiments. In the basic SM, these associated neutrinos are nominally massless (due to their chirality); however, the observation of neutrino oscillations has definitively shown that neutrinos are massive particles.

As with the quarks, individual leptons can decay spontaneously to lighter leptons through the weak force. These interactions preserve separately the total lepton number for each of the three lepton generations, $L_i = n_{\ell_i} + n_{\nu_i} - \left(n_{\bar{\ell}_i} + n_{\bar{\nu}_i}\right)$, so for example, the muon can only decay to a muon neutrino, electron, electron antineutrino, $\mu \rightarrow \nu_\mu + \mathrm{e} + \bar{\nu}_\mathrm{e}$. In neutrino oscillations, which are thought not to involve the weak force, the individual L_i are not preserved; instead, only the total lepton number, $\sum_{i=1}^{3} L_i$, is maintained.

2.2 Mathematical Structure of the SM: Particle Masses and Gauge Bosons

Both the observed particle masses as well as the interactions between the fermions and vector bosons arise naturally from the underlying mathematical structure of the SM, and thus discussions of either should invariably involve the other. The SM's formulation is that of a gauge quantum field theory: each "particle" is mathematically represented by a separate continuous field. The actual particles that we experimentally observe correspond to quantized excitations of these fields.

Lagrangian densities are used to describe the propagation of these particles. We can endow these particles with symmetries that act on internal degrees of freedom separate from the standard space-time degrees of freedom. The "gauge" aspect of the theory comes into play when we require these symmetries to be preserved under local transformations: the parameters of the transformation can have space-time dependence. Requiring that the associated Lagrangian densities respect these local symmetries naturally introduces into the Lagrangian interaction terms between the particle in question and a *gauge* field corresponding to a spin-1 *gauge* boson. The aforementioned spin-1 vector bosons of the SM are these gauge bosons. Depending upon the algebraic structure of the gauge group, these gauge bosons can also have interactions with one another.

From Noether's theorem, these conserved, local symmetries have an associated conserved charge, whose value is preserved in all interactions involving the gauge

boson. The strengths of these interactions are driven by the value of the conserved charge as well as an associated coupling constant.

Exact analytic solutions for QFTs are, in general, quite difficult, if not impossible. However, as long as the coupling strength for a given interaction is weak enough, perturbative calculations can yield extremely accurate results. A challenging aspect of these perturbative calculations is that they often contain divergent integrals. For a special class of QFTs, known as *renormalizable* QFTs, these divergences can be absorbed into the parameters of the Lagrangian, giving finite, testable predictions for the theory. A consequence of this renormalization and corresponding absorption of divergences is that the coupling "constants" of a renormalizable QFT are not actually constant, but instead change as a function of energy. This is known as the *running* of the coupling constant. We note here that all of the gauge QFTs associated with the SM are renormalizable.

Finally, gauge QFTs are strongly constrained: the only available choices are the specific symmetry group, the representation of the fermions in this symmetry group[1], and the specific values of the conserved charges. These choices subsequently define the exact nature and qualitative strength (i.e. running of the coupling) of all associated interactions and self-interactions in the theory.

The specific symmetry gauge group of the SM is the direct product of $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$.

Figure 2.1 contains a diagrammatic representation of the structure of the SM for the non-Higgs particles. All of the particles of the SM, except for the Higgs

[1]The gauge bosons are always in the adjoint representation of the gauge group

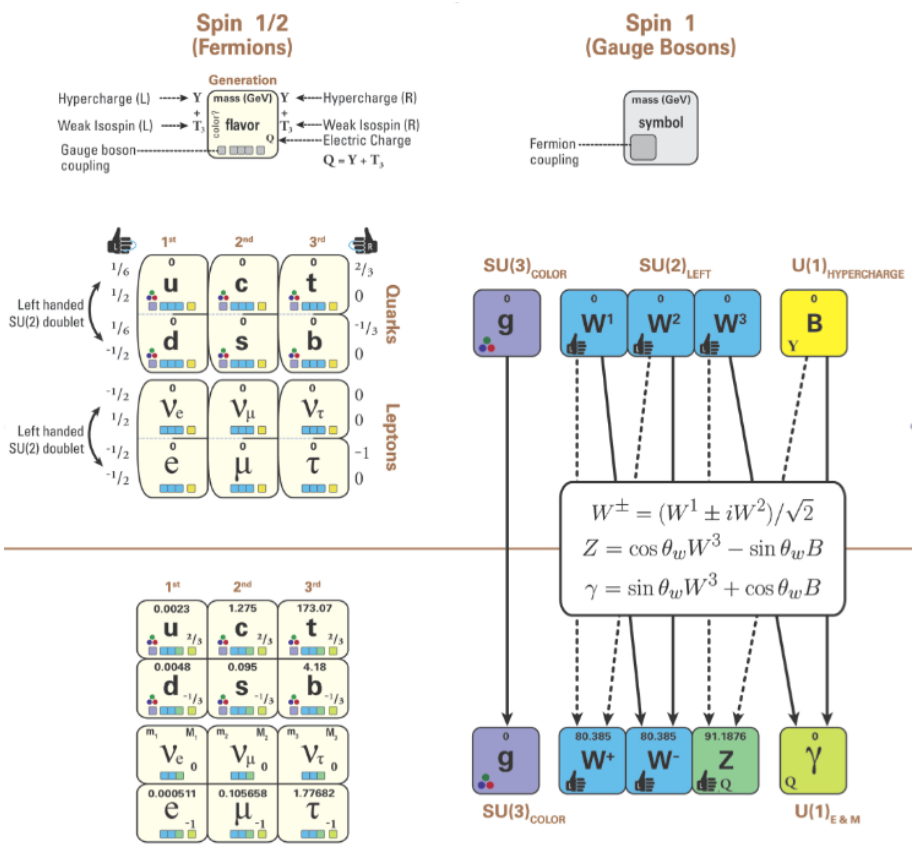

Figure 2.1: A diagram of the structure of the standard model (SM) of particle physics. All of the particles of the SM, except for the Higgs boson, are shown. The diagram also shows properties of each of these particles, including masses, spins, handedness, and so forth. These properties are shown both before and after the spontaneous breaking of the electroweak symmetry, represented by the brown horizontal line (see Fig. 2.2 or Sec. 2.2.2.1). Figure was created using parts of image from Ref. [4].

boson, are shown. The diagram also shows properties of each of these particles, including masses, spins, handedness, and so forth. These properties are shown both before and after the spontaneous breaking of the electroweak symmetry, represented by the brown horizontal line (see Fig. 2.2 or Sec. 2.2.2.1).

2.2.1 The Strong Nuclear Force: $SU(3)_C$

The strong nuclear force, associated with an $SU(3)_C$ gauge symmetry, can be described through the quantum chromodynamics (QCD) gauge theory. Associated with the non-Abelian $SU(3)_C$ symmetry group are three conserved charges, known as color and denoted by either (anti)red, (anti)green, or (anti)blue. The local $SU(3)$ gauge invariance introduces eight massless gauge bosons, the gluons, g, that mediate the strong force interactions. These gluons are bi-colored under the $SU(3)$ symmetry, carrying one unit of color and one unit of (a different) anti-color. Due to the $SU(3)_C$ group structure, the strength of α_S, the strong force coupling constant, decreases with increasing energy, which has important phenomenological consequences. At low energies, this α_S running manifests as *color confinement*. As two quarks travel apart from one another, the gluon field mediating their interactions forms tubes of color charge between them, in analogy with a taut elastic band. As the quarks continue to separate, these color tubes tighten further and eventually it becomes energetically favorable for a quark-antiquark pair to manifest from the vacuum at a point somewhere along this color tube. This process, known as hadronization, repeats recursively until, as mentioned above, all free color charge has been bound into color-neutral baryons or mesons. For highly energetic particles, hadronization results in a highly collimated spray of hadrons, commonly called jets [5].

In contrast with color confinement, at high energies, QCD exhibits *asymptotic freedom*, where the individual colored particles behave as free particles. The

relative transition between these two qualitative phenomena occurs at an energy scale $\Lambda_{\mathrm{QCD}} \sim 200$ MeV. This corresponds to a time scale for bare-color states of $\frac{\hbar}{\Lambda_{\mathrm{QCD}}} \sim \mathcal{O}\left(10^{-24}\mathrm{s}\right)$, or, alternatively, to distances of $\frac{\hbar c}{\Lambda_{\mathrm{QCD}}} \sim \mathcal{O}\left(10^{-15}\mathrm{m}\right)$. This effectively sets the maximum range of the strong force.

2.2.2 The Electroweak Force $SU(2)_L \otimes U(1)_Y$

At low energies, the electromagnetic and weak interactions appear to be separate interactions. This distinction only comes about from the spontaneous breaking of the electroweak $SU(2) \otimes U(1)$ symmetry by the Higgs mechanism, which will be discussed later in Sec. 2.2.2.1. At high energies, however, the electroweak gauge group symmetry is unbroken, and is separated into two components: weak isospin and weak hypercharge.

Weak Isospin

The $SU(2)_L$ gauge group describes weak isospin, with 3 generators corresponding to a triplet of massless gauge bosons, $\vec{W}$ or W_i, $i = 1, 2, 3$. The SM fermions fall into the simplest possible representations of $SU(2)$: left-handed (chirality) fermions are organized in the doublet representation with isospin, $T = \frac{1}{2}$, while right-handed fermions are organized into isospin singlets, $T = 0$, and thus do not interact with the W_i bosons.

In the unbroken electroweak theory, i.e. where the associated symmetries are still preserved, the different isospins for left- and right-handed fermions forbids the

insertion of explicit fermion mass terms, $m_f\bar{\psi}\psi$ into the SM Lagrangian, as these terms violate the isospin symmetry by mixing fermion fields of opposite chiralities.

Weak Hypercharge

Weak hypercharge is the conserved charge for the abelian group $U(1)_Y$. The neutral, massless B boson mediates interactions involving this charge. Nominally, since it is a $U(1)$ symmetry, the exact hypercharge values for particles can be arbitrary. However, at the quantum level, the $SU(2)\otimes U(1)$ electroweak theory possesses anomalies that do not cancel unless the $U(1)$ charges are exactly chosen based on the weak isospin values [6].

2.2.2.1 Electroweak Symmetry Breaking and the Higgs Mechanism

In the unbroken electroweak theory, all associated particles are massless. Empirically, however, we observe that all of the SM fermions, and three of the electroweak gauge bosons have mass. The naive, straightforward solution of directly inserting mass terms into the SM Lagrangian leads to issues, such as the possible violation of unitarity in electroweak gauge boson interactions at high enough energy scales [7, 8].

The answer to this problem is the Higgs mechanism [9, 10], a manifestation of a broader phenomenon known as spontaneous symmetry breaking. Spontaneous symmetry breaking arises in infinite dimensional, symmetric physical systems in which the minimum energy vacuum state is not symmetric. At low energies, the

symmetry of the system is spontaneously broken and massless, scalar bosons, known as *Goldstone bosons* appear in the theory, corresponding to excitations along the broken symmetry directions in the group space of the theory.

As an example, in the SM, the Higgs field, prior to EWSB, is a $SU(2)_L$ doublet, scalar, Φ, with weak hyper charge $Y = 1/2$. The potential for the complex Higgs field has a global, two-fold directional symmetry in the field components. In the vacuum state of this field, however, one of these directional symmetries is broken and four Goldstone bosons appear in the theory.

The Higgs field couples to the electroweak gauge bosons, and when the Higgs field symmetry is spontaneously broken, three of these four Goldstone bosons are absorbed as longitudinal polarization modes, and thus effective masses, for three of the electroweak gauge bosons. The remaining Goldstone boson appears as a massive, scalar particle, known as the Higgs boson.

Figure 2.2 contains a diagrammatic representation of the Higgs potential and the effect of the spontaneous electroweak symmetry breaking on the Higgs field.

Phenomenological Impact of EWSB in the SM

The photon (γ) mediates the electromagnetic force. The interactions of charged fermions with the photon manifest macroscopically as electromagnetism, but are described at a quantum level by quantum electrodynamics (QED), which is built on the Abelian $U(1)_{\mathrm{EM}}$ gauge symmetry. With this group structure, the photon is electrically neutral and α_{EM}, the coupling constant for QED, grows with increasing

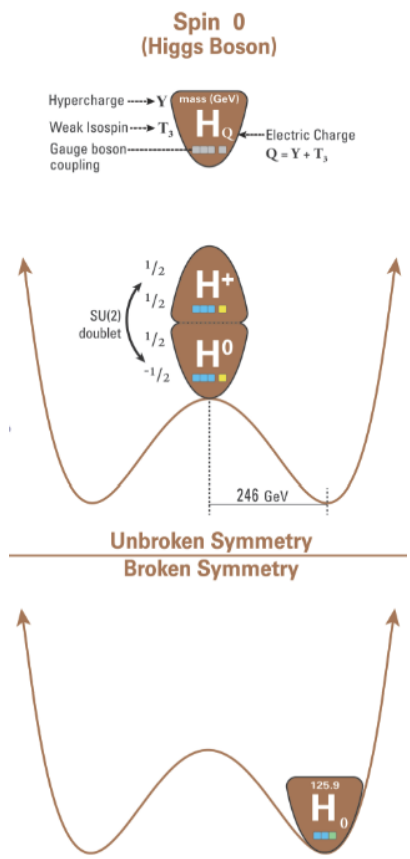

Figure 2.2: A diagram of the structure of Higgs sector of the standard model of particle physics. The diagram shows the properties of Higgs boson field, including its mass, spin, and gauge quantum numbers. The structure of the Higgs potential is also shown. The brown horizontal line represents the spontaneous breaking of electroweak symmetry, after which the Higgs boson particle manifests as the quantum excitation in this asymmetric vacuum ground state. Figure was created using parts of image from Ref. [4].

energy, in contrast with α_S. This means that, because the photon is massless, the electromagnetic force has infinite range.

The $W^{\pm}$ and Z bosons mediate the weak force, which manifests macroscopically primarily as radioactive decays of atomic nuclei. Because both the $W^{\pm}$ and Z bosons are massive, the macroscopic range of the weak force is extremely small, limited to approximately $\frac{\hbar c}{m_Z} \sim \mathcal{O}\left(10^{-18}\,\mathrm{m}\right)$. At a quantum level, the weak force interactions are related to the $SU(2)_L$ symmetry from the original, unbroken elec-

troweak group symmetry, $SU(2) \otimes U(1)$. Consequently, interactions involving the weak force preserve the third component of isospin. As well, even when the electroweak symmetry is broken, the distinction between fermion chiralities, due to their differing isospin values, remains to some degree. The charged current weak interactions ($W^{\pm}$) only involve left-handed fermions. The neutral current weak interactions (Z) involves both left- and right-handed fermions, but have different coupling strengths for each. Both the charged and neutral current weak interactions have a universal coupling strength parameter, α_{weak}. Similar to QED, this coupling strength grows with the energy scale of interactions.

Making Massive Fermions in the SM

As previously mentioned, the direct insertion of mass terms, $m_f \bar{\psi}\psi$, for the fermions is forbidden because these terms do not respect the isospin symmetry of the SM Lagrangian.

However, because the Higgs field is an $SU(2)_L$ doublet, Yukawa coupling terms like $\lambda_u u_L \Phi u_R$ respect the weak isospin symmetry and thus can be inserted into the SM Lagrangian. When the Higgs field acquires a VEV, in analogy with the gauge boson example above, these Yukawa coupling terms separate into a fermion mass term and an interaction term between the scalar Higgs boson and the fermion.

Chapter 3: Going beyond the SM

The SM is a remarkably successful effective field theory. In the subsequent decades since its invention, its theoretical predictions have been repeatedly confirmed through numerous experimental tests [1]. After the recent discovery of the SM Higgs boson, the SM is also a fully self-consistent theory.

3.1 Issues with the SM

That said, in its current form, there are a few problems with the SM. Some of these issues are only "aesthetic", and thus nominally do not require a direct answer. Others, however, represent distinct failures of the SM to describe the observable universe around us. In order for the SM, or some extension of it, to be a complete theory of Nature, this latter class of issues must be directly addressed. In this section, we discuss these issues with the SM. As well, we will describe a few "beyond the standard model" (BSM) theories for new physics that directly address these issues, both aesthetic and empirical, with the SM. Additional attention will be paid to one of these BSM models, supersymmetry, as it is the main motivation behind the searches described in Chapters 6 and 7.

3.1.1 Evidence for Issues from Observational Astrophysics

Observational astrophysics has provided several motivations, both empirical and aesthetic, for looking beyond the SM.

Dark Matter and Dark Energy

The most prominent of these is the observed presence of dark matter and dark energy in the universe around us. The first experimental evidence for dark matter came in the 1930's when Frist Zwicky studied the velocities of galaxies in the Coma cluster. He found that the average velocities of the galaxies were ≈ 160 times greater than would be expected from the observed luminosities [11]. The experimental evidence for dark matter has grown stronger in the subsequent decades.

Studies of the Bullet Cluster [12] – a merger of two separate galactic clusters – definitively show that there is some form of dark, non-baryonic matter in the universe. At the cosmological scale, cross-correlations between observed CMB (Cosmic Microwave Background) anisotropies and cosmological foregrounds corroborate the existence of dark matter [13, 14, 15]. In particular, maximum likelihood fits of the 6-parameter ΛCDM model of cosmology imply that the universe is composed of $\sim 5\%$ ordinary (e.g. baryonic) matter, $\sim 27\%$ dark matter, and $\sim 68\%$ "dark energy", a ubiquitous energy permeating the universe, repulsing against gravity and accelerating its expansion.

This implication of dark energy has been validated by direct comparisons across experiments [16]. As well, astrophysical studies of other cosmological ob-

jects, including BAO (baryon acoustic oscillations [1]) and SNe (Type 1a supernovae luminosities) further corroborate the CMB findings [17],

Neutrino Masses and Mixing

Ray Davis and John Bahcall's 1968 observed deficit of the solar neutrino flux provided one of the first pieces of evidence for neutrino flavor oscillations (and thus, massive neutrinos), which has since been confirmed by other experiments. This is at odds with the basic SM, as the basic SM does not contain right-handed neutrinos and thus predicts that neutrinos are massless.

Matter-Antimatter asymmetry

In addition to dark matter and dark energy, our observations of the universe around us have shown that there is a large matter-antimatter asymmetry. The CP violation in the weak sector of the basic SM is not strong enough to account for the magnitude of this asymmetry, and any dynamic explanations involving the early universe requires extensions to the SM [18].

3.1.2 SM Aesthetics

Beyond the direct empirical issues with the SM, there are also a number of "aesthetic" issues within the SM.

[1] the variations, with redshift z, in the density distributions of baryonic matter.

Fermion Generations and Yukawa Couplings

In the basic SM, the number of fermion generations is an arbitrary number[2]. Moreover, the Yukawa couplings of these fermions to the Higgs are nominally free, unpredicted parameters. The *flavor hierarchy*, i.e. the observed variation of these couplings from $\mathcal{O}(1)$ to $\mathcal{O}(10^{-6})$, like the number of generations, appears to be arbitrary.

Gauge Coupling Strengths, Unification, and Gravity

The difference in strength between the three gauge coupling parameters is also unexplained in the SM. Electroweak unification is an important and aesthetically pleasing component of the SM. It, together with charge quantization, suggest that the true theory of the universe might be a broader, Grand Unified Theory (GUT), that could unite all three fundamental forces [19].

Ideally, a GUT would also include the unification of gravity with the other 3 fundamental forces, as the basic SM does not include gravity. However, the graviton, the proposed spin-2 boson that mediates the gravitational force in quantum gravity theories, presents a particularly difficult challenge as it is difficult to construct renormalizable QFTs that include spin-2 particles.

Because gravity is $\mathcal{O}(10^{32})$ times weaker than the weak force, for most energy scales the lack of gravity in the SM does not matter. However, for energy scales larger than the Planck scale – $M_P \sim \mathcal{O}(10^{19}\text{ GeV})$ – gravity can longer be ignored

[2]Two interesting facts: 3 generations is the minimum number for CP violation in the SM. As well, the aforementioned fits to the ΛCDM model predict 3 neutrino flavors [13, 14].

in SM interactions. The Planck scale also plays a role in one the largest aesthetic issues in the SM.

The Higgs Mass, Fine Tuning, and the Gauge Hierarchy

Figure 3.1: One-loop diagrams that give corrections to the Higgs-boson mass parameter m_{H}^2 from a fermion f (left) and a scalar S (right) Reprinted from Fig. 1.1 of [20]..

The large splitting between the electroweak energy scale [$\mathcal{O}\left(10^2\text{ GeV}\right)$] and the Planck scale [$\mathcal{O}\left(10^{19}\text{ GeV}\right)$] presents one of the most challenging aesthetic issues in the SM. This is illustrated below.

Consider the interaction of a SM Dirac fermion f with the Higgs boson. When EWSB occurs and the Higgs field gains a VEV [Sec. 2.2.2.1], the strict distinction between fermion chiralities is broken by the spontaneous generation of the mass term from the Higgs-fermion Yukawa interaction, $-\lambda_f \bar{f} H f$. In addition to giving the fermion mass, this Yukawa interaction also contributes to quantum corrections to the Higgs mass parameter m_H.

For example, the one-loop Feynman Diagram on the left side of Fig. 3.1 yields a correction of the form

$$\Delta m_{\mathrm{H}}^2 = -\frac{|\lambda_f|^2}{8\pi^2}\Lambda_{\mathrm{UV}}^2 + \cdots\,, \tag{3.1}$$

where Λ_{UV} is the UV momentum cutoff used to regulate the loop integral. In practice, Λ_{UV} represents the energy scale at which new, BSM physics would enter, altering the high-energy theory and ostensibly removing the need for this integral regularization.

If Λ_{UV} is at the GUT scale [$\mathcal{O}\left(10^{16}\text{ GeV}\right)$] or Planck Scale [$\mathcal{O}\left(10^{19}\text{ GeV}\right)$], then this quantum correction to m_{H}^2 is $\mathcal{O}\left(10^{24}\right)$ – $\mathcal{O}\left(10^{30}\right)$ times larger than the physical value of the Higgs mass of ≈ 125 GeV [21].

It could be that the Lagrangian parameter m_{H}^2 is tuned in Nature so as to cancel the majority of Λ_{UV}^2, leaving behind a physical Higgs mass in the theory that agrees with current experimental results. However, this would require fine-tuning at the level of approximately 1 part in 10^{24} or greater. In a technical sense, nothing prevents this from occurring and this *gauge hierarchy problem* is seemingly just an issue of aesthetics but, to paraphrase one physicist, "...it would be like walking into a room to find a pencil standing perfectly balanced on its sharpened tip with no visible strings or other outside help. It technically could happen, but you would probably believe instead that there's some subtle, unknown thing working behind the scenes and keeping everything together."

One nominal solution to the gauge hierarchy problem is to pick a Λ_{UV} that is not too large, reducing the magnitude of parameter tuning required. There are a number of immediate concerns with this solution [20], but even in the case that it works, the m_{H}^2 will still receive corrections from any heavy, new particles.

For example, one could imagine there exists some heavy, complex scalar S, with mass m_S that couples to the Higgs with a Lagrangian term $-\lambda_S|H|^2|S|^2$. There

will be one-loop interaction diagrams between this new scalar and the Higgs that yield Feynman diagrams such as the right side of Fig. 3.1. This diagram yields a correction to m_{H}^2,

$$\Delta m_{\mathrm{H}}^2 = \frac{|\lambda_S|^2}{16\pi^2}\left[\Lambda_{\mathrm{UV}}^2 - 2m_S^2 \log\left(\Lambda_{\mathrm{UV}}/m_S\right) + \cdots\right]. \tag{3.2}$$

Although one can use dimensional regularization to remove the Λ_{UV}^2 term, it is much more difficult to remove the m_S^2 term. Similar arguments exists for other, arbitrary, new heavy particles, implying that the m_{H}^2 parameter is likely sensitive to the heaviest particles that the Higgs couples to. This issue persists even if the Higgs does not *directly* couple to heavy new particles. If the Higgs indirectly couples the heavy new particles, there are still cases where m_{H}^2 is quadratically sensitive to either Λ_{UV}, these new particles' masses, or both [20].

Looking at the rest of the SM, naively, one might worry that, like the Higgs, the masses of the chiral fermions and gauge bosons in the SM are also subject to these large, quadratic corrections. However, for both of these classes of particles, a form of symmetry in the theory serves to protect the masses.

For the chiral fermions, the explicit distinction, prior to EWSB, between chirality states (respecting the $SU(2)_L$ symmetry in the weak sector) prohibits the insertion of terms such as $\Lambda_{\mathrm{UV}}\bar{\psi}\psi$ into the Lagrangian. Consequently, Λ_{UV} can only enter into any fermion mass corrections through loop integrals, which, at worst, have logarithmic divergences for fermions, i.e. $\Delta m_f \sim \log\left(\Lambda_{\mathrm{UV}}/m_f\right)$.

For the gauge bosons, local gauge invariance plays a similar role, preventing the direct insertion of mass terms involving Λ_{UV}.

Using these examples as inspiration, for the scalar Higgs, some new kind of symmetry could be playing the role of the "subtle thing...keeping everything together". A number of BSM models propose new symmetries that mitigate this gauge hierarchy problem. Perhaps the most famous of these new symmetries is *supersymmetry* (SUSY), a proposed new symmetry between fermions and bosons.

3.2 Supersymmetry

In the simple example above of the gauge hierarchy problem, if the theory is made supersymmetric, then every fermion receives a scalar partner particle for each of its degrees of freedom (so 2) with the exact same Yukawa coupling to the Higgs boson, $\lambda_S = \lambda_f$. The total correction to the m_{H}^2 parameter is the net sum of the individual corrections; this leads to a cancellation of all terms with Λ_{UV}^2 in Eqs. (3.1) and (3.2). If SUSY is an unbroken symmetry of the theory, then these partner particles also share the same masses and the logarithmically divergent terms (and all other divergent terms) perfectly cancel as well.

Simple SUSY: the Wess-Zumino model

We can illustrate some basic, but notable aspects of SUSY through consideration of the Wess-Zumino Lagrangian [22],

$$\mathcal{L}_{\mathrm{free}} = -\partial^\mu \phi^* \partial_\mu \phi + i\psi^\dagger \bar{\sigma}^\mu \partial_\mu \psi + F^* F. \tag{3.3}$$

This Lagrangian describes a physical, massless, complex scalar ϕ, and a physical, massless, two-component Weyl fermion ψ. In the Lagrangian, there is also an "aux-

iliary" field F, whose sole purpose is to ensure that the Lagrangian respects the requirement of supersymmetry even when particles are propagating off-shell[3]

It can be shown that the Wess-Zumino Lagrangian is invariant under the supersymmetry transformation,

$$\begin{aligned} \delta\phi &= \epsilon\psi, & \delta\phi^* &= \epsilon^\dagger\psi^\dagger, \\ \delta(\psi)_\alpha &= -i(\sigma^\mu\epsilon^\dagger)_\alpha\,\partial_\mu\phi + \epsilon_\alpha F, & \delta(\psi^\dagger)_{\dot\alpha} &= i(\epsilon\sigma^\mu)_{\dot\alpha}\,\partial_\mu\phi^* + \epsilon^\dagger_{\dot\alpha}F^*, \\ \delta F &= -i\epsilon^\dagger\bar\sigma^\mu\partial_\mu\psi, & \delta F^* &= i\partial_\mu\psi^\dagger\bar\sigma^\mu\epsilon\,. \end{aligned} \tag{3.4}$$

where ϵ represents the infinitesimal, anti-commuting, two-component spinor that parametrizes this transformation.

Under these supersymmetry transformations, the boson field transforms into the fermion field and vice-versa. These particles are subsequently called *superpartners* of one another, and are commonly grouped into *supermultiplets*, irreducible representations of the underlying supersymmetry algebra governing the transformations. In the particular case of the Wess-Zumino Lagrangian, the supermultiplet of ϕ and ψ is known as a *chiral* supermultiplet, and is the simplest possible supermultiplet that can be constructed. If, instead of the scalar ϕ, we had paired ψ with a gauge boson, we would call this a *gauge* supermultiplet. Other supermultiplets can be constructed, e.g. pairing a massless spin-2 boson with a massless spin-3/2 fermion. Each always involves a boson-fermion pair or set of pairs where the partners differ in spin by 1/2.

Note that, for a chiral supermultiplet, although we did not do so for this

[3]Off-shell, ψ has 4 degrees of freedom, while on-shell it has only 2. On the other hand, ϕ always has 2 degrees of freedom, both on- and off-shell.

example, one can include explicit mass terms for either ϕ or ψ, and the Lagrangian will still respect the supersymmetry transformations.

By Noether's theorem, the Lagrangian's invariance under the continuous supersymmetry transformation signifies the existence of conserved currents that have associated conserved charges, Q_α. These Q_α are interesting because, as quantum mechanical operators, they are the generators of the supersymmetry transformation – i.e. $Q_\alpha |\phi\rangle \propto |\psi\rangle$. The α index identifies that they are spin-1/2 objects. They satisfy the following anticommutation relation,

$$\begin{aligned} \left\{Q_\alpha,\, Q_\beta^\dagger\right\} &\propto \delta_{\alpha\beta} P_\mu \\ \left\{Q_\alpha^\dagger,\, Q_\beta^\dagger\right\} &= 0 \\ \{Q_\alpha,\, Q_\beta\} &= 0 \end{aligned} \tag{3.5}$$

where P_μ is the generator of space-time translations.

A number of further interesting results immediately follow from the above relation. Supersymmetry is a space-time symmetry, meaning that supersymmetry automatically commutes with all symmetries related to a particle's internal degrees of freedom. As a result, superpartners fall into the same representations of gauge groups and share the same gauge couplings and quantum numbers (color, electric charge, etc.). Furthermore, the Q_α commute with the squared mass operator, $-P^\mu P_\mu$. Thus, in addition to gauge quantum numbers, superpartners also share the same masses.

These results play important roles in constructing supersymmetric extensions to the SM [23, 24]. Although many such extensions exist, it is most instructive to

Table 3.1: Chiral supermultiplets in the Minimal Supersymmetric Standard Model. The spin-0 fields are complex scalars, and the spin-1/2 fields are left-handed two-component Weyl fermions. Reprinted from Table 1.1 of [20].

Names		spin 0	spin 1/2	$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
squarks, quarks (×3 families)	Q_L	$\begin{pmatrix} \widetilde{\mathrm{u}}_{\mathrm{L}} \\ \widetilde{\mathrm{d}}_{\mathrm{L}} \end{pmatrix}$	$\begin{pmatrix} \mathrm{u}_L \\ \mathrm{d}_L \end{pmatrix}$	$(\mathbf{3} \otimes \mathbf{2} \otimes \frac{1}{6})$
	$\bar{\mathrm{u}}$	$\widetilde{\mathrm{u}}_{\mathrm{R}}$	u_R	$(\bar{\mathbf{3}} \otimes \mathbf{1} \otimes -\frac{2}{3})$
	$\bar{\mathrm{d}}$	$\widetilde{\mathrm{d}}_{\mathrm{R}}$	d_R	$(\bar{\mathbf{3}} \otimes \mathbf{1} \otimes \frac{1}{3})$
sleptons, leptons (×3 families)	L	$\begin{pmatrix} \widetilde{\nu}_L \\ \widetilde{\mathrm{e}}_{\mathrm{L}} \end{pmatrix}$	$\begin{pmatrix} \nu_L \\ \mathrm{e}_L \end{pmatrix}$	$(\mathbf{1} \otimes \mathbf{2} \otimes -\frac{1}{2})$
	$\bar{\mathrm{e}}$	$\widetilde{\mathrm{e}}_{\mathrm{R}}$	e_R	$(\mathbf{1} \otimes \mathbf{1} \otimes 1)$
Higgs, higgsinos	$\mathrm{H_u}$	$\begin{pmatrix} \mathrm{H_u^+} \\ \mathrm{H_u^0} \end{pmatrix}$	$\begin{pmatrix} \widetilde{\mathrm{H}}_{\mathrm{u}}^{+} \\ \widetilde{\mathrm{H}}_{\mathrm{u}}^{0} \end{pmatrix}$	$(\mathbf{1} \otimes \mathbf{2} \otimes \frac{1}{2})$
	$\mathrm{H_d}$	$\begin{pmatrix} \mathrm{H_d^0} \\ \mathrm{H_d^-} \end{pmatrix}$	$\begin{pmatrix} \widetilde{\mathrm{H}}_{\mathrm{d}}^{0} \\ \widetilde{\mathrm{H}}_{\mathrm{d}}^{-} \end{pmatrix}$	$(\mathbf{1} \otimes \mathbf{2} \otimes \frac{1}{2})$

consider the "minimally supersymmetric" SM (MSSM), which extends the SM with the minimal amount of additional particles required.

3.2.1 The MSSM

In the MSSM, every SM particle is paired with a new, distinct superpartner for each degree of freedom. All of these superpartners are symbolically denoted with tildes over the corresponding SM particle symbols, e.g. $\mathrm{g} \leftrightarrow \widetilde{\mathrm{g}}$.

Table 3.2: Gauge supermultiplets in the Minimal Supersymmetric Standard Model. Reprinted from Table 1.2 of [20].

Names	spin 1/2	spin 1	$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
gluino, gluon	$\widetilde{\mathrm{g}}$	g	$(\mathbf{8} \otimes \mathbf{1} \otimes 0)$
winos, W bosons	$\widetilde{\mathrm{W}}^{\pm}$, $\widetilde{\mathrm{W}}^{0}$	$\mathrm{W}^{\pm}$, W^{0}	$(\mathbf{1} \otimes \mathbf{3} \otimes 0)$
bino, B boson	$\widetilde{\mathrm{B}}^{0}$	B^{0}	$(\mathbf{1} \otimes \mathbf{1} \otimes 0)$

3.2.1.1 The Chiral Supermultiplets of the MSSM

Table 3.1 shows the chiral supermultiplets of the MSSM. Because the supersymmetric extension to the SM respects the preexisting gauge symmetries, distinct partners are added for both the left- and right-handed quarks and leptons separately. These bosonic, scalar superpartners are named after their corresponding SM fermions with the added prefix "s-", which stands for "scalar". Because they are scalars, the subscript on the "L" or "R" of the superpartners refers only to which corresponding SM fermion they are related to. They possess the same gauge couplings as their SM partners, i.e. of the two partners to the top quark, the "left-handed" top-squark, $\widetilde{\mathrm{t}}_L$, couples to the W boson, while the "right-handed" top-squark, $\widetilde{\mathrm{t}}_R$ does not.

The Higgs field, as a scalar field, receives a spin-1/2 partner field that has the same name but with the added suffix "-ino". There are a few interesting things to note in the supersymmetric extension to the Higgs sector. As a scalar, naively the Higgs might be placed, along with its fermionic "higgsino" superpartner, in a chiral supermultiplet with weak hyper charge either equal to $Y = 1/2$ or $Y = -1/2$. However, if only a single Higgs chiral supermultiplet is added to the theory, then

triangle gauge anomalies arise from these chiral higgsinos[4] and interfere with the consistency of the theory at a quantum level. Moreover, these hypercharge values also restrict the allowed Yukawa couplings between the Higgs boson and the SM fermions: only Higgs with $Y = 1/2$ can couple to the up-type quarks, while only Higgs with $Y = -1/2$ can couple to the down-type quarks and charged leptons.

The solution is quite straightforward: in the MSSM, two Higgs chiral supermultiplets are added, one with each weak hypercharge value, and labeled H_{u} and H_{d} for $Y = 1/2$ and $Y = -1/2$. Under EWSB, each Higgs doublet has its own VEV, labeled v_{u} and v_{d}, respectively. Each VEV is taken as a component of a single value, $v = v_{\mathrm{u}} \oplus v_{\mathrm{d}}$, and an angle β measures this relative difference via, $\tan(\beta) = v_{\mathrm{u}}/v_{\mathrm{d}}$. The presence of two Higgs doublets introduces two neutral scalar Higgs states, as well as a third, pseudoscalar Higgs state. The SM Higgs boson we observe is a mixture of the two neutral Higgs states, and is commonly assumed to be the lighter state, denoted with h^0; the heavier, mixed scalar state is denoted with H^0 and the pseudoscalar state is denoted by A^0. In parallel with this, a supersymmetric analogue to the Higgs mass, commonly denoted by μ, is added into the theory, via the Lagrangian term $\mu \mathrm{H_u H_d}$.

3.2.1.2 The Gauge Supermultiplets of the MSSM

Table 3.2 shows the gauge supermultiplets of the MSSM. Similar to the Higgs field, the fermionic superpartners to the SM gauge bosons are named after their corresponding SM bosons, but with the added suffix "-ino". As indicated by the

[4] In the basic SM, these gauge anomalies cancel somewhat coincidentally.

third column of Tab. 3.2, these partners fall into the same representations of the respective gauge groups.

3.2.1.3 R-parity in SUSY and Dark Matter

In the SM, both lepton and baryon number are conserved in all renormalizable interactions. In the MSSM, however, it is possible to write out renormalizable interaction terms involving squarks and sfermions that satisfy all SM gauge symmetries but violate lepton or baryon number. It is important to note that these terms are not required for a phenomenologically viable MSSM Lagrangian, and indeed, empirical constraints from experiments strongly implies the absence of such terms. For example, the interactions that would violate conservation of baryon number, for reasonable choices of SUSY parameters, lead to a proton lifetime of a fraction of a second, which is clearly refuted by the current observed lower bound of $\mathcal{O}\left(10^{32}\,\text{years}\right)$. The presence of these terms in the Lagrangian can be explicitly forbidden by adding in a new, discrete symmetry called R-parity,

$$R_p = (-1)^{3(B-L)+2s}, \tag{3.6}$$

where s is the spin of the field, and B and L are the baryon and lepton number respectively. The superpartners have $R_p = -1$ while the SM fields have $R_p = 1$. The conservation of R-parity has important phenomenological consequences: There must be an even number of sparticles in each interaction involving them, signifying that, in interactions at the LHC, they will be pair-produced at a minimum if at all. Furthermore, the lightest supersymmetric particle (LSP), even if it is heavier than

all other SM particles, is completely stable, as it cannot decay to SM particles alone. At the LHC, if produced, the LSP will escape the detector without interacting, just like neutrinos. This leads to an experimental signature of (typically) large amounts of *missing* energy. Moreover, the LSPs can, with the right mass, serve as a particle explanation for dark matter.

3.2.1.4 Running of Coupling Constants in the MSSM

The addition of the new superpartners and subsequent interactions alters the running of the coupling constants. Depending upon the exact spectrum of particle and sparticle masses, these changes can greatly aid grand unification, further motivating the existence of SUSY [25, 26, 27].

3.2.1.5 Sparticle Masses and the Lack of Discovery

Except for the last discussion point, we have made no mention thus far of the superpartner masses. If SUSY was an unbroken theory, for reasons discussed above, the superpartners would have the exact same masses as their SM counterparts. Empirically, this is obviously not the case, as we would have detected superpartners like the scalar electron (with a mass of only ~ 0.5 MeV) many years ago. Thus, SUSY must be a *broken* symmetry of nature.

3.2.2 Broken SUSY in the MSSM

The exact mechanism of SUSY breaking plays an important role in determining, among other things, what role SUSY plays in solving the gauge hierarchy problem.

In particular, the removal of quadratic divergences in the Higgs mass quantum corrections, will occur so long as the breaking mechanism maintains both the direct one-to-one correspondence between superpartners and the equivalence between fermion and scalar Yukawa couplings (within a given supermultiplet) to the Higgs.

This is generally written in the following fashion,

$$\mathcal{L} = \mathcal{L}_{\text{SUSY}} + \mathcal{L}_{\text{soft}}, \tag{3.7}$$

where $\mathcal{L}_{\text{SUSY}}$ consists of entirely SUSY-preserving terms, including couplings between the particles and the Higgs fields and the standard gauge interactions of the SM. $\mathcal{L}_{\text{soft}}$ is comprised of mass and coupling terms that violate SUSY and have positive mass dimensions. Associated with $\mathcal{L}_{\text{soft}}$ is a mass scale, m_{soft}, This mass scale is an approximate measure for the masses of all of the superpartners.

The effective result of soft-SUSY breaking is that the quadratic dependencies of m_{H}^2 on Λ_{UV} are still canceled, replaced instead by a quadratic dependence on m_{soft},

$$\Delta m_{\text{H}}^2 = m_{\text{soft}}^2 \left[\frac{\lambda}{16\pi^2} \log\left(\Lambda_{\text{UV}}/m_{\text{soft}}\right) + \cdots \right], \tag{3.8}$$

where λ is a generalized proxy for dimensionless couplings of the Higgs to various particles.

This may seem to have shifted the burden of fine-tuning from Λ_{UV} onto the m_{soft} parameter. However, so long as m_{soft} is $\mathcal{O}\,(1\ \mathrm{TeV})$, the overall degree of fine-tuning required is relatively low.

There are some other consequences of SUSY breaking. The soft SUSY breaking terms include interactions terms that lead to flavor-changing neutral currents (FCNC) and CP violation. Precision SM measurements severely constrain the magnitude both of these effects. A workaround is to assume that first- and second-generation sfermions have the same, flavor-blind masses, while the third-generation sfermions can have different masses. This removes the majority of flavor mixing through the superpartners and significantly reduces the free parameters of the MSSM.

3.2.2.1 Sparticle Mixing

Once electroweak symmetry breaking and SUSY breaking effects are included, the superpartners listed in Tabs. 3.1 and 3.2 are not necessarily the mass eigenstates of the theory. Specifically, all of the sparticles with the same spin and gauge quantum numbers can mix. In the sfermion sector, flavor-mixing between generations is constrained as per the prior discussion. Nevertheless, the two sfermion partners for each SM fermion will mix to form a lighter and heavier eigenstate.

Mixing of the $\widetilde{\mathrm{t}}$ Eigenstates

As an example, for the $\widetilde{\mathrm{t}}$ there can be significant mixing between the chiral eigenstates. The top squark mass matrix is written as follows [28]:

$$\mathbf{m}^{\mathbf{2}}_{\widetilde{\mathrm{t}}} = \begin{pmatrix} m^2_{\widetilde{\mathrm{t}}_{\mathrm{L}}} + m^2_{\mathrm{t}} + D^{\mathrm{t}}_L & m_{\mathrm{t}} X_{\mathrm{t}} \\ m_{\mathrm{t}} X_{\mathrm{t}} & m^2_{\widetilde{\mathrm{t}}_{\mathrm{R}}} + m^2_{\mathrm{t}} + D^{\mathrm{t}}_R \end{pmatrix} \tag{3.9}$$

where,

$$D^{\mathrm{t}}_L = \left(\frac{1}{2} - \frac{2}{3}s^2_w\right) m^2_{\mathrm{Z}} \cos 2\beta, \quad D^{\mathrm{t}}_R = \frac{2}{3}s^2_w m^2_{\mathrm{Z}} \cos 2\beta, \quad X_{\mathrm{t}} = A_{\mathrm{t}} - \frac{\mu}{\tan\beta}. \tag{3.10}$$

In the above equation, s_w is the sine of the Weinberg angle. When this mass matrix is diagonalized, a large mixing angle typically occurs because of the off-diagonal entries, which contain terms involving the large top Yukawa coupling (hidden within m_{t}), and soft coupling parameter, A_t. This means that one top squark mass eigenstate, $\widetilde{\mathrm{t}}_1$, will be lighter than the other, $\widetilde{\mathrm{t}}_2$, and in fact it likely will be the lightest squark.

The exact gauge coupling strengths for these sfermion mass eigenstates depends upon the magnitude of mixing. For example, the lighter top-squark eigenstate, $\widetilde{\mathrm{t}}_1$, will not couple at all to the charged W boson if it is entirely made up of the $\widetilde{\mathrm{t}}_{\mathrm{R}}$ eigenstate.

Mixing of the SM Boson Partners

In the SM boson sector, there are 8 degrees of freedom associated with the neutral SM bosons (3 for the Z, 2 for the γ, and 1 each for h^0, H^0, and A^0); Their superpartners (denoted generically by $\widetilde{N}_k$), the zino, photino, and higgsinos, can

mix to form 4, spin-1/2 neutralinos,

$$\widetilde{\chi}_j^0 = \sum_{k=1}^{4} N_{jk} \widetilde{N}_k, \tag{3.11}$$

where the mixing matrix, N_{jk}, is defined in Ref. [29].

Similarly, the charged SM bosons also have 8 degrees of freedom among them (2 for each charged W and 2 for each charged Higgs). Their superpartners (denoted generically by $\widetilde{C}_k^{\pm}$), the winos and charged higgsinos, can mix to form 2, spin-1/2 charginos with electric charges of ± 1,

$$\widetilde{\chi}_j^+ = \sum_{k=1}^{2} V_{jk} \widetilde{C}_k^+, \qquad \widetilde{\chi}_j^- = \sum_{k=1}^{2} U_{jk} \widetilde{C}_k^-, \tag{3.12}$$

where the mixing matrices, V_{jk} and U_{jk}, are defined in Ref. [29].

The couplings of these neutralinos and charginos to the fermions and sfermions are determined by their relative mixing of "gauge content". For example, if the lightest chargino, $\widetilde{\chi}_1^0$, is entirely wino-like, then only the $\widetilde{\mathrm{t}}_{\mathrm{L}}$ component of the $\widetilde{\mathrm{t}}_1$ will couple to it.

Beyond confirming the existence of these new particles, the main phenomenological tests of the MSSM are to measure the exact masses of the new particles as well as the relative mixings between the flavor and gauge eigenstates.

3.2.3 Phenomenological Summary of the MSSM

The salient details of the MSSM are the following. Every SM particle has a superpartner sparticle. From considerations of the hierarchy problem, these superpartner masses are expected to be around $\mathcal{O}\,(1\ \mathrm{TeV})$. These superpartners share the exact same gauge interactions as their SM counterparts. Thus, at the LHC, it is

Table 3.3: The undiscovered sparticles in the Minimal Supersymmetric Standard Model (with sfermion mixing for the first two families assumed to be negligible). Table adapted from Ref. [20].

Names	Spin	R_P	Gauge Eigenstates	Mass Eigenstates
squarks	0	-1	$\widetilde{\mathrm{u}}_{\mathrm{L}}\ \widetilde{\mathrm{u}}_{\mathrm{R}}\ \widetilde{\mathrm{d}}_{\mathrm{L}}\ \widetilde{\mathrm{d}}_{\mathrm{R}}$	(same)
			$\widetilde{\mathrm{s}}_{\mathrm{L}}\ \widetilde{\mathrm{s}}_{\mathrm{R}}\ \widetilde{\mathrm{c}}_{\mathrm{L}}\ \widetilde{\mathrm{c}}_{\mathrm{R}}$	(same)
			$\widetilde{\mathrm{t}}_{\mathrm{L}}\ \widetilde{\mathrm{t}}_{\mathrm{R}}\ \widetilde{\mathrm{b}}_{\mathrm{L}}\ \widetilde{\mathrm{t}}_{\mathrm{R}}$	$\widetilde{\mathrm{t}}_1\ \widetilde{\mathrm{t}}_2\ \widetilde{\mathrm{b}}_1\ \widetilde{\mathrm{b}}_2$
sleptons	0	-1	$\widetilde{\mathrm{e}}_{\mathrm{L}}\ \widetilde{\mathrm{e}}_{\mathrm{R}}\ \widetilde{\nu}_{\mathrm{e},L}$	(same)
			$\widetilde{\mu}_{\mathrm{L}}\ \widetilde{\mu}_{\mathrm{R}}\ \widetilde{\nu}_{\mu,L}$	(same)
			$\widetilde{\tau}_{\mathrm{L}}\ \widetilde{\tau}_{\mathrm{R}}\ \widetilde{\nu}_{\tau,L}$	$\widetilde{\tau}_1\ \widetilde{\tau}_2\ \widetilde{\nu}_{\tau,L}$
neutralinos	1/2	-1	$\widetilde{\mathrm{B}}^0\ \widetilde{\mathrm{W}}^0\ \widetilde{\mathrm{H}}_{\mathrm{u}}\ \widetilde{\mathrm{H}}_{\mathrm{d}}$	$\widetilde{\chi}_1^0\ \widetilde{\chi}_2^0\ \widetilde{\chi}_3^0\ \widetilde{\chi}_4^0$
charginos	1/2	-1	$\widetilde{\mathrm{W}}^{\pm}\ \widetilde{\mathrm{H}}_{\mathrm{u}}^{+}\ \widetilde{\mathrm{H}}_{\mathrm{d}}^{-}$	$\widetilde{\chi}_1^{\pm}\ \widetilde{\chi}_2^{\pm}$
gluino	1/2	-1	$\widetilde{\mathrm{g}}$	(same)

expected that the production cross-section for strongly-interacting sparticles should be notably larger than the corresponding electroweak-inos. If R-parity, a new discrete symmetry related to SUSY, is conserved, then, at the LHC, superpartners will be pair-produced. These superpartners will undergo cascade decays where, at each step, another, lighter superpartner is produced until finally, all that remains are SM particles and two of the lightest superpartner (LSP). These LSPs are completely stable and will be invisible to detection, giving an experimental signature of SM particles with missing energy.

Table 3.3 displays the full list of the undiscovered particles in the MSSM, including the distinctions between their gauge and mass eigenstates.

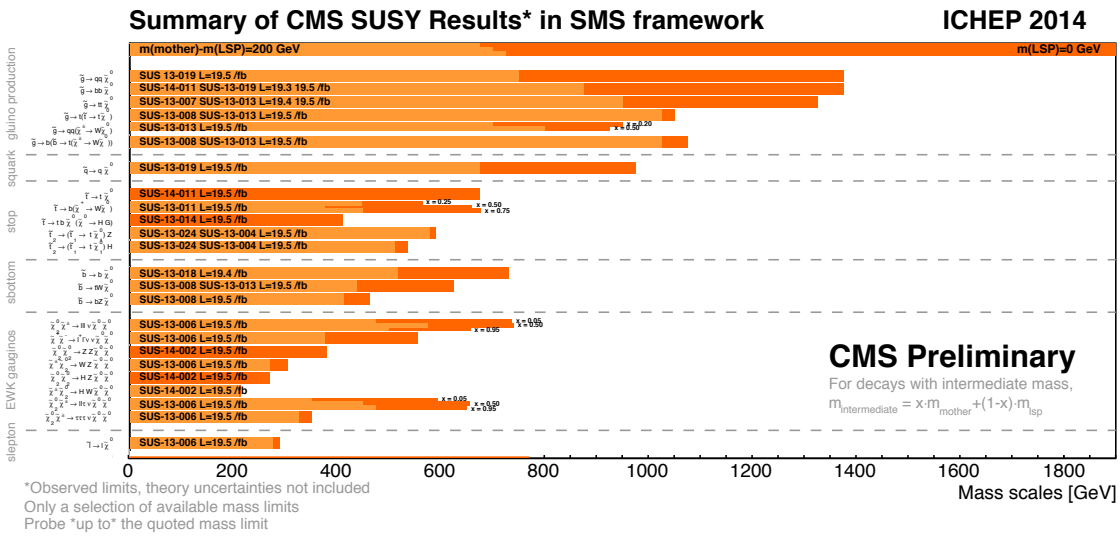

Figure 3.2: General summary (of a subset of results) of CMS searches (as of ICHEP 2014) for R-parity conserving supersymmetric particle production at the LHC (with results using $\sim 19.5\,\mathrm{fb}^{-1}$ of 8 TeV pp collision data). All limits are observed 95% CL lower-limits, Theoretical uncertainties were not taken into account for these limits. Figure built using information from from Ref. [30].

3.2.4 Current Experimental Limits on the MSSM

These details have driven the experimental searches for SUSY at numerous colliders, most recently with the ATLAS and CMS experiments.

Both of these experiments, in order to produce results with the broadest possible applicability, utilize simplified models (SMS), in which the SUSY particles not under direct consideration are assumed to have masses too large to be of relevance in any interactions.

Figure 3.2 summarizes results from a very wide range of SUSY searches done with the CMS experiment (Similar results from the ATLAS experiment can be found at Ref. [31]).

These results are shown as 95% confidence level observed upper limits on

the masses of specific sparticles. As can be seen, the masses for many of these sparticles are expected to be over 1 TeV. As we discussed earlier, however, there is a correlation between m_{soft}, the scale of SUSY breaking, and the sparticle masses. As well, there is an expectation from considerations of the gauge hierarchy problem that $m_{\text{soft}} \sim \mathcal{O}\,(1\text{ TeV})$. As a result, one might claim (and many scientists have) that SUSY is not experimentally viable as a "natural" solution to the hierarchy problem. Here, "natural" signifies that there is no major fine-tuning required of parameters in the theory.

This claim, however, ignores some of the more nuanced details of SUSY in relation to the hierarchy problem. In short, SUSY can still be a natural solution to the hierarchy problem if the superparters of the Higgs, top, and gluon have masses near the electroweak scale [32, 33, 34, 35, 36].

We turn now, to discuss some key aspects of "Natural" SUSY.

3.2.5 "Natural" SUSY

The requirement of naturalness in the MSSM can be generally summarized via the following relation that comes from evaluating the MSSM at the tree-level (i.e. leading-order approximation) [34],

$$-\frac{m_{\text{Z}}^2}{2} = |\mu|^2 + m_{H_u}^2. \tag{3.13}$$

In the above equation, m_{Z} is the experimentally measured mass of the Z boson, while μ and $m_{H_u}^2$ are, in some sense, tunable parameters of the theory.

The "naturalness" of the MSSM, or other supersymmetric extensions to the

SM that have similar such relations, is then measured by how tuned these parameters are relative to one another such that they achieve the experimentally measured m_{Z}.

In analogy with the initial discussion of the gauge hierarchy problem above, Eqs. (3.1) and (3.2), both of the parameters on the right side of Eq. (3.13) either directly predict sparticle masses or receive quantum corrections from sparticle interactions. Thus, by measuring the relevant sparticle masses and interactions, one can directly probe the "naturalness" of the MSSM.

For example, μ directly controls the mass of the Higgsinos, and so their masses should not be too far above the electroweak scale (i.e. m_{Z}). Defining a "fine-tuning" parameter, $\Delta \equiv 2\delta m_H^2/m_h^2$, where m_h is the physical Higgs mass and m_H^2 is the constant parameter for the quadratic part of the Higgs potential, this relation can be quantified [34],

$$\mu \lesssim 200 \text{ GeV} \left(\frac{m_h}{120 \text{ GeV}}\right) \sqrt{20\%\Delta}. \tag{3.14}$$

Similarly, the top-squark directly corrects $m_{H_u}^2$ in 1-loop diagrams, so the naturalness requirement imposes the following quantitative relation on the two top-squark masses [34],

$$\sqrt{m_{\tilde{\mathrm{t}}_1}^2 + m_{\tilde{\mathrm{t}}_2}^2} \lesssim 600 \text{ GeV} \frac{\sin\beta}{\sqrt{1+x_t^2}} \sqrt{\frac{3}{\log(m_{\mathrm{soft}}/\text{ TeV})}} \left(\frac{m_h}{120 \text{ GeV}}\right) \sqrt{20\%\Delta}, \tag{3.15}$$

where $x_t = A_t/\sqrt{m_{\tilde{\mathrm{t}}_1}^2 + m_{\tilde{\mathrm{t}}_2}^2}$ is a measure of the top-squark mixing in the soft SUSY breaking Lagrangian.

The other squarks, as well as the sleptons, give similar contributions to $m_{H_u}^2$, but their bounds are looser due to smaller Yukawa couplings to the Higgs.

The gluino corrects the top-squark mass at the 1-loop level, and thus corrects

$m_{H_u}^2$ at the 2-loop level. Requiring naturalness imposes the following relation [34], where M_3 represents the gluino mass,

$$\begin{aligned} M_3^{\text{Majoranna}} &\lesssim 900 \text{ GeV} \sin\beta \frac{3}{\log(m_{\text{soft}}/\text{ TeV})} \left(\frac{m_h}{120 \text{ GeV}}\right) \sqrt{20\%\Delta}, \\ M_3^{\text{Dirac}} &\lesssim 900 \text{ GeV} \sin\beta \sqrt{\frac{3}{\log(m_{\text{soft}}/\text{ TeV})}} \left(\frac{m_h}{120 \text{ GeV}}\right) \sqrt{20\%\Delta}, \end{aligned} \tag{3.16}$$

where the top (bottom) equation represents the scenario that the gluino is a Majoranna (Dirac) fermion.

There are analogous bounds on the other gauginos, the bino and wino, but these bounds are much looser relative to the aforementioned bounds.

The requirements in Eqs. (3.14), (3.15), and (3.16) define the most important SUSY searches to probe natural SUSY at the LHC. In particular, if SUSY is a natural solution to the gauge hierarchy problem, we expect to find, among other things, relatively light higgsinos (i.e. one or more of the charginos and neutralinos) as well as relatively light top-squarks.

3.2.6 Current Experimental Limits on Natural SUSY

Both the ATLAS and CMS experiments have performed extensive searches to either find, or constrain natural SUSY.

3.2.6.1 Gluino Pair-production

The experimental searches for gluino pair-production at ATLAS and CMS have primarily focused on the decay mode where the gluino decays through an off-shell $\widetilde{\mathrm{t}}$ into a pair of top quarks and the lightest neutralino, $\widetilde{\chi}_1^0$.

Summary plots from both experiments for the current experimental limits on this decay mode can be found in Refs. [30, 31]. Both experiments have excluded a wide swath of the low mass part of the 2D parameter space. Currently, approximately all of the region, $m_{\tilde{g}} \lesssim 1400$ GeV, $\tilde{\chi}_1^0 \lesssim 600$ GeV has been excluded at the 95% CL. These limits are already somewhat in tension with the requirements of natural SUSY. This tension will continue to grow as these analyses are performed again using the 13 TeV data currently being recorded by the ATLAS and CMS experiments.

3.2.6.2 Gaugino Pair-production

In the light gaugino sector, ATLAS and CMS have focused on the combined production of either $\tilde{\chi}_2^0$ or $\tilde{\chi}_1^\pm$ and the subsequent decay of these gauginos to $\tilde{\chi}_1^0$.

As with the gluino limits, summary plots of the current experimental limits can be found in Refs. [30, 31]. The exact regions excluded depend upon the exact combination of $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^\pm$ considered, as well as the particular details of the cascade decay chains.

Nevertheless, as with the gluinos, these limits are also somewhat in tension with the requirements of natural SUSY, but there are remaining, non-excluded parts of the "natural" SUSY parameter space.

3.2.6.3 Searches for Top-squark Pair-production

Because the search for top-squark pair-production is the focus of this dissertation, we will spend a bit more time discussing the current search results in the top-squark sector.

Final State Topologies for Top-squark Decays

The decays of the top-squark can be categorized in terms of separate n-body decay modes, where each has certain kinematic conditions – namely, the relative values of relevant particle masses – affecting its accessibility. If the kinematic conditions for two separate decay modes are both accessible in a given kinematic region, then both decay modes ostensibly occur.

Table 3.4: Description of the kinematic conditions of five decay modes of the top-squark. Reprinted from Table 1 of [37].

Kinematic conditions	n-body decays	Channel
$m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0} < m_{\mathrm{W}}$	2-body decays 4-body decays	$\widetilde{\mathrm{t}}_1 \to \mathrm{c}\widetilde{\chi}_1^0$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}ff'\widetilde{\chi}_1^0$
$m_{\mathrm{b}} + m_{\mathrm{W}} + m_{\widetilde{\chi}_1^0} \leq m_{\widetilde{\mathrm{t}}_1}$ and $m_{\widetilde{\mathrm{t}}_1} < m_{\mathrm{t}} + m_{\widetilde{\chi}_1^0}$	3-body decays	$\widetilde{\mathrm{t}}_1(\to \mathrm{t}^*\widetilde{\chi}_1^0) \to \mathrm{bW}^+\widetilde{\chi}_1^0$
$m_{\mathrm{t}} + m_{\widetilde{\chi}_1^0} \leq m_{\widetilde{\mathrm{t}}_1}$	2-body decays	$\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$, and $\mathrm{t} \to \mathrm{bW}^+$
$m_{\mathrm{b}} + m_{\mathrm{W}} + m_{\widetilde{\chi}_1^0} \leq m_{\widetilde{\mathrm{t}}_1}$ and $m_{\widetilde{\chi}_1^0} < m_{\widetilde{\chi}_1^+} < m_{\widetilde{\mathrm{t}}_1} - m_{\mathrm{b}}$	2-body decays	$\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}^{\pm}$, and $\widetilde{\chi}_1^{\pm} \to \mathrm{W}^{\pm(*)}\widetilde{\chi}_1^0$

These kinematic conditions are summarized in Tab. 3.4.

Polarization of Daughter Particles in Top-squark Decays

One of the more interesting aspects of the top-squark sector of the MSSM is the possibly large relative mixing between the two chiral top-squark eigenstates [38],

$$\begin{pmatrix} \widetilde{\mathrm{t}}_1 \\ \\ \widetilde{\mathrm{t}}_2 \end{pmatrix} = \begin{pmatrix} \cos\theta_t & \sin\theta_t \\ \\ -\sin\theta_t & \cos\theta_t \end{pmatrix} \begin{pmatrix} \widetilde{\mathrm{t}}_{\mathrm{L}} \\ \\ \widetilde{\mathrm{t}}_{\mathrm{R}} \end{pmatrix}. \tag{3.17}$$

where

$$\sin 2\theta_t = \frac{2m_{\mathrm{t}}X_{\mathrm{t}}}{m^2_{\widetilde{\mathrm{t}}_1} - m^2_{\widetilde{\mathrm{t}}_2}}, \tag{3.18}$$

$$\cos 2\theta_t = \frac{m^2_{\widetilde{\mathrm{t}}_{\mathrm{L}}} + D^{\mathrm{t}}_L - m^2_{\widetilde{\mathrm{t}}_{\mathrm{R}}} - D^{\mathrm{t}}_R}{m^2_{\widetilde{\mathrm{t}}_1} - m^2_{\widetilde{\mathrm{t}}_2}}. \tag{3.19}$$

For both the $\widetilde{\mathrm{t}} \to \mathrm{t}\widetilde{\chi}^0_1$ and $\widetilde{\mathrm{t}} \to \mathrm{b}\widetilde{\chi}^+_1$ decay modes of the top-squark, this mixing can lead to non-trivial polarizations of the daughter particles. For example, in the $\widetilde{\mathrm{t}} \to \mathrm{t}\widetilde{\chi}^0_1$ decay mode, there is one relevant interaction vertex between the mother $\widetilde{\mathrm{t}}$ and its daughters [39],

$$g^{\mathrm{t}}_{\mathrm{eff}} \left(\sin\theta^{\mathrm{t}}_{\mathrm{eff}} P_L + \cos\theta^{\mathrm{t}}_{\mathrm{eff}} P_R\right) \widetilde{\mathrm{t}}_1 \bar{\mathrm{t}} \widetilde{\chi}^0_1, \tag{3.20}$$

where $g^{\mathrm{t}}_{\mathrm{eff}}$ is the effective coupling constant for this interaction and $\theta^{\mathrm{t}}_{\mathrm{eff}}$, the effective mixing angle between top quark polarization states, is defined by,

$$\tan\theta^{\mathrm{t}}_{\mathrm{eff}} = \frac{y_{\mathrm{t}} N_{14} \cos\theta_{\mathrm{t}} - \frac{2\sqrt{2}}{3} g' N_{11} \sin\theta_{\mathrm{t}}}{\sqrt{2}\left(\frac{g}{2} N_{12} + \frac{g'}{6} N_{11}\right)\cos\theta_{\mathrm{t}} + y_{\mathrm{t}} N_{14} \sin\theta_{\mathrm{t}}}. \tag{3.21}$$

Similarly, in the $\widetilde{\mathrm{t}} \to \mathrm{b}\widetilde{\chi}^+_1$ decay mode, the relevant $\widetilde{\mathrm{t}}$, $\widetilde{\chi}^+$ vertex is,

$$g^{\widetilde{\chi}^+_1}_{\mathrm{eff}} \left(\sin\theta^{\widetilde{\chi}^+_1}_{\mathrm{eff}} P_L + \cos\theta^{\widetilde{\chi}^+_1}_{\mathrm{eff}} P_R\right) \bar{\mathrm{b}}\widetilde{\mathrm{t}}_1 \widetilde{\chi}^+_1, \tag{3.22}$$

where $\theta_{\text{eff}}^{\widetilde{\chi}_1^+}$ is defined by,

$$\tan\theta_{\text{eff}}^{\widetilde{\chi}_1^+} = \frac{y_{\text{b}} U_{12}^* \cos\theta_{\text{t}}}{-g V_{11}\cos\theta_{\text{t}} + y_{\text{t}} V_{12}\sin\theta_{\text{t}}}. \tag{3.23}$$

To note, for the $\widetilde{\text{t}} \to \text{b}\widetilde{\chi}_1^+$ decay mode, the relative mixing of the gauginos can also lead to non-trivial polarizations for the daughter W, as the relevant vertex between the $\widetilde{\chi}^+$ and the W is,

$$g_{\text{eff}}^{\text{W}^-}\left(\sin\theta_{\text{eff}}^{\text{W}^-} P_L + \cos\theta_{\text{eff}}^{\text{W}^-} P_R\right) \text{W}_\mu^- \gamma^\mu \bar{\widetilde{\chi}}_1^0 \widetilde{\chi}_1^+, \tag{3.24}$$

where $\theta_{\text{eff}}^{\text{W}^-}$ is defined by,

$$\tan\theta_{\text{eff}}^{\text{W}^-} = \frac{-N_{14}V_{12}^* + \sqrt{2}N_{12}V_{11}^*}{N_{13}^* U_{12} + \sqrt{2}N_{12}^* U_{11}}. \tag{3.25}$$

These non-trivial polarizations are both relevant and exciting from an experimental perspective. Considering a hypothetical top-squark discovery, once there are a sufficient amount of available statistics, we will be able to measure the polarizations of the top-squark daughter particles. For example, in the $\widetilde{\text{t}} \to \text{t}\widetilde{\chi}_1^0$ decay mode, the angular distribution of the lepton in the rest frame of its mother top is [38, 39],

$$\frac{d\sigma}{d\cos\hat{\theta}_\ell} \propto E_{\widetilde{\chi}_1^0} + \sin 2\theta_{\text{eff}}^{\text{t}} m_{\widetilde{\chi}_1^0} + p_{\widetilde{\chi}_1^0}\cos 2\theta_{\text{eff}}^{\text{t}}\cos\hat{\theta}_\ell, \tag{3.26}$$

which, upon comparison with the standard angular distribution for the child lepton in the top quark rest frame [40],

$$\frac{1}{N}\frac{dN}{d\cos\Theta^*} = \frac{1}{2}\left(1 + P\cos\Theta^*\right), \tag{3.27}$$

directly yields a relation between the top quark polarization and $\theta_{\text{eff}}^{\text{t}}$,

$$\mathcal{P}_{\text{t}} = \frac{p_{\widetilde{\chi}_1^0}\cos(2\theta_{\text{eff}}^{\text{t}})}{E_{\widetilde{\chi}_1^0} + m_{\widetilde{\chi}_1^0}\sin(2\theta_{\text{eff}}^{\text{t}})}, \tag{3.28}$$

where $p_{\widetilde{\chi}_1^0} = (\sqrt{(m_{\text{t}}^2 + m_{\widetilde{\chi}_1^0}^2 - m_{\widetilde{\text{t}}_1}^2)^2/4 - m_{\text{t}}^2 m_{\widetilde{\chi}_1^0}^2})/m_{\widetilde{\text{t}}_1}$ and $E_{\widetilde{\chi}_1^0} = \sqrt{p_{\widetilde{\chi}_1^0}^2 + m_{\widetilde{\chi}_1^0}^2}$. A

similar relation can be derived for the polarization of the chargino and the polarization of the W in the $\widetilde{\mathrm{t}} \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode. As an example, the polarization of the chargino is given by [38],

$$\mathcal{P}_{\widetilde{\chi}_1^+} = \frac{p_{\mathrm{b}} \cos(2\theta_{\mathrm{eff}}^{\widetilde{\chi}_1^+})}{E_b + m_b \sin(2\theta_{\mathrm{eff}}^{\widetilde{\chi}_1^+})}. \tag{3.29}$$

Post-discovery, measurements of the daughter particle polarizations in top-squark decays will subsequently constrain the relative mixing between the top-squark chiral eigenstates and possibly the gaugino eigenstates as well (if the top-squark decay is observed in the $\widetilde{\mathrm{t}} \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode).

Prior to discovery, these non-trivial polarizations are still relevant, however, as they can strongly affect the sensitivity of analyses looking for top-squark pair-production. For example, the leptonic decays of right-handed top quarks lead to notably harder momentum leptons (because the lepton's momentum is aligned with the top quark's momentum), thereby increasing the object selection efficiency of analyses investigating lepton top-squark decays. The impact of top-squark daughter particle polarizations on experimental sensitivity are discussed in more detail in Chapters 6 and 7 as well as Appendix D. We do note, however, that for the large majority of the current searches for top-squark pair-production at the LHC, ATLAS and CMS utilized different polarization scenarios; we discuss this in more detail in the next section.

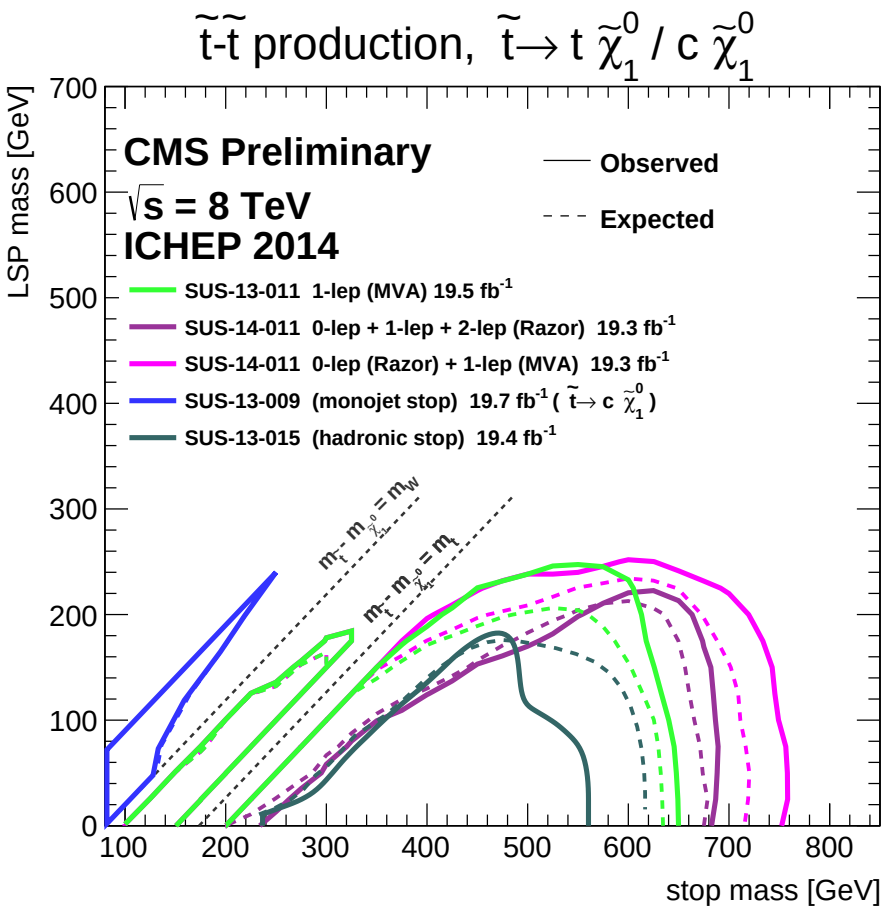

Figure 3.3: Summary of the currently published CMS Run 1 ($20\,\mathrm{fb}^{-1}$ of 8 TeV pp collision data) searches for direct top-squark pair-production where no supersymmetric particle other than the $\tilde{t}_1$ and $\tilde{\chi}_1^0$ are involved in the $\tilde{t}_1$ decay. Both the expected (dashed) and observed (solid) limits at the 95% CL are shown. Three different $\tilde{t}_1$ decay modes (c.f. Table 3.4) are considered individually with assumed branching ratios of 100%. Figure taken from [30].

Summary of Experimental Limits on Top-squark Pair-production

Due to issues of kinematic accessibility for some of the other decay modes, the experiments at LEP and the Tevatron heavily explored the $\tilde{t}_1 \to c\tilde{\chi}_1^0$ decay channel [42, 43, 44, 45, 46, 47]. As well, ATLAS has recently provided more recent experimental limits on the sensitivity to this decay channel [48].

Underlying theoretical motivations, however, have shown that this decay chan-

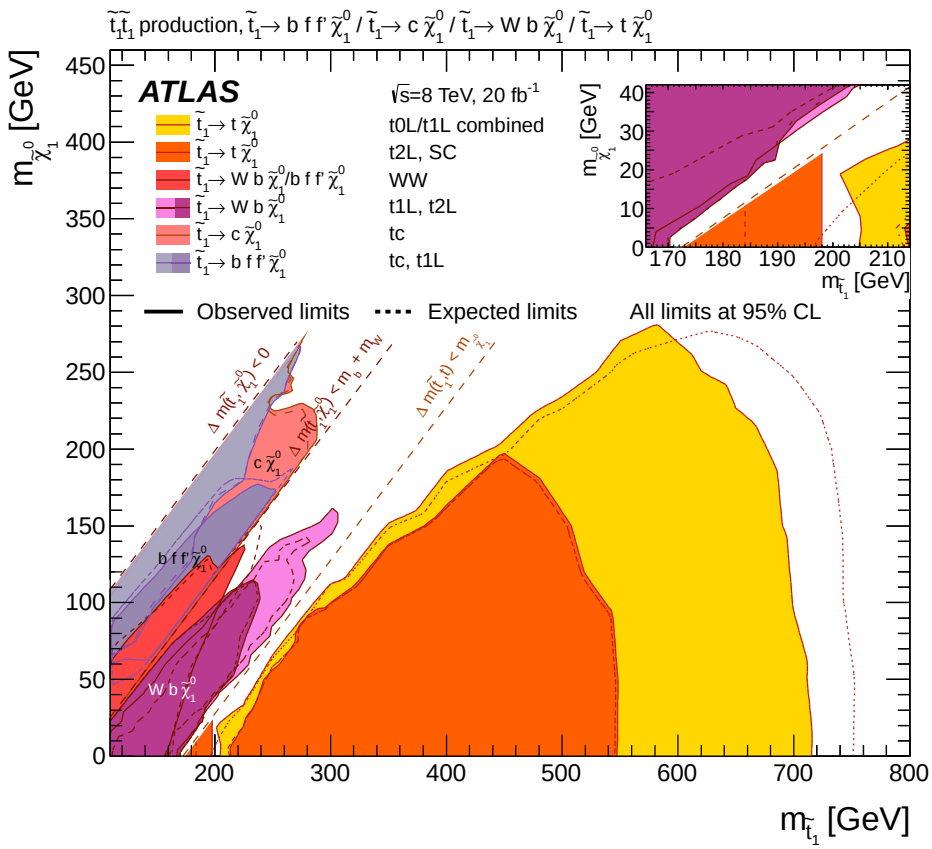

Figure 3.4: Summary of the ATLAS Run 1 ($20\,\mathrm{fb}^{-1}$ of 8 TeV pp collision data) searches for direct top-squark pair-production where no supersymmetric particle other than the $\widetilde{\mathrm{t}}_1$ and $\widetilde{\chi}_1^0$ are involved in the $\widetilde{\mathrm{t}}_1$ decay. Both the expected (dashed) and observed (solid) limits at the 95% CL are shown. Four different $\widetilde{\mathrm{t}}_1$ decay modes (c.f. Table 3.4) are considered individually with assumed branching ratios of 100%. Reprinted from Fig. 4 of [41].

nel might not be the dominant one [49, 50], even for the lower $\sqrt{s}$ experiments.

At the LHC, due to the relatively large $\sqrt{s}$, the focus has shifted instead to the $\widetilde{\mathrm{t}}_1(\rightarrow \mathrm{t}^*\widetilde{\chi}_1^0) \rightarrow \mathrm{bW}^+\widetilde{\chi}_1^0$ and $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{b}\widetilde{\chi}^\pm$, $\widetilde{\chi}_1^\pm \rightarrow \mathrm{W}^{\pm(*)}\widetilde{\chi}_1^0$ decay channels.

ATLAS has published the results of searches across multiple final states for top-squark decays into these channels with both their 7 TeV dataset [51,52,53,54,55] and their 8 TeV dataset [56, 57, 58, 59].

CMS has published multiple searches for top-squark pair-production. The

most recent published results include searches using the 8 TeV data in the all-hadronic and semi-leptonic final states [60, 61, 62].

Figures 3.3 and 3.4 show the current experimental limits from ATLAS and CMS on top-squark pair-production in the $\widetilde{\mathrm{t}} \rightarrow \mathrm{t}\widetilde{\chi}_1^0$ decay mode. In this decay mode, the experimental signature is $\mathrm{t}\bar{\mathrm{t}}$ with extra, invisible particles (the $\widetilde{\chi}_1^0$). As noted above, ATLAS and CMS utilized different polarization scenarios when investigating the sensitivity of their respective analyses to top-squark pair-production. ATLAS's polarization scenario corresponded to an approximate top quark polarization of 95% right-handed, while CMS's polarization scenario corresponded to a unpolarized top quark polarization. All other things being equal, this leads to higher sensitivity for ATLAS, purely for reasons of this arbitrary choice alone.

In the region, $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0} \approx m_{\mathrm{t}}$, the top-squark signal looks very similar to $\mathrm{t}\bar{\mathrm{t}}$, as the $\widetilde{\chi}_1^0$ are produced at rest. This is why this region is, in general, not yet excluded for both plots. The ATLAS analyses that are able to exclude this region somewhat, i.e. the orange triangle shown in the sub-figure in the upper right of Fig. 3.4, accomplish this feat by both exploiting differences in decay product correlations as well as measuring the $\mathrm{t}\bar{\mathrm{t}}$ pair-production cross-section and comparing this to theory.

This is one of the most interesting regions from a theoretical perspective. If the two $\widetilde{\mathrm{t}}$ particles exist and one of their masses is in this range, then this will strongly support the validity of SUSY as a natural solution to the gauge hierarchy problem.

The search for top-squark pair-production, in particular the search in this difficult but theoretically interesting region, motivates the analyses described in

Chapters 6 and 7.

Chapter 4: The Large Hadron Collider and the CMS Detector

In the previous chapter, we discussed the Standard Model, making note of current aesthetic and empirical problems present in the SM and possible extensions to the SM that solve these issues, including a more detailed focus on one such extension: supersymmetry.

Since its invention in the 1960s, there has been a concerted effort on the part of experimentalists to both verify the SM and to test proposed extensions to it. One of the main methods for testing the SM, notably for testing extensions that predict new, very massive particles, is to attempt to directly make these very massive particles through collisions of highly-energetic, lower mass particles. This approach – directly exploring the "energy frontier" – has yielded numerous successes, including proving the existence of the W and Z bosons through direct production at the LEP collider. Later in the 20th century, the discovery of the top quark was a hallmark success of the CDF and D0 experiments at the Tevatron [63,64]. Currently, the energy frontier is being explored at the Large Hadron Collider (LHC), a high-energy hadron collider located in Europe.

The entirety of my work presented in this dissertation was performed using data from proton-proton collisions provided by the LHC and measured by the CMS

detector, a detector designed to comprehensively study the high-energy hadron collisions delivered by the LHC. I will devote the rest of this chapter to discussing the experimental design (e.g. technologies, operating principles, etc.) of both the LHC and the CMS detector. I begin in Section 4.1 with a description of the LHC. Following this, in Section 4.2, I discuss the CMS detector.

4.1 The LHC

The information from this section primarily comes from Ref. [65].

The LHC is the world's largest machine, a circular particle accelerator that collides beams of hadrons at the highest controlled energies and instantaneous beam luminosities in history. It is located underground at the CERN laboratory located northwest of Geneva, Switzerland, near the border between France and Switzerland. The LHC was installed in the same 27 km tunnel used by the LEP experiments some decades earlier.

4.1.1 General Goals of the LHC

The design of the LHC was set with several key physics goals in mind:

- Search for the source of electroweak symmetry breaking, the SM Higgs boson, including measurements of its mass, spin, and couplings.
- Precision measurements of the SM such as,
  - Precision tests of QCD, including measuring jet production cross sections as well as α_S at Q^2 values not previously attained.

  - Precision electroweak measurements, including further reducing the uncertainty on the W boson mass.
  - Precision measurements of the properties of the top quark, as the production cross-section for top quarks at the LHC is much larger than at the Tevatron.

- Search for new physics beyond the SM, with particular focus paid towards supersymmetry and other highly motivated BSM theories at the TeV scale.

- Precision flavor-physics using B-hadrons, including further measurements of CP-violation, rare decays, and B^0 oscillations.

- Heavy-ion physics, with a focus towards improving the understanding of the phase transition, in QCD, of colored particles from bound hadrons to the quark-gluon plasma.

4.1.2 LHC Layout and Design

The overall layout of the LHC ring is shown in Fig. 4.1.

The LHC uses superconducting dipole magnets to steer two beams of hadrons around the LHC ring and collide these beams at one of four interaction points (IPs). A separate, independent particle physics experiment operates at each of these interaction points. The high-luminosity ATLAS and CMS experiments operate at Points 1 and 5, directly across the ring from one another (so as to ensure approximately equivalent delivered beam luminosity) while the heavy-ion experiment ALICE op-

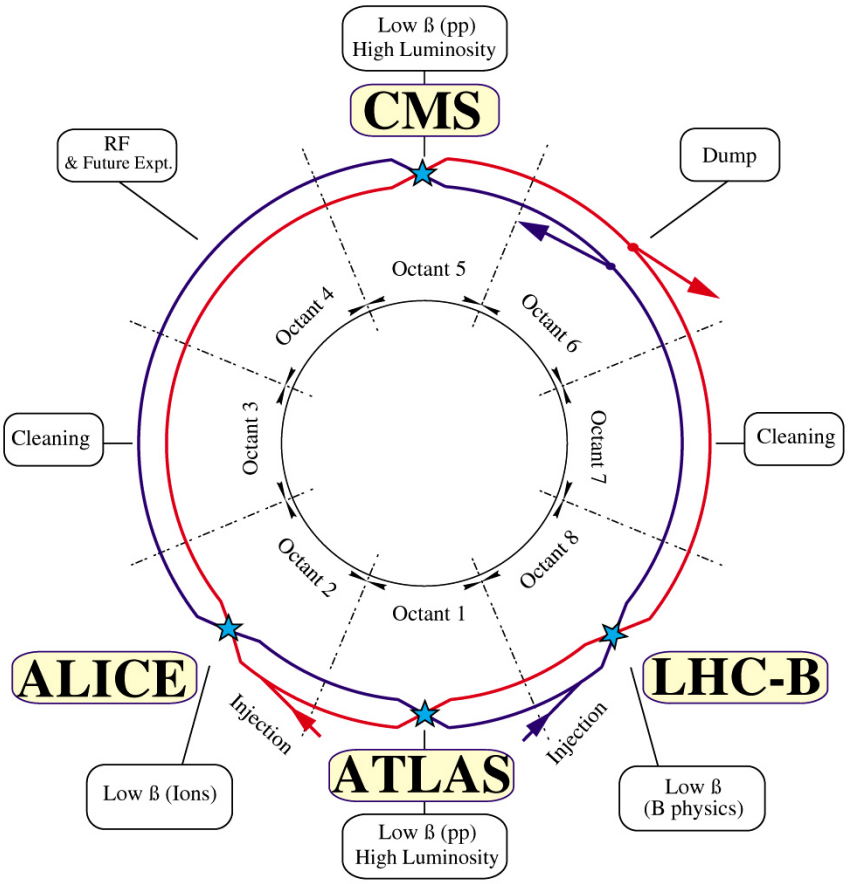

Figure 4.1: An overhead diagram of the LHC; the four major experiments are labeled, along with coloring to indicate the hadron beam directions (red — clockwise, blue — counter-clockwise). Reprinted from Fig. 2.1 of [65].

erates at Point 2 and the precision b-physics experiment LHCb operates at Point 8.

Figure 4.2 shows the LHC injection chain system used to accelerate hadrons from rest to their final beam energies. Two kinds of hadron beams are run in the LHC: beams of single protons and beams of ionized lead nuclei[1]. For proton beams, the injection chain begins with a bottle of hydrogen gas attached to a duoplasmatron.

In this duoplasmatron, this hydrogen is ionized by firing electrons into it,

[1]The lead beams are used for studying heavy-ion physics; consequently, we will not discuss them further here.

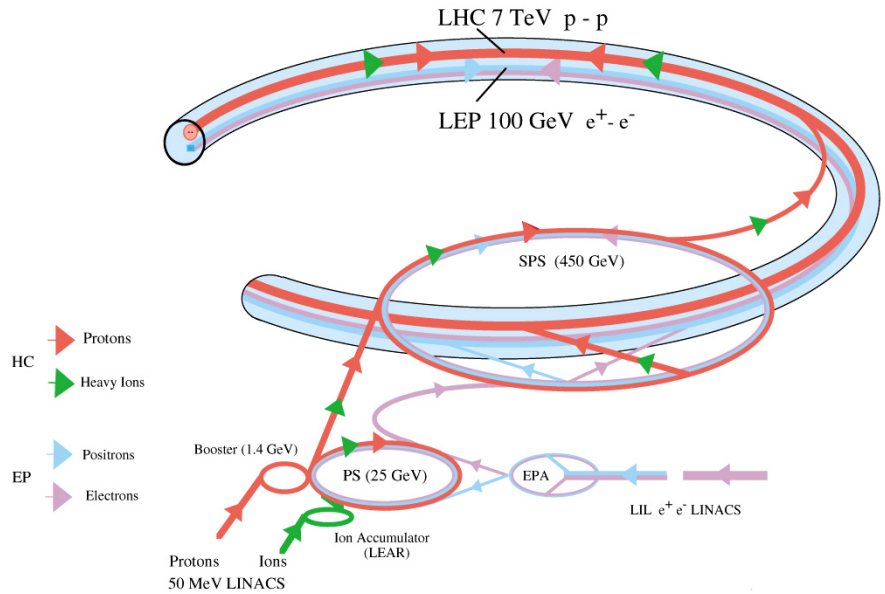

Figure 4.2: Diagram detailing the relative placement and final beam energies of the CERN accelerators that comprise the LHC acceleration chain [66].

forming a proton plasma. This proton plasma is then accelerated via charged grids and shaped by RF quadrupole magnets into a beam, leaving the duoplasmatron source to enter the Linac2 linear accelerator with a per-proton energy of 100 keV.

The Linac2, a multichamber RF cavity, applies an AC voltage to accelerate the proton beams. Drift tubes placed throughout the Linac2 serve to rectify the applied AC voltage so that the protons only experience a net acceleration. The Linac2 accelerates the proton beams up to 50 MeV before injecting them into the Proton Synchrotron (PS) Complex, consisting of the PS Booster followed by the PS.

As the protons leave the Linac2, a system of magnets focus (20 quadrupole magnets) and direct (2 bending and 8 steering magnets) the beams into the PS Booster. The PS Booster, which consists of 4 rings each comprised by 16 individual RF cavity sections, accelerates the proton beams up to 1.4 GeV and also splits them into individual bunches that are spaced, for the most part, in 25 ns intervals. The

PS Booster then injects the protons into the PS where the bunches are accelerated up to 25 GeV and further split by a factor of 3. Once this is accomplished, the beams are injected into the Super Proton Synchrotron (SPS). In the SPS, which uses the same kind of RF cavity system as the PS Booster and PS, the proton beams are accelerated up to an energy of 450 GeV before they are finally injected into the main LHC ring.

4.1.2.1 Proton Acceleration in the Main LHC Ring

The main LHC ring has 1,232 twin-bore dipole magnets that steer the high-energy beams of hadrons through their circular orbits. Each magnet consists of two beam pipes (one beam pipe for each proton beam traveling in opposite directions) surrounded by superconducting niobium-titanium (NbTi — critical temperature 9.2 °K) wires, all sitting within an iron yoke that provides a temperature-stabilizing mass. The yoke plus beam pipe assembly is cooled by a surrounding vessel of liquid helium, maintained at 1.9°K, that is itself thermally insulated from the rest of the LHC-ring cavern by a vacuum chamber.

In addition to the dipole magnets that steer the beams through the LHC ring, the LHC also uses quadrupole magnets to control the focus of the particle beams and higher-pole magnets (sextupole, decapole) to correct and to stabilize the trajectory of particles in order to improve the beam lifetime.

4.1.3 Performance Limitations of the LHC

The LHC was designed to study very high-energy, often rare, processes.

The design energy is 7 TeV per beam; in order to reach this, the main dipole magnets must provide relatively large magnetic fields of 8.33 T.

At high energies, charged particles in a circular accelerator lose energy through synchrotron radiation; this radiation loss has an inverse quartic dependence on the particle mass, $E_{\text{loss}} \sim m^{-4}$, meaning that electrons lose 10^{13} more energy than protons per turn with all other parameters kept constant. At the LHC design energy of 7 TeV, protons only lose about 7 keV/ turn, but a 7 TeV electron would lose $\sim 7 \times 10^4$ TeV / turn!

The delivered instantaneous luminosity of the LHC, $\mathcal{L}_{\text{inst.}}$, which is a direct measure of how often the LHC produces collisions, can be written entirely in terms of parameters of the LHC beams [65],

$$\mathcal{L}_{\text{inst.}} = \frac{n_b N_p^2 f_{rev} \gamma_r}{4\pi \epsilon_n \beta *} F \tag{4.1}$$

where n_b is the number of colliding bunches, N_p is the number of protons per bunch, f_{rev} is the revolution frequency, γ_r is the relativistic gamma factor, ϵ_n is the transverse beam emittance, $\beta*$ is the beta function at the collision point, and F is a geometrical factor arising from the non-zero crossing angle of the beams at the IP.

The transverse beam emittance, ϵ_n, is essentially the transverse (relative to beam direction) width of the beam at the collision point. The beta function, β, characterizes the bunch shape profile along the beam direction. Both ϵ_n and β are

carefully tuned in order to avoid resonant feedback loops that would destroy the beams – ϵ_n is typically $\mathcal{O}\,(10\,\mu\text{m})$ while $\beta*$ is typically 0.5 m [65].

The proton bunch density is limited by nonlinear beam-beam interactions; these can be measured by the linear tune shift,

$$\xi = \frac{N_p r_p}{4\pi\epsilon_n} \tag{4.2}$$

where r_p is the classical Bohr radius of the proton, $r_p = e^2/4\pi\epsilon_0 m_p c^2$. Previous hadron colliders have found that the net linear tune shift, which is linearly additive over all IPs requiring head-on proton collisions, should not exceed 0.015 [65]. This leads to an upper bound on the number of circulating protons, $N_p < 1.15 \times 10^{15}$ protons / bunch.

At the nominal LHC bunch spacing, 25 ns, $f_{rev} = 40\,\text{MHz}$.

As noted above, the geometric factor, F, arises from the non-zero crossing angle of the beams at the IP; this angle helps to mitigate long range beam-beam interactions that would otherwise threaten the beam stability. Assuming round, equally-sized beams with a typical bunch width much smaller than $\beta*$, F can be written as [65],

$$F = \left(1 + \left(\frac{\theta_c \sigma_z}{2\sigma*}\right)^2\right)^{-1/2} \tag{4.3}$$

In Eq. (4.3), θ_c is the maximum crossing angle at the IP (300 μ rad), σ_z is the RMS bunch length (7.55 cm), and $\sigma*$ is the transverse RMS beam size at the IP (16.6 μm), leading to $F \approx 83\%$.

Combining the above values yields a nominal instantaneous luminosity of 10 nb^{-1}, or $\mathcal{L} = 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$. This luminosity is not constant over a given fill

of the LHC, however, as the beam intensities and emittances change due to interactions. The degradation of beam intensities and emittances naturally leads to a decay of the instantaneous luminosity, which can be characterized by the time constant τ,

$$\tau = \frac{N_{p,\,0}}{\mathcal{L}_0 \sigma_{\text{tot}} k} \tag{4.4}$$

where $N_{p,\,0}$ is the initial proton bunch density, $\mathcal{L}_0$ is the initial beam luminosity, σ_{tot} is the total hadron interaction cross-section (10^{25} cm^{-2} at $\sqrt{s} = 14$ TeV), and k is the number of IPs (4). Plugging in the nominal values yields a decay time (for the 14 TeV LHC) of 29 hours; additional effects not accounted for in Eq. (4.4) further reduce this decay time to 14.9 hours.

Overall, the LHC, thanks to the high density of protons, the $\mathcal{O}$(half-day) of stable beams per fill, and the high beam energy, is the most powerful particle accelerator in the world.

4.2 The CMS Detector

The majority of information in this section comes from Ref. [67].

The idea for the CMS detector was originally conceived of in the early 1990s. After general approval by CERN management in the late 1990s, construction began and the detector was fully commissioned by the time the LHC first collided proton beams on September 10th, 2008.

Overview of CMS

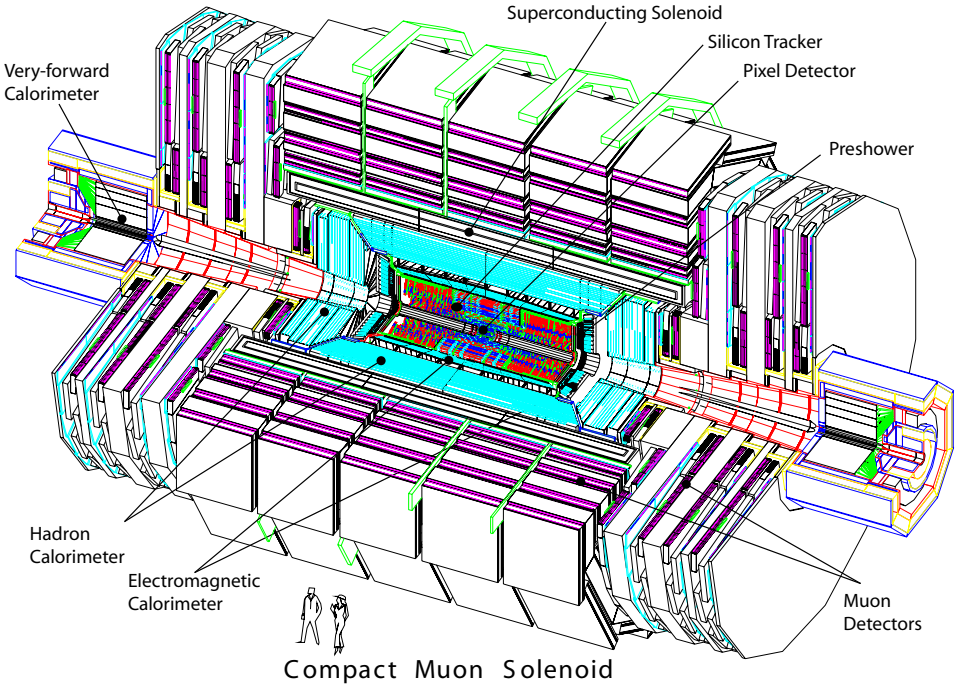

Figure 4.3: Inner view of the CMS detector. Reprinted from Fig. 1.1 of [67].

We provide here a quick overview of the CMS detector, with more extensive details included in individual sections in this chapter.

Figure 4.3 gives an inner view into the CMS detector. In some sense, it can

be thought of as a cylindrical onion, with a length of 22 m and a radius of 7 m, where two beams of hadrons enter from opposite sides of the detector and collide in the detector, typically within $\mathcal{O}$ (10 cm) of the exact detector center. The products from these collisions — e.g. electrons, hadrons, or γs, interact with the various sub-detectors located inside CMS (the layers of the onion), leaving behind charge and energy deposits that are subsequently read out and algorithmically combined to reconstruct a comprehensive picture of the original collision.

The central feature of CMS is a large, superconducting solenoid (Section 4.2.2), with a radius of 3 m, that provides a uniform, 3.8 T, axial magnetic field to bend the trajectories of charged particles as they leave the IP. Immediately surrounding the IP is a silicon tracking system (Section 4.2.3) that detects and measures these charged-particle trajectories. One of the subsystems of the tracking detector is used in the default method for measuring the instantaneous luminosity of hadron collisions (Section 4.2.7) being delivered to CMS by the LHC.

Also located in the solenoid field volume are a homogeneous electromagnetic calorimeter (Section 4.2.4) utilizing lead-tungstate crystals that measures deposits from (primarily) electrons and photons, as well as a brass, hadronic sampling-calorimeter (Section 4.2.5) that measures the energy from hadrons.

Outside of the solenoid's magnetic field volume, in the far-forward regions along the beam pipes, is an extensive forward calorimetry system (Section 4.2.5.5). These forward calorimeters significantly improve the hermeticity of CMS, especially in regards to the reconstruction of missing transverse energy ($\not{E}_{\mathrm{T}}$), by measuring the energy of high-energy, far-forward particles. The forward calorimeters are also

used to cross-check the instantaneous luminosity measurements performed with the tracking system.

Muons are measured partly in the silicon tracking system. High-p_{T} muons ($p_{\mathrm{T}} \gtrsim 5$ GeV) escape through the solenoid, where they are also measured by gas-ionization chambers (Section 4.2.6) located in the solenoid's iron return yoke.

In order to read out and record these particle interactions with the detector, CMS utilizes a sophisticated trigger and data-acquisition system (Section 4.2.8). The trigger system is a two-level system; first, custom-made programmable electronics reduce the effective rate of data-taking from the nominal LHC bunch spacing of 20–40 MHz down to 100 kHz; immediately thereafter, a processor farm running custom reconstruction software and algorithms performs a second, sequential reduction down to ~ 400 Hz.

4.2.1 Coordinate System for CMS

Due to its design geometry, CMS uses a cylindrical coordinate system, where z represents the coordinate along the beam axis, ϕ represents the azimuthal angle, calculated relative to the positive x-axis that points toward the center of the LHC ring, and θ is the polar angle calculated relative to the positive z-axis. The polar angle is commonly transformed into the pseudorapidity variable, $\eta = -\ln\left[\tan(\theta/2)\right]$, as differences in pseudorapidity are Lorentz-invariant. Pseudorapidity is an approximation of the standard rapidity variable, $y = \ln\left(\frac{E+p_z}{E-p_z}\right)$.

Design Goals for the CMS experiment

As with the LHC, the design of CMS was guided by several underlying physics goals. Because one of the main goals (arguably **the** main goal) of the LHC was to explore the origin of electroweak symmetry, most of the original physics goals for the CMS experiment were related to physics performance in the SM Higgs-boson analyses. Specifically, the CMS experiment was designed to have the following qualities [67]:

- An overall relative electromagnetic energy resolution of $< 0.5\%$ to ensure a precise measurement of the mass of the (any) light Higgs boson(s) in the $\gamma\gamma$ final state.

- Effective and robust rejection of π^0 and efficient photon and lepton isolation at high luminosities.

- A relative muon momentum resolution better than 6–17% for muon momenta $p = 1$ TeV (for reconstructing Z' masses) and 1–1.5% for muon momenta $p = 10$ GeV (for precision measurements of the Higgs boson mass in the $\mathrm{H} \rightarrow \mathrm{ZZ} \rightarrow 4\ell$ final state).

- Unambiguous determination ($< 1\%$ mis-assignment rate) of the relative sign of the charge of muons for muon momenta $p < 1$ TeV.

- A strong efficiency for spatially distinguishing interaction vertices close to the beam line and interaction points (using the transverse position to tag the

decays of heavy-flavor quarks and the longitudinal vertex position to tag pileup interactions).

- A relatively good measurement precision for dijet-mass and $\not{E}_{\mathrm{T}}$ reconstruction, requiring hermetic hadron calorimeters with lateral segmentation.

- Fast-response trigger systems and detector electronics to quickly read out particle energies from hadron collisions spaced as little as 25 ns apart.

- Fine-resolution detector granularity to minimize channel occupancy during future, planned high-luminosity runs of the LHC.

- Radiation-resistant detector materials, especially in the far-forward regions, in order to maintain optimal physics performance in the face of radiation damage from the copious amounts of high energy particles (e.g. neutrons) produced in LHC collisions.

These physics goals outlined above served as guiding principles in such things as detector material choices, inter- and intra-sub-detector layouts, electronics readout systems, and reconstruction algorithms.

We now discuss the individual components of the CMS detector.

4.2.2 The CMS Magnet

The central feature of the CMS detector is a superconducting magnet [67,68]. The magnet is a cylindrical solenoid with a radius $r = 3$ m and a length of 12.5 m. It is designed to generate a uniform magnetic field of 3.8 T within its free bore. In

order to do this, four winding layers of NdTi conductors[2] provide the necessary (and large) number of amperes per turn required (4.2×10^7 amp/turn). These winding layers are cooled (via saturated helium) to 4.5 °K and are split into five rings of equal length.

These conductors are mechanically reinforced with aluminum alloy; in spite of this reinforcement, the large amount of stored energy within the magnet, 2.7 GJ, with respect to the relatively small cold mass of the magnet, 220 metric tonnes, leads to large mechanical deformations of $\mathcal{O}\,(0.15\%)$ [67]. In comparison with other particle physics experiments, the CMS solenoid has the largest stored energy-to-mass ratio, as can be seen in Fig. 4.4.

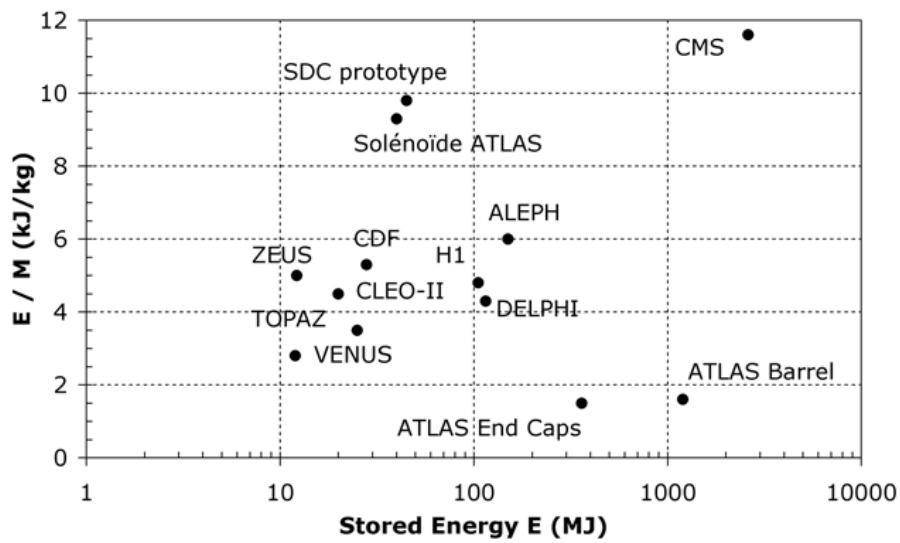

Figure 4.4: A comparison scatter plot of the stored energy E to mass M ratio, $\frac{E}{M}$, versus stored energy for a number of detector magnets from various particle-physics experiments. Reprinted from Fig. 2.3 of [67].

The return field of the CMS solenoid is directed through a 1.5 m, 10 kiloton iron yoke. Although the field inside the solenoid free bore is a uniform 3.8 T, the

[2] in contrast with the typical one used in other large solenoids.

uniformity and field strength outside of this region vary [69, 70], as can be seen in Fig 4.5.

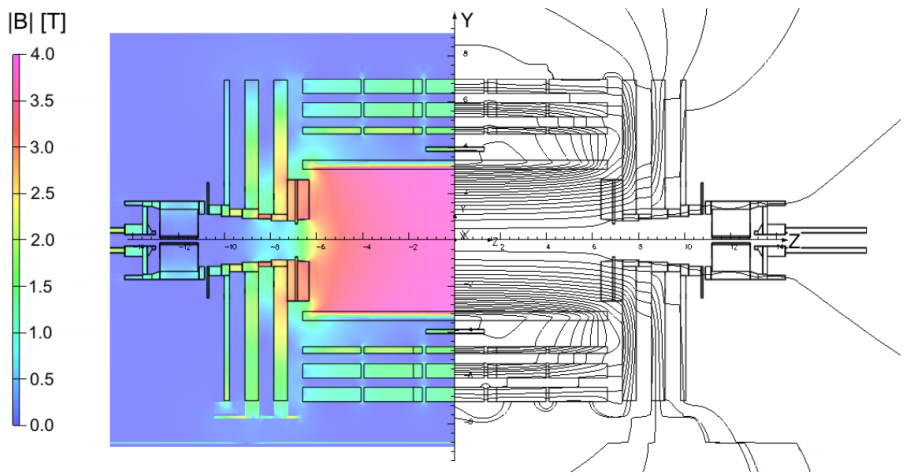

Figure 4.5: Two different diagrams of a longitudinal section of the magnetic field generated by the CMS solenoid. (Left) A color-map of the B field strength. (Right) A contour map of the B field lines, with each line representing a 6 Wb flux increase. Reprinted from Fig. 5 of [69].

4.2.3 Charged-particle Tracks at CMS

The detection and measurement of charged-particle tracks has played a crucial role in particle physics experiments, including as far back as the discovery of the positron in 1932. The innermost CMS sub-detector, the CMS tracker, is dedicated to the measurement and detection of charged-particle tracks.

There are many notable aspects of CMS analyses that depend crucially upon the robust performance of this inner detector. For analyses that use charged leptons (such as the search described in Chapter 6) the CMS tracker, in conjunction with the ECAL (Section 4.2.4) for electrons and the muon system (Section 4.2.6) for muons, provides high precision measurements of the charged-lepton momenta. Detailed

properties of the charged-lepton tracks, such as their transverse and longitudinal impact parameters with respect to the measured IP, enable discrimination between prompt charged leptons produced directly from W and Z decays and charged leptons produced through other means, e.g. electrons produced from photon conversions or muons produced through (semi-)leptonic decays of pions and kaons.

The CMS tracker also enables robust reconstruction of charged hadrons, improving the momentum measurement precision for hadronic jets. Furthermore, the reconstruction of charged hadrons facilitates the application of relative isolation requirements to selected leptons and photons in physics analyses, improving the rejection of "fake" leptons and photons – i.e. either jets spuriously reconstructed as leptons or photons or real leptons and photons produced during hadronic decays and fragmentation.

It is also worth noting that charged hadron reconstruction, impact parameter measurements, and secondary vertex identification comprise the main key ingredients for effective tagging (c.f. Section 5.6.2) of jets from heavy-flavor (b- and c-quarks) decays, which are a characteristic signature of the decays of supersymmetric partners of the third-generation quarks.

Finally, the accurate and precise reconstruction of individual charged particle kinematics plays a **crucial** role in the calculation of $\not{E}_{\mathrm{T}}$, a variable that is central to the searches described in Chapter 6 as well as detector performance studies [71]. In particular, the association of production vertices with individual charged particles enables the association of particles with pileup interactions, thereby enabling the mitigation of the negative impact of pileup on aspects of the $\not{E}_{\mathrm{T}}$ reconstruction

(Section 5.7.3).

In order to contextualize the design choices used in the CMS tracker, it is prudent to first discuss charged-particle tracking in general.

Measuring Charged-particle Momenta

The measurement of charged-particle momenta begins with the Lorentz force law,

$$\frac{d\vec{p}}{dt} = q\vec{v} \times \vec{B}, \tag{4.5}$$

where p is the particle momentum, q is the particle's electric charge, v is the particle's velocity, and B is the magnetic field the charged particle is traversing through. If the B field is uniform, Eq. (4.5) leads to helical trajectories for charged particles. Assuming, without loss of generality, a uniform B-field completely oriented in the z direction (with strength B_z), one can determine the radius r, relative to the z-axis, of the helical trajectory of the charged particle,

$$\frac{1}{r} = \frac{qB_z}{p_{\mathrm{T}}}, \tag{4.6}$$

Eq. (4.6) provides a direct means for calculating the particle's transverse momentum p_{T}, given a measured q, B_z, and r. The precision of the measurement of p_{T} thus naturally depends on the corresponding measurement precision for the other parameters. In particular the p_{T} measurement precision depends most strongly upon the determination of the radius of curvature r^{-1} of the particle track. One method to reconstruct r^{-1} is to sample the particle's helical trajectory at N points. In that

situation, this component of the resolution can be written as [72, 73],

$$\sigma_{p_{\mathrm{T}}} \sim \frac{1}{N} \frac{\sigma_x p_{\mathrm{T}}}{0.3 B_z R^2}, \tag{4.7}$$

where R is the radial distance (i.e. transverse to the z-axis) from the first hit to the last hit and σ_x is the hit position resolution in the transverse plane. In the context of tracker design, minimizing the p_{T} resolution requires a relatively "deep" (large R), high granularity (large N) detector with good transverse-hit position resolution. The technology and overall layout of the CMS tracking detector was chosen in order to optimize this momentum resolution.

The CMS Charged-particle Tracker

As noted above, the CMS tracker is the first sub-detector surrounding the IP. It is 5.8 m in length and 2.5 m in diameter. It is completely hermetic in ϕ and has a pseudorapidity coverage, $|\eta| < 2.5$.

The tracker is comprised of two subsystems: the pixel detector and the strip detector. Both utilize silicon semi-conductor p-n junctions, albeit in notably different geometric layouts, to record and measure the interactions of charged particles with the detector material. These p-n junctions are reverse biased with a large enough voltage such that no free charge is present in the bulk material (depleted bulk). When charged particles pass through the silicon, they ionize the tracker material, and these holes and electrons drift to edge sensors where they are read out.

In the pixel detector, the junction is designed by placing n^+-doped pixels on

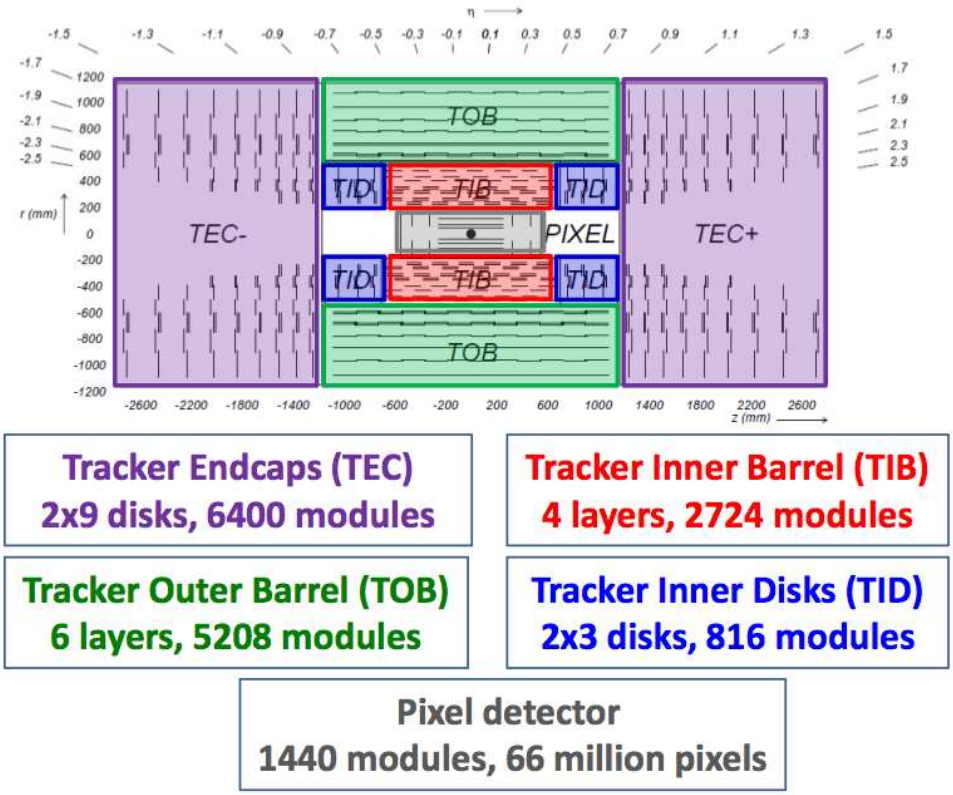

Figure 4.6: The schematic layout of the CMS tracker, including r, η positions and labels for the subsystems.

an overall p-doped bulk, while in the strip detector, the junction is designed by placing p^+-doped (SiO2) strips on an n-doped bulk [67].

Thin aluminum sensors along the strips/pixels, in conjunction with amplifiers and read-out systems, record the charge depositions, while aluminum plates provide the bias voltage that generates the bulk depletion region.

The inner tracking detector is closest to the beam line, and hence experiences large dosages of high-energy radiation.This radiation degrades the silicon through ionization damage (oxide deposits) and non-ionization damage, inducing defects in the silicon. This radiation damage reduces the overall signal efficiency of the tracker by both increasing the noise current and reducing the size of the overall

bulk depletion region.

In order to minimize the adverse effects of this radiation damage, the doping of the silicon was performed such that the depletion voltage decreases over time, leading to a lower required bias voltage (and hence, smaller noise current) to generate the same signal strength [67, 74]. Furthermore, during beam collisions, the silicon is maintained at a temperature of $-10\,^{\circ}\mathrm{C}$ to reduce leakage current through the semiconductor junctions [67]. This last step is particularly important because the leakage current is not only correlated with temperature, but also (for obvious reasons) heats up the silicon, which can lead to a run-away positive feedback loop.

We provide below additional details on the design and layout for each of the two tracker sub detectors.

4.2.3.1 The CMS Pixel Detector

The pixel detector handles charged track reconstruction in the closest region to the IP, $r \lesssim 10\,\mathrm{cm}$. The pixel detector consists of 48 million pixels spread across three barrel layers – designated collectively as the BPIX and covering up to $|\eta| < 1.5$ in pseudorapidity – and 18 million pixels arranged in two endcap layers on each side of the IP – known collectively as the FPIX and covering the pseudorapidity range, $1.5 < |\eta| < 2.5$. The "pixels" are hybrid silicon detector cells with dimensions $100 \times 150\,\mu\mathrm{m}^2$.

4.2.3.2 The CMS Strip Detector

The silicon strip detector handles charged track reconstruction in the region $20\,\text{cm} \lesssim r \lesssim 110\,\text{cm}$. It is comprised of four subsystems. The tracker inner barrel (TIB – labeled in red in Fig. 4.6) covers the regions, 25 cm $< r <$ 50 cm and 70 cm $< |z| <$ 70 cm. It consists of four layers, each using $320\,\mu$m thick silicon strips.

The tracker inner disks (TID – labeled in blue in Fig. 4.6) cover the regions, 20 cm $< r <$ 50 cm and 80 cm $< |z| <$ 90 cm. The TID consists of three layers. Like the TIB, the TID uses $320\,\mu$m thick silicon strips.

The tracker outer barrel (TOB – labeled in green in Fig. 4.6) covers the regions, 55 cm $< r <$ 116 cm and 00 cm $< |z| <$ 00 cm. The TOB consists of six layers. The first four of these use $500\,\mu$m thick strips; the last two use $122\,\mu$m thick strips.

The tracker endcaps (TEC – labeled in purple in Fig. 4.6) cover the regions, 22 cm $< r <$ 114 cm and 124 cm $< |z| <$ 280 cm. Each of the TECs have nine disks, using $320\,\mu$m ($500\,\mu$m) thick strips for the first four (last three) rings.

In total, there are 9.3 million silicon strips utilized across the four silicon strip subsystems.

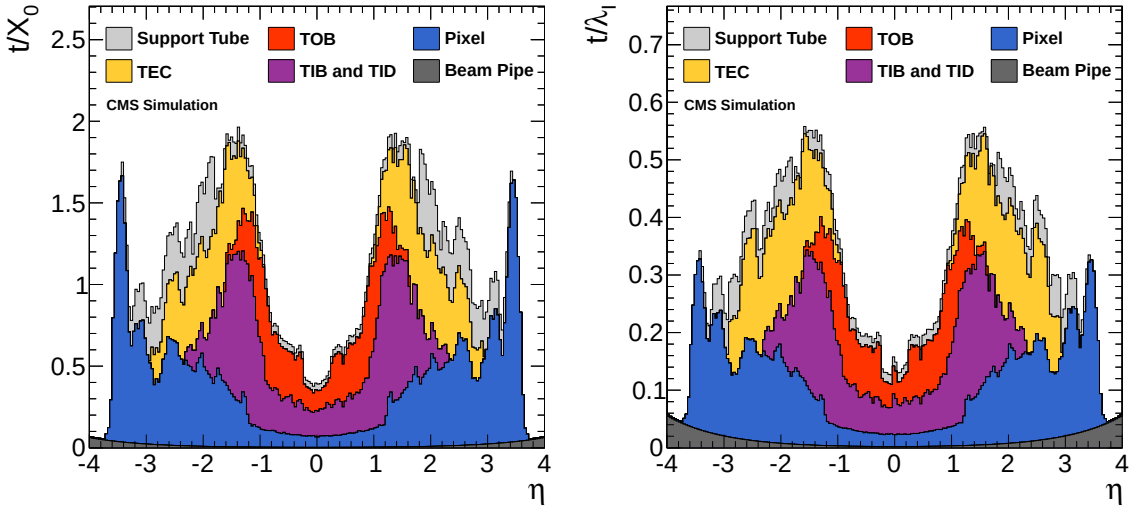

(a) CMS tracker material in units of radiation length, X_0.

(b) CMS tracker material in units of nuclear interaction length, λ_0.

Figure 4.7: The material budget for the CMS tracker, calculated using the latest simulation geometry as of June 2012. Reprinted from Fig. 2 of [75].

4.2.3.3 Tracker Summary

Tracker performance

Thanks to the design choices, in particular the small pixel size, the pixel detector has excellent track position resolution: track positions can be resolved to 10 μm in the r-ϕ direction and 20 μm in the z direction.

At design specifications, the pixel detector experiences an average hit occupancy per bunch crossing of 10^{-4} hits / pixel, while the strip detector experiences an average hit occupancy per bunch crossing of 10^{-2} hits / strip.

During the 2012 run of the LHC, the tracker maintained excellent performance. In the pixel sub-detector, 97.7% (92.8%) of the channels were operational in the BPIX (FPIX). For the silicon strip sub-detectors, 97.5% of the channels remained

active. Both the pixel and silicon strip sub-detectors maintained greater than 99% hit reconstruction efficiencies throughout the entirety of the 2012 run of the LHC. Consequently, the momentum resolution for charged particle tracks in the 2012 CMS dataset is quite excellent, as can be seen from Fig. 4.16, where the blue curve indicates the observed p_{T} resolution for muons when using just the inner tracker system.

4.2.4 Electromagnetic Calorimeter

The CMS electromagnetic calorimeter (ECAL) is the primary detector for photons at CMS. In particular, the accurate reconstruction of the diphoton final state of the decay, H$\rightarrow \gamma\gamma$, was one of the key driving factors in the design of the ECAL.

In addition, although the CMS tracking detector provides high precision measurements of electrons, the ECAL provides important, complementary measurements of electrons; in particular, the ECAL provides important additional measurements of high momentum electrons, where the tracker's measurement precision – see Eq. (4.7) – suffers notably, as well as capturing and measuring bremsstrahlung photons radiated by electrons in the tracker material (c.f. Fig. 4.7a).

The ECAL is a homogenous calorimeter comprised of 75,848 lead tungstate ($PbWO_4$) crystals. These crystal possess a number of properties that make them useful for electromagnetic particle detection at CMS.

4.2.4.1 EM Properties of $PbWO_4$ Crystals

Electrons and photons that pass through these crystals induce electromagnetic showers that create a large number – $\mathcal{O}(10^5)$ – of very low-energy photons. The $PbWO_4$ crystals are transparent to these photons, so the photons travel to the ends of the crystal where they are then measured by photodetectors (the exact photodetector technology depends upon the ECAL detector region – see Sec. 4.2.4.2 below).

Because $PbWO_4$ crystals have a high-density – 8.28 g cm^{-3} – and short radiation length – 0.89 cm – the crystals can be designed to be relatively compact along the r direction. Furthermore, $PbWO_4$ has a relatively small Moliére radius, $R_M = 2.2$ cm[3]; this allows for relatively fine detector granularity, facilitating the measurement of the direction and shower shape of electrons and photons, and in turn both improving the proper assignment of the photon vertex as well as helping the analysis-level discrimination between genuine and "fake" electrons and photons.

The scintillation time-scale of $PbWO_4$ is fast – approximately 80% of the scintillation light is emitted within the nominal LHC bunch spacing of 25 ns. This is important for minimizing the impact of out-of-time pileup interactions on particle measurements.

The spectrum of the scintillation light released in electromagnetic showers in $PbWO_4$ crystals is broadly centered around ≈440 nm (blue-green) [77, 78].

[3]The Moliére radius represents the typical transverse (relative to the incident particle direction) spread of the electromagnetic shower – on average 90% of the energy is contained within 1 R_M and 99% lies within 3.5 R_M [76]

4.2.4.2 Layout and instrumentation of the ECAL

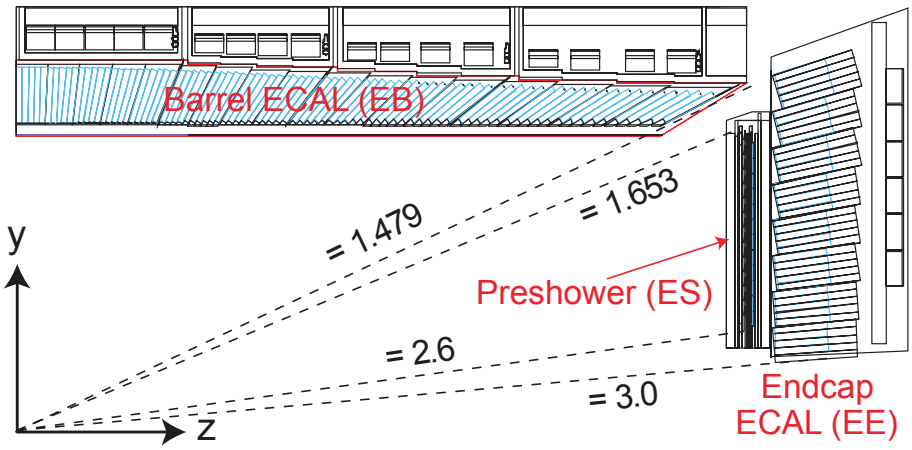

Figure 4.8: Longitudinal cross-section of the CMS ECAL, showing the relative $|\eta|$ placement of the ECAL sub-detectors.

Figure 4.8 shows the layout of the ECAL as a function of η.

The geometric layout of the ECAL is split into two primary sub-detector regions each covering 2π in ϕ and complementary regions in $|\eta|$. There is a barrel detector (ECAL Barrel, or EB) that covers up to $|\eta| = 1.479$, and two endcap wheel-shaped detectors (ECAL Endcap, or EE) that cover the pseudorapidity range, $1.479 < |\eta| < 3.0$. Although there is contiguous pseudorapidity coverage between the EB and EE, the energy resolution is significantly degraded in the region, $1.4442 < |\eta| < 1.566$, due to a splitting of coverage between the EB and EE at the point, $|\eta| = 1.479$.

4.2.4.3 ECAL Crystal Geometry

In the EB, the aforementioned $PbWO_4$ crystals are arranged into supermodules (SM), consisting of 1700 crystals each, arranged in a 20×85 grid in $\phi \times \eta$. Within an SM, the crystals are arranged in a slightly non-projective geometry to ensure complementary coverage of the cracks in between crystals – there is a 3° relative angle between the crystal central axes and the nominal interaction vertex [67]. To cover a single ϕ block in the EB, two SMs are laid end-to-end; the complete 2π coverage in ϕ requires 18 SMs in total. The $PbWO_4$ crystals in the EB have dimensions $2.2\,\text{cm} \times 2.2\,\text{cm} \times 23\,\text{cm}$ at the front face and $2.6\,\text{cm} \times 2.6\,\text{cm} \times 23\,\text{cm}$ at the back face. This corresponds to a radiation-length thickness of 25.8 X_0 and a crystal granularity in $\eta - \phi$ space of approximately 0.0174×0.0174.

In the EE, the $PbWO_4$ crystals are arranged into half-disks consisting of 3,662 tapered crystals, where each crystal has dimensions $2.86\,\text{cm} \times 2.86\,\text{cm} \times 22\,\text{cm}$ on the front face. and $3.0\,\text{cm} \times 3.0\,\text{cm} \times 22\,\text{cm}$ at the back face. With these crystal dimensions, the radiation-length thickness is 24.7 X_0 and the granularity in $\eta - \phi$ ranges from 0.0174×0.0174 at low η all the way up to 0.05×0.05 at high η. As with the EB, the crystals within each half-disk are arranged in a non-projective fashion – the crystals point at a point 1.3 m beyond the nominal interaction point, leading to relative axis-IP angles of 2° to 8°.

The Moliére radius for $PbWO_4$ (c.f. Sec. 4.2.4.1 above) is small enough that, in both the EB and EE, the EM shower generated from an electron or photon is typically contained within one $PbWO_4$ crystal. As noted before, the crystals are

transparent to the scintillation light produced in the EM showers – scintillation light produced during electron/photon EM showers travels through the entire crystal length before being read out by avalanche photodiodes (APDs) in the EB or vacuum phototriodes (VPTs) in the EE. Both types of photodetectors are run at $-18\,^{\circ}\mathrm{C}$; at this temperature both APDs and VPTs tend to measure approximately 4.5 photoelectrons per MeV.

ECAL photodetectors and radiation damage

Both APDs and VPTs are sensitive to radiation damage. APDs are sensitive to damage from neutrons; incident neutrons create defects in the silicon lattice, leading to an increase in the dark current. When the neutron dose is stopped (i.e. in-between LHC fills), these defects in the APDs undergo some degree of self-annealing. The short time component of this annealing is 20 days, however, meaning that appreciable reductions of dark current only occur during long-term technical stops of the LHC.

Figure 4.9a shows plots of the EB's APD dark currents as a function of time in terms of MeV. The self-annealing of APD defects is visible in this plot, in particular around January 2012 (the annual winter shutdown of the LHC). The neutron fluence in the EB is expected to be larger by a factor of 2 for high η relative to $\eta = 0$; the dark-current noise scales somewhat proportionally to this.

VPTs, like APDs, are also sensitive to damage from neutrons, with a similar net effect (increase in dark current). Figure 4.9b shows the dark current of the EE

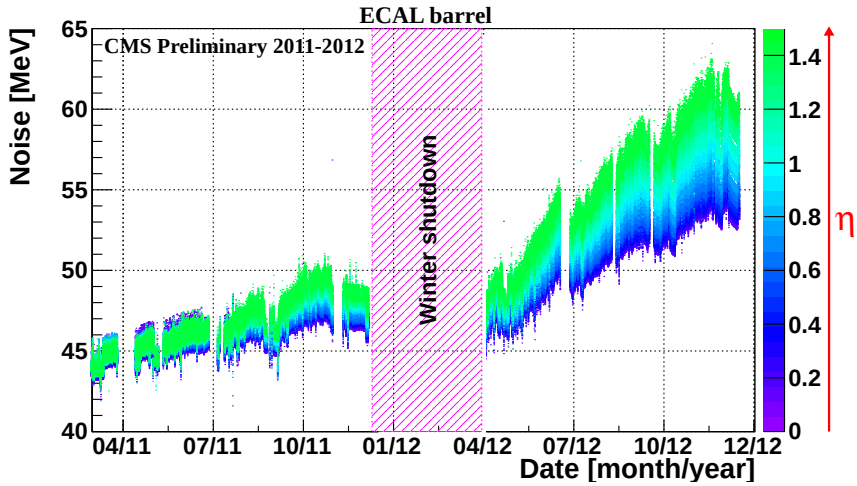

(a) Plot of the $|\eta|$ dependence of the average dark current measured per channel in the EB in units of MeV as a function of time.

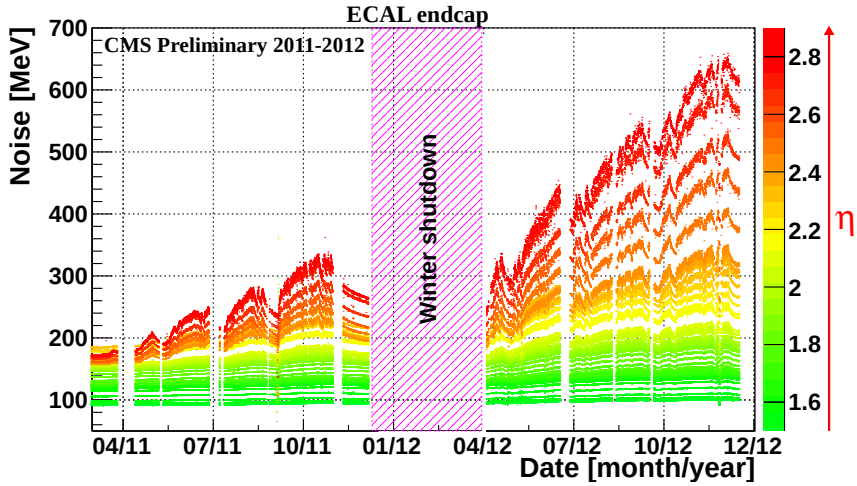

(b) Plot of the $|\eta|$ dependence of the average dark current measured per channel in the EE in units of MeV as a function of time.

Figure 4.9: Two plots showing the effects of radiation damage on the ECAL photodetectors [79].

VPTs as a function of time. As with the APDs, the defects created in the VPTs self-anneal, reducing the dark current, as can be seen around January 2012. The VPTs, being located in the EE, are subject to even larger neutron fluences than the APDs in the EB (and also larger variations of said fluences with η); this can correspondingly be seen in the relative growth of the dark current over the course of the 2012 run. Independent of radiation damage, VPT gains are known to fluctuate, depending upon incident light frequency and amplitude [80].

These gain fluctuations, as well as the radiation damage to the $PbWO_4$ crystals and photodetectors, threaten the accuracy and precision of the ECAL particle measurements. In order to correct for these effects, the CMS ECAL utilizes a calibration monitoring system consisting of 447 nm lasers as well as blue (450 nm) and orange (617 nm) LEDs. The laser measurements form the primary calibration tool. The blue LEDs are used to complement the laser measurements. Orange light is transparent to the $PbWO_4$ crystals, so the orange LEDs enable direct monitoring of the VPT gain fluctuations.

4.2.4.4 ECAL Pre-shower Detector

The relatively coarser η coverage of the EE at high-η subsequently leads to worse position and energy resolution for charged particles at high-η. In particular, the separation of individual photon showers, e.g. from the decay, $\pi^0 \to \gamma\gamma$, becomes difficult. To aid this differentiation, as well as improve the vertex position reconstruction for high-η individual photons, in front of each EE wheel, there is a

pre-shower detector (ES), covering the pseudorapidity range, $1.653 < |\eta| < 2.6$.

The ES detectors are sampling calorimeters consisting of two plates of silicon sensors – each containing 1,072 silicon strip sensors – interleaved with two planes of lead absorbers (with thicknesses of 2 and 1 X_0 for the front and back layers respectively); the lead blocks induce electromagnetic showers and the silicon strips measure the shower energy and position. Each silicon sensor has dimensions $6.3\,\text{cm} \times 6.3\,\text{cm} \times 0.032\,\text{cm}$. In practice, the pre-shower can resolve photons separated by a minimum of 2 mm.

4.2.4.5 Performance of the ECAL

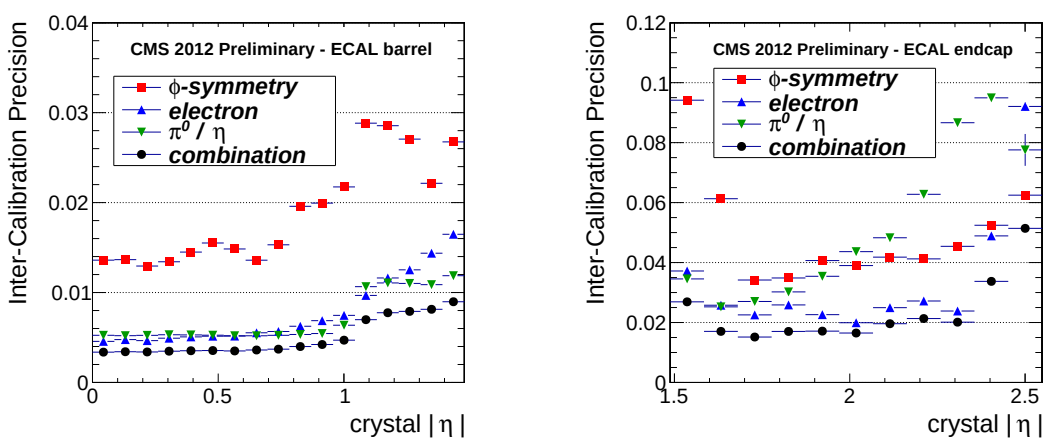

(a) The performance of the ECAL channel intercalibration in the EB.

(b) The performance of the ECAL channel intercalibration in the EE.

Figure 4.10: Precision of the ECAL channel inter-calibrations, using energy deposits, as a function of η [79]. The individual results from three methods [81, 82] — ϕ-symmetry, reconstruction of $\pi_0 \to \gamma\gamma$ and $\eta \to \gamma\gamma$ decays, and the measurement of the ratio E/p from W and Z decays — are shown, along with the weighted average of the three.

The overall ECAL energy resolution can be represented by the three-term

equation [83],

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2. \tag{4.8}$$

In Eq. (4.8), the first term (the stochastic term) represents the stochastic component of the measurement, driven by Poisson fluctuations in the number of photons produced in a given shower; the second term (the noise term) accounts for noise in the detector readout technology as well as photon excitations in the crystal; the third and final term (the constant term) accounts for imperfect inter-crystal calibrations (c.f. Fig. 4.10), non-uniform crystal performance, and the degree of shower non containment within one crystal.

Using test beam studies, S, N, and C were measured to be approximately 2.8%, 12%, and 0.3% respectively [83]. However, the *in-situ* precision is a little worse than the test beam would suggest. The biggest reason for this is radiation damage darkening the $PbWO_4$ crystals. The second largest effect worsening the *in-situ* resolution are crystal-to-crystal variations in energy resolution and transparency. Finally, although this is a very minor effect, temperature variations throughout the detector lead to gain variations that further worsen the resolution.

The ECAL crystals and photodetectors suffered radiation damage during the 2012 run; the average signal loss in the EB (EE) was $\sim 5\%$ ($\sim 18\%$). Nevertheless, the ECAL calibration system worked extremely well, enabling a robust recovery of the ECAL energy scale with minimal degradation of energy resolution. This can be seen clearly in Fig. 4.11 that shows the time dependence of the ratio of electron energy E as measured in the ECAL, against the electron momentum p using $\mathrm{W} \to \mathrm{e}\nu$

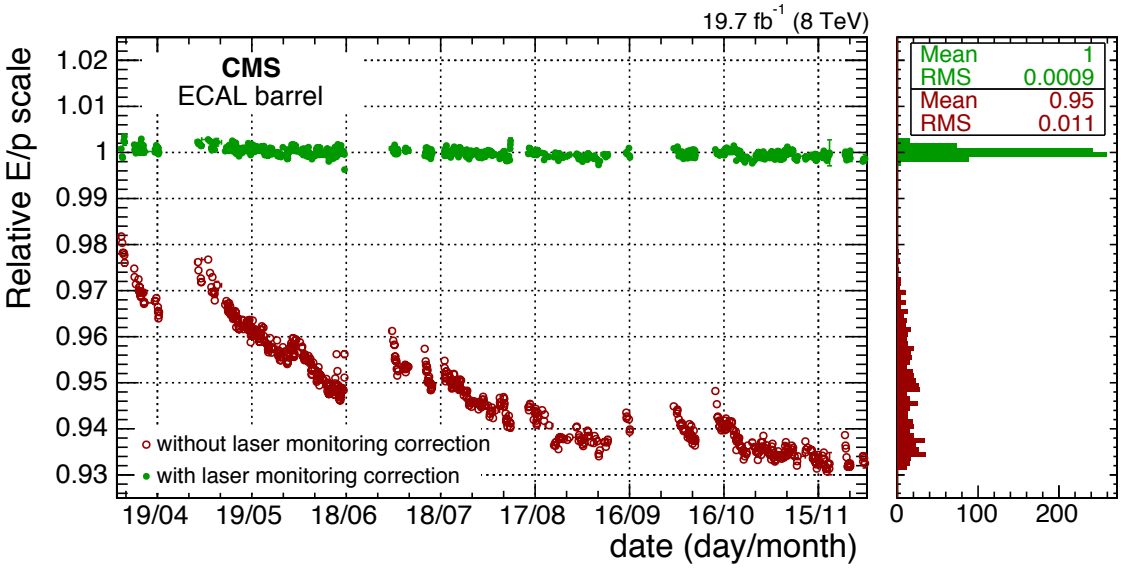

(a) The time-dependence of the E/p electron ratio in the EB.

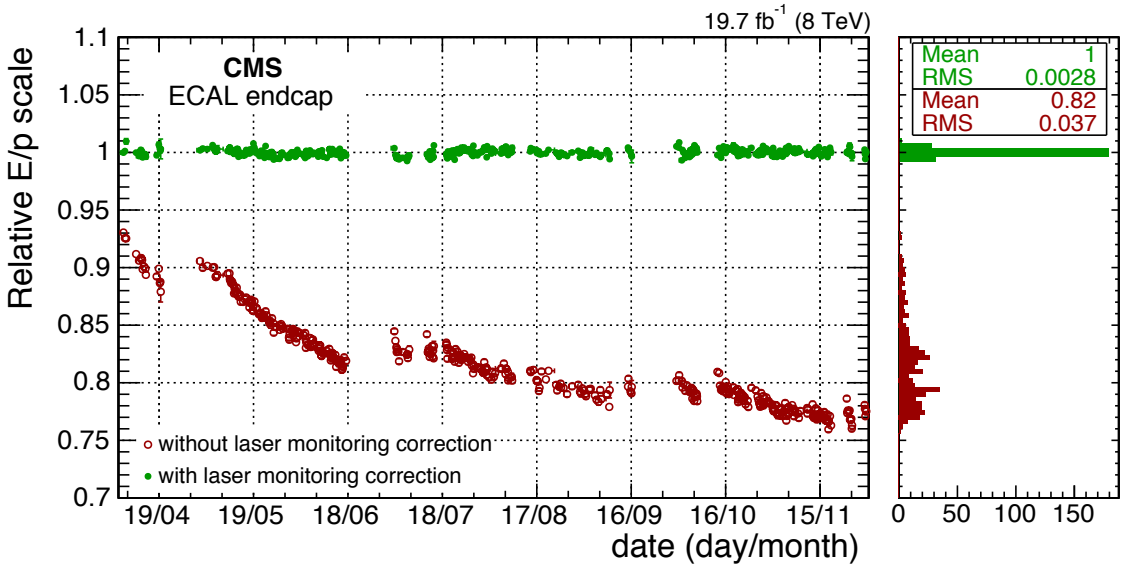

(b) The time-dependence of the E/p electron ratio in the EE.

Figure 4.11: Ratio comparison of electron energy, E, as measured in the EB (EE), against the electron momentum, p, as measured in the tracker. Each point is computed from 2×10^4 (1×10^4) W $\rightarrow$ eν events. Red points represent the ratio with no correction for crystal transparency loss, while green points show the recovery of a stable energy scale after the application of laser calibration corrections. Reprinted from Fig. 3 of [84].

events. The green points, which represent this ratio after the application of laser calibration corrections, demonstrate the recovery of the ECAL energy scale; the overall stability RMS introduced from these corrections is 0.09% (0.28%) for the EB (EE).

During the 2012 run, there were very few dead channels in the ECAL: For the EB, EE, and ES, the fraction of dead channels was approximately 0.9%, 1.6%, and 3.2% respectively.

4.2.5 Hadron Calorimeter

The CMS hadron calorimeter (HCAL) is designed to detect and measure hadrons, providing both additional energy precision for charged hadrons and, since they leave no deposits in the tracker, the **primary** energy and position measurements for neutral hadrons. Furthermore, the forward hadronic calorimeters (HF) provide additional high-η coverage, which is crucial for reconstructing certain interesting physics scenarios, such as the vector-boson fusion production mode for the Higgs boson, as well as greatly improving the hermeticity of CMS, which is crucial for the robust reconstruction of $\not{E}_{\mathrm{T}}$ (increasing the $|\eta|$ coverage from 3 to 5 improves $\not{E}_{\mathrm{T}}$ resolution by a factor of 3).

4.2.5.1 HCAL Layout and Instrumentation

The HCAL is divided into four sub-detectors that in total, provide η coverage up to $|\eta| < 5$.

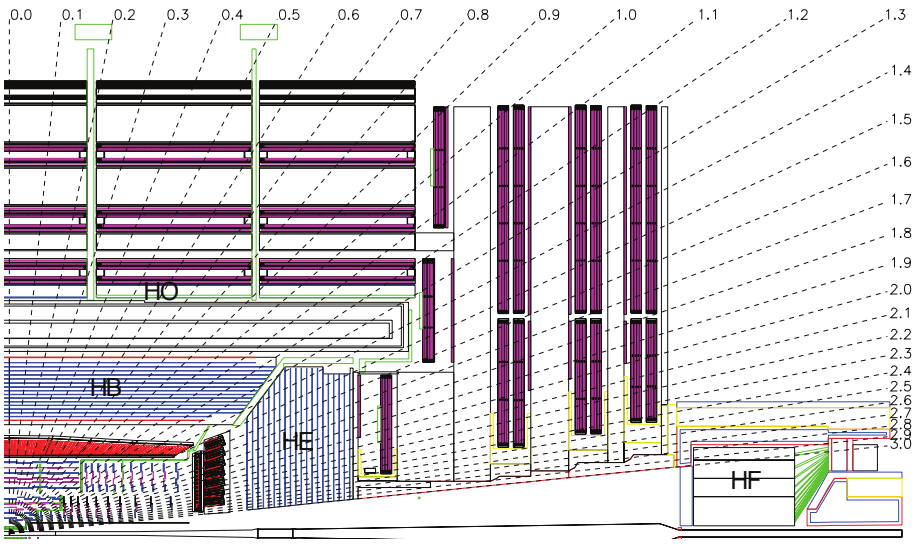

Figure 4.12: Longitudinal view of the CMS HCAL including demarcations in $|\eta|$; the blue layers designate the HB and HE; extending radially outward from the HB is the solenoid followed immediately by the HO; finally, the HF is located in the far-forward (large $|\eta|$) region on the right.

The HCAL barrel (HB) covers up to $|\eta| < 1.3$. the outer hadronic calorimeters (HO) cover the pseudo rapidity range, $0 < |\eta| < 1.2$; the HCAL endcaps (HE) cover the pseudo rapidity range, $1.3 < |\eta| < 3.0$; and finally, the forward hadronic calorimeters (HF) cover the pseudo rapidity range, $3 < |\eta| < 5$;

The first three sub-detectors are sampling calorimeters, which work in a general sense as follows: hadronic particles interact with absorbing layers and generate showers of pions. Any neutral pions, π^0, produced in these interactions decay promptly to photons, $\pi^0 \rightarrow \gamma\gamma$, that subsequently start electromagnetic showers. On the other hand, any charged pions, $\pi^{\pm}$ that are produced continue traveling through the material, generating their own hadronic showers from nuclear interactions. This process continues recursively until the hadronic particles being produced no longer have enough energy to start their own hadronic showers. At this point, they get stopped by the absorbing material until they decay, primarily to very low-energy

muons.

4.2.5.2 Hadronic Barrel

The hadronic barrel (HB) has 16 absorbing layers. Except for the first and last absorbing layers, the absorbing layers of the HB are 56.5 mm thick plates made of C26000 cartridge brass (70% copper, 30% zinc), which has a density of $8.53\,\mathrm{g/\,cm^3}$, a radiation length of 1.49 cm, and a nuclear interaction length, λ_0, of 16.42 cm. The first and last absorbing layers, for structural stability reasons, are 40 mm- and 75 mm-thick steel plates ($X_0 = 1.68$ cm, $\lambda_0 = 16.8$ cm). In total, the thickness, in units of λ_0, of the HB ranges from 5.82 at $|\eta| = 0$ up to 10.6 at $|\eta| = 1.3$.

It is worth noting that the EB, which radially lies in front of the HB, has a thickness of $1.1\lambda_0$, meaning that the initial electromagnetic energy generated in hadronic showers is often measured in the EB.

Most of the scintillating readout layers (layers 1 - 15) are 3.7 mm-thick Kuraray SCSN81 plastic scintillator, a polystyrene base doped with fluors. Layer 0, placed before the first steel absorbing layer, is 9 mm-thick Bicron BC408 plastic scintillator, a polyvinyltoluene base doped with fluors; Layer 0 measures any energy deposited by hadronic showers in the dead material between the EB and HB. Layer 16, like layers 1 - 15, is made of the Kuraray scintillator, but is 9 mm-thick in order capture the tails of hadronic showers that start in the later absorbing layers.

The scintillators are placed in tiles where the scintillation light is wavelength-shifted (WLS) downward by Kuraray Y-11 green WLS fibers before being read out by

hybrid photodiodes (HPDs). There are 70,000 tiles arranged in a radially projective geometry across 16 η divisions and 36 ϕ divisions, such that each scintillator tile covers a range in $\eta - \phi$ of 0.087×0.087.

4.2.5.3 Outer Hadronic Calorimeters

As noted above, at low $|\eta|$, the HB is relatively thin, which means that a non-trivial energy fraction of hadronic showers can escape detection in the HB. As an example, about 5% of all hadrons above 100 GeV fail to deposit all of their energy in the HB [85].

The outer hadronic calorimeters (HO) act as an extension to the HB to capture this escaping shower energy. The absorbing layer of the HO is the 19.5 cm-thick iron solenoid itself, leading to a minimum thickness for the combined HB + HO of 11.8 λ_0. The HO is divided into five rings with z-width 2.536 m, labeled Rings 0, ± 1, and ± 2. Ring 0, the central ring, covers up to $|\eta| < 0.35$; Rings ± 1, where the sign refers to $\pm\eta$, cover the pseudorapidity range, $0.35 < |\eta| < 0.87$. In a similar fashion, Rings ± 2 cover the pseudorapidity range, $0.87 < |\eta| < 1.2$. All 5 rings have a scintillating layer located outside the solenoid, but Ring 0 has an additional scintillating layer located inside the solenoid.

The HO's scintillation layers use the same technology as the HB (scintillator material, WLS fibers, and readout devices), and possess the same effective granularity in $\eta - \phi$. The HPD readout devices used in the HB and HO rely on relative alignment with any external magnetic fields – within the uniform solenoidal field

(from Fig. 4.5, all of the HB and the first layer of Ring 0 of the HO), this alignment is straightforward. Outside of the solenoid, however, misalignments (as large as 40°) for the HPDs in the HO lead to random discharges [85, 86].

4.2.5.4 Hadronic Endcaps

The hadronic endcaps (HE) have 17 absorbing layers, each of which is 79 mm-thick and is made of the same cartridge brass as the majority of the absorbing layers in the HB. There are 20,916 scintillator tiles in the HE, each of which uses the same technology as the HB and the HO, including an analogous Layer 0 scintillator layer designed to measure showers that start in the dead material between the EE and HE. For $|\eta| < 1.74$, these tiles have the same granularity in $\eta - \phi$ as the tiles in the HB and the HO – 0.087×0.087. For $|\eta| > 1.74$, the η segmentation ranges from 0.09 up to 0.35 and the ϕ segmentation is 0.175.

Including the EE, the HE calorimeter system is approximately 10 λ_0 thick.

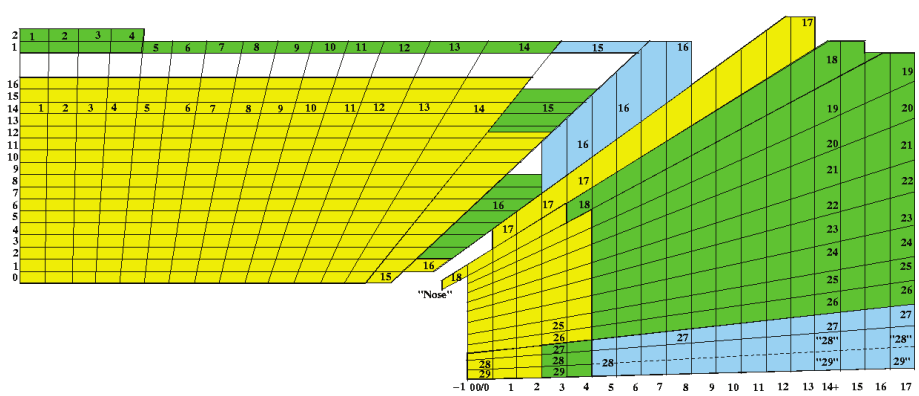

Figure 4.13: Longitudinal view of the CMS HCAL including demarcations in $|\eta|$; the blue layers designate the HB and HE; extending radially outward from the HB is the solenoid followed immediately by the HO; finally, the HF is located in the far-forward (η) region on the right. Figure adapted from Ref. [67]

In contrast to the HB and HO, the scintillating layers in the HE are split into multiple depths before being read out by HPDs; because the HE experiences a higher radiation dose than the HB, this additional segmentation enables a more precise recalibration of individual HE towers to mitigate the impact of radiation damage. Figure 4.13 displays a diagram of this depth segmentation.

4.2.5.5 Forward Hadronic Calorimeters

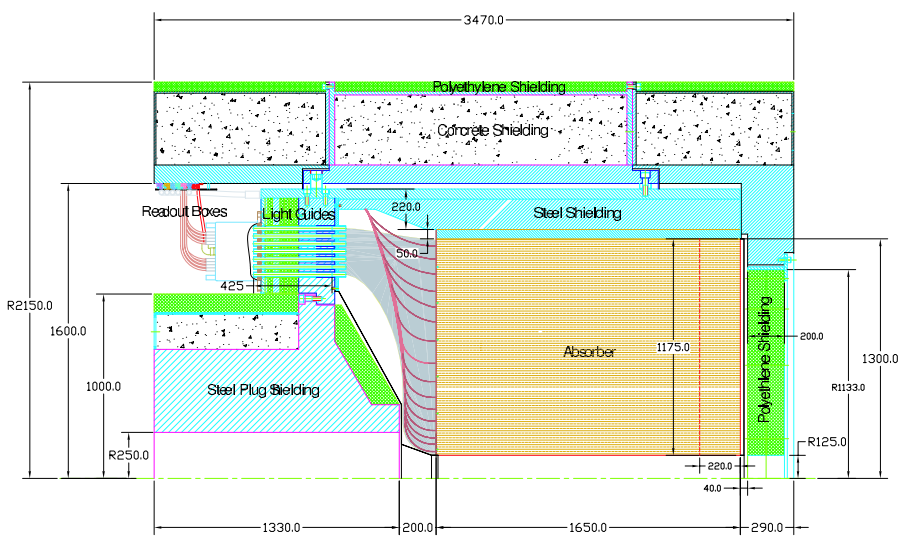

Figure 4.14: Cross-sectional (r-z) view of one of the HF calorimeters. The IP is off-screen to the right. Dimensions are in mm. Reprinted from Fig. 5.28 of [67].

Figure 4.14 shows a diagram of the forward hadronic calorimeters (HF). The HF calorimeters are built from 165 cm thick steel absorbers (10 λ_0). Embedded within the steel are 800 μm thick polymer-cladded quartz fibers, which are bundled together into 13 towers arranged in a non-projective geometry. For $|\eta| < 4.7$, the HF tower granularity in $\eta - \phi$ is 0.175×0.175. For $|\eta| > 4.7$, the HF tower granularity in $\eta - \phi$ is 0.175×0.35. These fibers measure particle showers using Cherenkov

radiation that is read out by PMTs located at the end of the HF. Half of these fibers run the full 165 cm length of the detector, while the other half start 22 cm in relative to the IP. Electromagnetic showers deposit most of their energy in the initial first 22 cm; thus, separate readouts for the two lengths of fibers can be used to distinguish between shower (and thus particle) types.

4.2.5.6 HCAL Performance

The energy resolution of the HCAL for single pions can be represented by,

$$\left(\frac{\sigma_E}{E}\right)^2 = \left(\frac{A}{\sqrt{E}}\right)^2 + B^2. \tag{4.9}$$

where A represents the stochastic term and B represents the constant noise term. For the HE, pion test beam studies in 2007 measured $A = 113\%$ and $B = 3\%$.

These numbers differ for hadronic jets, and depend upon the particular quark and gluon content of jets. Furthermore, the resolutions are often strongly improved by including tracker information into the jet reconstruction, see Sec. 5.6.

A number of individual channels in the HCAL experienced performance issues during the 2012 run of the LHC. This only strongly affected the HO, which was only partly operational. For the HB (HE), 99.80% (99.97%) of the channels were functional. The HF maintained 100% channel functionality.

4.2.6 Muon System

Muons interact in the CMS tracker, but deposit at most ~ 3 GeV of energy in the ECAL and HCAL systems. Because of this, CMS has an additional sub-

detector located outside of the solenoid that is dedicated primarily to the detection and additional measurement of muons produced in the hadron collisions.

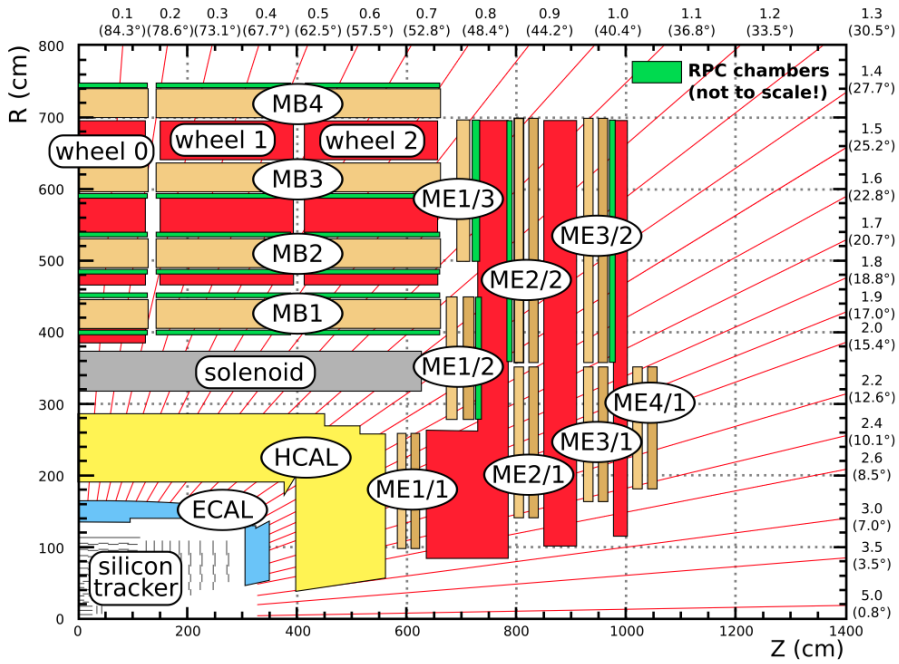

Figure 4.15: Longitudinal (r-z) cross section of the CMS detector, with a relative focus on the muon chambers. For geometric context the other subdetectors as well as the solenoid are labeled. The three gas-based muon detection technologies — RPC, DT, and CSC — are respectively denoted by green squares, the gold blocks labeled (MBn), and the gold blocks labeled (MEn/m). Included on the upper x-axis and right y-axis are values for η and θ. Reprinted from Fig. 3.1 of [87].

Figure 4.15 shows a longitudinal ($r-z$) view of the CMS muon system relative to the other CMS sub-detectors (e.g. tracker, ECAL, HCAL).

The CMS muon system is comprised of three subsystems, each using a different gas-based detection technology. These subsystems are built into the solenoid's return yoke — 10,000 tons of iron split across five barrel rings and six endcap disks that shape the magnetic return flux of the solenoid and absorb high-energy hadrons escaping the HCAL.

The three subsystems are the drift tubes (DTs), the cathode strip chambers (CSCs) and the resistive plate chambers (RPCs).

4.2.6.1 Drift Tubes

Drift tubes (DTs) are used in the muon system's barrel region, $|\eta| < 1.3$. They consist of anode wires placed between cathode strips within a tube-like structure. The entire tube is filled with a gas mixture of 85% Ar and 15% CO_2. A muon passing through the chamger ionizes the gas; the electrons produced from this ionization travel to the anode wire where they are measured. The net drift time and position of the wire constrain the location of the ionization deposition to an annulus around the anode wire. Three-dimensional position information for muons can be extracted by combining multiple DTs, oriented in orthogonal directions. Each DT station has a muon position resolution of $100\,\mu$m in the $r - \phi$ direction and $150\,\mu$m in the z direction with an overall $\sim 1\,$mrad precision on the muon direction.

4.2.6.2 Cathode Strip Chambers

Drift tubes are unusable in the endcap region, $0.9 < |\eta| < 2.4$, due to the highly non-uniform magnetic fields there, as well as the higher flux of hadronic punch-through and radiation. Cathode strip chambers (CSCs), which are more robust against the aforementioned effects, are used instead in the endcap regions. Anode wires are placed in between two cathodes, where one of the cathode planes is non-segmented and the other is segmented into strips oriented orthogonally to the

anode wires. The entire volume is filled with a gas mixture of 40% Ar, 50% CO_2, and 10% CF_4. Muons passing through the gas mixture leave ionization deposits; the electrons produced drift to the anode wires, where they produce a charge avalanche that is read out, giving information on the $r - \theta$ coordinate based on which wire. At the same time, the positive ions produced from the muon's interactions with the gas drift onto the cathodes. The induced charge distribution on the segmented cathode strips is fit with a four bit Gatti function [67,88], providing a measurement of the $r - \phi$ coordinate. By combining the information from multiple strips and wires, an overall 3D image can be reconstructed of the muon's track through the CSC chambers. Each CSC station has a position resolution of 75–150 μm in the $r - \phi$ direction and 200 μm in the z direction.

4.2.6.3 Resistive-plate Chambers

One issue with the DTs and CSCs is the relatively long rise time in their signal – typically $\mathcal{O}$ (10 ns) but as long as 400 ns for the DTs and 60 ns for the CSCs.

DTs and CSCs thus cannot be used for fast muon trigger decisions; furthermore, there can be concerns with properly assigning the correct LHC bunch crossing for a given muon detected in the DT and CSC. In order to have a dedicated fast muon p_{T} trigger, resistive-plate chambers (RPCs) are used in the pseudorapidity range, $|\eta| < 2.1$. RPCs consist of two layers of Bakelite, a highly resistive material ($\rho \sim 10^{10} \Omega\,\mathrm{cm}$), separated by a 2 mm gap filled with a gas mixture of 95.2% freon ($C_2H_2F_4$), 4.5% isobutane ($\mathrm{iso}C_4H_{10}$), and 0.3% sulphur hexafluoride (SF_6). When

an ionizing muon passes through this gap, a large voltage differential (9.6 kV) across this gap induces a discharge that is semi-proportional to the ionization deposit. This discharge develops extremely rapidly, i.e. faster than 3 ns, but is large geometrically, limiting the position resolution to a few cm.

4.2.6.4 Muon System Performance

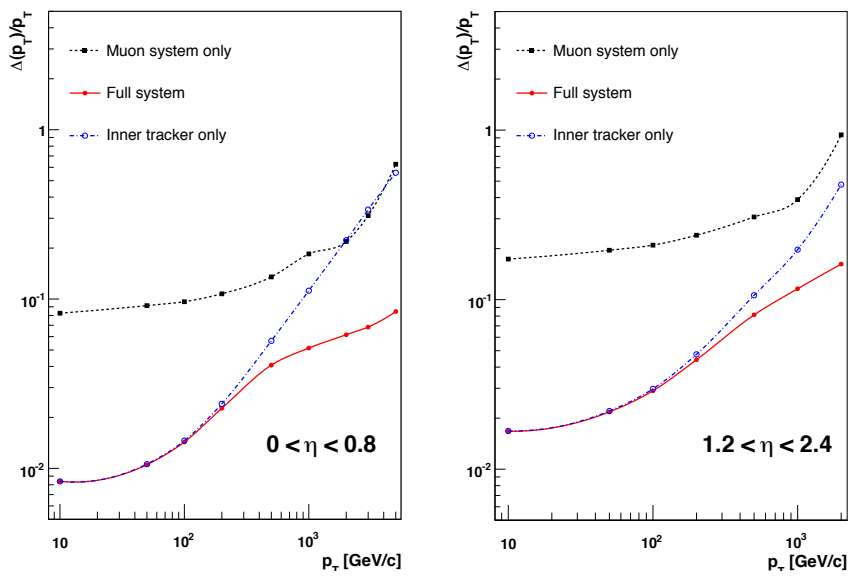

Figure 4.16: The dependence of the muon p_{T} resolution on muon p_{T}. The differently colored curves indicate the p_{T} resolution when using just the tracker information (blue), muon system information (black), or the combination of the two (red). Reprinted from Fig. 1.2 of [67].

The reconstruction of muons, see Sec. 5.3, can be performed using either the inner tracker, the outer muon systems, or a combination of both. Figure 4.16 shows the dependence of the reconstructed muon's relative p_{T} resolution on the muon p_{T}. The relatively large distance from the collision vertex hurts the muon reconstruction when using only the outer muon systems. Nevertheless, these outer muon systems provide useful, complementary information about the momentum of high-

p_{T} muons, yielding significant improvements over the muon reconstruction using the inner tracker alone.

4.2.7 Luminosity Measurement with CMS

The instantaneous luminosity is a measure of how much physics "data" is produced in a given beam crossing; higher instantaneous luminosities are directly correlated with the probability of a given event (e.g. production of a Higgs boson) occurring in a given proton-proton interaction

Instantaneous luminosities are often measured in units of $\mathrm{barns}^{-1}/\,\mathrm{s}$, where a barn (shorthand — b) is a unit of area comparable to the size of a heavy atomic nucleus, $10^{-24}\,\mathrm{cm}^2$. Calibrating luminosity in these units simplifies the discussion and comparison of different physics processes.

Measuring the instantaneous luminosity, and by extension, the total recorded luminosity integrated across an entire LHC fill, is crucial for physics analyses in order to properly normalize the expected number of events for the various physics processes considered. The luminosity can be expressed simply as [89, 90],

$$\mathcal{L} = \frac{\nu \langle n_{\mathrm{int}} \rangle}{\sigma_{\mathrm{int}}}, \tag{4.10}$$

where ν is the beam-crossing frequency of the bunches, $\langle n_{\mathrm{int}} \rangle$ is the average number of collisions per beam-crossing, and σ_{int} is the effective cross section for proton-proton interactions in a given beam-crossing. Determining $\mathcal{L}$ then boils down to determining $\langle n_{\mathrm{int}} \rangle$ and σ_{int}.

4.2.7.1 Measuring the Average Number of Interactions, $\langle n_{\text{int}} \rangle$

CMS uses three methods to measure the average number of interactions; one of these involves pixel counting in the pixel detector, while the other two involve using information from the HF calorimeter.

Pixel Counting

There are approximately 66 million pixels in the pixel detector (c.f. Sec. 4.2.3). At the design luminosity of the LHC, $\mathcal{L} = 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$, this leads to a smaller than 0.1% detector occupancy. Thus, a given pixel should be activated by at most one charged particle (and thus a single interaction) per bunch crossing. Consequently, the number of pixel hits is linearly correlated with $\langle n_{\text{int}} \rangle$ [89, 90]. Moreover, the pixel detector tends to be extremely stable during the entirety of a given LHC fill, and furthermore, the pixel detector will still maintain this relatively low occupancy even at the higher luminosity conditions expected in Run 2 and beyond, where 100 pileup events per bunch crossing will become the norm [90].

Measuring $\langle n_{\text{int}} \rangle$ with the HF

One disadvantage of the pixel counting method is that the pixel detector requires stable beams to be safely run, meaning that the measurement must be performed offline [90]. In contrast, the HF calorimeters (Sec. 4.2.5) can be safely run during unstable beams. CMS uses the HF calorimeters to measure the instantaneous luminosity using two separate methods.

Zero-counting

One of the two methods, the "zero-counting" technique relates the average fraction of empty HF towers to $\langle n_{\text{int}} \rangle$. Assuming the number of interactions in a given HF tower follow a Poisson distribution and that we can measure the probability of there being 0 interactions in a bunch crossing (in a given HF tower), we can directly calculate $\langle n_{\text{int}} \rangle$,

$$p(n_{\text{int}} = n) = \frac{\langle n_{\text{int}} \rangle^n \, e^{-\langle n_{\text{int}} \rangle}}{n!} \tag{4.11}$$

$$\rightarrow \langle n_{\text{int}} \rangle = -\ln(p(n_{\text{int}} = 0)), \tag{4.12}$$

Average E_{T} / Tower

The second method that utilizes the HF calorimeter calculates the average transverse energy, E_{T}, measured per tower. This is directly and linearly related to the luminosity; This method, however, suffers from slight instabilities in the gain of the PMTs used in the HF. Moreover the pileup conditions during the 2012 run caused the HF to operate in a nonlinear regime of its response, further adversely affecting the robustness of this measurement.

Summary of $\langle n_{\text{int}} \rangle$ Measurement Methods

Through a combination of the aforementioned two methods, the HF calorimeter can be used to measure the average instantaneous luminosity to a statistical uncertainty of 1% in under 1 s [89, 90].

Due to the aforementioned systematic effects that limit the accuracy of the HF-based methods, though, the pixel counting method is used as the default method for estimating $\langle n_{\text{int}} \rangle$.

4.2.7.2 Measuring the Beam Interaction Cross Section, σ_{int}

In order to measure, σ_{int}, CMS uses what is known as a Van De Meer scan [91]. The two LHC beams of protons are slowly scanned across one-another in the transverse plane. The relative interaction rate is measured for each point of the scan. The integral of these relative rates across the $x - y$ plane is then directly related to σ_{int}.

4.2.7.3 Summary of the Luminosity Measurement Method

Through a combination of the two measurements from Secs. 4.2.7.1 and 4.2.7.2, CMS can currently measure the instantaneous luminosity per bunch crossing with an overall precision of 2.5% (syst.) $\pm$ 0.5% (stat.) = 2.6%.

4.2.8 Trigger and Data-acquisition Systems

Under nominal running conditions, the LHC produces hadron collision data at a rate of 20–40 MHz, although accounting for the total proton-proton interaction cross-section (100 mb), this rate is closer to 1 MHz; given a typical event size of 1.25 MB, this leads to a byte production rate of 1.25 TB/ s.

It is completely unfeasible to fully reconstruct events in this time-frame, much

less physically record this data. Thus, the onus is to select promptly the events most likely to contain interesting physics and to record only these events. CMS has a system of hardware and software dedicated to this, known as the CMS trigger system. It is divided into two systems that sequentially throttle the data-taking rate down to feasible levels: a hardware-level trigger known as the level 1 (L1) trigger that reduces the rate down to at least 100 kHz, and a software-level trigger, known as the high-level trigger (HLT) that further reduces the data output rate to $\mathcal{O}$ (100 Hz).

Level 1 Trigger System

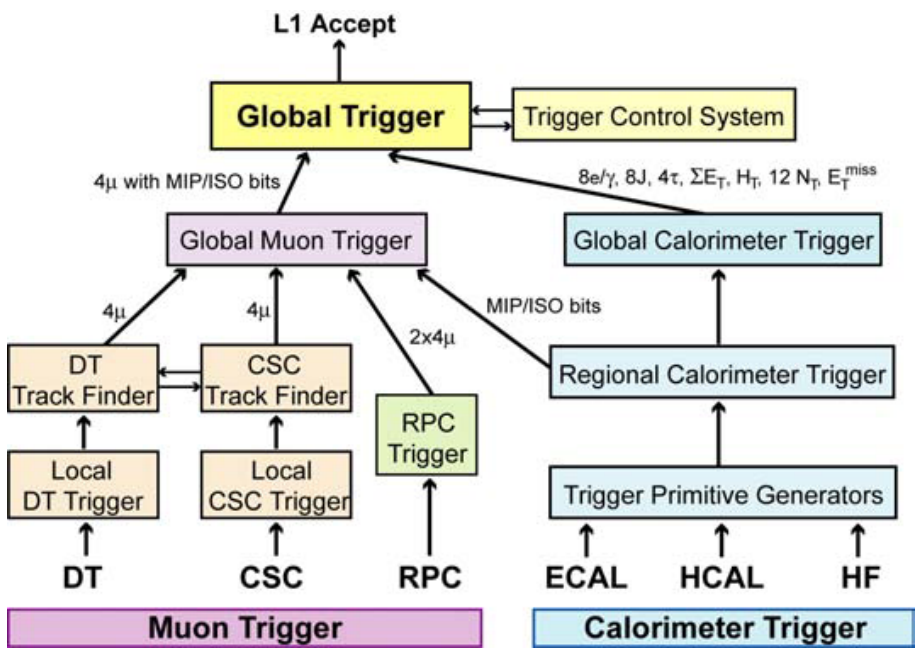

Figure 4.17: A flow schematic of the L1 Trigger. Reprinted from Fig. 8.1 of [67].

The L1 trigger is divided into individual detector triggers:

- A fast muon track reconstruction trigger.

- An electron and γ trigger based on reconstructed EM deposits.
- A fast hadronic jet reconstruction trigger
- A $\not{E}_{\mathrm{T}}$ trigger based on the vector sum of calorimeter deposits.
- An H_{T} trigger based on the scalar sums of calorimeter deposits.

These individual L1 sub-triggers have a small latency window in which to work, as the readout buffer can store at most 200 events for processing. The L1 trigger system as a whole can use individual sub-triggers by themselves or in concert with one another to make fast trigger decisions — 3.2 μs per event at most.

It accomplishes this through the use of extremely high-speed, custom-designed hardwired electronics, including field-programmable gate arrays (FPGAs), memory lookup tables (LUTs), and application-specific integrated circuits (ASICs).

Figure 4.17 displays a flow schematic for the L1 trigger system. As can be seen from this diagram, the muon triggers largely operate independently of the calorimeter triggers, although there is some information sharing between the two to create the global muon trigger, due to, for example, minimum ionizing particles (MIPs). The global muon and calorimeter triggers are combined into an overall global trigger that assigns the L1 acceptance value. Events that are accepted by the L1 trigger decision are passed along to the HLT trigger system.

HLT Trigger System

The HLT system is completely software based and is implemented on a large computing farm (over 13,000 CPUs divided across 720 computing nodes [92, 93]). The HLT system emulates the full offline reconstruction (c.f. Chapter 5). Successive stages of reconstruction are checked against sets of trigger conditions known as a trigger path. As an example, one trigger path might require at least one reconstructed muon candidate with $p_{\mathrm{T}} > 28$ GeV, while another one might require two reconstructed muon candidates, one with $p_{\mathrm{T}} > 17$ GeV, and another with $p_{\mathrm{T}} > 8$ GeV.

The HLT is designed to be as computationally expedient as possible; this is accomplished through several optimizations. Within a given trigger path, the reconstruction algorithms are grouped so that the fastest reconstructions are run first; if these reconstructed objects fail to pass the trigger path's object-specific quality filter, the rest of the path (and thus, rest of the reconstruction) is skipped. Moreover, the reconstruction algorithms themselves are seeded with candidate objects from the L1 trigger system; when performing the object reconstruction, the algorithms only consider a small area around these L1 seed objects. Finally, every trigger path has a characteristic known as a prescale. A prescale of N means that only every Nth event to pass a trigger path is recorded.

Most trigger paths, especially those involving e and μ candidates, have prescales of 1, so all events that pass the trigger selection are recorded. Prescales greater than 1 are primarily designed to further throttle the data-taking rate for events that are likely to be "uninteresting" (for example, most events that pass a simple require-

ment of two reconstructed jets with $p_{\mathrm{T}} > 30$ GeV are far more likely to be QCD multijet events rather than some new, exotic di-jet resonance).

The HLT system utilizes different sets of trigger paths known as trigger menus. This enables the HLT system to accommodate different data-taking conditions. For example, as the instantaneous luminosity decays during the course of an LHC fill, the prescale of a given trigger path is often lowered in order to maintain approximately the same data-taking rate.

4.2.8.1 Data Acquisition and Storage

After the events pass the HLT selections, a custom data-acquisition system (DAQ) records the information from each sub-detector. This information is sent to a dedicated computer processing farm located in the same building as the main detector. These linked CPUs perform full reconstructions of the event data. After this prompt, full reconstruction of the data, it is then disseminated worldwide into institution data centers using a tiered storage structure.

CERN serves as Tier 0, and has both analog and digital copies of the complete sets of data recorded by the CMS DAQ. Additional digital copies of these data are also stored at 13 Tier 1 sites located around the world, including Fermilab and DESY. The Tier 1 institutions are also the primary institutions for large-scale reprocessing of the full CMS dataset.

After Tier 1, there are the Tier 2 institutions (approximately 155 worldwide) and Tier 3 institutions ($>$ 300 institutions worldwide). Both the Tier 2 and Tier

3 institutions store digital copies of parts of the full CMS dataset. These digital copies are apportioned and duplicated in such a way that a complete copy of the full, recorded CMS dataset can be created without needing all of the Tier 2 and Tier 3 institutions to be accessible. In addition to data storage, the Tier 2 institutions are primarily responsible for the generation, reprocessing, and storage of simulated physics events (c.f. Section 4.2.9.2).

4.2.8.2 Run 1 Performance of the Trigger and DAQ

The CMS trigger system had excellent performance metrics during the 2012 LHC run. The L1 Trigger system maintained rates up to 100 kHz with only 3% dead time at worst [94]. The HLT system maintained rates up to 1 kHz, utilizing only 200 ms on average to process events [93].

The overall data-taking efficiency by CMS increased from 91% in 2010 to 93% in 2011 and 95% in 2012. The majority of this dead time came from either detector infrastructure issues or the sub-detector DAQs, as the central DAQ availability exceeded 99.6% during all three years of data taking [95].

4.2.9 Event Generation and CMS Detector Simulations

The majority of CMS analyses compare recorded collision events against simulated proton-proton collision events. These simulated events are built primarily from a combination of two separate simulation processes: the generation of the proton-proton collision events and the subsequent simulation of the detector response to

the collision products in each event.

4.2.9.1 Event Generation

Protons are composite particles, meaning that an individual proton's momentum is fractionally split amongst its comprising partons. The distributions of these parton momentum fractions, known as parton distribution functions (PDFs), depend both upon the type of parton considered as well as the energy scale, Q^2, that the proton is being probed at. Several groups measure these PDFs using experimental data, with the exact analysis approaches and empirical proton behavior models varying from group to group. CMS primarily uses PDFs provided by Martin-Stirling-Thorne-Watt (MSTW) and the Coordinated Theoretical-Experimental Project on QCD (CTEQ). At the LHC, which probes notably large Q^2, for most values of the momentum fraction, x, up to around $x \approx 0.1$, the gluon PDFs overwhelmingly dominate all other parton PDFs. The PDFs for the proton valence quarks, the two up quarks and the one down quark, dominate all other parton PDFs and also dominate the gluon PDF for values of $x \gtrsim 0.1$.

Under the QCD factorization theorem [96], large logarithmic factors in QCD scattering calculations, such as the collinear emission of gluons from a quark, can be included in the definition of the PDFs for all hard-scattering processes. Thus, one can calculate the total cross-section, $\sigma_{AB\to X}(s)$, for the scattering process $a+b \to X$ at a center of mass energy $\sqrt{s}$, where a and b are partons from their respective parent

hadrons A and B, and X in an arbitrary final state, as,

$$\sigma_{AB\to X}(s) = \int \mathrm{dx}_a \mathrm{dx}_b f_a(x_a, \mu_F^2) f_b(x_b, \mu_F^2) \hat{\sigma}_{ab\to X}(\hat{s}, \mu_R^2). \tag{4.13}$$

In Eq. (4.13), x_i refers to the momentum fraction of parton i; $f_i(x_i, \mu_F^2)$ is the PDF of parton i at the factorization scale, μ_F, the scale that separate long-range and short-range physics in the interaction; $\hat{\sigma}_{ab\to X}(\hat{s}, \mu_R^2)$ is the cross-section for the scattering process $a + b \to X$; and μ_R is the renormalization scale of the QCD coupling parameter. Typically, when simulating these parton scattering processes, μ_F and μ_R are set to be equivalent to one another. As well, the final state X can include additional radiation of gluons or photons from the initial or final charged particles (ISR and FSR respectively); however, this is often handled empirically through parton-showering processes.

Finally, the remaining partons that did not participate in the primary hard scattering often undergo soft interactions with one another; these soft interactions comprise the "underlying event", a sea of low-energy quarks and gluons that accompany most hard scattering events.

Different computer programs are often used to simulate the various components of the proton-proton interactions. The hard-scattering cross-section, $\hat{\sigma}_{ab\to X}(\hat{s}, \mu_R^2)$, is often simulated using $2 \to n$ matrix element calculations. Programs such as PYTHIA [97] or MADGRAPH [98] simulate the primary hard-scattering at LO in the matrix element and PDFs. Notably, MADGRAPH simulates ISR and FSR through consideration of additional LO diagrams where the ISR or FSR have been explicitly included in the diagram. Other programs, such as POWHEG [99], utilize NLO matrix

elements and PDFs.

Commonly, the hadronization of partons produced in the hard scatter, as well as the parton-showering from soft or collinear radiation, are handled through empirical models implemented in PYTHIA.

When programs other than PYTHIA are used to simulate the primary hard-scattering, the linking between PYTHIA and the hard-scattering calculation are performed so as to avoid double counting.

Occasionally, dedicated programs are used to simulate the decays of specific particles. For example, the program TAUOLA [100] provides one of the most accurate simulations of tau lepton decays.

4.2.9.2 CMS Detector Simulation

After event generation, the list of particles produced in the event are fed into a simulation of the CMS detector in order to model the expected detector response to these particles.

There are two primary versions of the CMS detector simulation:

1. CMS FullSim: A full simulation based on GEANT4 [101, 102] that includes modeling of the full geometry of the CMS detectors (including malfunctioning sub-detector elements), the effects of materials and detector electronics, and simulations of the reconstruction of sub-detector hits.

2. CMS FastSim: A fast simulation [103, 104] that requires a factor of ~ 100 less time, compared to the FullSim, to simulate and reconstruct events. The

FastSim achieves this increase in speed by utilizing simplified models of the CMS detector geometry and by parametrizing material and electronics effects at the level of reconstructed hits. The FastSim can model most physics objects with an accuracy comparable to the FullSim [103, 105].

For the purposes of accuracy, the CMS FullSim is used to generate most simulated event samples, especially centrally produced background samples used in multiple analyses (e.g. SM $Z \to \ell^+\ell^-$ or $t\bar{t}$ events). The CMS FastSim is used when it is infeasible, from a computational resource standpoint, to generate sufficient statistics for the desired event sample. For example, most SUSY searches scan through $\mathcal{O}(300)$ points in a 2D plane of unknown SUSY particle masses, where each point requires its own statistically orthogonal event sample.

Chapter 5: Object/Event Reconstruction at CMS

The different particles produced in the hadronic collisions delivered by the LHC interact with detector materials in various ways. As discussed in Section 4.2, the CMS detector utilizes a wide range of detector technologies across multiple sub-detectors to measure and record the kinematic qualities of these particles.

In order to use these collisions for physics analyses, the information from these sub-detectors has to be properly combined and reconstructed. Multiple kinds of particles can interact with a given sub-detector — e.g. any particle with electric charge will interact with the CMS tracker. Fortunately, however, the CMS sub-detectors provide complementary information; algorithms that combine this information can provide accurate particle measurement and identification.

5.1 Particle Flow Reconstruction

CMS applies this paradigm — combining information from multiple sub-detectors to construct a global picture of an event — in what is known as the "particle-flow" (PF) approach to object and event reconstruction [106, 107]. Extensive work has been performed to test and commission the PF algorithm — using early 7 TeV CMS data using minimum-bias and jet events [108] as well as leptons from J/Ψ and W

decays [109] and using 8 TeV CMS data in a wide variety of events [110]. The PF algorithm has been found to give significantly better performance, particularly in the areas of jet and $\not{E}_{\mathrm{T}}$ reconstruction, when compared against simpler reconstruction algorithms.

In the remainder of this section, we will detail the basic principles of this approach, as well as the particular details of its implementation at CMS.

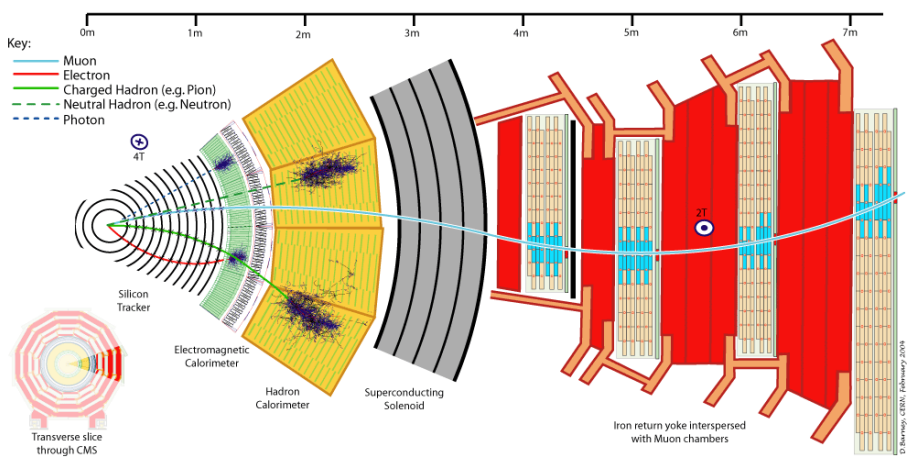

Figure 5.1: A transverse wedge cross-section of the CMS detector with the various sub-detectors labeled. Also shown are typical visible signatures left by the various kinds of particles CMS can detect.

To see an example of how this approach works, consider the passage of various particles through CMS, as shown in Fig. 5.1. Muons, electrons, and charged hadrons create tracks in the tracker, while photons and neutral hadrons do not. Electrons and photons deposit all of their (remaining) energy in the ECAL, while charged and neutral hadrons deposit the bulk of their energy in the HCAL. Muons lose very little energy interacting with the ECAL and HCAL and high-energy muons escape all the way through to the muon systems[1]. Finally, neutrinos (e.g. from the leptonic

[1] $dE/dx \sim 0.26$ GeV$/\lambda_0 \rightarrow$ all muons with momentum greater than 5.2 GeV (1.9 GeV) should, on average, make it to fourth muon station in the barrel (endcap)

decay of a W) do not interact with any of the sub-detectors, and thus are "invisible" to CMS; their presence must be inferred through special variables such as missing transverse energy ($\not{E}_{\mathrm{T}}$– c.f. Section 5.7).

The first stage to PF reconstruction is building the individual elements that are later reconstructed into individual candidate particles

5.1.1 The Basic Building Blocks of the PF Algorithm

The PF reconstruction algorithm requires two basic kinds of input ingredients: charged particle tracks and energy clusters. The creation of these basic input ingredients begins with the reconstruction of particle hits ("RecHits") for each subsystem of each subdetector [111]. These RecHits include appropriate information depending upon which subsystem they are created by: tracker (muon) system RecHits include position information of clusters — adjoint strips and pixels (chambers) with recorded signals above threshold — along with the energy deposition information; RecHits in the ECAL and HCAL contain magnitude, position, and timing information for the energy deposits.

In the second step, global reconstruction, these RecHits from the various subsystems are combined and further processed. Charged-particle tracks are created in both the inner tracker and muon systems by using iterative algorithms to link together the RecHits for these subdetectors. For the ECAL and HCAL, seed cells are created from the calorimeter cells in which the recorded energy deposits are local maxima above a certain energy threshold. Clusters are built around these seed cells,

consisting of all neighboring cells above a certain energy threshold. "CaloTowers" are built from these clusters by combining both ECAL and HCAL clusters that match in a projective $\eta - \phi$ geometry.

As part of the second step, rejection algorithms are run on both the charged-particle tracks and CaloTowers to select only tracks and calorimeter deposits likely coming from genuine particles, as opposed to spurious fits of the tracking algorithm or noise signals in the calorimeters.

These calorimeter clusters and charged-particle tracks are then linked together and classified to create PF candidates — candidate "particles" used for higher-level event reconstruction and data analysis.

The PF candidate linking algorithm evaluates possible links between either a track and a cluster, two clusters, or an inner-tracker track and a muon-system track. When linking a charged-particle track and a calorimeter cluster, or an inner-tracker track and a muon-system track, the algorithm accounts for the propagation of charged particles in the magnetic field. When linking an inner-tracker track and a muon-system track, the algorithm attempts a global fit of the whole track, and utilizes the resulting χ^2 as a "distance"; for the other linking cases (e.g. involving clusters), the η and ϕ distances are used. Detector "blocks" are created from the sets of linked elements by minimizing the link distances.

These blocks are then classified as one of five different types of PF candidates. This classification is done iteratively by category so that double-counting is avoided and energy deposits can be properly associated with each particle candidate.

Muons are the first type of PF candidate that is classified. PF muons are

built from the global muon-track fits that have consistent momenta for each section of the track (inner tracker and muon system). The expected deposited energy of this muons (from minimum ionization) are subtracted from the relevant calorimeter clusters.

The remaining blocks with associated charged particle tracks are then tested to see if they're consistent with electrons. This testing accounts for expected energy loss due to bremsstrahlung.

After removing the blocks corresponding to electrons, charged-hadron candidates are built by testing, within each block, that the linked calorimeter cluster's total energy is similar in magnitude but smaller than the charged-particle track. If it is too small, the track is either assumed to come from a low-energy muon, or is classified as a spurious track.

At this stage of classification, the only remaining unclassified blocks are calorimeter clusters. Clusters where most of the energy was deposited in the ECAL are classified as photons. Any remaining clusters are then classified as neutral hadrons.

We will now turn to specific discussions of the reconstruction algorithms used to recreate the various "high-level" objects used in analyses.

5.2 Reconstruction of Charged-particle Tracks and Interaction Vertices

The reconstruction of charged-particle tracks using hits from the tracker subsystems is based on the Combinatorial Track Finder (CTF) algorithm [75], which is

an adaptation of the combinatorial Kalman filter [112].

There is a tremendous amount of combinatorial complexity due to the large number of individual hits in the tracker subsystems. To combat this, the CTF algorithm follows an iterative procedure. It first identifies the tracks that are easiest to find, e.g. with large p_{T} and numerous hits in the pixel system, and then gradually relaxes the requirements from there in order to continue to reconstruct additional tracks.

Each iteration utilizes the same 4-step procedure:

1. Seeds are generated, using very few (2 to 3) hits in order to provide initial track candidates, giving initial estimates of the track trajectory and its associated uncertainties.

2. Additional hits are added into the track candidate via extrapolation along the expected flight path.

3. The overall collection of hits for a given track candidate is then fit with a Kalman filter and smoother; this provides the best estimate of the parameters of the track.

4. Selection criteria are applied to reject track candidates that do not pass the requirements.

The salient differences between each iteration are in the specific configurations of the seeds used to generate the initial track candidates as well as the final track selection criteria.

5.2.1 Seed Choices for Iterations

Initial iterations will use as seeds either pairs or triplets of pixel hits. These seeds contain the majority of tracks produced by prompt particle decays (e.g. W/Z bosons).

A subsequent iteration requires at least 1 pixel hit and < 3 strip hits in its initial seeds. This iteration tends to find slightly displaced tracks from heavy-flavor (c, b) hadron decays, nuclear interactions with tracker material, or photon pair-conversions.

Finally, the last set of iterations use two hits from the strip detectors, which accounts for charged particles that did not enter the pixel detector.

Table 5.1 shows the specific seeding strategies and additional selection requirements used for each tracking iteration. The selection criteria for each iteration are based on the transverse impact parameter d_0, as well as either the absolute value of z_0, the z-position of the track origin relative to the IP or the significance of z_0, determined by the number of σ away from the center of the IP, based on a Gaussian fit.

5.2.2 Vertex Reconstruction

The set of reconstructed tracks are then used to reconstruct the primary interaction vertices in the event [75]. This includes both the main hard-scatter vertex as well as the additional vertices corresponding to pileup collisions.

As a pre-step, the set of reconstructed tracks are run through additional se-

Table 5.1: The sequence of tracking iterations used during the 2012 run. The specifications on the initial seeding are shown in the first three columns. The additional selection requirements applied are shown in the last two columns. Reprinted from Table 3 of [113].

step	seed type	seed subdetectors	p_{T} [GeV]	d_0 [cm]	$\lvert z_0 \rvert$
0	triplet	pixel	>0.6	<0.02	<4.0σ
1	triplet	pixel	>0.2	<0.02	<4.0σ
2	pair	pixel	>0.6	<0.015	<0.09 cm
3	triplet	pixel	>0.3	<1.5	<2.5σ
4	triplet	pixel/TIB/TID/TEC	>0.5 − 0.6	<1.5	<10.0 cm
5	pair	TIB/TID/TEC	>0.6	<2.0	<10.0 cm
6	pair	TOB/TEC	>0.6	<2.0	<30.0 cm

lection requirements. The maximum value of the d_0 significance must be < 5. The number of pixel and strips hits associated with a track must be ≥ 2 and ≥ 5 respectively. Finally, the normalized χ^2 from the Kalman fitting of the track-trajectory must be < 20.

The tracks that pass these requirements are then clustered together into vertices based on a deterministic annealing algorithm. Each track is assigned a probability to be associated with each individual vertex. A "free-energy" is then constructed out of this system, where each vertex χ^2 serves the role of "energy" for that vertex. This free-energy is minimized, in analogue with a macroscopic system that has attained thermodynamic equilibrium.

The initial number of vertices that are fed into the algorithm can be arbitrarily large. Nevertheless, after the deterministic annealing has been performed, many of these vertices will overlap with one another, leading to a much reduced set of "effective vertices".

The set of initial vertices is built from weighted effective vertices. These weighted effective vertices will get split into individual vertices as part of the minimization procedure, leading to a final set of reconstructed vertex candidates located at various z throughout the beam axis.

An adaptive vertex fitter is applied to each of these vertex candidates. This fitter utilizes information from the nominal initial vertex position and a weighted combination of the information for each track associated with this vertex.

After the adaptive vertex fitting, a final set of quality criteria are applied in order to select vertices that are most likely represent true proton-proton interactions [114]. The number of degrees of freedom for a vertex's adaptive fit, $n_{\mathrm{dof}} = -3 + 2\sum w_i$, where w_i is the weight for the ith track in the vertex, must be larger than 4. The vertex must be within $|z| = 24\,\mathrm{mm}$ of the IP center. Finally, the transverse position of the vertex must be less than $2\,\mathrm{mm}$ from the IP center.

Because every vertex contains a unique, non-overlapping set of associated tracks, every single reconstructed charged particle can be assigned a unique vertex. We will discuss this further in the subsequent sections.

5.3 Muon Reconstruction

The track reconstruction principles used at CMS, see Sec. 5.2, can also be applied to the muon-track reconstruction in the muon systems located outside of the solenoid [115]. The tracks that result from this reconstruction are known as "stand-alone muon tracks"

These reconstructed muon-system tracks are combined with reconstructed tracks in the tracker subsystem ("inner-tracker tracks") with either an inside-out or outside-in approach.

"Global muons" are constructed with the outside-in approach. The reconstruction tries to link stand-alone muon tracks with inner-tracker tracks, accounting for the changes in the magnetic field and expected muon energy losses in the intervening material. The stand-alone muon tracks must have at least two associated muon subsystem segment hits. This method has the strongest momentum resolution for high momentum ($p_{\mathrm{T}} \gtrsim 200$ GeV) muons [116, 117]

"Tracker muons" are constructed with the inside-out approach. Reconstructed inner-tracker tracks are extrapolated to the muon system, again accounting for expected changes in the magnetic field and expected muon energy losses in the intervening material. The inner-tracker track is considered a tracker muon track so long as the extrapolated position of the inner-tracker track is consistent with at least one of the segments of the muon subsystems. This method yields stronger momentum resolution for low momentum muons and, for very low momentum muons ($p < 5$ GeV), this method tends to be much more efficient than the global method, because it only requires that the tracker track be matched with a single muon system segment.

A collection of "stand-alone" muons are built from the stand-alone tracks that are not part of either the global muon or tracker muon collections.

5.4 Electron Reconstruction

The majority of the information for this section is paraphrased from Ref. [118].

Electrons at CMS are built from a combination of momentum information that is measured in the tracker and energy information that is measured in the ECAL.

One of the biggest challenges in electron reconstruction at CMS is the often non-trivial amount of bremsstrahlung radiation generated by electrons as they pass through the tracker material. An average of anywhere from 33% - 86% of an electron's initial energy is radiated before the electron reaches the ECAL, depending upon how much detector material the electron travels through [118].

Thus, it is crucial to capture the energy of the bremsstrahlung-generated photons in order to measure the initial energy and momentum of electrons accurately. This energy primarily spreads in the ϕ direction due to the bending of the electrons' trajectories in the magnetic field (both the initial electron and the e^+e^- pairs produced from the bremsstrahlung photon conversions); the spread in the η direction is only notable for low-p_T ($p_T \lesssim 5$ GeV).

Two separate algorithms are used to cluster the electron candidate's deposited energy in the ECAL: a "hybrid" algorithm in the EB, and a "multi-5x5" algorithm in the EE.

5.4.1 Hybrid Electron Reconstruction Algorithm

The hybrid algorithm starts with a seed crystal in the EB, where, in a given geometric region being considered, a seed crystal is defined as the crystal containing

the most deposited energy of all crystals in the geometric region; this seed crystal must have measured E_{T} above a threshold value, $E^{\mathrm{min}}_{\mathrm{T,\,seed}}$ (1 GeV). The algorithm then conditionally adds crystal arrays of $(\eta \times \phi)$ size 5×1 to this seed crystal, where the arrays are only added if they're within N_{steps} (17) in the ϕ direction and if their energies are above a minimum, $E^{\mathrm{min}}_{\mathrm{array}}$ (0.1 GeV). Any contiguous arrays (ignoring the initial seed crystal) are grouped into clusters, where one of the arrays in a given cluster must have energy greater than a threshold $E^{\mathrm{min}}_{\mathrm{seed\text{-}array}}$ (0.35 GeV). The clusters that pass these requirements, along with the initial seed crystal, are then grouped together into a final cluster, known as a supercluster (SC).

5.4.2 Multi-5×5 Electron Reconstruction Algorithm

The multi-5×5 algorithm, like the hybrid algorithm, starts with a seed crystal. This seed crystal is defined as the crystal with the most energy relative to its four direct neighbors (as a reminder, in the EE, the crystals are not arranged in a projective $\eta \times \phi$ array geometry). As with the hybrid algorithm, the seed crystals must have measured E_{T} bigger than a threshold value (0.18 GeV). Using these seeds, sorted by E_{T}, as reference points, the multi-5×5 algorithm builds clusters of 5×5 crystals (these clusters can partially overlap with one another). All clusters within a fixed η range (±0.07) and fixed ϕ (±0.3 rad) range of one another are then grouped together into a SC if their total E_{T} is greater than a threshold, $E^{\mathrm{min}}_{\mathrm{T,\,cluster}}$ (1 GeV). In order to account for energy deposited into the preshower, for a given SC, the energy-weighted positions of its constituent clusters are extrapolated to planes of

the preshower. For a given SC, the clustering range for cells in the pre-shower is equal to the total range of the SC in $\eta - \phi$, but extended by ± 0.15 in both coordinate directions. These energies are then added into the SC energy.

The values listed in parentheses for the various parameters of both algorithms were tuned to optimize the ECAL-energy resolution over a wide range of electron p_{T} values [118].

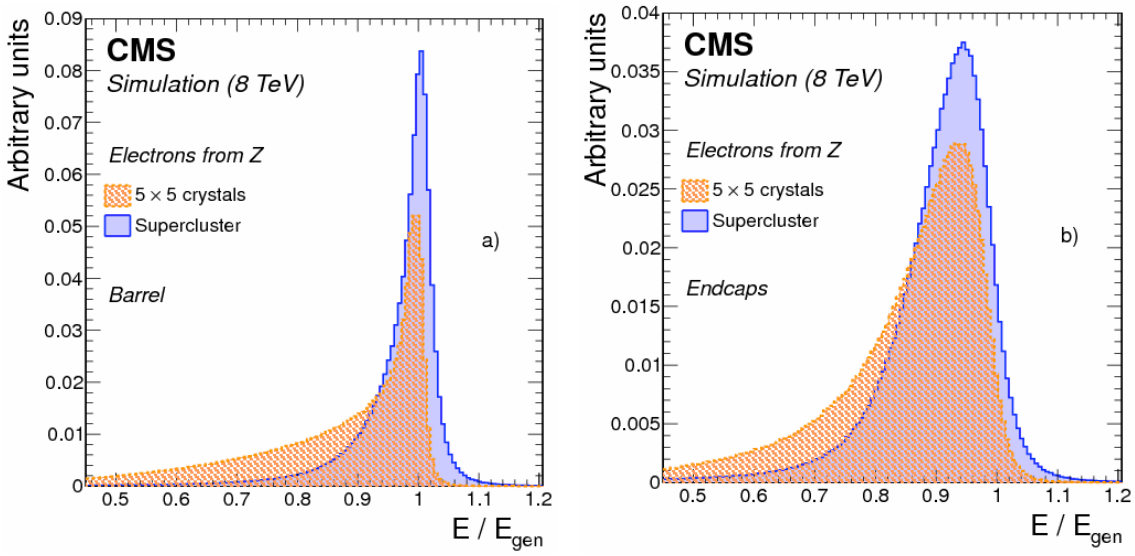

(a) (b)

Figure 5.2: Distributions of the ratio of reconstructed over generated electron energies for electrons coming from Z-boson decays in a) the ECAL barrel and b) the ECAL endcap. The yellow, dashed histograms show the distribution for electrons reconstructed using 5 x 5 crystal arrays, while the solid blue histogram shows the distribution for electrons reconstructed using superclustering algorithms, see Section 5.4. Reprinted from Fig. 3 of [118].

Although these algorithms are somewhat detailed, they bring notable gains in reconstruction performance. Figure 5.2 shows the distributions of the ratio of reconstructed over generated electron energies for electrons coming from Z-boson decays in the barrel and the endcap. The yellow, dashed histograms show the distribution for electrons reconstructed using 5 x 5 crystals, while the solid blue

histogram shows the distribution for electrons reconstructed using the previously mentioned SC-based algorithms. As can clearly be seen, the long energy-loss tail is significantly reduced by performing superclustering. Although a slight amount of spurious additional energy can be added into the reconstructed SC (the tail extending to the right of 1), this is a very small effect.

5.4.3 Linking the ECAL Clusters with Tracks: The Gaussian-Sum Filter Algorithm

The standard Kalman-filter track reconstruction algorithm, see Sec. 5.2, does not perform well with electron tracks. This is primarily due to the energy losses from bremsstrahlung radiation. The Kalman filter is designed to account for Gaussian energy losses, and bremsstrahlung radiation leads to very non-Gaussian energy losses.

A dedicated track-reconstruction procedure is used to reconstruct electron tracks. It is significantly more time-consuming than the normal procedure. To help combat this, two different, dedicated seeding algorithms provide a reduced set of initial track candidates that are very likely to be coming from electrons.

One of these algorithms, the ECAL-based seeding algorithm, extrapolates backwards from SCs in the ECAL to match one of the collision vertices. The other, the tracker-based seeding algorithm, attempts to match, through extrapolation, the tracks reconstructed during the normal track reconstruction with SCs in the ECAL. These tracks are often poorly reconstructed due to possible bremsstrahlung radia-

tion. Thus, a dedicated multivariate analysis tool utilizes quality variables associated with these tracks, included energy information from their matched ECAL SCs, to best identify the initial tracker seed. The two algorithms, working in concert with one another, can achieve a seeding efficiency of 95% [118].

From these seeds, a variant of the Kalman-filter reconstruction is applied in order to reconstruct the associated electron tracks. The electron energy loss due to bremsstrahlung is modeled by a Bethe-Heitler function. The standard tracking requirements, in particular the requirements on predicted versus found hits in each tracker layer, are relaxed to increase the efficiency of the track reconstruction.

Once the set of hits associated with an electron track have been collected, a Gaussian-sum Filter fit is used to estimate the global track parameters. The energy loss in each layer is approximated by a mixture of Gaussian distributions. The overall mode of this sum distribution is used to characterize all the associated electron track parameters.

5.5 Photon Reconstruction

The reconstruction of photons at CMS is somewhat similar to the reconstruction of electrons, c.f. Section 5.4 or Refs. [84, 118]. The same SC-based approach used to reconstruct the energy deposits of electrons in the ECAL can also be applied to reconstruct the deposits of photons. Anywhere from $\sim$ 20–60% (depending upon the photon η) of photons produced at collisions in CMS convert to electron pairs before reaching the ECAL. Just as the electron reconstruction algorithms have to ac-

count for bremsstrahlung radiation, photon reconstruction algorithms must account for these electron-pair conversions.

The R_9 variable, the energy sum of the 3×3 matrix of ECAL crystals centered on the most energetic crystal in the SC divided by the total SC energy, is a strong indicator of whether a given photon converted before reaching the ECAL. For EB (EE) photons with $R_9 < 0.94\,(< 0.95)$, the simplest estimation of the candidate photon energy is the SC energy. For photons with larger R_9 values close to unity, the total energy of the 5×5 crystal matrix centered on the most energetic crystal in the SC provides the simplest energy estimation.

Dedicated calibration algorithms based on, among other variables, the photon candidate p_{T}, η, and R_9 values, provide accurate recalibrations of the reconstructed photon energies [84]. These algorithms, along with more general photon ID algorithms, were tuned to maximize their performance in $\mathrm{H} \to \gamma\gamma$ events, but their performance is still excellent for other classes of photon events [84].

5.6 Jet Reconstruction

As noted in Section 2.2.1, every single (high-)momentum quark or gluon, except for top quarks, produced in pp collisions at the LHC eventually becomes part of a (highly-)collimated spray of color-neutral hadrons known as a hadronic jet.

By properly clustering the individual hadrons and reconstructing these jets, we are able to get the only available visible information corresponding to the original parton(s) kinematics.

There has been a tremendous amount of theoretical and experimental work across multiple decades devoted to developing jet clustering algorithms [5]. Naive algorithms based more or less on the construction and clustering of objects falling within simple cones in $\eta - \phi$ space suffer from collinear[2] and infrared[3] instabilities. These naive algorithms, are thus unsuitable for perturbative QCD calculations.

CMS uses the infrared-safe, collinear-safe anti-k_{T} algorithm [119] as its primary jet clustering algorithm. Jet substructure studies often utilize the Cambridge-Aachen (CA) algorithm for clustering sub-jets. Furthermore, jet corrections, as part of the removal of contributions from pileup interactions, utilize the k_{T} algorithm.

The k_{T}, anti-k_{T} and CA algorithms are part of a more general family of iterative-cone jet clustering algorithms, based on the following equation,

$$\begin{aligned} d_{ij} &= \min(p_{\mathrm{T}i}^{2p}, p_{\mathrm{T}j}^{2p}) \frac{\Delta R_{ij}^2}{R^2}, \\ \Delta R_{ij}^2 &= (\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2, \\ d_{iB} &= p_{\mathrm{T}i}^{2p}. \end{aligned} \tag{5.1}$$

The above equation contains two distance variables: the distance between particles i and j, d_{ij}, and the distance between particle i and the beam, d_{iB}. These distances are weighted by the p_{T} of the particles as indicated, with R serving as a size parameter. Every iteration of the clustering algorithm involves calculating these distance variables for all remaining particles in the event. These distance variables are then sorted. If the minimum distance is d_{iB}, then particle i is marked as a fully-clustered jet and is removed from the list of particles. Otherwise, one of the d_{ij} must be the

[2]Splitting a single hard (high p_{T}) particle into two softer ones can change the final clusterings.
[3]The addition of soft seed particles, e.g. soft gluon emission, can change the final clusterings.

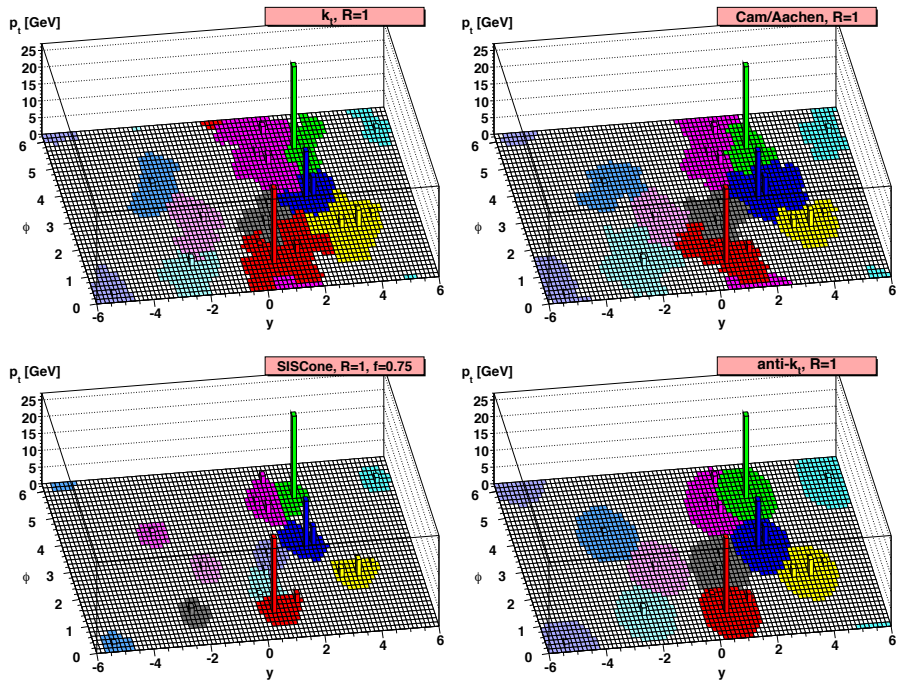

Figure 5.3: A sample event clustered with four different algorithms (algorithms described in [5]). Numerous soft, ghost particles have been added; all four algorithms demonstrate collinear and infrared stability, but the jet "area" varies notably for each. In particular, for the k_{T} and CA algorithms, the detailed jet shapes strongly depend upon the distribution and magnitude of the ghosts. Reprinted from Fig. 7 of [5].

minimum distance. These two particles (i and j) are grouped together into a single particle. The algorithm continues to iterate until there are no remaining particles to cluster.

Figure 5.3 shows the clustered jets that result from applying Eq. (5.1) for the k_{T}, anti-k_{T}, and CA algorithms, as well as a simple cone-based algorithm (which we will not discuss further here).

The anti-k_{T} algorithm uses Eq. (5.1) with $p = -1$. This leads to the algorithm favoring clustering high-p_{T} particles first. As can be seen from Fig. 5.3, this algorithm tends to lead to relatively circular jets whose exact shapes are relatively independent of additional ghost (i.e. soft) particles. This significantly aids in the calibration of jet energy responses in the detectors. For jet substructure calculations, this algorithm's performance suffers because it tends to preferentially cluster elements from the individual sub jets (e.g. the two quarks in a hadronic W or Z decay) together rather than keeping them separate.

On the other hand, the CA algorithm uses $p = 0$ and thus only clusters based on $\eta - \phi$ distances, with no preferences based on particle p_{T}. This enables it to maintain proper separation between the individual sub-jet elements in jet substructure calculations.

Finally, the k_{T} algorithm uses $p = 1$, causing it to favor clustering low-momentum particles together. This makes it a relatively robust algorithm for properly clustering jets coming from pileup interactions, which tend to be very wide and relatively non-collinear.

During the 2012 run, CMS used distance parameter values of $R = 0.5$, $R = 0.8$,

and $R = 0.6$ for the anti-k_{T}, CA, and k_{T} algorithms respectively. In Run 2 and onwards, it is planned that CMS will use $R = 0.4$ for the anti-k_{T} algorithm to help mitigate the adverse effects of additional pileup.

Finally, we note that these algorithms are typically run using input PF candidates to seed the initial particle list. The resulting jets produced in this fashion are known as PF jets.

Alternatively, tracker-jets or Calo-jets can be constructed using either tracks or CaloTowers as the input particle seeds. The performance of these jets suffers notably in relation to PF jets, however. [106, 108].

The relative gains in performance for PF jets can be traced back to the relative fraction of jet constituents. Typically, 60% of jet energy goes into charged particles, 25% goes into photons, and 15% goes into neutral hadrons [120].

Track-based jets have strong energy resolution, due to the excellent momentum resolution of the tracker, but suffer from a relatively lower energy response due to missing particles. Calo-jets, in principle, are built from all visible particles, and thus have relatively acceptable energy response, but suffer from the relatively poor energy resolution of the hadronic calorimeters. Using PF candidates ensures that all visible particles are input into the jet clustering with the optimal energy measurements.

5.6.1 Jet Corrections

The calibration of the resolution and response for jet energies is one of the more important aspects of CMS event reconstruction.

5.6.1.1 Pileup Corrections

Pileup interactions tend to lead a relatively constant, diffuse amount of energy spread approximately evenly throughout the detector. A large fraction ($\sim 65\%$) of this energy is contained in charged hadrons, while the remaining fraction is contained in neutral particles [120].

The undesired contributions of these charged hadrons to the jet clustering can be removed through a process known as *charged-hadron subtraction* (CHS). Prior to the jet-clustering, charged hadrons that are associated with a interaction vertex that is not the hard-scatter vertex are removed from the list of input particle candidates.

The contributions from the neutral-pileup hadrons, as well as any charged-pileup hadrons that were improperly assigned to the hard scatter vertex, are removed via an offset correction, known as the L1 correction.

The L1 Jet Correction

The overall average energy deposited by the sum total of all pileup interactions (ignoring the charged hadrons already removed via CHS) is estimated by clustering all the PF candidates using the k_{T} algorithm, including a large number of random *ghost* particles. The variable $\rho = \mathrm{median}(p_{\mathrm{T,i}}/A_i)$ is then calculated from the jets created by this clustering. The use of the median makes ρ insensitive to both hard scatter jets as well as jets composed entirely of ghost particles.

The jets j that resulted from the anti-k_{T} clustering are then corrected by removing an amount of energy from the jet equal to $\rho * A_j$. This correction, on

average, removes the contribution of pileup from each jet, leading to better overall jet energy response and slightly improved jet energy resolution.

The L2 Jet Correction

After the L1 jet correction is applied, the jet energy response is made uniform in η via correction factors derived from simulated dijet events where two hadronic jets balance each other in the transverse plane.

The L3 Jet Correction

The final correction applied to both simulation and collision data events is the multiplicative absolute correction. This correction is designed to make the jet response uniform in p_{T}. The correction factors are derived from $\gamma +$ jets events, as these events provide an extremely well-measured probe object (the γ) that facilitates the direct calibration of the recoiling jet.

The Residual L3 Jet Correction

All of the above corrections are derived from simulated events in order to both minimize statistical uncertainties for relatively rare classes of events as well as to avoid systematic biases from data-based calibration methods.

Nevertheless, correcting only from simulation introduces residual systematic biases of its own due to differences between data and the simulation. Additional corrections serve to account for these residual differences. QCD dijet events are used

to correct for remaining residual differences in the η distributions of calibrated jets, while γ + jets, $Z \to \ell^+\ell^-$ + jet, and QCD multijet events are used to correct for residual differences in the p_T distributions of calibrated jets.

Results of the JEC

The end result of the JEC is a near-unity jet-energy response regardless of jet p_T, jet η, and pileup conditions. This can be seen in Fig. 5.4c, where the "MC truth ratio", the measured jet p_T divided by the particle-level jet p_T, is shown versus particle-level jet-p_T. The different colored points represent sub-samples of events with different ranges of μ, the average number of pileup interactions per bunch crossing. All JECs have been applied to arrive at the MC truth ratio in this figure. This can be compared against Fig. 5.4a, where no JECs have been applied, and Fig. 5.4b, where the pileup offset corrections have been applied.

5.6.1.2 Jet Energy Resolution

The energy resolution of jets (JER) at CMS is much poorer in comparison to other physics objects. This is partly due to the nature of hadronic showers, which tend to be highly stochastic, non-linear processes. As well, the fact that the CMS HCAL is primarily a sampling calorimeter also adversely affects the JER at CMS.

Accurate modeling of JER is particularly important for low-p_T jet physics and thus, $\not{E}_T$ reconstruction. In order to understand and quantify the performance of this modeling, CMS has executed a number of detailed studies [120].

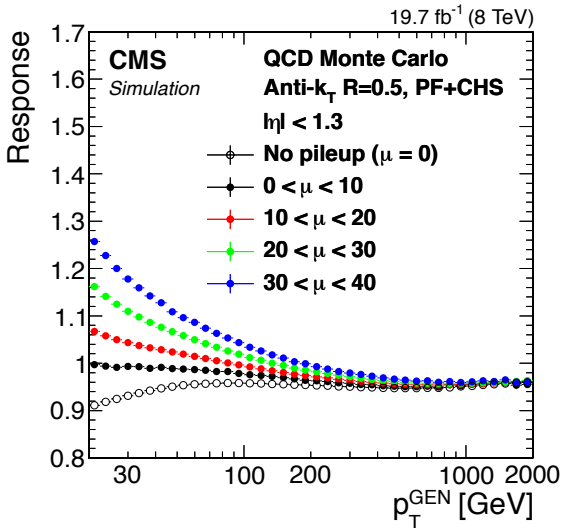

(a) The MC truth ratio before any JEC corrections have been applied.

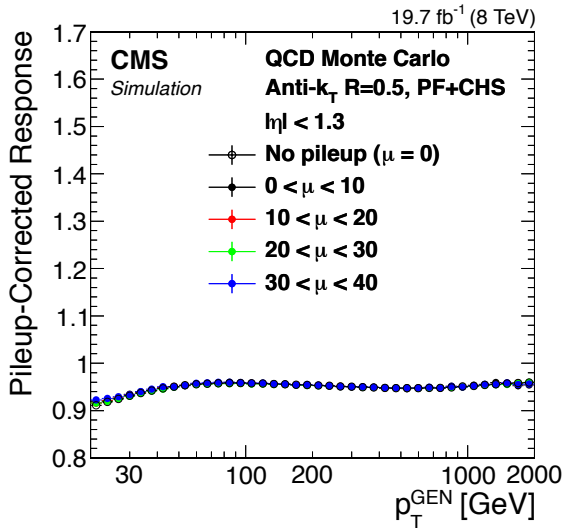

(b) The MC truth ratio after the application of pileup offset corrections.

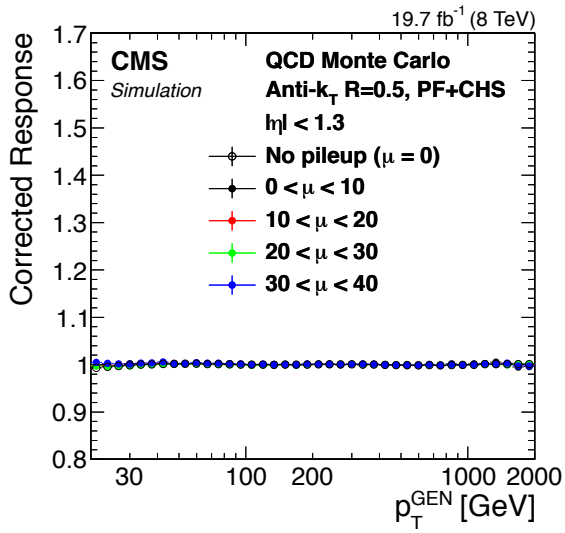

(c) The MC truth ratio after the application of all JECs.

Figure 5.4: The "MC truth ratio", the measured jet p_{T} divided by the particle-level jet p_{T}, in QCD multijet simulation events at various levels of the jet-energy calibrations. The different colored points represent the sub-samples with different ranges of μ, the average number of pileup interactions per bunch crossing. Reprinted from Fig. 1 of [120].

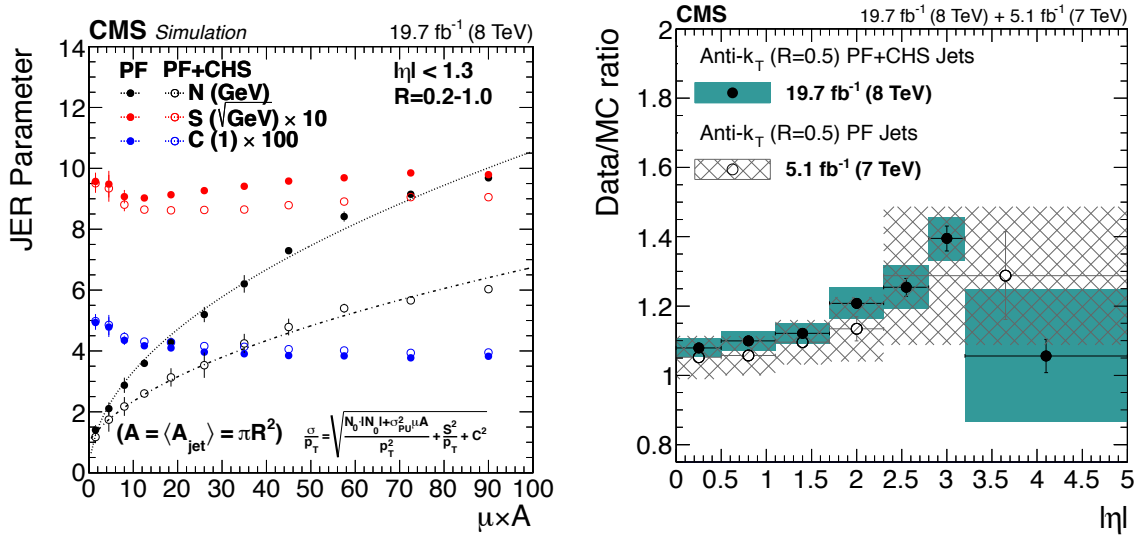

(a) Measured parameters of the JER, Eq. (5.3), fit in bins of μ, the number of pileup interactions, and then averaged over $\mu \times A$, where A is the jet area. Comparisons are made between jets constructed using the PF algorithm with and without charged hadron subtraction (CHS). Reprinted from Fig. 35 of [120].

(b) The data/simulation scale factor for the JER, based on measurements of the Gaussian core of a dijet p_{T} asymmetry distribution, Eq. (5.2). Reprinted from Fig. 39 of [120].

Figure 5.5: Two CMS performance plots related to JER.

One of the general ways of quantifying the JER is to measure the asymmetry distribution in QCD dijet events, where the dijet asymmetry is defined by,

$$\mathcal{A} = \frac{p_{\mathrm{T,\ 1st\ jet}} - p_{\mathrm{T,\ 2nd\ jet}}}{p_{\mathrm{T,\ 1st\ jet}} + p_{\mathrm{T,\ 2nd\ jet}}}. \tag{5.2}$$

The core of the distribution of $\mathcal{A}$ is then fit with a Gaussian and the asymmetry resolution, $\sigma_{\mathcal{A}}$ is extracted.

The relative dependence of this resolution on the jet p_{T} can be parameterized in terms of three quantities,

$$\frac{\sigma_{p_{\mathrm{T}}}}{p_{\mathrm{T}}} = \sqrt{\frac{\mathrm{sign}(N)N^2}{p_{\mathrm{T}}^2} + \frac{S^2}{p_{\mathrm{T}}} + C^2}, \tag{5.3}$$

where N is for noise and pileup contributions, S accounts for stochastic fluctuations in the jet showers, and C represents constant terms that do not scale with p_{T}, such as interdetector calibrations.

Figure 5.5a shows the dependence of the three jet-p_{T} resolution parameters on the number of pileup interactions, μ. The parameters were measured in a QCD multijet simulation for relatively central jets ($|\eta| < 1.3$). The two sets of points for each parameter represent the measured values for the cases where the PF jets are constructed with or without CHS.

It has been known for some time (early 2011) that the JER is underestimated by the simulation. Comparisons between data and simulation of $\sigma_{\mathcal{A}}$ can be used to derive scale factors to correct the modeling of the JER in the simulation. These scale factors are empirically found to have little p_{T} dependence, but are η-dependent.

Figure 5.5b shows the measured values of these scale factors for the full 7 TeV and 8 TeV CMS collision datasets. The 8 TeV scale factors only recently became

available; thus, the majority of Run 1 CMS analyses, including the analyses described in Chapters 6 and 7, used the 7 TeV JER scale factors, which have notably larger uncertainties.

5.6.2 Tagging Bottom Quark Jets

The information in this section primarily is paraphrased from Refs. [121, 122].

Bottom quarks are interesting objects from a particle physics standpoint. They are associated with many interesting and important physical processes, most notably the decay of heavy SM particles like the top quark and Higgs boson, as well as their expected production in many BSM theories such as SUSY.

In addition, during the hadronization process, bottom quarks form hadrons with relatively long lifetimes and relatively large masses, which subsequently decay to high-momentum daughter particles. These properties facilitate the efficient and accurate identification of the jets produced during the bottom quark hadronization (b-jets as they're colloquially known).

CMS has a number of dedicated algorithms that utilize these discriminating features of b-jets in order to identify them in the reconstructed data. All of these algorithms utilize reconstructed hadronic jets; as well, these algorithms capitalize on additional information provided by reconstructed tracks, vertices, and identified leptons. Consequently, b-tagging occurs at the later stages of event reconstruction, after the aforementioned input objects have been reconstructed and calibrated.

In order to improve the purity of the selection of input tracks, additional

requirements are placed beyond those listed in Section 5.2:

- Track $p_{\mathrm{T}} > 1$ GeV;
- The track must have an angular distance, $\Delta R < 0.3$, from its associated jet;
- The track must be within 700 μm of the jet axis at its point of closest approach to the jet axis (axis defined relative to the PV);
- The normalized χ^2 of the track reconstruction must be less than 5;
- The decay length of the track must be less than 5 cm, where the decay length is the distance from the PV to the track at its point of closest approach to the jet axis;
- The number of associated hits in the tracker system (pixel alone) for the track must be ≥ 8 (≥ 2);
- The distance in the transverse plane (along z) between the PV and the track at its point of closest approach to the PV must be less than 0.2 cm (17 cm).

Quantities Used in b-tagging: 3D Impact Parameter Significance

A useful quantity for b-tagging is the 3D impact parameter (3D IP), the distance from the PV to the track its point of closest approach to the PV. In particular, the 3D IP significance, $\mathcal{S}_{\mathrm{3D\ IP}}$, the ratio of the 3D IP to its estimated uncertainty, is the basic discriminating variable used by the *Track Counting* (TC), *Jet Probability* (JP), and *Jet B Probability* (JPB) algorithms.

The TC algorithm sorts the tracks in a jet by their $\mathcal{S}_{\text{3D IP}}$ values and uses the $\mathcal{S}_{\text{3D IP}}$ of the second and third ranked tracks as the discriminating value. The JP and JPB algorithms combine the 3D IP information of all tracks in the jet, although the JPB algorithm assigns additional weight to the tracks with the highest $\mathcal{S}_{\text{3D IP}}$.

We do not use the aforementioned b-tagging algorithms in this dissertation, and thus will not discuss them further from this point. Additional, detailed descriptions of them, however, can be found in Ref. [121].

Quantities Used in b-tagging: Secondary Vertices

Additional vertices located within the projected cone of a reconstructed jet, a.k.a. secondary vertices (SV), and their associated kinematic variables are another useful tool for b-jet discrimination.

In the CMS event reconstruction, these secondary vertices are constructed using techniques similar to the primary vertex construction techniques described in Section 5.2. For the purposes of b-tagging, the SVs must meet the following additional requirements

- SVs must share less than 65% of their associated tracks with the PV;
- The significance (defined similarly as $\mathcal{S}_{\text{3D IP}}$ above) of the radial distance between the SV and PV must exceed 3σ;
- In order to reduce the contamination from both long-lived mesons such as K^0s as well as interactions of particles with the detector material, SV candidates that have a radial distance (from the PV) of more than 2.5 cm or SV candidates

that have an invariant mass either compatible with M_{K^0} or exceeding 6.5 GeV are rejected;

- The SV candidate and its associated jet must satisfy $\Delta R < 0.5$.

The *Simple Secondary Vertex* (SSV) algorithm uses the significance of the flight distance as a discriminating variable. However, this algorithms efficiency is limited by the SV reconstruction efficiency, 65%, and consequently, it is not commonly used at CMS.

A more complicated b-tagging algorithm that utilizes the available SV information is the *Combined Secondary Vertex* (CSV) algorithm. The CSV algorithm combines the SV information with track-based lifetime information and other discriminating variables in order to recover tagging efficiency even when there are no "real" SV candidates in a given jet. It primarily accomplishes this by reconstructing "pseudo-vertices", effective SV candidates built from tracks with $\mathcal{S}_{\mathrm{3D\ IP}} > 2$, although it still can provide b-jet discrimination when no additional vertices, real or pseudo, are present in a jet.

To be more specific, the CSV algorithm takes as input the following set of variables, chosen for their high discriminating power and relatively low inter-correlations (note that if the SV is non-existent, i.e. neither a real nor pseudo-SV candidates exists in the jet, only the last two variables in this list are used):

- The SV vertex category — real, pseudo, or non-existent;

- The SV candidate's flight distance significance in the transverse plane ("2D");

- The SV candidate's invariant mass;
- The SV candidate's track multiplicity;
- The ratio of the energy carried by the SV candidate's tracks relative to all tracks associated with the jet containing the SV candidate;
- The η values for the SV candidate's tracks relative to its containing jet's axis;
- The 2D IP significance of the first track in the SV that raises the SV invariant mass above the charm quark threshold (1.5 GeV), where the SV invariant mass is calculated iteratively by adding in the SV candidate's individual tracks, and the tracks have been sorted by 2D IP significance in decreasing order;
- The total number of track in the containing jet;
- The 3D IP significance of each track in the containing jet.

Using the full list of aforementioned discriminating variables, the CSV algorithm calculates two likelihood ratios. The first of these ratios discriminates between b- and c-jets, while the second preferentially selects jets coming from heavy-flavor quark hadronization (c and b quarks) relative to those coming from light-parton hadronization (u, d, s, and gluons). A weighted combination of the two likelihood ratios (where the two respective weights are 0.25 and 0.75) is then combined into an overall discriminator. This discriminator takes values between 0 and 1, with values close to 1 signifying that the jet is likely coming from the hadronization of a heavy-flavor quark (c- or b-jet).

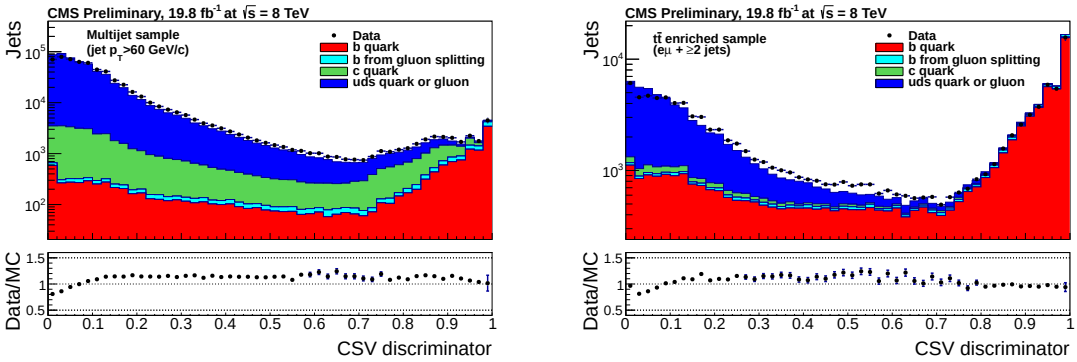

(a) The output CSV discriminator value in a sample dominated by QCD multijet events.

(b) The output CSV discriminator value in a sample dominated by $t\bar{t}$ events.

Figure 5.6: The distributions, for two separate event samples, of the output CSV discriminator value. See Ref. [122] for exact definitions of these event samples. Reprinted from Fig. 6 of [122].

Figure 5.6 shows the distributions of CSV discriminator values in a sample of events dominated by QCD multijet events, and a sample of events dominated by $t\bar{t}$ events[4]. The power of the CSV discriminator to effectively select c- and b-jets is quite clear. In this dissertation, we utilize the "medium" working point of the CSV algorithm, where we tag a given jet as a b-jet if its CSV discriminator value is 0.679 or higher. At this working point, the light-flavor mistag rate (i.e. the probability to incorrectly identify light-flavor jets as c- or b-jets) is $\sim 1\%$ and the heavy-flavor tagging efficiency is $\sim 70\%$ for jets with transverse momentum 80 GeV $< p_{\mathrm{T}} <$ 120 GeV.

5.6.2.1 Correcting the Modeling of the CSV Discriminator

Although it has strong discriminating power between light- and heavy-flavor hadronic jets, the CSV discriminator is not modeled perfectly by the CMS simula-

[4]See Ref. [122] for exact definitions of these event samples.

tion, as can be seen by the ratios comparing the data and simulation in Fig. 5.6. Thus, in addition to evaluating the performance of the CSV discriminator, the authors of Ref. [122] also derived scale factors to correct the mismodeling in the simulation of both the light-flavor mistag rate and the heavy-flavor tagging efficiency. Systematic uncertainties on these scale factors were also derived; these systematic uncertainties account for possible systematic biases in the modeling of relevant physics processes, including the impact of pileup interactions, $\mathrm{b\bar{b}}$ production by gluons, and the fragmentation of b-quarks.

These scale factors, and associated systematic uncertainties, (primarily) depend upon both the jet p_{T} and jet η. The average scale factor to correct the light flavor mis-tag rate (heavy-flavor tagging efficiency) is ~ 1.15 (~ 0.95), with a typical total systematic uncertainty on the scale factor of ~ 0.10 (~ 0.02).

Figure 5.7 shows the b-tagging efficiency dependence upon jet p_{T} and η, when using the medium working point of the CSV b-tagging algorithm, in a sample of events with $\mathrm{t\bar{t}}$-like topologies. Included in this figure are scale factors calculated from comparisons between data and simulation as well as associated systematic uncertainties on these scale factors. Figure 5.8 shows the analogous results as Fig. 5.7a, but for the light-flavor mistag rate.

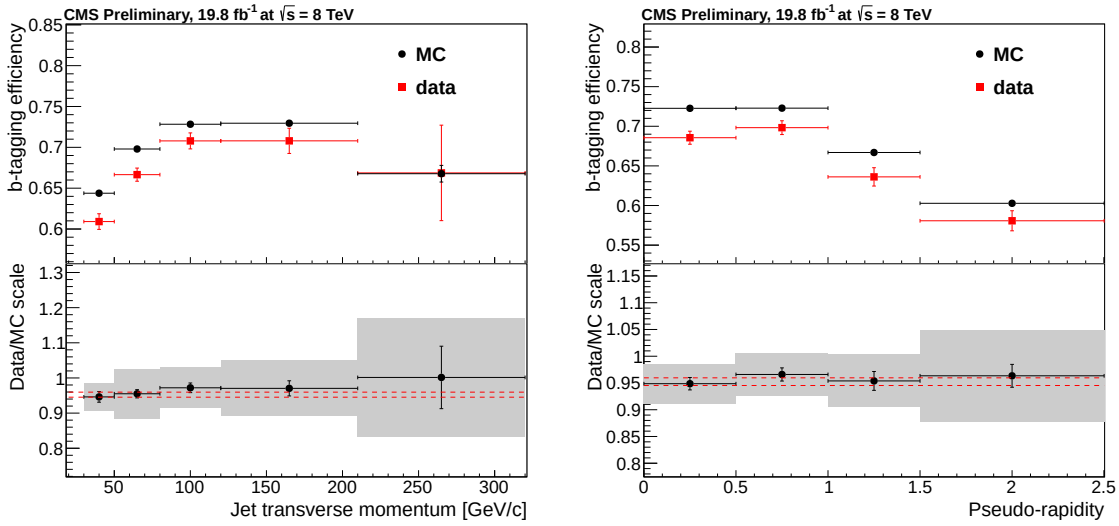

(a) The dependence on the b-tagging efficiency on jet p_{T}.

(b) The dependence on the b-tagging efficiency on jet η.

Figure 5.7: The b-tagging efficiency dependence upon jet p_{T} and η, when using the medium working point of the CSV b-tagging algorithm, in a sample of events with $\mathrm{t\bar{t}}$-like topologies. Scale factors, designed to correct the modeling of these quantities in simulation, are calculated from direct comparisons of the performance in data and simulation and are shown in the bottom panels of each sub-figure. The error bars in the top panels of each sub-figure are statistical only, while the gray filled areas in the bottom panels represent the combined statistical and systematic uncertainties. Reprinted from Fig. 15 of [122].

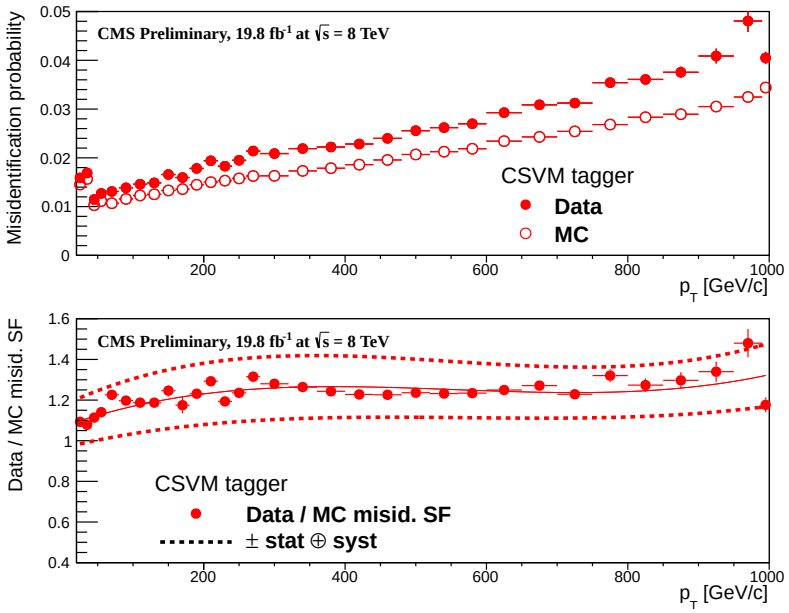

Figure 5.8: (top panel): The light-flavor mistag rate dependence on jet p_{T} for the medium working point of the CSV b-tagger. (bottom panel): The scale-factor for the light-flavor mistag rate calculated from comparisons of data and simulation in the top panel. The dashed lines represent the combined statistical and systematic uncertainties. Reprinted from Fig. 9 of [122].

5.7 Reconstructing the Missing Transverse Energy, $\not{E}_{\mathrm{T}}$

The main part of my PhD disseration work consisted of studies related to $\not{E}_{\mathrm{T}}$, both in the context of characterizing and optimizing the performance of the $\not{E}_{\mathrm{T}}$ reconstruction[5], and utilizing $\not{E}_{\mathrm{T}}$ as part of a search for signatures of R-parity conserving top-squark pair-production, Chapters 6 and 7. As such, I will devote a notable amount of additional discussion to the subject of $\not{E}_{\mathrm{T}}$ reconstruction.

Missing transverse energy, commonly denoted with the symbol $\not{E}_{\mathrm{T}}$ or the acronym MET, is used to estimate the missing energy (and momentum) carried away by undetected particles at particle physics experiments. Because there is a conservation of net momentum in the plane transverse to the beam axis, the transverse momenta of particles created and/or scattered by the primary hard interaction must balance[6]. When a particle is not detected, this balance of momentum in the transverse plane naturally does not occur. At the CMS detector, the only undetectable SM particles are the neutrino partners for the three charged leptons – the electron, muon, and tau neutrinos.

Many analyses utilize $\not{E}_{\mathrm{T}}$ as a relevant variable. For example, R-parity conserving SUSY, see Section 3.2, along with many other models for new physics beyond the SM, predict new, weakly-interacting particles that are stable on the distance scale of the CMS detector (or longer) and thus are effectively invisible. The search for signatures of top-squark pair-production described in Chapters 6 and 7 capi-

[5] I contributed the $\gamma + \text{jets}$ results to CMS's official $\not{E}_{\mathrm{T}}$ performance study paper [71].

[6] Momentum is also conserved along the beam axis. However, the center of mass frame and the lab frame are not necessarily equivalent due to the differing parton momenta. Furthermore, there is of course the technical difficulty of detecting particles that scatter down the beam pipe.

talizes on this fact and utilizes a variable derived from $\not{E}_{\mathrm{T}}$ as a means to separate the top-squark signal from the main $\mathrm{t\bar{t}}$ background. The $\not{E}_{\mathrm{T}}$ variable also played a crucial role in key analyses contributing to the discovery of the Higgs boson, in particular in the analyses examining the $\mathrm{WW} \to \ell\ell\nu\nu$, $\mathrm{ZZ} \to \ell\ell\nu\nu$, where ℓ is e or μ, and $\mathrm{H} \to \tau\tau$ final states [123]. Finally, the precise measurement of $\vec{\not{E}}_{\mathrm{T}}$ is a critical component of precision SM physics involving W bosons and top quarks.

The standard method to calculate $\not{E}_{\mathrm{T}}$ in an event is to compute the negative vector sum of visible transverse momenta of objects in the event,

$$\vec{\not{E}}_{\mathrm{T}} = -\sum_{i \in \mathrm{vis.}} \vec{p}_{\mathrm{T},\,i}. \tag{5.4}$$

The challenging aspect of $\not{E}_{\mathrm{T}}$ reconstruction comes in determining the correct momenta to use in Eq. (5.4). As of the writing of this dissertation, the canonical choice at CMS is to utilize the PF candidates reconstructed by the PF algorithm after corrections and calibrations have been applied as necessary.

5.7.1 $\not{E}_{\mathrm{T}}$ Corrections

Nonzero $\not{E}_{\mathrm{T}}$ values can be measured even in events that, nominally, should have no inherent $\not{E}_{\mathrm{T}}$ (i.e. no neutrinos or other weakly interacting particles). For example, the p_{T} thresholds in the tracker, minimum energy thresholds in the calorimeters, and the calorimeter system's non-linear response to hadronic particles can induce biases in the measurement of individual particle momenta. Most, if not all, of these biases are removed by dedicated energy calibrations for individual particles. Even after re-calibration, however, imperfect object energy resolutions lead to

stochastically-driven reconstructed momentum imbalances in events.

More generally, detector and collider conditions also affect the $\not{E}_{\mathrm{T}}$ performance. Malfunctioning detector elements (hot or cold detector cells, biased readout devices, etc.) can easily induce momentum imbalance and subsequent $\not{E}_{\mathrm{T}}$ into events. As well, visible particles that travel outside of the fiducial geometry of the detectors are not included in the sum in Eq. (5.4), again leading to a (spurious) measured imbalance in the event.

The additional (pileup) interactions per beam crossing that are a hallmark of high-luminosity hadron colliders also play a significant role in $\not{E}_{\mathrm{T}}$ measurements. Just as with the primary hard-scatter, the transverse momentum is balanced in these additional pileup interactions. Furthermore, these pileup interactions tend to contain small amounts, if any, of prompt high-p_{T} neutrino production and subsequently contain little to no true $\not{E}_{\mathrm{T}}$. Nevertheless, due to the aforementioned imperfect detector energy resolutions, every individual pileup interaction adds more energy into the event that can subsequently be mismeasured.

5.7.1.1 Type-0 $\not{E}_{\mathrm{T}}$ Correction

As noted above, pileup interactions are expected to contribute little to no genuine $\vec{\not{E}}_{\mathrm{T}}$ to events. This means that the vector sum of the $\vec{p}_{\mathrm{T}}$ of charged particles from pileup is expected to balance the analogous sum for neutral particles coming pileup,

$$\sum_{i \in \mathrm{neuPU}} \vec{p}_{\mathrm{T}i}^{\,\mathrm{true}} + \sum_{i \in \mathrm{chPU}} \vec{p}_{\mathrm{T}i}^{\,\mathrm{true}} = 0. \tag{5.5}$$

The (typically low-p_{T}) charged particles from pileup are well-measured by the tracker detector,

$$\sum_{i \in \mathrm{chPU}} \vec{p}_{\mathrm{T}i} = \sum_{i \in \mathrm{chPU}} \vec{p}_{\mathrm{T}i}^{\,\mathrm{true}}. \tag{5.6}$$

For the neutral pileup particles, however, nonlinearities in the calorimeters as well as minimum energy thresholds can significantly bias the energy measurements,

$$\sum_{i \in \mathrm{neuPU}} \vec{p}_{\mathrm{T}i} = R^0 \sum_{i \in \mathrm{neuPU}} \vec{p}_{\mathrm{T}i}^{\,\mathrm{true}}. \tag{5.7}$$

In Eq. (5.7), $R^0 < 1$ on average, causing the $\vec{\not{E}}_{\mathrm{T}}$ from each pileup interaction to point on average in the direction of the net vector sum of the neutral pileup particles' $\vec{p}_{\mathrm{T}}$.

The "Type-0" $\not{E}_{\mathrm{T}}$ correction works to remove this effect by using $\vec{v}$, the net $\vec{p}_{\mathrm{T}}$ of charged pileup particles, as an estimator of the induced $\vec{\not{E}}_{\mathrm{T}}$ from the neutral pileup. The functional dependence of R^0 on $\vec{v}$ is estimated by fitting the correlation between the $\vec{\not{E}}_{\mathrm{T}}$ component parallel to the direction of $\vec{v}$ and the magnitude of $\vec{v}$,

$$f(\vec{v}) = c_1 \left[1 + \mathrm{erf}\left(-c_2 |\vec{v}|^{c_3}\right)\right]. \tag{5.8}$$

This fit is performed in a sample of simulated minimum bias events with exactly one generated pp interaction. An example of this correlation and resulting fit can be found in Fig. 5.9. For the 2012 8 TeV CMS dataset, this fit yielded the values, $c_1 = -0.71$, $c_2 = 0.09$, and $c_3 = 0.62$ for the coefficients.

To apply this correction, the factor $f(\vec{v})\vec{v}$, which gives the expected total induced $\vec{\not{E}}_{\mathrm{T}}$ for each pileup interaction, is calculated for each pileup vertex and the

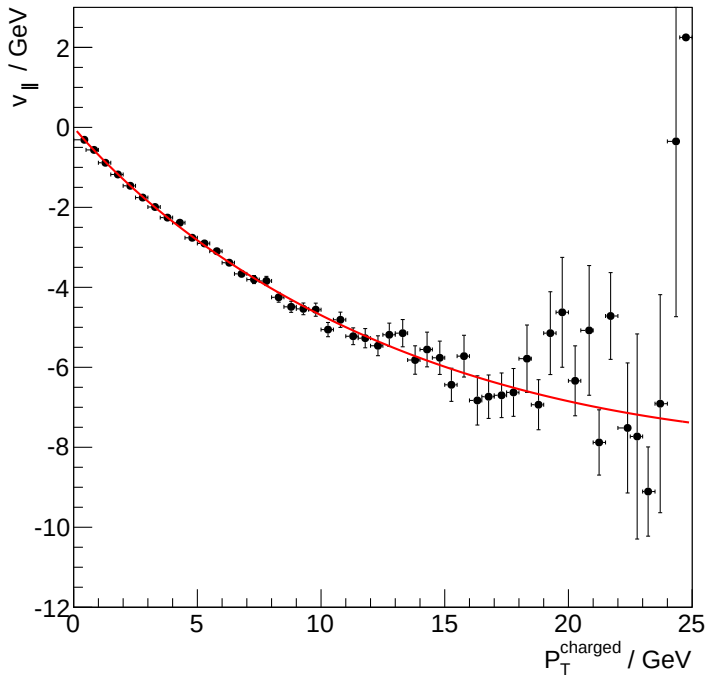

Figure 5.9: The correlation, in simulated minimum bias events, between the net $\vec{p}_{\mathrm{T}}$ of charged pileup particles and the reconstructed $\vec{\not{E}}_{\mathrm{T}}$ component parallel to this direction. This correlation has been fit by Eq. (5.8). Reprinted from Fig. 1 of [124].

total sum of these individual factors is subtracted from the uncorrected $\vec{\not{E}}_{\mathrm{T}}$,

$$\begin{aligned} \vec{C}_{\mathrm{T}}^{\text{Type-0}} &= -\sum_{i \in \mathrm{chPU}} f(\vec{v}_i)\vec{v}_i \\ \rightarrow \vec{\not{E}}_{\mathrm{T}}^{\text{Type-0}} &= \vec{\not{E}}_{\mathrm{T}}^{\text{raw}} + \vec{C}_{\mathrm{T}}^{\text{Type-0}}. \end{aligned} \tag{5.9}$$

5.7.1.2 Type-1 $\not{E}_{\mathrm{T}}$ Correction

The bulk of measurement biases in the $\not{E}_{\mathrm{T}}$ reconstruction are removed by correcting the p_{T} of jets back to their original, particle-level p_{T} using jet energy

corrections, c.f. Section 5.6.1 or Refs. [71, 125],

$$\begin{aligned}\vec{\not{E}}_{\mathrm{T}}^{\mathrm{corr}} &= \vec{\not{E}}_{\mathrm{T}}^{\mathrm{raw}} - \vec{\Delta}_{\mathrm{jets}} \\ &= \vec{\not{E}}_{\mathrm{T}}^{\mathrm{raw}} - \sum_{i \in \mathrm{jets}} \left(\vec{p}_{\mathrm{T},i}^{\mathrm{corr}} - \vec{p}_{\mathrm{T},i} \right),\end{aligned} \tag{5.10}$$

where the superscript "corr" refers to the corrected values.

5.7.1.3 $\vec{\not{E}}_{\mathrm{T}}\,\phi$ Asymmetry Correction

On average, particles are produced uniformly in ϕ. However, CMS analyses have observed an overall ϕ asymmetry in the $\vec{p}_{\mathrm{T}}$ sums of various visible objects (e.g. calorimeter energy deposits, tracks) leading to an overall ϕ asymmetry in the reconstructed $\vec{\not{E}}_{\mathrm{T}}$. This ϕ asymmetry is roughly sinusoidal and has been observed in both the data and the simulation. The primary sources of this asymmetry have been identified to be detector inefficiencies, imperfect detector alighnment, a shift between the center of the detector and the beamline, and an overall residual ϕ dependence of particle calibrations [126].

This observed ϕ asymmetry in the $\vec{\not{E}}_{\mathrm{T}}$ stems from overall shifts in the x and y components of the $\vec{\not{E}}_{\mathrm{T}}$. These shifts are correlated with $\mathrm{N}_{\mathrm{vtx}}^{\mathrm{reco}}$, the number of reconstructed interaction vertices, and this correlation is exploited in order to correct the asymmetry. Linear functions are fit to the correlation of $\not{E}_{\mathrm{x}}$ and $\not{E}_{\mathrm{y}}$ with $\mathrm{N}_{\mathrm{vtx}}^{\mathrm{reco}}$,

$$\begin{aligned}\langle \not{E}_{\mathrm{x}} \rangle &= c_{x_o} + c_{x_s} \mathrm{N}_{\mathrm{vtx}}^{\mathrm{reco}}, \\ \langle \not{E}_{\mathrm{y}} \rangle &= c_{y_o} + c_{y_s} \mathrm{N}_{\mathrm{vtx}}^{\mathrm{reco}}.\end{aligned} \tag{5.11}$$

The linear dependence of $\langle \not{E}_{\mathrm{x}} \rangle$ and $\langle \not{E}_{\mathrm{y}} \rangle$ on $\mathrm{N}_{\mathrm{vtx}}^{\mathrm{reco}}$ is used to correct $\vec{\not{E}}_{\mathrm{T}}$ on an

event-by-event basis as,

$$\begin{aligned} \not{E}_{\mathrm{x}}^{\mathrm{corr}} &= \not{E}_{\mathrm{x}} - \langle \not{E}_{\mathrm{x}} \rangle = \not{E}_{\mathrm{x}} - (c_{x_0} + c_{x_s} N_{\mathrm{vtx}}), \\ \not{E}_{\mathrm{y}}^{\mathrm{corr}} &= \not{E}_{\mathrm{y}} - \langle \not{E}_{\mathrm{y}} \rangle = \not{E}_{\mathrm{y}} - (c_{y_0} + c_{y_s} N_{\mathrm{vtx}}). \end{aligned} \tag{5.12}$$

The coefficients c_{x_0}, c_{x_s}, c_{y_0}, and c_{y_s} depend slightly upon the event sample considered, and also vary between data and simulation. Table 5.2 shows the calculated coefficients in an example sample dominated by $\mathrm{Z} \to \mu^+\mu^-$ events.

Table 5.2: The parameters for the PF $\vec{\not{E}}_{\mathrm{T}}$ ϕ-asymmetry corrections for data and simulation. As the detector alignment and ϕ-intercalibrations are different between data and simulation, the values of the respective parameters are expected to be different. Reprinted from Table 1 of [71].

	c_{x_0} (GeV)	c_{x_s} (GeV)	c_{y_0} (GeV)	c_{y_s} (GeV)
Data	0.05	0.25	−0.15	−0.08
Simulation	0.16	−0.24	0.36	−0.13

5.7.1.4 Correcting the $\not{E}_{\mathrm{T}}$ Resolution Modeling

The CMS simulation does not model the energy resolutions of jets with full accuracy, c.f. Section 5.6.1.2 or Ref. [127]. Although most jet-based analyses, especially those analyses looking at high-p_{T} jets, are not severely affected by this mismodeling, the underestimation of jet resolution in the simulation notably impacts the modeling of $\not{E}_{\mathrm{T}}$ resolution.

In order to improve the $\not{E}_{\mathrm{T}}$ resolution modeling, as part of the $\not{E}_{\mathrm{T}}$ performance studies detailed in Refs. [71,128], a smearing procedure was developed to individually smear the energies of reconstructed jets in simulated events.

The Smearing Procedure

For the first stage of this procedure, a collection of reconstructed jets (reco-jets) are built using all PF candidates as input. As well, a collection of generator-level jets (gen-jets) are built using the same clustering algorithm as that used to build the reco-jets (e.g. anti-k_{T} with distance parameter 0.5), but using all stable generator-level particles as input instead.

The second stage of the procedure is matching each reco-jet with a gen-jet, where this matching will be defined later. If a matching gen-jet is found, then the energy of the reco-jet is set to be,

$$E_{\mathrm{jet}}^{\mathrm{reco}} \mapsto E_{\mathrm{jet}}^{\mathrm{reco}} + (k_{\mathrm{jet}} - 1) \cdot \left(E_{\mathrm{jet}}^{\mathrm{reco}} - E_{\mathrm{jet}}^{\mathrm{gen}}\right).$$

Otherwise, it is set to,

$$E_{\mathrm{jet}}^{\mathrm{reco}} \mapsto E_{\mathrm{jet}}^{\mathrm{reco}} \cdot \left(1 + \frac{\mathcal{G}\left(\mu = 0,\, \left(k_{\mathrm{jet}}^{2} - 1\right) \cdot \sigma_{\mathrm{MC}}\right)}{E_{\mathrm{jet}}^{\mathrm{reco}}}\right),$$

where k_{jet} is an η dependent scale factor measured from comparisons between data and simulation (see Table 5.3); $\mathcal{G}(\mu, \sigma)$ refers to a random number drawn from a Gaussian distribution with mean μ and standard deviation σ; and σ_{MC} refers to the individual reco-jet's energy resolution as simulated by the MC.

Regardless of whether a matching gen-jet was found, the reco-jet's $\eta_{\mathrm{jet}}^{\mathrm{reco}}$, $\phi_{\mathrm{jet}}^{\mathrm{reco}}$ and the ratio of jet mass to energy, $m_{\mathrm{jet}}^{\mathrm{reco}}/E_{\mathrm{jet}}^{\mathrm{reco}}$, are all kept constant.

Table 5.3: Measured jet resolution correction factors k_{jet} used to smear the jets energies in Monte Carlo simulated events. Scale factors were calculated using 7 TeV dijet data, see Fig. 5.5b.

$\lvert\eta_{\text{jet}}^{\text{reco}}\rvert$ range	k_{jet}
$\lvert\eta_{\text{jet}}^{\text{reco}}\rvert < 0.5$	1.052 ± 0.065
$0.5 < \lvert\eta_{\text{jet}}^{\text{reco}}\rvert < 1.1$	1.057 ± 0.059
$1.1 < \lvert\eta_{\text{jet}}^{\text{reco}}\rvert < 1.7$	1.096 ± 0.070
$1.7 < \lvert\eta_{\text{jet}}^{\text{reco}}\rvert < 2.3$	1.134 ± 0.102
$\lvert\eta_{\text{jet}}^{\text{reco}}\rvert > 2.3$	1.288 ± 0.222

Matching Reco-jets with Gen-jets

Reco-jets are matched with gen-jets by searching in a cone of size $\Delta R < \Delta R_{\text{match}}$ around the reco-jet, where,

$$\Delta R_{\text{match}} = \min\left[0.5,\, 0.1 + 0.3 \cdot \exp\left(-0.05 \cdot \left(p_{\text{T, jet}}^{\text{gen}} - 10.0\right)\right)\right] \tag{5.13}$$

In Eq. (5.13), The dependence of ΔR_{match} on the reco-jet p_{T} helps to reduce matches stemming from pileup that have large differences between $E_{\text{jet}}^{\text{reco}}$ and $E_{\text{jet}}^{\text{gen}}$[7]. If there are multiple gen-jets geometrically closer than ΔR_{match}, the gen-jet with the smallest ΔR is chosen as the matched gen-jet; Even if a nominal match is found, if the difference of $E_{\text{jet}}^{\text{reco}}$ and $E_{\text{jet}}^{\text{gen}}$ is found to be larger than $3\sigma_{\text{MC}}$, then the jet is designated as "not-matched". This additional check serves to further reduce the number of jets where an excessive amount of jet-energy smearing would be applied due to fluctuations in the L1 pileup jet-energy correction.

[7] $E_{\text{jet}}^{\text{gen}}$ does not include energy from pileup particles, while $E_{\text{jet}}^{\text{reco}}$ can; it is true that the L1 pileup jet-energy corrections do, on average, remove the contributions of pileup to $E_{\text{jet}}^{\text{reco}}$, but large stochastic fluctuations in the amount of energy contributed to a jet from pileup can lead to too much additional smearing for simulated jets.

Impact of the Jet-energy Smearing on $\not\!\!E_{\mathrm{T}}$ Resolution

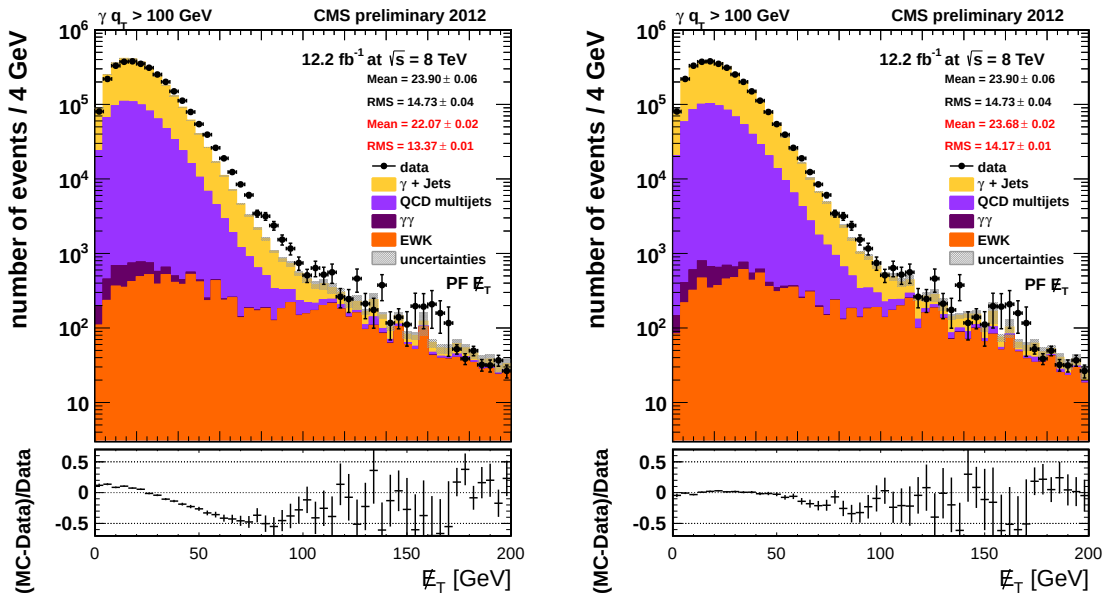

(a) The PF $\not\!\!E_{\mathrm{T}}$ distribution without jet-energy smearing.

(b) The PF $\not\!\!E_{\mathrm{T}}$ distribution with jet-energy smearing.

Figure 5.10: A demonstration of the impact of jet-energy smearing (Sec. 5.7.1.4) on the modeling of $\not\!\!E_{\mathrm{T}}$ resolution. Both sub-figures show the PF $\not\!\!E_{\mathrm{T}}$ distribution in an event sample of $\gamma +$ jets events with high-q_{T} (p_{T}^{γ}). Details on the sample can be found in Refs. [129, 130]. The left sub-figure shows the distribution without the application of jet-energy smearing, while the right sub-figure shows the distribution after the application of jet-energy smearing.

Figure 5.10 shows how this jet-energy smearing impacts the modeling of $\not\!\!E_{\mathrm{T}}$-resolution. In this figure, the PF $\not\!\!E_{\mathrm{T}}$ distribution in a sample of $\gamma +$jets events with high-q_{T} (i.e. p_{T}^{γ}) is shown with and without jet-energy smearing (more details on this event sample can be found in Refs. [71, 128, 129, 130]). As can be seen, without the application of the jet-energy smearing the $\not\!\!E_{\mathrm{T}}$-resolution in the simulation is clearly underestimated, but after smearing the jet-energies, the agreement between data and simulation is notably improved.

5.7.2 Large $\not{E}_{\mathrm{T}}$ due to Misreconstruction

The reconstruction of $\vec{\not{E}}_{\mathrm{T}}$ in an event can be substantially biased by detector and reconstruction issues, leading to spuriously high-$\not{E}_{\mathrm{T}}$ events in the collision data that must be identified and suppressed. There are a number of known, specific sources for these fake $\not{E}_{\mathrm{T}}$ events. Dedicated filters have been developed [71,131,132] in order to either remove the spurious signals in the detector, or to filter the event entirely if the spurious signal removal is not an option or is not applicable.

These known sources of fake $\not{E}_{\mathrm{T}}$ include,

- Particles at the radial fringes of the proton beams can interact with the LHC beam pipe material upstream of CMS (beam-scraping), leading to high-occupancy events;
- The HCAL hybrid photodiodes and readout boxes will occasionally generate spurious, large spikes of electronics noise; the individual spurious signals are filtered using shape-based filters;
- Direct particle interactions with the light guides and photomultiplier tubes in the HF can lead to spurious particle reconstruction and subsequent fake $\not{E}_{\mathrm{T}}$;
- Events where a substantial fraction of the energy is deposited in dead cells in the ECAL are removed by using ECAL trigger primitives;
- The ECAL and HCAL laser calibration systems can occasionally misfire, producing false signals in nearly all channels in a subdetector. These misfires can

overlap with bunch crossings, resulting in a trigger and a subsequent large, fake reconstructed $\not{E}_{\mathrm{T}}$. These events can easily be identified by the hit occupancies in the subdetector channels used for signal and calibration readout;

- A few channels in the ECAL endcaps occasionally produce high-amplitude anomalous pulses; these events can identified by the total energy and number of low-quality hits within the same reconstructed supercluster;

- The silicon strip tracker will occasionally be affected by electronic noise that is coherent across the entire subdetector, leading subsequently to $\mathcal{O}\left(10^4\right)$ clusters widely distributed in the silicon detectors. The online trigger reconstruction vetoes the majority of these events, but it is not fully efficient. When these events are reconstructed, the transverse momentum of the resulting spurious tracks can easily exceed 100 GeV, resulting in ultra high-p_{T} jets (with subsequent large, fake $\not{E}_{\mathrm{T}}$). These events are identified by the cluster multiplicity in the silicon strip and pixel detectors.

Figure 5.11 demonstrates the performance of the $\not{E}_{\mathrm{T}}$-filters in an event sample dominated by high-p_{T} QCD dijet events (c.f. Ref. [71] for an exact definition of this event sample). One can see that after the application of the $\not{E}_{\mathrm{T}}$-filters, the high-$\not{E}_{\mathrm{T}}$ tail is substantially reduced and agrees well with the expected distribution from simulation for PF $\not{E}_{\mathrm{T}}$ above 500 GeV.

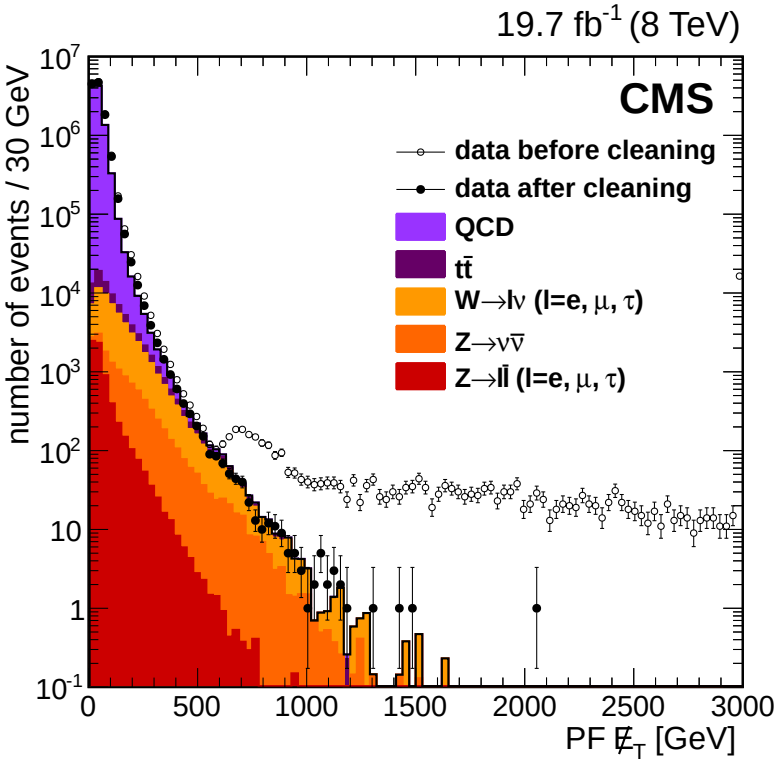

Figure 5.11: The distribution of PF $\vec{\not{E}}_{\mathrm{T}}$ in an event sample dominated by high-p_{T} QCD dijet events. The filled histograms represent the distribution in simulated events, while the filled (open) markers represent the distribution in collision data events with(out) the application of $\not{E}_{\mathrm{T}}$-cleaning algorithms. Reprinted from Fig. 4 of [71].

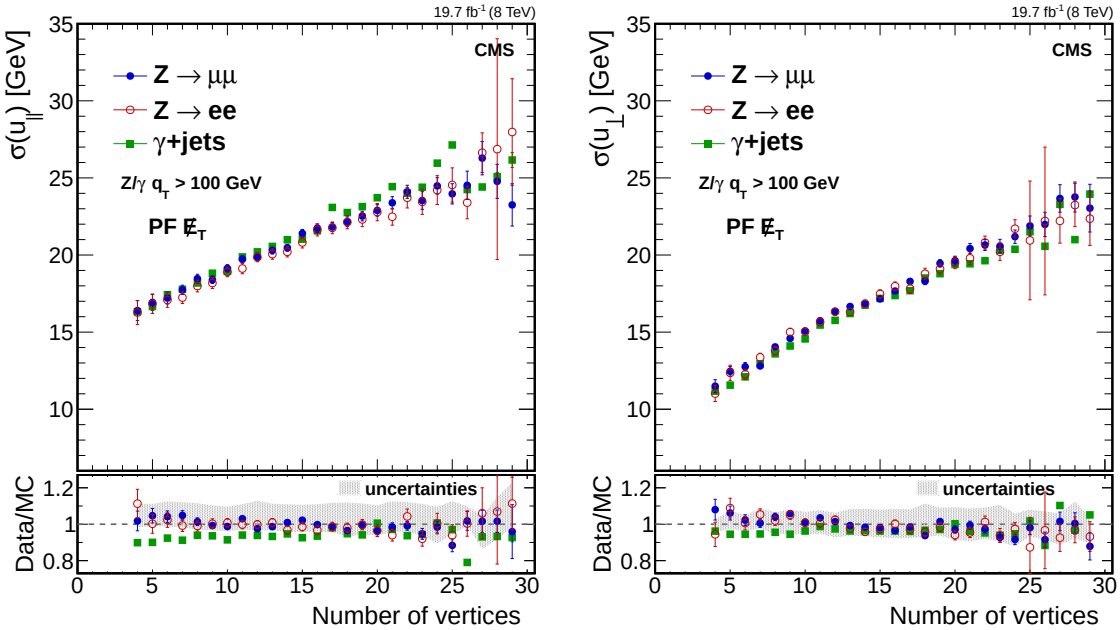

(a) (b)

Figure 5.12: The dependence of the energy resolution of $u_\parallel$ and $u_\perp$ on the number of reconstructed vertices in events with a Z-boson or γ. Note that the energy resolution of these variables is directly related to the $\not{E}_{\mathrm{T}}$ energy resolution (see Section 6.1 of Ref. [71] for more discussion on this relation). Results are shown for $\mathrm{Z} \to \mu^+\mu^-$ events (full blue circles), $\mathrm{Z} \to \mathrm{e}^+\mathrm{e}^-$ events (open red circles), and direct-photon events (full green squares). The upper frame of each figure shows the resolution in data; the lower frame shows the ratio of data to simulation with the grey error band displaying the systematic uncertainty of the simulation, estimated as the maximum of each channel's systematic uncertainty. Reprinted from Fig. 11 of [71].

5.7.3 Pileup Mitigation in $\not{E}_{\mathrm{T}}$ Reconstruction

As discussed above, pileup interactions have adverse effects of the $\not{E}_{\mathrm{T}}$ reconstruction. Most notably, they degrade the energy resolution of the $\not{E}_{\mathrm{T}}$ reconstruction. This resolution degradation can be quantified by looking at the dependence of the energy resolution of $u_{\parallel}$ and $u_{\perp}$ on the number of reconstructed vertices, as the energy resolution of these variables is directly related to the $\not{E}_{\mathrm{T}}$ energy resolution (see Section 6.1 of Ref. [71] for more discussion on this relation).

Figure 5.12 displays these resolution curves for $\mathrm{Z} \to \mu^{+}\mu^{-}$ events (full blue circles), $\mathrm{Z} \to \mathrm{e}^{+}\mathrm{e}^{-}$ events (open red circles), and direct-photon events (full green squares). From these resolution curves, one can quantify the degradation in $\not{E}_{\mathrm{T}}$ resolution due to pileup: each additional pileup interaction degrades the $\not{E}_{\mathrm{T}}$ energy resolution by $\sim$ 3.3–3.6 GeV, added in quadrature.

CMS has developed several improved $\not{E}_{\mathrm{T}}$ reconstruction algorithms that help mitigate the adverse effects of pileup interactions on the $\not{E}_{\mathrm{T}}$ resolution. These algorithms each begin by dividing a given event into two components: the particles that are likely to have come from the primary hard-scatter (HS), and the particles that are likely to have come from additional pileup interactions (PU).

One of these algorithms, the NoPU PF $\not{E}_{\mathrm{T}}$ algorithm, reconstructs the $\vec{\not{E}}_{\mathrm{T}}$ by separately weighting the contributions of the HS and PU particles. The contribution of the HS particles to the $\vec{\not{E}}_{\mathrm{T}}$ sum is not changed, but the PU particles' contribution is relatively down-weighted in order to reduce their overall impact on the $\vec{\not{E}}_{\mathrm{T}}$ reconstruction.

The other algorithm, the MVA PF $\not{E}_{\mathrm{T}}$ algorithm, utilizes a set of multivariate regressions. Five different $\vec{\not{E}}_{\mathrm{T}}$ variants are constructed, each utilizing different components (e.g. HS, PU) of the event. The MVA PF $\not{E}_{\mathrm{T}}$ algorithm utilizes all of these $\vec{\not{E}}_{\mathrm{T}}$ variants in order to calculate an overall correction to the observed PF $\vec{\not{E}}_{\mathrm{T}}$. There are currently two separate trainings of the MVA PF $\not{E}_{\mathrm{T}}$, one that is designed to achieve the optimal $\not{E}_{\mathrm{T}}$ resolution, and one that is designed to achieve a unity $\not{E}_{\mathrm{T}}$ response (i.e. no systematic bias in the $\not{E}_{\mathrm{T}}$ reconstruction) while still reducing the impact of pileup interactions on the $\not{E}_{\mathrm{T}}$ resolution.

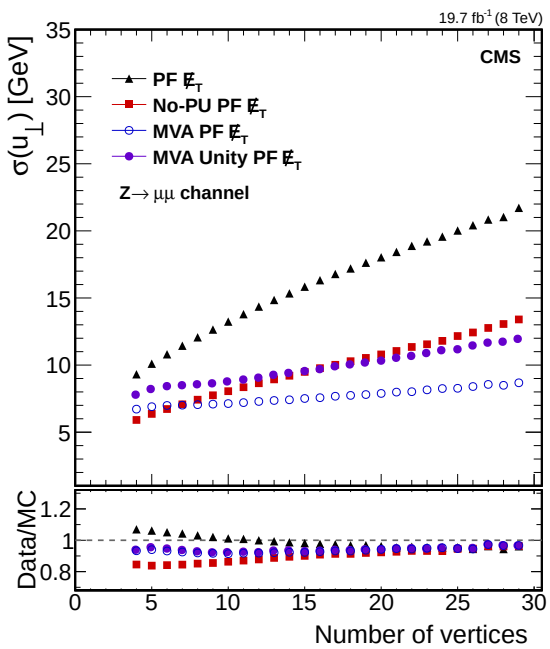

Figure 5.13: The pileup dependence of the $u_{\perp}$ resolution, for PF $\vec{\not{E}}_{\mathrm{T}}$ (black triangles), NoPU PF $\vec{\not{E}}_{\mathrm{T}}$ (red squares), MVA PF $\vec{\not{E}}_{\mathrm{T}}$ (blue open circles), and MVA PF Unity $\vec{\not{E}}_{\mathrm{T}}$ (violet full circles) in $\mathrm{Z} \to \mu^{+}\mu^{-}$ events. The upper frame of each figure shows the resolution in data; the lower frame shows the ratio of data to simulation. Reprinted from Fig. 21 of [71].

Figure 5.13 shows the dependence of the $u_{\perp}$ energy resolution on the number of reconstructed vertices in $\mathrm{Z} \to \mu^{+}\mu^{-}$ events, for four separate PF $\not{E}_{\mathrm{T}}$ reconstruc-

tion algorithms: the basic PF (black triangles), the NoPU PF (red squares), MVA PF (blue open circles), and MVA Unity PF $\vec{\not{E}}_{\mathrm{T}}$ (violet full circles). As can be clearly seen, in the pileup-mitigating $\not{E}_{\mathrm{T}}$ algorithms, the negative impact of pileup on the $\not{E}_{\mathrm{T}}$ energy resolution has been reduced; the degradation per additional pileup interaction is reduced, relative to the basic PF $\not{E}_{\mathrm{T}}$, by approximately a factor of two.

Chapter 6: Search for Top-squarks in the Dileptonic Final State using $M_{\mathrm{T2}}(\ell\ell)$

In the previous chapter of this dissertation, we outlined supersymmetry as a possible solution to current issues in the Standard Model, most notably the gauge hierarchy problem.

As noted in Sections 3.2.4 and 3.2.5, current LHC data has provided strict limits on the production of generic colored superpartners. However, SUSY can still be a natural solution to the gauge hierarchy problem – i.e. no major fine tuning required for model parameters – if the superparters of the Higgs, top, and gluon have masses near the electroweak scale. In particular, the experimental search for top-squark pair-production is one of the most important tests of "natural" SUSY.

In this chapter, we describe a search with the CMS detector for signatures of top-squark pair-production. We begin by providing some general context for this search.

6.1 General Context for the Top-squark Search

In searches for top-squark pair-production at the LHC, in order to provide the best experimental limits possible, it is important to directly target the expected

final state topology. For top-squark pair-production, the pair production of SM top quarks represents one of the major backgrounds for the majority of the top-squark decay modes considered.

This can quickly be seen by comparing the production and subsequent R-parity conserving decay of a top-squark pair into the dileptonic final state via an intermediate top quark,

$$\mathrm{pp} \to \widetilde{\mathrm{t}}_1 + \widetilde{\mathrm{t}}_1^* \to \widetilde{\chi}_1^{0*}\mathrm{t} + \widetilde{\chi}_1^0\bar{\mathrm{t}} \to \mathrm{bW}^+\widetilde{\chi}_1^{0*} + \bar{\mathrm{b}}\mathrm{W}^-\widetilde{\chi}_1^0 \to \mathrm{b}\ell^+\bar{\nu}_\ell\widetilde{\chi}_1^{0*} + \bar{\mathrm{b}}\ell^-\nu_\ell\widetilde{\chi}_1^0,$$

against the production and dileptonic decay chain of SM $\mathrm{t}\bar{\mathrm{t}}$ production,

$$\mathrm{pp} \to \mathrm{t} + \bar{\mathrm{t}} \to \mathrm{bW}^+ + \bar{\mathrm{b}}\mathrm{W}^- \to \mathrm{b}\ell^+\bar{\nu}_\ell + \bar{\mathrm{b}}\ell^-\nu_\ell.$$

The compositions of the final states of SM $\mathrm{t}\bar{\mathrm{t}}$ and top-squark pair events are identical except for the additional invisible particles, the $\widetilde{\chi}_1^0$.

A baseline selection for a top-squark search is to treat it as a $\mathrm{t}\bar{\mathrm{t}}$ search with additional $\not{E}_\mathrm{T}$. This would mean selecting for b-jets (all final states), high-p_T leptons (the semi-leptonic or dileptonic final states), and multiple jets, possibly including substructure information to tag jets coming from the top quark and W-boson decays (the semi-leptonic and fully hadronic final states).

6.1.1 Discriminating Variables in Top-squark Searches

After performing these basic selections, the key to top-squark searches is then properly distinguishing between the large $\mathrm{t}\bar{\mathrm{t}}$ background and the top-squark signal. To that end, the presence of the additional invisible particles naturally leads to using $\not{E}_\mathrm{T}$ or $\not{E}_\mathrm{T}$-related quantities as key discriminating variables.

For top-squark searches in the all-hadronic final state, the basic $\not{E}_{\mathrm{T}}$ is often used, although more sophisticated variables, such as $\not{E}_{\mathrm{T}}$ significance (defined in [71]), can provide additional discrimination. For top-squark searches in the semi-leptonic (dileptonic) final state, the power of $\not{E}_{\mathrm{T}}$ as a discriminating variable is hindered by the presence of genuine $\not{E}_{\mathrm{T}}$ stemming from the neutrino(s) produced in the leptonic decays of the W bosons. Rather than use $\not{E}_{\mathrm{T}}$ by itself, analyses searching for top-squark signatures in these final states must exploit the correlations between the reconstructed $\not{E}_{\mathrm{T}}$ and the observed visible objects in the events. As an example, for the semi-leptonic final state, M_{T}, the transverse mass constructed with the lepton and the $\not{E}_{\mathrm{T}}$, is often used instead of $\not{E}_{\mathrm{T}}$.

In the case where a mother particle, X, decays to two daughter particles a and b, M_{T} can be defined as follows,

$$\begin{aligned} M^2_{\mathrm{T},\,X}(\vec{p}_{\mathrm{T},\,a},\, \vec{p}_{\mathrm{T},\,b},\, m_a,\, m_b) &\equiv m_a^2 + m_b^2 + 2\left[E_{\mathrm{T},\,a}E_{\mathrm{T},\,b} - \vec{p}_{\mathrm{T},\,a}\cdot\vec{p}_{\mathrm{T},\,b}\right] \\ &= 2|\vec{p}_{\mathrm{T},\,a}||\vec{p}_{\mathrm{T},\,b}|\left(1-\cos(\Delta\phi_{a,\,b})\right), \end{aligned} \tag{6.1}$$

where the second line of Eq. (6.1) shows the definition of M_{T} in the limit that the daughter particle masses can be neglected.

The expression for M_{T} can be compared against an analogous expression for the reconstructed invariant mass of X,

$$m_X^2 = m_a^2 + m_b^2 + 2\left[E_{\mathrm{T},\,a}E_{\mathrm{T},\,b}\cosh(\Delta\eta) - \vec{p}_{\mathrm{T},\,a}\cdot\vec{p}_{\mathrm{T},\,b}\right], \tag{6.2}$$

where $\cosh(\Delta\eta)$ is the hyperbolic cosine of the difference in pseudorapidity between a and b. Since $\cosh(x) \geq 1$, the transverse mass has the convenient property that when the two input objects come from the decay of a single particle, the distribution

of M_{T} has a kinematic endpoint at the mother particle's mass. Analyses at the Tevatron utilized this property of M_{T} to measure the W-boson mass [133, 134, 135]. In semi-leptonic top-squark searches, the distribution of M_{T} for SM $\mathrm{t\bar{t}}$ events, under the assumption that the observed $\not{E}_{\mathrm{T}}$ stems from a single invisible particle (i.e. the neutrino from the leptonic W-decay), has a kinematic edge at $m_{\mathrm{W}} \sim 80\,\mathrm{GeV}$. The presence of the two additional, invisible LSPs in top-squark pair events breaks the assumption and subsequently, the distribution of M_{T} for top-squark pair events often has a significant tail for values of $M_{\mathrm{T}} > m_{\mathrm{W}}$.

In dileptonic top-squark searches, M_{T} no longer makes sense as a discriminating variable; not only is there the ambiguity of selecting which of two leptons to pair the reconstructed $\vec{\not{E}}_{\mathrm{T}}$ with, but also, the observed $\vec{\not{E}}_{\mathrm{T}}$ no longer stems from a single invisible particle, even for the SM $\mathrm{t\bar{t}}$ background. A separate, related kinematic variable, M_{T2} [2, 3], can be used instead. This variable M_{T2}, or the "stransverse mass", is a generalization of M_{T} to a system of pair-produced particles that decay semi-invisibly. In particular, M_{T2} constructed with the two leptons, $M_{\mathrm{T2}}(\ell\ell)$, has the same property in SM dileptonic $\mathrm{t\bar{t}}$ events that M_{T} has in SM semi-leptonic $\mathrm{t\bar{t}}$ events – the distribution of $M_{\mathrm{T2}}(\ell\ell)$ has a kinematic edge at m_{W}. We devote additional discussion to $M_{\mathrm{T2}}(\ell\ell)$ in Section 6.2.5.

6.1.2 Overview of the Analysis

For this search, we use the dilepton final state (where lepton signifies e or μ). Dilepton events are straightforward to trigger on, the position and momentum

measurements for the leptons are precise, and many of the backgrounds can be estimated using robust, data-driven methods. In order to suppress most of the SM backgrounds in the dilepton channel, we use events with two opposite-charge, high-p_{T}, isolated leptons, and at least two jets with at least one b-tagged jet. This results in an event sample where the dominant background is SM dileptonic $\mathrm{t\bar{t}}$ events.

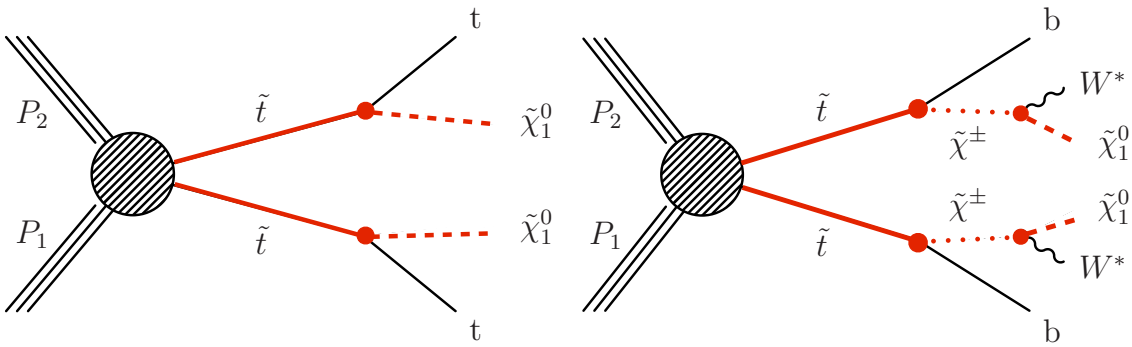

(a) Diagram for the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0 \to \mathrm{bW}^+\tilde{\chi}_1^0$ (T2tt) decay mode.

(b) Diagram for the $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+ \to \mathrm{bW}^+\tilde{\chi}_1^0$ (T2bw) decay mode.

Figure 6.1: Feynman diagram representations of the symmetric decay modes of top-squark pair-production at the LHC. Propagators and vertices for supersymmetric particles are colored in red.

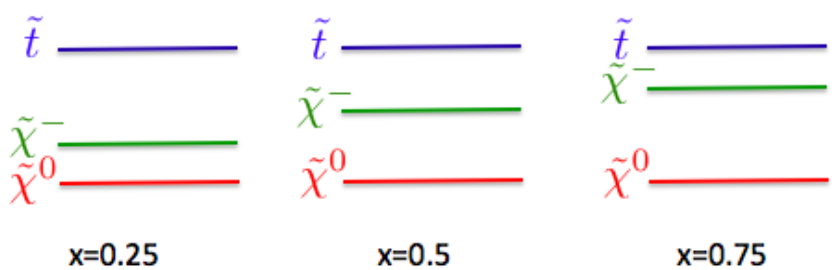

Figure 6.2: The three relative mass spectra of the $\tilde{\mathrm{t}}_1$, $\tilde{\chi}_1^\pm$, and $\tilde{\chi}_1^0$ in the $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode that are considered for this analysis. Each is demarcated by the value of the chargino mass-splitting parameter, $x = \frac{m_{\tilde{\chi}_1^+} - m_{\tilde{\chi}_1^0}}{m_{\tilde{\mathrm{t}}_1} - m_{\tilde{\chi}_1^0}}$, which measures the relative mass-splitting between the three supersymmetric particles in the $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay chain.

In line with the other searches for top-squark pair-production at the LHC (Section 3.2.6.3), we considered two top-squark decay modes in this search, one

where each top-squark decays directly to a SM top quark and the lightest neutralino, i.e. $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0 \to \mathrm{bW}^+\widetilde{\chi}_1^0$ (T2tt), and one where each top-squark decays to a b quark and the lightest chargino, i.e. $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+ \to \mathrm{bW}^+\widetilde{\chi}_1^0$ (T2bw).

These are hereafter primarily referenced by their CMS simplified model acronyms: T2tt and T2bw. Under the CMS simplified model paradigm, all SUSY particles, except the ones directly involved in the decay mode, are assumed to have large enough masses to be irrelevant in the top-squark decay. This is why the decay chains discussed above involve the lighter top-squark $\widetilde{\mathrm{t}}_1$. Moreover, the relevant SUSY couplings are chosen such that they uniformly represent all possible chirality or handedness scenarios (e.g. in the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode, the $\widetilde{\mathrm{t}}_1$ decays with equal probability to left- and right-handed top quarks). The sensitivity for specific SUSY scenarios can then be investigated by reweighting or combining different simplified model results.

Figure 6.1 shows general Feynman diagram representations for these two decay modes. Due to an assumption of R-parity conservation, the $\widetilde{\mathrm{t}}_1$ particles are pair produced and furthermore, the lightest neutralino, $\widetilde{\chi}_1^0$, is treated as the stable LSP, escaping the detector and adding to the $\vec{\not{E}}_{\mathrm{T}}$ of the event.

In order to ensure that the search is as model-independent as possible, for both decay modes we explore the experimental sensitivity in the full range of kinematically accessible regions for the unknown masses of the $\widetilde{\mathrm{t}}_1$ and $\widetilde{\chi}_1^0$. For the T2bw decay mode, the mass of the $\widetilde{\chi}_1^\pm$ is a third unknown. Rather than independently scan through all kinematically allowed $m_{\widetilde{\chi}_1^\pm}$, we fix the relative mass-splitting between the three supersymmetric particles with a parameter x that is set by the relation,

$x = \frac{m_{\tilde{\chi}_1^+} - m_{\tilde{\chi}_1^0}}{m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}}$. In the rest of this dissertation, we will refer to x in this context as the chargino mass-splitting parameter. We consider three values for x, $x = 0.25$, $x = 0.50$, and $x = 0.75$, leading to three relative mass-spectra, as shown in Fig. 6.2.

We utilize the dileptonic stransverse mass, $M_{\mathrm{T2}}(\ell\ell)$, to discriminate between our main background, SM dileptonic $\mathrm{t\bar{t}}$ events, and the top-squark signal. Specifically, using a sample of events constructed from the aforementioned selections, we perform a "cut-and-count" experiment in the general signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$. The exact $M_{\mathrm{T2}}(\ell\ell)$ threshold we cut on is determined on a case-by-case basis. For each point in the 2D SUSY mass plane $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ and for each considered top-squark decay mode, we calculate the "data-blind" median expected upper limit (c.f. Appendix G) on the top-squark pair production cross-section, $\sigma^{\mathrm{UL,\,exp.}}_{\tilde{t}_1\tilde{t}_1^*}$, for 5 different $M_{\mathrm{T2}}(\ell\ell)$ thresholds (iterating from 80 GeV to 120 GeV in 10 GeV steps). For calculating our final values of $\sigma^{\mathrm{UL}}_{\tilde{t}_1\tilde{t}_1^*}$, both the expected and observed versions, we utilize the threshold that yielded the optimal data-blind median expected limit.

With that general overview complete, we now provide the roadmap for the rest of this chapter. In Section 6.2 we provide details on:

- The samples of simulation and collision data used (Section 6.2.1).
- The object and event selections applied to extract our final event samples (Sections 6.2.2 and 6.2.3).
- Details on reweighting/corrections applied to the simulated data (Section 6.2.4).
- Additional discussion on $M_{\mathrm{T2}}(\ell\ell)$ and the signal region definition (Section 6.2.5).

In Section 6.3 we provide details on the estimation of our major backgrounds, including detailed discussions on the data-driven estimation methods used for major backgrounds.

In Section 6.4, we provide details on the notable systematic uncertainties affecting our analysis and how we estimate their magnitude and effect.

In Section 6.5, we discuss the multitude of checks we performed to both develop our general understanding of $M_{\mathrm{T2}}(\ell\ell)$ and to validate the accuracy of our simulation's modeling of $M_{\mathrm{T2}}(\ell\ell)$.

In Section 6.6, we discuss the aspects of this analysis directly related to our top-squark signal.

In Sections 6.7 and 6.8, we provide the results of this search, including comparisons between data and simulation in our $M_{\mathrm{T2}}(\ell\ell)$ signal region as well as exclusion contours for regions of the 2D SUSY mass plane $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$.

6.2 Selection

6.2.1 Data Samples and Trigger Selection

Selection of Collision Datasets

The data used in this analysis is 8 TeV pp collision data taken by the CMS detector during the 2012 run of the LHC. Several HLT paths were used to select the collision data events used in this analysis. Dilepton triggers provided the main samples of collision data. These triggers require a combination of at least two re-

constructed HLT muon and electron candidates, where the (sub-)leading candidate must have $p_{\mathrm{T}} > (8)\,17\,\mathrm{GeV}$. Reconstructed HLT electrons are subject to additonal detector isolation requirements in order to alleviate trigger pressure from the mis-reconstruction of hadronic jets as electrons. The total integrated luminosity of this dilepton-triggered dataset is $19.66 \pm 0.51\ \mathrm{fb}^{-1}$.

In order to estimate the efficiency for leptons to pass the trigger selection requirements (as described in Section 6.2.4.1), we also use a selection of collision data events selected by their HLT-reconstructed $\not{E}_{\mathrm{T}}$ values. These events are provided by various $\not{E}_{\mathrm{T}}$-based triggers, chosen so as to minimize statistical correlations between the $\not{E}_{\mathrm{T}}$-based triggers and our main dilepton triggers.

In order to estimate the efficiency with which lepton candidates pass the isolation and identification requirements (As described in Section 6.2.4.1), we also used a sample of events selected by single-lepton triggers; these triggers require the presence of a reconstructed electron or muon candidate with high p_{T}: electron (muon) candidates must have $p_{\mathrm{T}} > 27\,(24)$ GeV. As well, muons must have $|\eta| < 2.1$ and electrons must pass the medium working point of the electron selection requirements – c.f. Section 6.2.2 for details on these requirements.

A sample of single-lepton triggered events are also used to both estimate the contribution of "fake" leptons (c.f. Section 6.3.3) to our signal regions, and to investigate the modeling of the $M_{\mathrm{T2}}\,(\ell\ell)$ shape. For these triggers, the reconstructed lepton must satisfy $p_{\mathrm{T}} > 8\,\mathrm{GeV}$. As with the single-lepton trigger used in the isolation and identification efficiency measurements, electron candidates are required to pass the medium working point of the electron selection requirements.

Simulation Samples

MC simulations were used to study both the expected analysis selection efficiency for top-squark events as well as the expected contribution of SM background processes to our final selected data sample. Several MC generators are used to model the expected SM background processes:

- The POWHEG 1.0 r138 [99, 136, 137] generator was used to model both the main background of $t\bar{t}$ events as well as single top quark production events.

- The MadGraph v5.1.3.30 generator [98] was used to model the Z + jets events (Drell-Yan, or DY) as well as provide samples of more rare background processes, such as triple vector boson (e.g. ZZZ, WWZ, etc.) or the production of vector bosons in association with a $t\bar{t}$ pair ($t\bar{t}$ + Z, $t\bar{t}$ + WW, etc.). The MadGraph generator is also used to provide a cross-check sample of both the $t\bar{t}$ and WW backgrounds. For the modeling of Z + jets production, two kinds of samples were generated, one allowing any number of additional hadronic jets ("jet-inclusive") and one requiring an explicit number (either 1, 2, 3, or 4) of additional jets at the generator level.

- The MC@NLO generator [138] was used to provide a SM $t\bar{t}$ cross-check sample.

- The PYTHIA v6.4.24 generator [97] was used to model diboson production — e.g. WW, WZ, ZZ, Wγ, Zγ, e.t.c.

Signal Simulation Samples

The MadGraph generator was utilized to generate the samples of top-squark events that were used to estimate the signal efficiency. These samples include interactions diagrams with up to two additional partons in the final state.

Each of these samples consisted of scans in fixed intervals over a chosen region in the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane. The interval sizes used in each sample were either 25 GeV or 12.5 GeV with the same interval size used in both the $m_{\widetilde{\mathrm{t}}_1}$ and $m_{\widetilde{\chi}_1^0}$ directions for a given sample. For each point in the scan of the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane, $\mathcal{O}(10^5)$ events were generated. The available $\widetilde{\mathrm{t}}_1\widetilde{\mathrm{t}}_1^*$ final states also depended upon the sample.

In general, these top-squark event samples can be separated into three categories: "inclusive" samples, "ℓ-filtered" samples, and "ℓ-filtered" samples with tight intervals. These ℓ-filtered samples were created during the progression of earlier CMS analyses looking for top-squark pair-production in the semi-leptonic final state. Their purpose was to provide additional statistics in interesting regions of the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane for the two decay modes considered. For the T2tt decay mode, the interesting region is the area with $\Delta\mathrm{M} \approx m_{\mathrm{t}}$ and $m_{\widetilde{\chi}_1^0} \lesssim 150$ GeV, as this area is not only relatively well-motivated by natural SUSY, but it also quite experimentally difficult, see Figs. 3.3 and 3.4. For the T2bw decay modes, the interesting region is the area where the $\widetilde{\mathrm{t}}_1$ is not excessively heavy, $m_{\widetilde{\mathrm{t}}_1} \lesssim 600$ GeV and where the mass splitting between the $\widetilde{\mathrm{t}}_1$ and $\widetilde{\chi}_1^0$ is not excessively large, $\Delta\mathrm{M} \lesssim 300$ GeV.

Tables 6.1 and 6.2 contain the relevant details for these three categories.

Table 6.1: Details on the generation of the $\widetilde{t}_1\widetilde{t}_1^*$ event samples for the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode. All samples were generated such that nominally all top-squark final states were available; however, the generation of the ℓ-filtered samples included the the requirement that each event have at least one generator-level lepton with $p_T \geq 10\,\text{GeV}$.

Sample	interval size [GeV]	$m_{\widetilde{t}_1}$ range [GeV]	$m_{\widetilde{\chi}_1^0}$ range [GeV]	$\Delta M = m_{\widetilde{t}_1} - m_{\widetilde{\chi}_1^0}$ range [GeV]	gen. no. of events
Inclusive	25	100:1000	0:900	100:1000	1.40×10^5 ($m_{\widetilde{\chi}_1^0} = 0$ GeV)
					1.15×10^5 ($m_{\widetilde{\chi}_1^0} \neq 0$ GeV)
ℓ-filtered	25	100:350	-	100:200	3.70×10^5
ℓ-filtered tight interval	12.5	150:400	-	150:200	4.10×10^5

Table 6.2: Details on the generation of the $\widetilde{t}_1\widetilde{t}_1^*$ event samples for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode. All samples were generated such that nominally all top-squark final states were available; however, the generation of the ℓ-filtered samples included the the requirement that each event have at least one generator-level lepton with $p_T \geq 10\,\text{GeV}$.

Sample	interval size [GeV]	$m_{\widetilde{t}_1}$ range [GeV]	$m_{\widetilde{\chi}_1^0}$ range [GeV]	$\Delta M = m_{\widetilde{t}_1} - m_{\widetilde{\chi}_1^0}$ range [GeV]	gen. no. of events
Inclusive	25	100:800	0:800	100:700	1.50×10^5 ($m_{\widetilde{\chi}_1^0} = 0$ GeV)
					1.75×10^4 ($m_{\widetilde{\chi}_1^0} \neq 0$ GeV)
ℓ-filtered	25	100:650	-	100:350	7.25×10^4

Events with τ Leptons

For all events involving the production of τ leptons, the decay of the τ leptons are handled via the TAUOLA software package [100].

Simulation of the CMS Detector

After their creation, all of the generated events for the simulated SM background processes are run through the CMS FullSim, c.f. Section 4.2.9.2. Due to the large number of signal models and mass points considered, it is infeasible to run each generated top-squark signal event through the CMS FullSim; instead the top-squark signal events are run through the CMS FastSim, c.f. Section 4.2.9.2.

Normalization of the Simulated Samples

We normalize the MC simulation event samples to their expected contributions by scaling each process with its own respective cross section. For the background, these cross sections are calculated at various orders of perturbation theory:

- The cross sections for $t\bar{t}$ [139] and single top quark processes [140, 141, 142], are calculated at next-to-next-to-leading order (NNLO).
- The cross sections for both the jet-inclusive and individual jet-exclusive Z + jets processes are calculated at NNLO [143].
- The cross sections for the VV processes as well as the production of $t\bar{t}$ pairs with associated EWK boson production are calculated at NLO [144].

- The cross sections for the triple EWK boson production processes are calculated at NLO [138].

The dependence of the 8 TeV top-squark pair-production cross section on $m_{\tilde{t}_1}$ was calculated at NLO in the strong coupling constant, with the inclusion of the resummation of soft gluon emission at next-to-leading logarithmic accuracy [145,146]. Figure 6.3 shows the dependence of this cross section on $m_{\tilde{t}_1}$. Standard theoretical uncertainties on this cross section stemming from PDF eigenvector set uncertainties as well as renormalization scale variations (a factor of 2 in both directions) have been added in quadrature and are represented by the blue band.

6.2.2 Object Selection

In this section, we provide verbal descriptions of the selection requirements we apply to our visible objects. Summary tables of these requirements are provided for reference in Appendix A.

Lepton Selection

Common Lepton Requirements

There are a number of requirements common to both flavors of lepton candidate used in this analysis.

The two hardest lepton candidates are required to have $p_{\mathrm{T}} > 20\,(10)$ GeV for the harder (softer) candidate.

Selected lepton candidates must also be isolated: The total transverse momen-

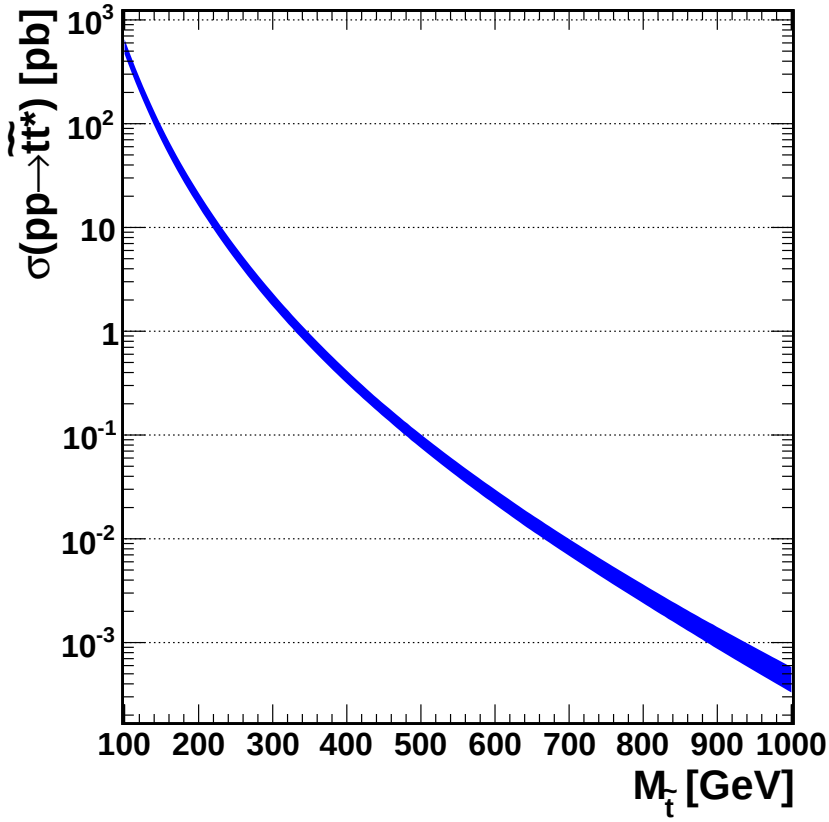

Figure 6.3: The NLO cross-section for direct top-squark pair-production at the LHC at $\sqrt{s} = 8\,\text{TeV}$. The blue-band represents the uncertainties on the cross-section corresponding to PDF uncertainties added in quadrature with renormalization scale variations.

tum of PF candidates in a cone of radius, $\Delta\mathrm{R} = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 0.3$, around a given lepton candidate (also known as the PF isolation of the lepton) is required to be less than 15% of the lepton candidate's p_{T}.

For electrons, the PF isolation of the electron, $\mathrm{Iso_e}$, is calculated as the sum of the total transverse momentum sum of charged hadrons, $\mathrm{Iso_{h^\pm}}$; neutral hadrons, $\mathrm{Iso_{h^0}}$; and photons, Iso_γ, corrected for the effect of pileup using the electron's effective area in the (ϕ, η) plane, A_{e}, and the average energy density of pileup, ρ,

$$\mathrm{Iso_e} = \mathrm{Iso^e_{h^\pm}} + \max\left[0, \left(\mathrm{Iso^e_{h^0}} + \mathrm{Iso^e_\gamma} - (\rho A_{\mathrm{e}})\right)\right].$$

For muons, the PF isolation of the muon, Iso_μ, like the analogous isolation for the electrons, depends upon the individual momentum sums of charged hadrons, neutral hadrons, and photons; unlike the electrons, though, muon candidates' PF isolations are corrected for pileup based on the total sum of charged PF candidates identified with a pileup vertex, $p_{\mathrm{T}}^{\mathrm{PU}\pm}$,

$$\mathrm{Iso}_\mu = \mathrm{Iso^\mu_{h^\pm}} + \max\left[0, \left(\mathrm{Iso^\mu_{h^0}} + \mathrm{Iso^\mu_\gamma} - 0.5 p_{\mathrm{T}}^{\mathrm{PU}\pm}\right)\right].$$

where the factor of 0.5 in front of $p_{\mathrm{T}}^{\mathrm{PU}\pm}$ accounts for the empirically observed relative ratio of charged to neutral particles produced by pileup interactions [108].

Lepton candidates must have a position consistent with the primary hard-scatter vertex (the vertex with highest $\sum p_{\mathrm{T}}^2$ of tracks associated with the vertex); for muon candidates, we require the reconstructed transverse (longitudinal) impact parameter with respect to the primary hard-scatter vertex, d_0 (d_z), to be less than 2 (5) mm. For electron candidates, we require these numbers to be less than 0.02

and 0.1 mm respectively.

In addition to coming from the primary hard-scatter vertex, lepton candidates are also required to be central; for muons we require $|\eta| < 2.4$; for electrons, we require $|\eta| < 2.5$, with an additional veto on the boundary region between the electromagnetic barrel and endcap calorimeters ($1.4442 < |\eta| < 1.566$).

Both muon and electron candidates must have a consistent p_{T} value as measured by the basic reconstruction and more sophisticated PF reconstruction; The two values must differ by less than 5 (10) GeV for muon (electron) candidates.

Muon-specific Requirements

Muons must be reconstructed as a tight, global, prompt muon. As well, they must be successfully reconstructed as PF muons.

The reconstructed global track associated with the muon candidate must pass a few specific requirements:

- To suppress hadronic punch-through and muons from hadronic decays in flight, the normalized χ^2 for the muon's global track fit must be less than 10.
- To further suppress muons from hadronic decays in flight, the overall global muon track reconstruction must include at least 1 of the layers of the pixel detector.
- To ensure a robust measurement of the muon p_{T}, as well as to further suppress muons from hadronic decays in flight, the muon candidate must have associated hits in at least 6 tracker layers.

- To suppress hadronic punch-throughs into the muon system as well as accidental matches between inner tracks and segments in the muon system, we require at least 2 matched stations in the muon system.

The performance of these, and related, selection criteria was studied in detail in Refs. [116, 117, 147].

Electron-specific Requirements

Electron candidates are subject a number of stringent selection requirements in order to reject both hadronic jets misreconstructed as electrons as well as photons that converted into electron pairs.

- $\Delta\eta$ (SC, trk), the difference in η of the electron's associated supercluster and its track associated with the primary hard-scatter vertex, must be less than 0.004 (0.007) for electrons in the EB (EE).

- $\Delta\phi$ (SC, trk), the difference in ϕ of the electron's associated supercluster and its track associated with the primary hard-scatter vertex, must be less than 0.06 (0.03) radians for electrons in the EB (EE).

- $\sigma_{i\eta i\eta}$, the second moment of the log(energy)-weighted electron shower shape along the η direction, must be less than 0.01 (0.03) for electrons in the EB (EE).

- H/E, the ratio of the electron's energy measured in the HCAL over that measured in the ECAL, must be less than 0.12 (0.10) for electrons in the EB

(EE).

- $E^{-1} - p^{-1}$, the difference of the inverse of the electron's reconstructed energy and momentum, must be less than 0.05 for all selected electrons.

- n_{miss}, the number of missing hits in the tracker for the track associated with the electron, must be less than 1 for all selected electrons.

We also apply the standard veto on conversion electrons that rejects photons that underwent an e^+e^- conversion in the tracker and were reconstructed incorrectly as electrons. This conversion veto exploits differences in the patterns of track hits between electrons produced from EWK boson decays and those produced from photon conversions [118]. It introduces an $\mathcal{O}(1\%)$ inefficiency for prompt electrons reconstruction, but has a strong, $\mathcal{O}(45\%)$ rejection efficiency for electrons coming from photon conversions.

The performance of these, and related, selection criteria was studied in detail using both 7 TeV and 8 TeV data in Refs. [118, 148].

Jet Selection

Jets are built using the anti-k_{T} clustering algorithm with a distance parameter of 0.5, after first removing all charged hadrons coming from pileup interactions. The individual energies of jets are rescaled using simulation-derived corrections (c.f. Section 5.6.1).

Jets are required to satisfy $p_{\text{T}} > 30\,\text{GeV}$ and fall within $|\eta| < 2.4$. To prevent the misidentification of leptons as jets, we require that there be no "tight" lepton

(lepton that satisfied the previously mentioned lepton selection requirements) with $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ within a cone of $\Delta\mathrm{R} < 0.4$ around the jet.

In addition, jets are required to pass the loose working point of the PF Jet ID [149, 150]. The exact list of variables and their cut values can be found in Appendix A.1.

Jets originating from the hadronization of b quarks are an expected signature of top-squark pair production. Each reconstructed jet is considered b-tagged if it passes the medium working point of the CSV b-tagging algorithm (Section 5.6.2). As a reminder, the medium working point has a $\sim 70\%$ tagging efficiency for heavy flavor (c, b) jets and a $\sim 1\%$ mistag rate for light flavor (u, d, s, and gluon) jets.

$\not{E}_{\mathrm{T}}$ Selection

We use the basic PF $\vec{\not{E}}_{\mathrm{T}}$ (Section 5.7), which is computed as the negative vector sum of the $\vec{p}_{\mathrm{T}}$ of all PF candidates. We apply the Type-1 (Section 5.7.1.2) and $\not{E}_{\mathrm{T}}\,\phi$ asymmetry (Section 5.7.1.3) corrections to the $\vec{\not{E}}_{\mathrm{T}}$ to account for observed asymmetries in the $\not{E}_{\mathrm{T}}\,\phi$ distribution and the jet energy recalibrations. Note that we did not apply the Type-0 $\not{E}_{\mathrm{T}}$ correction, Section 5.7.1.1, in order to remain consistent with a preliminary version of this analysis.

In order to account for the inaccurate modeling of jet energy resolutions in the simulation, most notably their impact on $\not{E}_{\mathrm{T}}$ resolution, c.f. Section 5.7.1.4, we apply a corrective "smearing" to the magnitude and direction of the $\vec{\not{E}}_{\mathrm{T}}$ in order to improve the agreement between data and simulation. Section 6.2.4.3 contains the

Table 6.3: Event selection defining the full preselection sample.

Object	Selection
e, μ	At least two, oppositely charged
	Highest sum-p_{T} pair used
M_{ee} or $M_{\mu\mu}$	$M_{\ell^+\ell^-} < 76$ GeV $\cup$ $M_{\ell^+\ell^-} > 106$ GeV
$\not{E}_{\mathrm{T}}$	≥ 40 GeV in ee, $\mu\mu$ channels
$M_{\ell^+\ell^-}$	> 20 GeV (all flavors)
N_{jets}	≥ 2
N_{b}	≥ 1

details of this smearining.

6.2.3 Event Selection

We require at least two oppositely-charged leptons with an invariant mass larger that 20 GeV (to reject low-mass dilepton resonances). If there are two or more opposite-charge lepton pairs that satisfy the mass requirement, we select the pair that has the largest scalar sum of the individual lepton p_{T}. Finally, once we have selected an opposite-charge lepton pair, the event itself must have passed the dilepton HLT path associated with the flavor of the selected lepton pair (e.g. if we select an $\mathrm{e}^+\mathrm{e}^-$ pair, the event must have passed the dielectron HLT path). If the event does not meet this requirement, it is discarded. For the simulation samples, we emulate and apply the online dilepton trigger requirement using a simulation of the CMS dilepton triggers.

We also require at least 2 jets, where at least one of the jets is b-tagged.

To further reduce the background contribution from SM Drell-Yan events, we

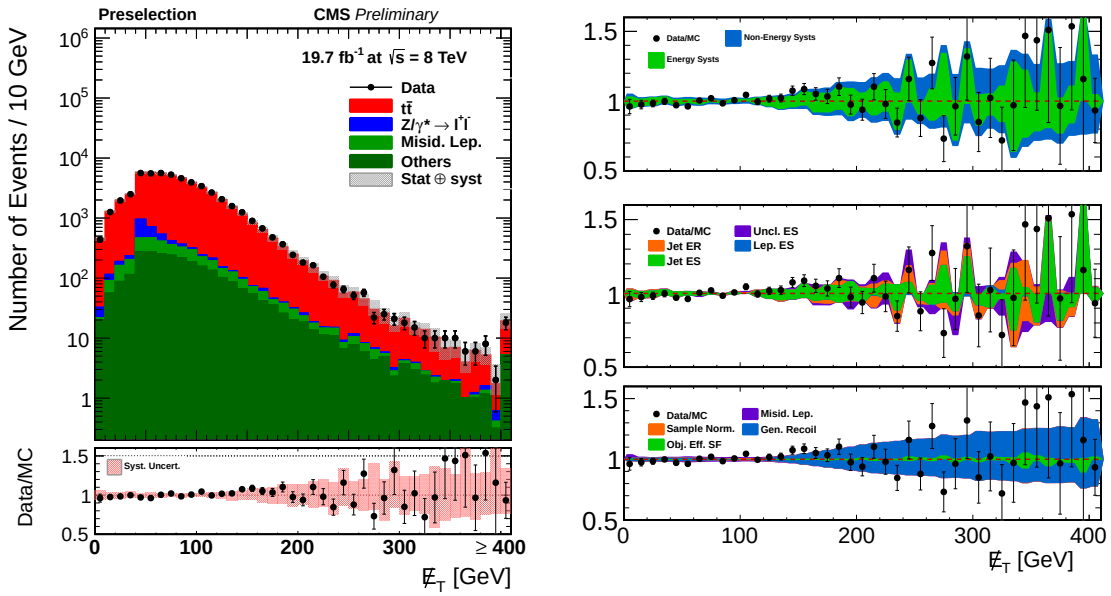

(a) (b)

Figure 6.4: The $\not{E}_T$ distribution in the sample of events passing the basic preselection. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

apply a few additional requirements when the selected leptons in the lepton pair are the same flavor. Specifically, we veto the event if the invariant mass of the leptons, $M_{\ell^+\ell^-}$, falls inside the Z-veto region, $|m_Z - M_{\ell^+\ell^-}| < 25$ GeV, or if the reconstructed $\not{E}_T$ is less than 40 GeV.

Table 6.3 summarizes these requirements, which define our analysis preselection.

Figure 6.4a shows the distribution of $\not{E}_T$ at the preselection level. It provides two forms of comparison between data and simulation.

The top panel shows the $\not{E}_T$ distribution in the the measured data (black circles) as well as the expected $\not{E}_T$ distribution from each individual background (colored histograms). These individual backgrounds are discussed in more detail

in Section 6.3, but to clarify which colored histograms refer to which, the $t\bar{t}$ refers to the expected contribution from direct production of $t\bar{t}$ pairs; the $Z/\gamma^* \to \ell^+\ell^-$ refers to the expected contribution from the dileptonic Drell-Yan background; The "Misid. Lep." refers to the expected contribution from events with non-prompt or misreconstructed leptons; finally, "Others" refers to other backgrounds considered in this analysis.

The bottom panel shows a direct ratio (black circles) between the individual histogram bins in the measured data and the total sum of the expected background contributions. The statistical uncertainties on this ratio are shown as the error bars on the individual points, while the overall expected impact of the simulation's systematic uncertainties (c.f. Section 6.4) on this ratio is shown as the red band.

The second sub-figure, Fig. 6.4b, shows a itemized breakdowns of the expected impact of specific sources of systematic uncertainties on the aforementioned ratio between data and simulation.

In this sub-figure, the top panel compares the relative impact of systematic uncertainties involving individual object energy scales and energy resolutions ("Energy Systs") and all other systematic uncertainties ("Non-Energy Systs").

The middle panel provides a comparison between the various individual components of the "Energy Systs" band in the top panel; this includes: Jet energy scale ("Jet ES"); the impact of Jet energy resolution smearing on the $\not{E}_{\mathrm{T}}$ ("Jet ER"); the energy scale of low-p_{T} unclustered objects ("Uncl. ES"); and the energy scale of the reconstructed lepton candidates ("Lep. ES").

The bottom panel provides a comparison between the systematic uncertainty

on the data-driven estimate of the "misidentified lepton" background ("Misid. Lep."); the reweightings applied to the simulation based on generator-level object p_{T} ("Gen. Recoil"); the normalization of several backgrounds ("Sample Norm."); and the individual object efficiency (e.g. b-tagging, lepton reconstruction, etc.) scale factors that are applied to the simulation ("Obj Eff. SF").

From the top panel of Fig. 6.4a, it is clear that the production of dileptonic $\mathrm{t\bar{t}}$ pairs comprise the vast majority of events at the preselection level. Moreover, comparisons between simulation and data in the bottom panel of Fig. 6.4a show that the overall modeling of $\not{E}_{\mathrm{T}}$ is quite good at the preselection level, as data and simulation easily agree within the systematic uncertainties on the simulation.

From the top panel of Fig. 6.4b, it is clear that the object-energy systematics dominate the overall systematic uncertainty on the modeling of $\not{E}_{\mathrm{T}}$ in the simulation; within this subset of systematic uncertainties, the middle panel of the same figure shows that the systematic uncertainty on the energy scale of reconstructed jets has the largest impact on $\not{E}_{\mathrm{T}}$ reconstruction. Finally, for the non-object-energy systematics, the bottom panel shows that at high-$\not{E}_{\mathrm{T}}$, the systematic uncertainty on the generator-level recoil reweighting clearly dominates the subset of non-object-energy systematics.

6.2.4 Reweighting the Simulation

We apply a number of corrections to the simulation in order to improve its accuracy in modeling the collision data.

6.2.4.1 Lepton Reconstruction and Trigger Efficiences

The CMS simulation tends to model the reconstruction, identification, and isolation of lepton candidates with fairly reasonable accuracy. Nevertheless, analyses often find some degree of disagreement between data and the predictions made by simulation.

In this section we detail how we calculate scale factors to correct the lepton efficiencies — the relative rate at which leptons pass trigger, identification, and isolation requirements — in the simulation to match those observed in data.

Lepton Trigger Efficiencies

The general principle for estimating the dilepton trigger efficiency is to identify a cross-trigger that is both weakly-correlated with the dilepton trigger and has sufficient statistics. After applying the basic dilepton selection with no trigger requirement, one can then estimate the dilepton trigger efficiency by calculating the ratio of events passing both triggers, $N^{\ell^+\ell^-}_{\ell\ell\text{-trig.},\,\text{x-trig.}}$, divided by the number of events passing just the cross-trigger, $N^{\ell^+\ell^-}_{\text{x-trig.}}$,

$$\varepsilon_{\text{trig.}} = \frac{N^{\ell^+\ell^-}_{\ell\ell\text{-trig.},\,\text{x-trig.}}}{N^{\ell^+\ell^-}_{\text{x-trig.}}} \tag{6.3}$$

As noted above, the cross-trigger should both have sufficient statistics and be weakly-correlated with the dilepton triggers. So, for example, using low p_{T} jet triggers would not be wise, as the reconstructed L1/HLT electron candidates could also easily be the L1/HLT "jets" triggering the acceptance for the jet triggers.

We chose $\not{E}_{\mathrm{T}}$-based HLT paths for the cross-triggers, as the associated $\not{E}_{\mathrm{T}}$-triggered datasets were found to have enough events in order to maintain a statistical uncertainty below 1%. As well, they were found to have a low statistical correlation with the dilepton triggers. The correlation was evaluated using simulated $\mathrm{t\bar{t}}$ events by calculating the parameter α,

$$\alpha = \frac{\varepsilon^{\mathrm{MC}}_{\ell\ell-\mathrm{trig.}} \times \varepsilon^{\mathrm{MC}}_{\not{E}_{\mathrm{T}}-\mathrm{trig.}}}{\varepsilon^{\mathrm{MC}}_{\ell\ell-\mathrm{trig.},\not{E}_{\mathrm{T}}-\mathrm{trig.}}} \tag{6.4}$$

where $\alpha = 1$ would indicate no correlation, $\alpha > 1$ would indicate anti-correlation, and $\alpha < 1$ would indicate correlation between the two triggers. We found that on average, α was 1.006, 0.999, and 0.995, for the $\mathrm{e^+e^-}$, $\mu^+\mu^-$, and $\mathrm{e}^{\pm}\mu^{\mp}$ channels respectively. Note that all three measurements were found to be consistent with 1 within the statistical precision of the measurements.

We calculated the efficiencies, i.e. Eq. (6.3), for both the collision data and MC simulation; we then derived p_{T}- and η-dependent scale factors to apply to the MC by taking the ratio of these efficiencies, $\mathrm{SF}_{\mathrm{trig.}} = \varepsilon^{\mathrm{data}}_{\mathrm{trig.}}/\varepsilon^{\mathrm{MC}}_{\mathrm{trig.}}$.

In addition to the statistical uncertainty of the calculated scale factors, we also assigned an additional, conservative systematic uncertainty of 1% to account for possible biases, such as the (weak) correlations between the two sets of triggers or any differing kinematics between the events used to derive the scale factors ($\mathrm{t\bar{t}}$) and events where they are applied ($\mathrm{Z}/\gamma^* \to \ell^+\ell^-$, EWK boson events).

Note that since the signal simulation samples were run using the CMS FastSim, as opposed to the FullSim used for the main background simulation samples, we calculated separate efficiencies and subsequent scale factors for signal simulation

samples.

The trigger efficiency scale factors for the FullSim (FastSim) samples range from 0.94 to 1.00 (0.98 to 1.10) depending on the specific channel as well as the lepton candidates' η. The net uncertainties on these scale factors are typically around 0.013, but for electrons in the endcap, whether in the e^+e^- channel or $e^\pm\mu^\mp$ channel, these uncertainties are higher, with values around 0.023 - 0.030.

Lepton Identification and Isolation Efficiencies

We utilized single lepton-triggered collision data samples (Section 6.2.1) to measure the lepton identification and isolation efficiencies. This choice was made in order to minimize any correlations with the measurement of the dilepton trigger efficiency. We utilized a standard tag-and-probe method[1] in order to estimate the lepton identification and isolation efficiencies. In addition to the relative strong HLT requirement on the HLT lepton p_{T}, we also require the reconstructed "tag" lepton to be associated with the HLT lepton that seeded the acceptance of the HLT path.

The identification efficiency corresponds to the number of probe lepton candidates passing the complete lepton selection criteria (Section 6.2.2 or Tables A.1 and A.2), **except** for the isolation requirement, divided by the total number of probe lepton candidates. Similarly, the isolation efficiency is defined as the ratio of the number of probe lepton candidates passing the complete lepton selection criteria including the isolation requirement divided by the number of probe lepton candidates

[1]See Section 6.1 of Ref. [151] for a concise pedagogical discussion on the tag-and-probe technique.

passing the complete lepton selection criteria without any requirement on lepton isolation.

We estimated systematic uncertainties for the identification and isolation efficiencies by varying the invariant mass window as well as the tag lepton selection requirements and then recalculating the scale factors after reapplying the tag-and-probe method. The largest variation of the scale factors w.r.t. their calculated central values was found to be $\approx 0.5\%$. We thus assigned a conservative systematic uncertainty of 1% — in addition to the intrinsic statistical uncertainty of the method — in order to cover this 0.5% variation as well as account for any additional possible systematic biases such as notable differences between the Z events from which the scale factors were derived and the primarily $\mathrm{t\bar{t}}$ events where the scale factors are applied.

As with the lepton trigger efficiencies, we calculated separate scale factors for the FullSim and FastSim simulation samples.

The muon identification and isolation efficiency scale factors for the FullSim (FastSim) samples range from 0.97 to 1.00 (0.96 to 0.98) depending on the muon candidate's p_{T} and $|\eta|$.

For the FullSim samples, the electron identification and isolation efficiency scale factors are typically around 0.96 to 0.98, although these scale factors dip down to values around 0.93 for low-p_{T} electrons $p_{\mathrm{T}} < 30\,\mathrm{GeV}$ around $|\eta| \approx 1.75$.

For the electron candidates in the FastSim samples, these scale factors are typically around 0.93 - 0.96, but for electrons with $|\eta| \approx 1.75$, these scale factors range from 0.88 ($p_{\mathrm{T}} > 50\,\mathrm{GeV}$) to 0.75 $p_{\mathrm{T}} < 30\,\mathrm{GeV}$)

For both lepton candidate types, in both the FullSim and FastSim samples, the uncertainty on these scale factors is dominated by the systematic uncertainty of 1%, as the statistical uncertainty on these scale factors is miniscule.

6.2.4.2 b-tagging Efficiencies

We apply corrective scale factors (Section 5.6.2.1) to improve the simulation's modeling of both the light-flavor mistag rate and the efficiency for tagging heavy-flavor decays.

As a reminder, the average (with respect to jet p_{T} and jet η) scale factor to correct the light flavor mis-tag rate (heavy flavor tagging efficiency) is approximately 1.15 (0.95), with a total systematic uncertainty on the scale factor of approximately 0.10 (0.02).

6.2.4.3 Smearing of the $\vec{\not{E}}_{\mathrm{T}}$ Resolution in Simulation

As noted in Section 6.2.2 (as well as Section 5.7.1.4), in order for the simulation to properly model the $\vec{\not{E}}_{\mathrm{T}}$ energy resolution, we have to apply corrective smearing in the simulation. In this section, we provide additional details on that smearing.

Using a separate set of simulation samples with representative event topologies — $\mathrm{t\bar{t}}$, DY + jets, etc. — where the individual jets in each event have been smeared (by the procedure described in Section 5.7.1.4) and subsequently propagated back into the $\vec{\not{E}}_{\mathrm{T}}$ calculation, we derived template smearing functions to correct the $\vec{\not{E}}_{\mathrm{T}}$ magnitude and direction in our simulation samples to better match that expected

in data. Figure 6.5 shows examples of these template distributions. The template distributions shown are 2D distributions of the observed difference between the smeared and unsmeared $\vec{\not{E}}_{\mathrm{T}}$ magnitude and direction for a given unsmeared $\not{E}_{\mathrm{T}}$. For an event with a given unsmeared $\not{E}_{\mathrm{T}}$, these 2D templates are then projected into associated 1D templates, from which corrective factors for both the $\not{E}_{\mathrm{T}}$ and $\not{E}_{\mathrm{T}}\,\phi$ are stochastically sampled.

The impact of this smearing on the simulation's modeling of $\not{E}_{\mathrm{T}}$ is shown in Figure 6.6. It is clear that without the smearing, the agreement between data and simulation is rather poor. After smearing, the overall agreement is acceptable within systematic uncertainties. Using the hadronic recoil method [71], where the hadronic recoil, u_{T}, is estimated using the $\vec{\not{E}}_{\mathrm{T}}$ and the $\vec{p}_{\mathrm{T}}$ of the dilepton system, q_{T}, and then subsequently broken down into its components perpendicular and parallel to $\vec{q}_{\mathrm{T}}$, we can directly probe the impact of this smearing on the parallel and perpendicular components of the $\vec{\not{E}}_{\mathrm{T}}$. Figures 6.7 and 6.8 show this comparison. Again, it is clear from direct comparison of the before and after that the overall agreement post-smearing is acceptable within systematic uncertainties.

6.2.4.4 Pileup Reweighting

The number of pileup interactions per event affects the analysis in several ways. Pileup reduces the probability of identifying the correct primary vertex in the event. It worsens the energy resolution for the selected objects (especially jets and $\not{E}_{\mathrm{T}}$), and makes lepton identification more difficult by putting additional energy into

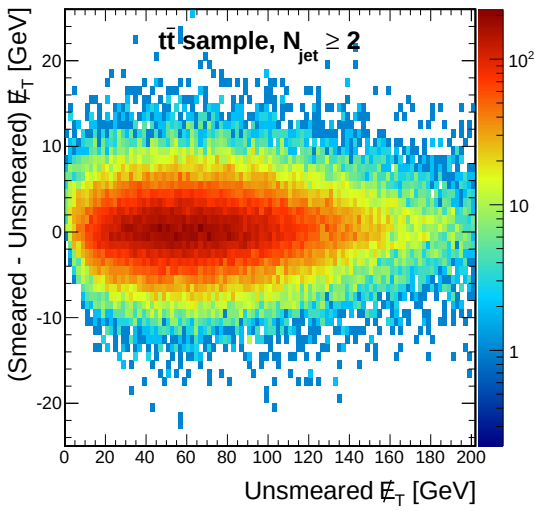

(a) The distribution of (Smeared $\not{E}_T$ - Unsmeared $\not{E}_T$) versus Unsmeared $\not{E}_T$ in a sample of dileptonic $t\bar{t}$ events.

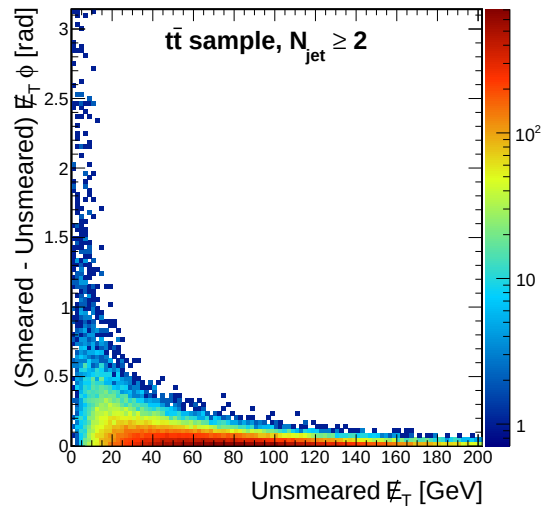

(b) The distribution of (Smeared $\not{E}_T$ ϕ - Unsmeared $\not{E}_T$ ϕ) versus Unsmeared $\not{E}_T$ in a sample of dileptonic $t\bar{t}$ events.

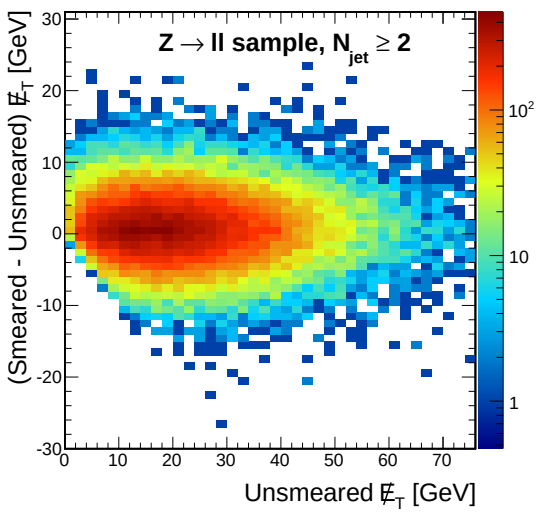

(c) The distribution of (Smeared $\not{E}_T$- Unsmeared $\not{E}_T$) versus Unsmeared $\not{E}_T$ in a sample of $Z \rightarrow \ell^+\ell^-$ events.

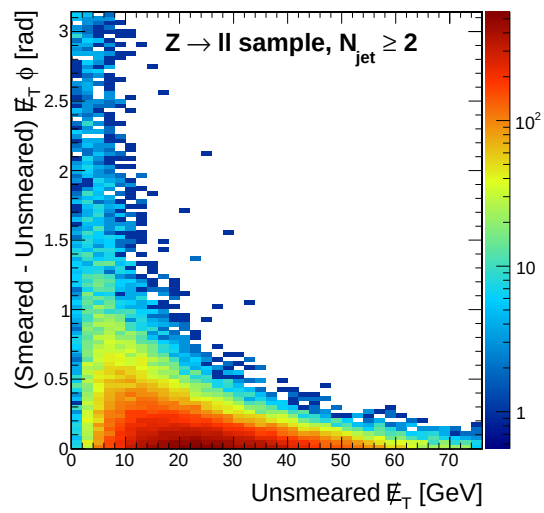

(d) The distribution of (Smeared $\not{E}_T$ ϕ - Unsmeared $\not{E}_T$ ϕ) versus Unsmeared $\not{E}_T$ in a sample of $Z \rightarrow \ell^+\ell^-$ events.

Figure 6.5: Example template distributions used for smearing the magnitude and direction of the $\vec{\not{E}}_T$ in simulation. For a given value of the initial reconstructed $\not{E}_T$ in a simulation event, the corrections to the calculated $\not{E}_T$ and $\not{E}_T$ ϕ are stochastically sampled from 1D projections of the corresponding 2D template distributions.

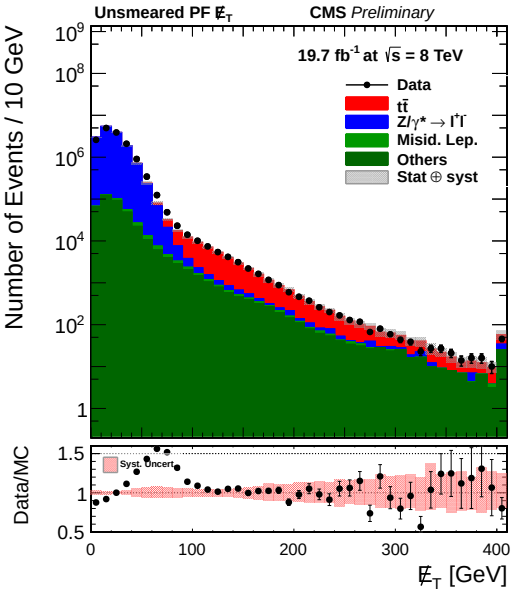

(a) Distribution of $\not\!\!E_T$ without jet energy resolution smearing.

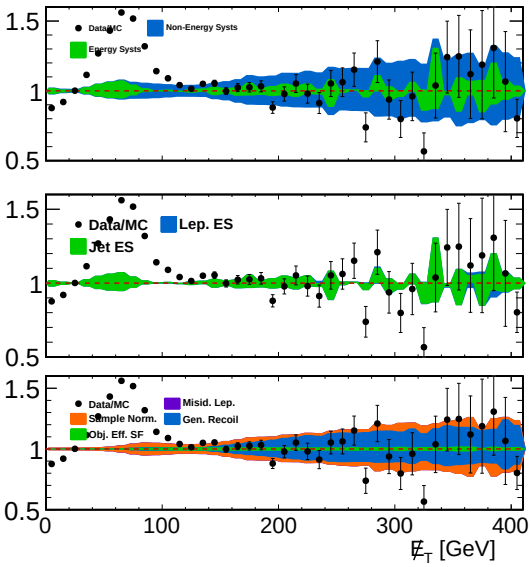

(b) Systematic uncertainties on the overall level of agreement between data and simulation

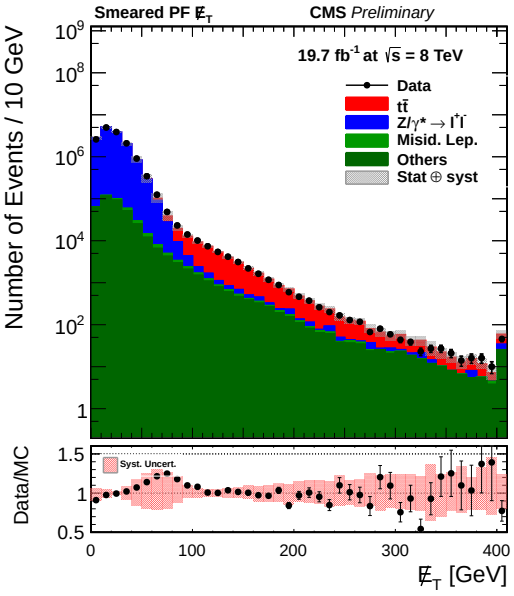

(c) Distribution of $\not\!\!E_T$ with jet energy resolution smearing.

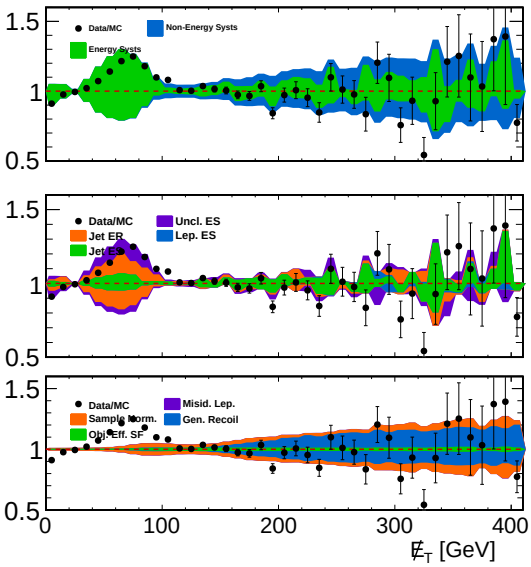

(d) Systematic uncertainties on the overall level of agreement between data and simulation

Figure 6.6: Distributions of $\not\!\!E_T$ with and without jet energy resolution smearing in a sample of events passing just the basic dilepton selection. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

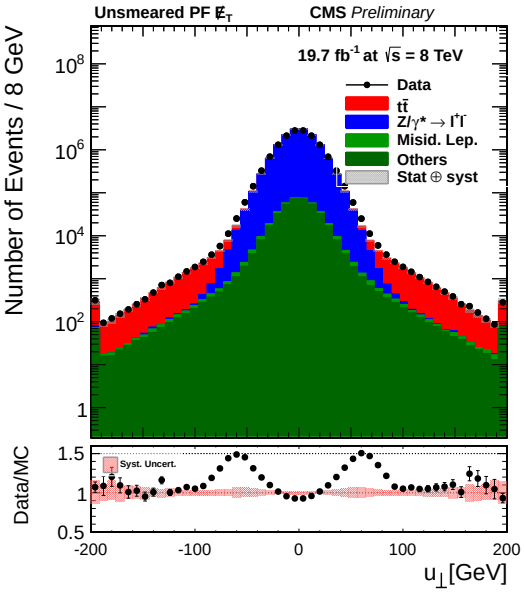

(a) Distribution of $u_\perp$ without jet energy resolution smearing.

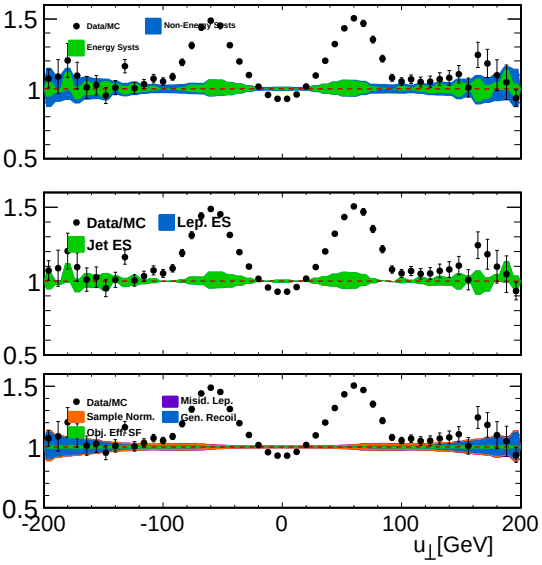

(b) Systematic uncertainties on the overall level of agreement between data and simulation

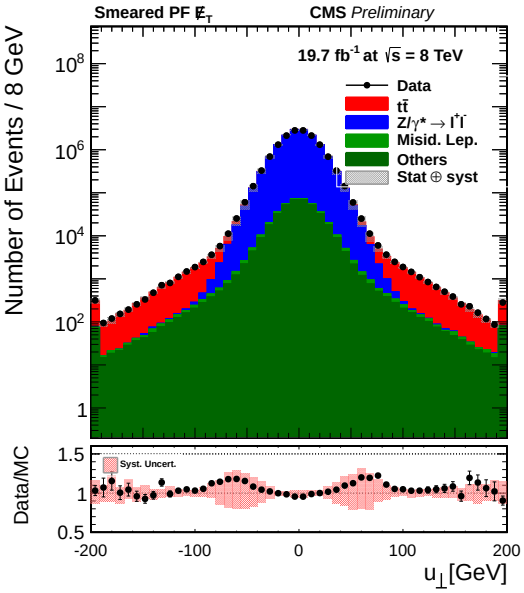

(c) Distribution of $u_\perp$ with jet energy resolution smearing.

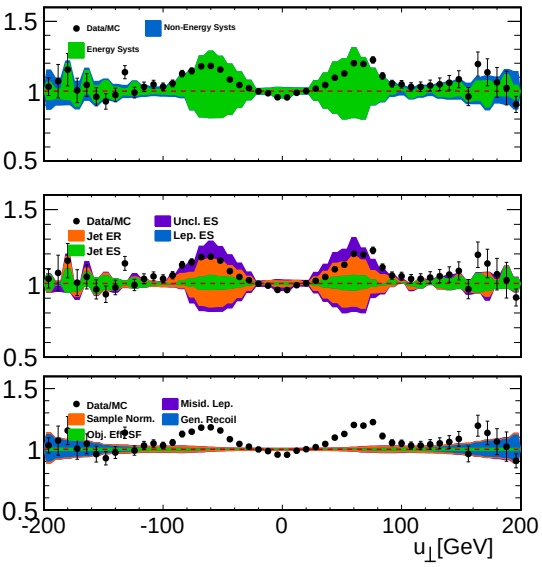

(d) Systematic uncertainties on the overall level of agreement between data and simulation

Figure 6.7: Distributions of the variable, $u_\perp$, with and without jet energy resolution smearing in a sample of events passing just the basic dilepton selection. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

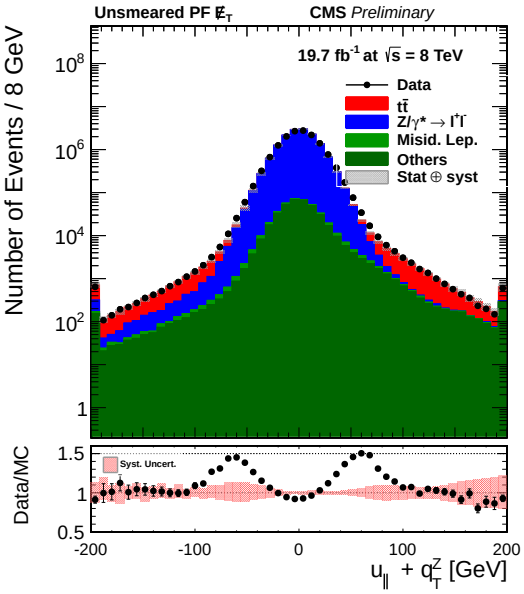

(a) Distribution of $u_{\|}+q_{\mathrm{T}}$ without jet energy resolution smearing.

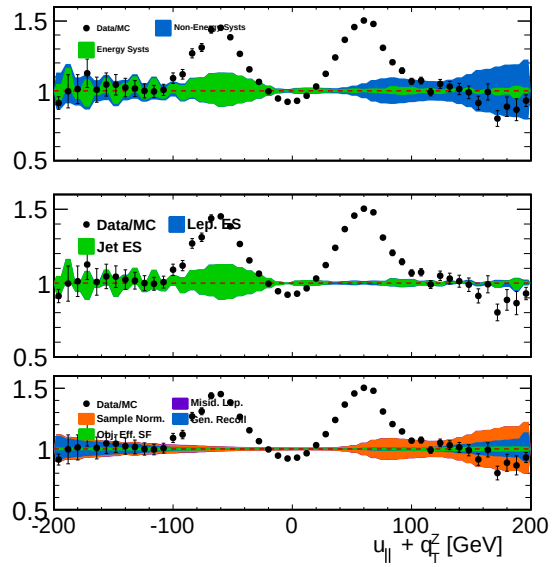

(b) Systematic uncertainties on the overall level of agreement between data and simulation

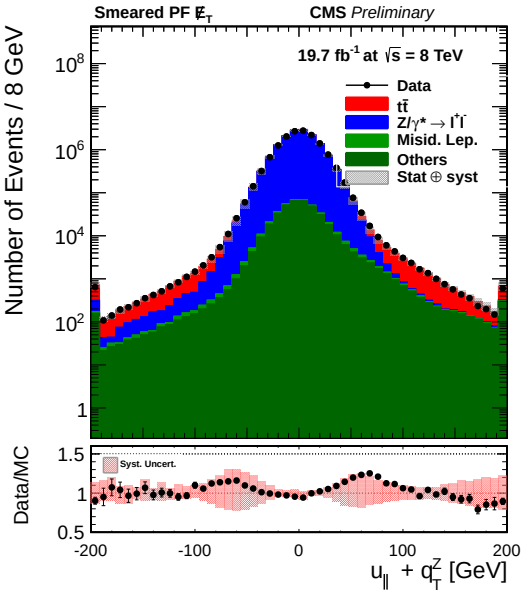

(c) Distribution of $u_{\|} + q_{\mathrm{T}}$ with jet energy resolution smearing.

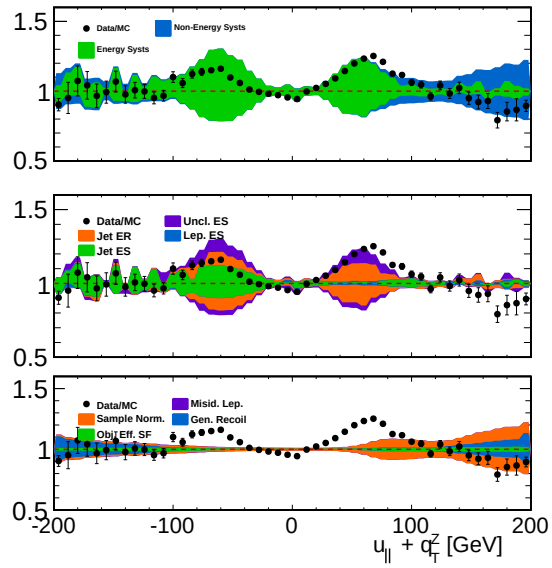

(d) Systematic uncertainties on the overall level of agreement between data and simulation

Figure 6.8: Distributions of the variable, $u_{\|}+q_{\mathrm{T}}$, with and without jet energy resolution smearing in a sample of events passing just the basic dilepton selection. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

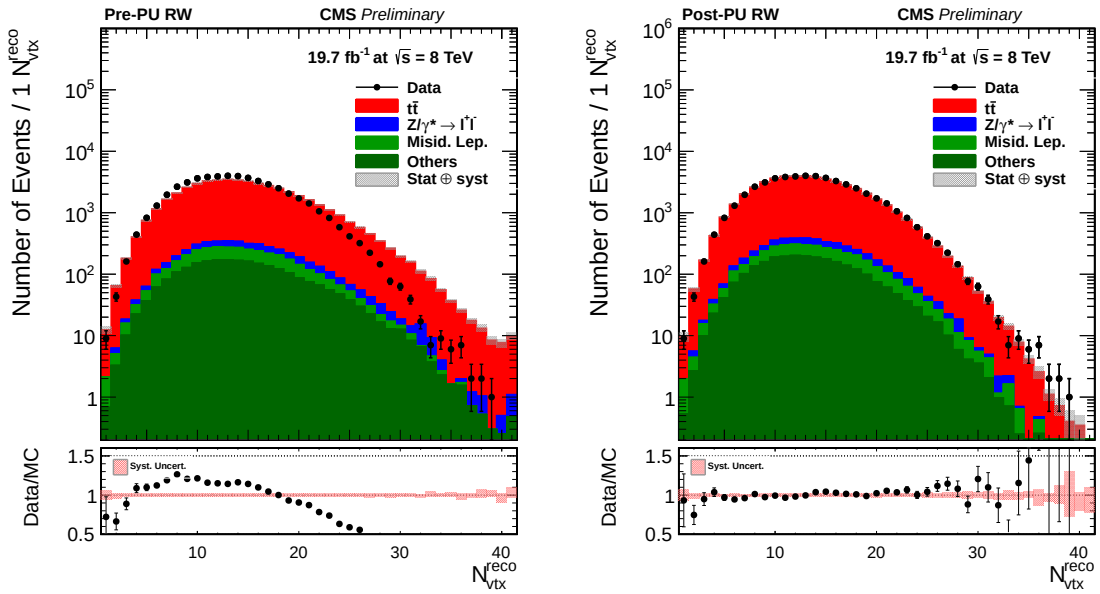

(a) Distribution of N_{vtx}^{reco} before reweighting the simulation.

(b) Distribution of N_{vtx}^{reco} after reweighting the simulation.

Figure 6.9: Distributions (at the full preselection level) of the number of reconstructed vertices, N_{vtx}^{reco}, before and after reweighting the simulation to match the expected distribution in data.

the isolation cones of lepton candidates. It has similar effects at the trigger level. For all these reasons, we reweight the simulation events to have the same pileup distribution as in data. We determine event weights using the Poisson mean for the true number of pileup vertices in the event. Figure 6.9 compares the distributions in data and simulation of N_{vtx}^{reco}, the number of reconstructed primary vertices in an event, at the preselection level. As you can see, the pileup reweighting leads to good agreement between data and simulation.

6.2.4.5 Top p_T Reweighting and Recoil Reweighting

Analyses at CMS that measured the $t\bar{t}$ production cross section [152, 153] found that, in order to have acceptable agreement between data and simulation,

most notably in the high-$p_{\mathrm{T}}^{\mathrm{t\bar{t}}}$ tail, the $p_{\mathrm{T}}^{\mathrm{t\bar{t}}}$ spectrum in simulation must be reweighted. For this analysis, we found that this reweighting was necessary in order to achieve acceptable data and simulation agreement.

A global fit to the semi-leptonic and dileptonic channels with $12.2\,\mathrm{fb}^{-1}$ of available data yielded the following reweighting formula that is calculated individually for both the t and $\bar{\mathrm{t}}$,

$$w(p_{\mathrm{T}}) = e^{0.156-0.00137p_{\mathrm{T}}}. \tag{6.5}$$

The event weight is the geometric average of the calculated t and $\bar{\mathrm{t}}$ weights for a given event. The overall effect of this reweighting is to decrease the number of high-$p_{\mathrm{T}}^{\mathrm{t\bar{t}}}$ events in the simulation. We note that this correction is rather small for typical $p_{\mathrm{T}}^{\mathrm{t\bar{t}}}$ values. The exact origin of this difference is still not known at this time. Some possible explanations are higher-order QCD corrections, EWK corrections, non-resonant production of $\mathrm{t\bar{t}}$-like final states, or something else entirely. Thus, in order to conservatively cover this mismodeling by the simulation, we also associate with this correction a 100% systematic uncertainty (i.e. equal in magnitude to the correction itself).

Drell-Yan Recoil Reweighting

Similar conclusions were found by an analysis [62, 154], comparing the high Z-boson p_{T} spectrum in data and simulation when using the MadGraph generator. A hypothesis for the source of this disagreement is a mismodeling of ISR by MadGraph. From this comparison, the authors of Refs. [62,154] derived corrective

shape scale factors: events where the generator-system p_{T} (e.g. the Z-boson p_{T} in DY events) is larger than 150 (250) GeV are scaled down by 10% (20%).

Because the top-squark simulation samples were also generated using MadGraph, analyses that utilized these samples were advised to apply this reweighting. We did not see this same disagreement when comparing data and simulation. Thus, we did not apply it to our DY background. However, to maintain consistency with other analyses searching for top-squarks, we chose to apply this shape reweighting to our top-squark simulation samples. Note that the initial state for top-squark pair-production events is different from DY events (gg versus qg). However, the authors of Refs. [62, 154] performed analogous studies of the modeling of high-p_{T} tails in $\mathrm{t\bar{t}}$ events and found that similar scale factors were needed to achieve agreement between data and simulation.

As with the $\mathrm{t\bar{t}}$-p_{T} reweighting, we associated a 100% systematic uncertainty to the correction because the exact source of the disagreement is still unknown. It should be noted that the effect of this shape reweighting is to scale down the relative contribution of top-squark events into our signal region. It thus leads to more statistically conservative final results.

6.2.5 Signal Region Selection and Analysis Strategy

We utilize the dileptonic stransverse mass $M_{\mathrm{T2}}(\ell\ell)$ to separate SM $\mathrm{t\bar{t}}$ production from our top-squark signal. We provide here a quick overview of how to construct $M_{\mathrm{T2}}(\ell\ell)$ and some important properties.

Under the assumption that we have two identical mother particles each decaying to a lepton and neutrino, we can define $M_{\mathrm{T2}}(\ell\ell)$ as,

$$M_{\mathrm{T2}}(\ell\ell) \equiv \min_{\vec{p}_{\mathrm{T}}^{\widehat{\nu}_1}+\vec{p}_{\mathrm{T}}^{\widehat{\nu}_2}=\vec{\not{E}}_{\mathrm{T}}} \left(\max\left[M_{\mathrm{T}}(\vec{p}_{\mathrm{T}}^{\ell_1}, \vec{p}_{\mathrm{T}}^{\widehat{\nu}_1}),\ M_{\mathrm{T}}(\vec{p}_{\mathrm{T}}^{\ell_2}, \vec{p}_{\mathrm{T}}^{\widehat{\nu}_2})\right]\right). \tag{6.6}$$

To construct $M_{\mathrm{T2}}(\ell\ell)$, we begin with the two selected lepton candidates, ℓ_1 and ℓ_2. Under the assumption that the $\vec{\not{E}}_{\mathrm{T}}$ is only comprised of contributions from the two neutrinos, we partition the $\vec{\not{E}}_{\mathrm{T}}$ into two hypothetical neutrinos, $\widehat{\nu}_1$, $\widehat{\nu}_2$ with transverse momenta $\vec{p}_{\mathrm{T}}^{\widehat{\nu}_1}$ and $\vec{p}_{\mathrm{T}}^{\widehat{\nu}_2}$, chosen so that their transverse vector sum equals the observed $\vec{\not{E}}_{\mathrm{T}}$, i.e. $\vec{p}_{\mathrm{T}}^{\widehat{\nu}_1} + \vec{p}_{\mathrm{T}}^{\widehat{\nu}_2} = \vec{\not{E}}_{\mathrm{T}}$. We construct lepton-neutrino pairs (i.e. ℓ_1 with $\widehat{\nu}_1$ and ℓ_2 with $\widehat{\nu}_2$), calculate the transverse mass M_{T} for both pairs, and record the **maximum** of these two M_{T}. We calculate this maximum M_{T} for other viable 2-neutrino partitions of the $\vec{\not{E}}_{\mathrm{T}}$, iterating through all viable partitions until we have found the *minimum* of these maximum M_{T}. This minimal maximum M_{T} is the event's $M_{\mathrm{T2}}(\ell\ell)$. The actual implementation of this calculation in the analysis code is slightly different than the above description, but achieves the same qualitative and quantiative results; see Ref. [155] for the technical details of these differences.

To understand why $M_{\mathrm{T2}}(\ell\ell)$ can be a useful variable, consider a dileptonic $\mathrm{t\bar{t}}$ event. One of the viable 2-neutrino partitions of the $\vec{\not{E}}_{\mathrm{T}}$ in this event is the "true" partition. In the true partition, each of the hypothetical neutrino $\vec{p}_{\mathrm{T}}$ are correctly assigned, both in terms of the magnitude $p_{\mathrm{T}}^{\widehat{\nu}}$ as well as the associated ℓ. Thus, for this true partition, the maximum M_{T} of the two possible M_{T} values will, by construction, have an upper bound at m_{W}. The minimum maximum M_{T} across all viable partitions — i.e. $M_{\mathrm{T2}}(\ell\ell)$— consequently also has an upper bound at m_{W}.

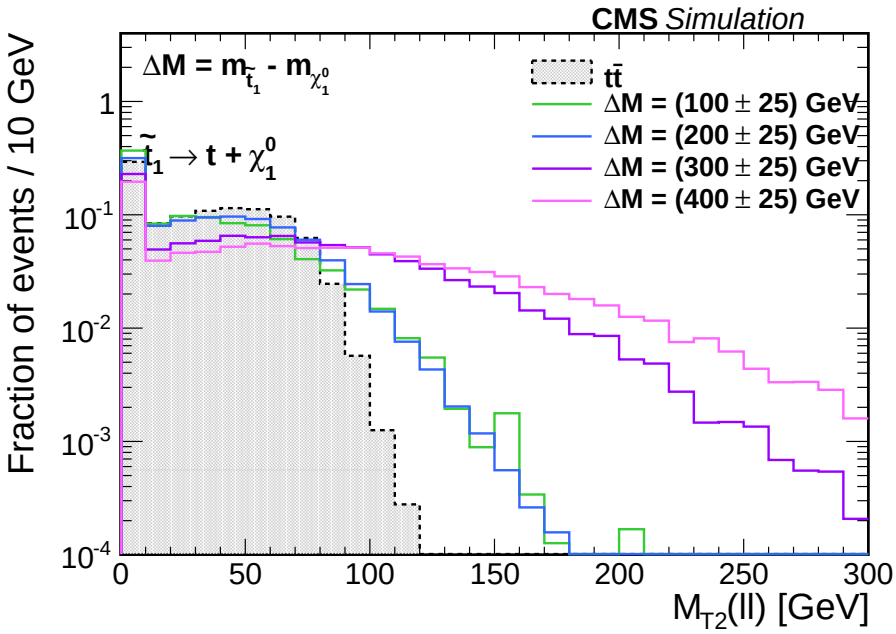

Figure 6.10: $M_{\mathrm{T2}}(\ell\ell)$ distribution for the $\mathrm{t\bar{t}}$ background and different signal mass points of the T2tt decay-mode regrouped in constant ΔM bands; distributions are normalized to the same area.

To constrast this, for a dileptonic top-squark pair event, discussing a "true" partition is no longer sensical, as the additional invisible particles (e.g. the $\widetilde{\chi}^0_1$) break the assumption that the observed $\vec{\not{E}}_{\mathrm{T}}$ only stems from 2 neutrinos. Constructing $M_{\mathrm{T2}}(\ell\ell)$ as per the method discussed above is still feasible, but it is no longer guaranteed that the upper bound of $M_{\mathrm{T2}}(\ell\ell)$ will be at m_{W}.

In particular, the additional $\vec{\not{E}}_{\mathrm{T}}$ coming from the $\widetilde{\chi}^0_1$ tends to lead to larger values of maximum M_{T} on average, and thus larger $M_{\mathrm{T2}}(\ell\ell)$. The distribution of $M_{\mathrm{T2}}(\ell\ell)$ in top-squark pair events is particularly sensitive to the p_{T} distribution of the $\widetilde{\chi}^0_1$ produced in the event as these are what directly add to the $\vec{\not{E}}_{\mathrm{T}}$. These $\widetilde{\chi}^0_1$ momentum distributions themselves are dependent upon the overall mass-splitting between their mother particle and themselves; as an example, for the T2tt decay mode, this is $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}^0_1}$. Figure 6.10 compares the normalized $M_{\mathrm{T2}}(\ell\ell)$

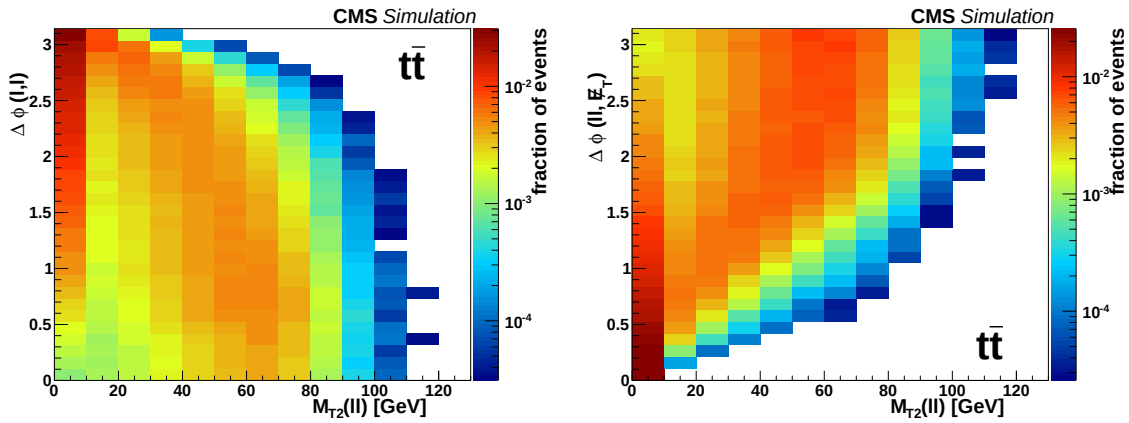

(a) $M_{\mathrm{T2}}(\ell\ell)$ vs. $\Delta\phi(\ell^+, \ell^-)$.

(b) $M_{\mathrm{T2}}(\ell\ell)$ vs. $\Delta\phi(\ell^+\ell^-, \not{E}_{\mathrm{T}})$.

Figure 6.11: Normalized distributions for the $t\bar{t}$ background comparing the dependence of $M_{\mathrm{T2}}(\ell\ell)$ on the relative angular correlations of its input objects.

distribution in $t\bar{t}$ events against the $M_{\mathrm{T2}}(\ell\ell)$ distribution for various ΔM in the T2tt decay mode. For larger ΔM, the tail of the $M_{\mathrm{T2}}(\ell\ell)$ distribution extends much further beyond m_{W}. The strong kinematic endpoint at m_{W} in the $t\bar{t}$ background dictates the primary demarcation between the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, and our general signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$.

It is important to also discuss the properties of the lower bound of $M_{\mathrm{T2}}(\ell\ell)$. Solutions where the maximum M_{T} is 0 for a given 2-neutrino $\vec{\not{E}}_{\mathrm{T}}$ partition will be kept if found, thus yielding $M_{\mathrm{T2}}(\ell\ell) = 0$. A particular class of events where this will always be the case is when the $\vec{\not{E}}_{\mathrm{T}}$ falls within the smaller of the two transverse plane opening angles defined by the two lepton $\vec{p}_{\mathrm{T}}$. For these events, it is always possible to partition $\vec{\not{E}}_{\mathrm{T}}$ into two hypothetical neutrino $\vec{p}_{\mathrm{T}}$ that lie along each respective lepton $\vec{p}_{\mathrm{T}}$, subsequently yielding $M_{\mathrm{T}} = 0$ for each lepton-neutrino pairing, and consequently $M_{\mathrm{T2}}(\ell\ell) = 0$. In Fig. 6.10, the large spike in the distribution of events near $M_{\mathrm{T2}}(\ell\ell) = 0$ is comprised primarily of this class of events.

The shape of $M_{\mathrm{T2}}(\ell\ell)$ is thus highly depedent upon the relative angular correlations between its input objects. This can be clearly seen in Fig. 6.11, which compares the dependence of $M_{\mathrm{T2}}(\ell\ell)$ on the relative angular correlations of its input objects using normalized distributions for the $\mathrm{t\bar{t}}$ background. When the opening angle between the two input leptons (Fig. 6.11a) is large, it is far more likely that the $\vec{\not{E}}_{\mathrm{T}}$ falls within this opening angle, yielding $M_{\mathrm{T2}}(\ell\ell) = 0$ as per the argument above. An analogous situation arises when the opening angle between the $\vec{\not{E}}_{\mathrm{T}}$ and the dilepton system (Fig. 6.11b) is small.

6.3 Backgrounds

The object and event selection described in Section 6.2.3 rejects most SM backgrounds; only final states that contain two high-p_{T}, opposite charge lepton candidates along with two or more jets, at least one of which must be b-tagged, contribute to the background.

For the $M_{\mathrm{T2}}(\ell\ell)$ signal regions that we use, the notable backgrounds are $\mathrm{t\bar{t}}$; Drell-Yan (DY); "misidentified lepton" events — single-lepton events with an non-prompt lepton (e.g. semi-leptonic $\mathrm{t\bar{t}}$ or W+jets); and other, less common, processes. These less common processes include: single top quarks produced in association with a W boson (tW); WW, WZ, and ZZ production (VV); W or Z production with an associated photon (VG); and "Rare" backgrounds, which include triple EWK vector boson production and $\mathrm{t\bar{t}}$ production in association with one or two EWK bosons.

The normalization of the Drell-Yan and $\mathrm{t\bar{t}}$, and the shape and normalization

of the misidentified lepton backgrounds are evaluated from data using data-driven estimation methods applied to control samples (c.f. Sections 6.3.1, 6.3.2 and 6.3.3).

The shapes of the Drell-Yan and $t\bar{t}$, and the normalization and shapes of the less common processes are all estimated from the simulation.

When displaying the expected contributions of these individual backgrounds, the $t\bar{t}$, DY, and misidentified lepton backgrounds are individually shown, while other backgrounds are grouped together with the name "Other".

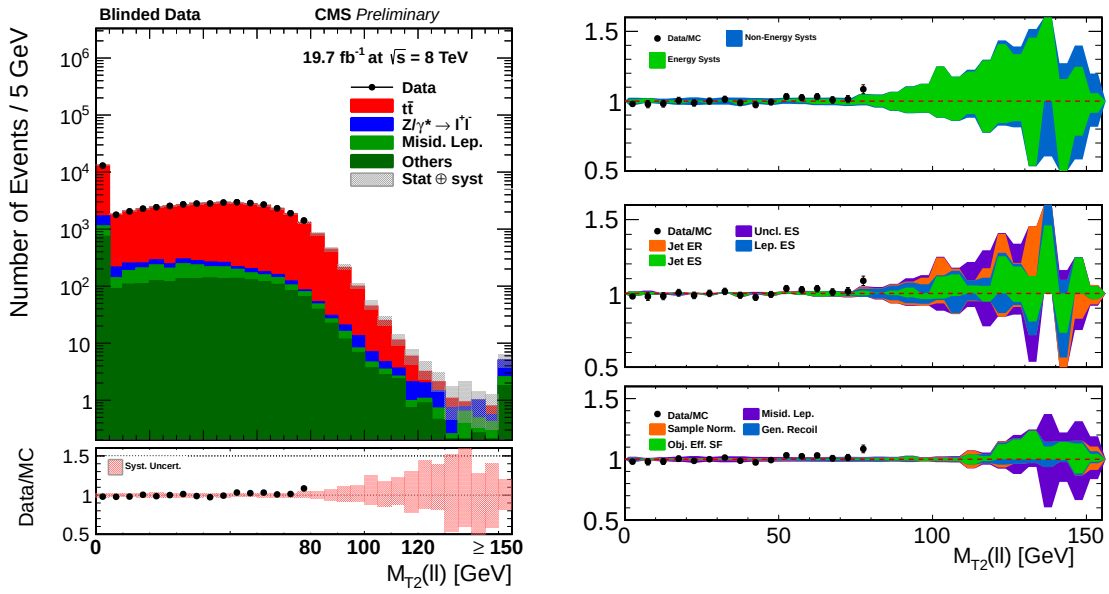

(a) (b)

Figure 6.12: The $M_{\mathrm{T2}}(\ell\ell)$ distribution in the sample of events passing the basic preselection where events in data with $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$ have been blinded (removed); as can be clearly seen, $t\bar{t}$ events comprise the large majority of events in the control region $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, enabling a data-driven normalization of the $t\bar{t}$ contribution (c.f. Section 6.3.1). (right) An itemized breakdown of the contributions from individual sources of systematic uncertainty

6.3.1 $t\bar{t}$ Estimation

Figure 6.12 shows the $M_{\mathrm{T2}}(\ell\ell)$ distribution at the preselection level, where we have blinded the signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$. As stated earlier, dileptonic SM $t\bar{t}$ events comprise a large majority (over 90%) of the events in our control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$. We can thus use this $M_{\mathrm{T2}}(\ell\ell)$ control region to set the normalization of our expected SM $t\bar{t}$ contribution in the $M_{\mathrm{T2}}(\ell\ell)$ signal region. We first count the number of events in the $M_{\mathrm{T2}}(\ell\ell)$ control region in data and subtract the expected contributions of all non-$t\bar{t}$ backgrounds (including the data-driven background estimations for the DY and misidentified lepton backgrounds)

We derive a normalization scale factor by comparing our predicted $t\bar{t}$ contribution in the $M_{T2}(\ell\ell)$ control region with this modified collision data count,

$$\mathrm{SF}_{\mathrm{DD}\,t\bar{t}} = \frac{N^{\mathrm{ctrl}}_{\mathrm{Data}} - N^{\mathrm{ctrl}}_{\mathrm{Non\text{-}t\bar{t}\ Bkg.}}}{N^{\mathrm{ctrl}}_{t\bar{t}\ \mathrm{MC}}}. \tag{6.7}$$

The total number of $t\bar{t}$ events in the signal region is then,

$$N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{DD}} = \mathrm{SF}_{\mathrm{DD}\,t\bar{t}} N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{MC}}. \tag{6.8}$$

The calculated scale factors, as well as the estimated $t\bar{t}$ contribution in the control region, $M_{T2}(\ell\ell) < 80\,\mathrm{GeV}$, shown for the individual channels as well as their weighted average, are contained in Table 6.4. As can be seen from this table, the calculated scale factor for the e^+e^- channel is different from the other channels at the level of approximately 2–3σ. We do not have a direct explanation for this, but cross-checks, such as comparisons of the $M_{T2}(\ell\ell)$ shape in each of the individual dilepton channels, which is shown in Fig. 7.18, did not indicate any clear systematic biases in the simulation or backgrounds that were not accounted for.

This method implicitly assumes that the $M_{T2}(\ell\ell)$ shape in the simulation is well-modeled. Section 6.5 contains details on the studies we performed to to confirm this assumption.

6.3.1.1 Systematic Uncertainties in the $t\bar{t}$ Estimation

Systematic uncertainties like the jet energy scale calibration uncertainty affect both the expected number of $t\bar{t}$ events passing the final preselection (i.e. the overall normalization of the $t\bar{t}$ contribution to the $M_{T2}(\ell\ell)$ distribution) as well as their overall shape (i.e. the relative fraction of $t\bar{t}$ events falling into the region,

Table 6.4: Results of the calculation of (top) the $t\bar{t}$ data-driven normalization scale factor and (bottom) the estimated $t\bar{t}$ contribution to the control region, $M_{T2}(\ell\ell) < 80\,\text{GeV}$, in a sample of events passing the full preselection. Also shown are the statistical and systematic uncertainties on these calculated quantities, using absolute uncertainty magnitude for the normalization scale factor, and relative (%) uncertainty magnitude for the estimated $t\bar{t}$ contribution. The scale factor and estimated yield for $t\bar{t}$ were calculated using events that fell into the control region, $M_{T2}(\ell\ell) < 80\,\text{GeV}$.

Event Flavor	$t\bar{t}$ DD Scale Factor
	$M_{T2}(\ell\ell) < 80\,\text{GeV}$
$\mu^+\mu^-$	0.98 ± 0.01 (stat.) $^{+0.08}_{-0.07}$ (syst.)
e^+e^-	1.13 ± 0.02 (stat.) $^{+0.08}_{-0.07}$ (syst.)
$e^\pm\mu^\mp$	1.02 ± 0.01 (stat.) $^{+0.06}_{-0.06}$ (syst.)
Weighted Ave.	1.02 ± 0.01 (stat.) $^{+0.07}_{-0.06}$ (syst.)

Event Flavor	$t\bar{t}$ Estimated Yield
	$M_{T2}(\ell\ell) < 80\,\text{GeV}$
$\mu^+\mu^-$	$10{,}689.16 \pm 1.16\%$ (stat.) $^{+2.86\%}_{-2.86\%}$ (syst.)
e^+e^-	$6{,}937.14 \pm 1.43\%$ (stat.) $^{+0.78\%}_{-0.78\%}$ (syst.)
$e^\pm\mu^\mp$	$26{,}713.67 \pm 0.73\%$ (stat.) $^{+1.42\%}_{-1.42\%}$ (syst.)
Weighted Ave.	$44{,}340.41 \pm 0.57\%$ (stat.) $^{+1.65\%}_{-1.65\%}$ (syst.)

$M_{T2}(\ell\ell) < 80\,\text{GeV}$). When we calculate the impact of systematic uncertainties on the $t\bar{t}$ estimation by using versions of the simulation where we have systematically varied some component of the simulation (see Section 6.4), we remove the first of these two effects (the change in the normalization of the expected $t\bar{t}$ contribution to the control region) by recalculating this normalizing scale factor individually for all of the systematic variations of the simulation. We then add an additional systematic uncertainty to account for the statistical uncertainty of the normalization; this uncertainty naturally tends to be quite small as there are over 4×10^4 events

used to normalize the $t\bar{t}$ in the region, $M_{T2}(\ell\ell) < 80\,\text{GeV}$.

6.3.1.2 Signal Contamination in the $t\bar{t}$ Estimation

Although the extra $\vec{\not{E}}_T$ in top-squark events tends to push the $M_{T2}(\ell\ell)$ values past the kinematic edge at $M_{T2}(\ell\ell) = m_W$, it is clear from Fig. 6.10 that a non-trivial fraction of top-squark events, particularly for lower values of ΔM, can populate the $M_{T2}(\ell\ell)$ control region used for normalizing our expected $t\bar{t}$ contribution to the $M_{T2}(\ell\ell)$ signal region.

Under the hypothesis that the signal exists, the normalization described in Eq. (6.7) is technically incorrect, as it fails to account for the presence of the signal events in the data control region used. We account for this "signal contamination" effect when we compute the upper limits on the top-squark pair-production cross section in Section 6.8.

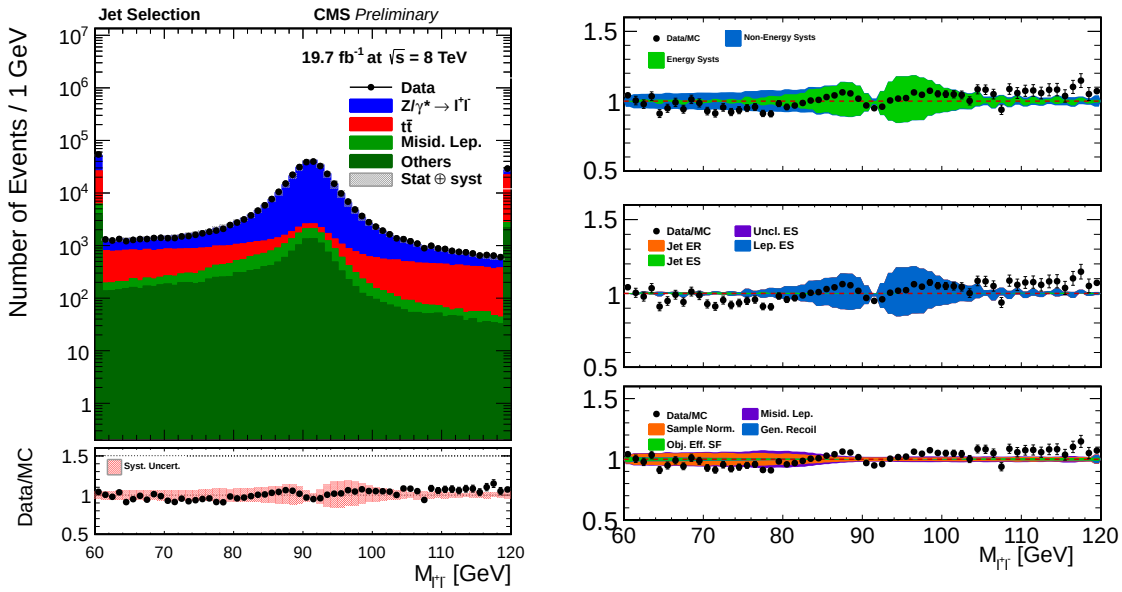

(a) (b)

Figure 6.13: (left) The distribution of the dilepton invariant mass for events passing a basic jet selection (2+ jets required). (right) An itemized breakdown of the contributions from individual sources of systematic uncertainty

6.3.2 Estimation of the Drell-Yan background

The analysis makes use of a data-driven DY estimation method described in Refs. [156, 157].

The presence of the well-defined Z-boson mass resonance in the $M_{\ell^+\ell^-}$ distribution for opposite charge, same-flavor dilepton events enables using collision data to measure the total contribution of DY events in our final selected sample of events.

We calculate a normalization scale factor for our simulated DY by comparing the expected yields in the Z-mass window,

$$\mathrm{SF}^{\ell^+\ell^-}_{\mathrm{DD\,DY}} = \frac{N^{\ell^+\ell^-}_{\mathrm{in}} - 0.5 N^{\mathrm{e}^\pm\mu^\mp}_{\mathrm{in}} k_{\ell^+\ell^-}}{N^{\mathrm{DY\ MC}}_{\mathrm{in}}}. \tag{6.9}$$

where we remove the contamination from non-DY processes, such as $\mathrm{t\bar{t}}$ or the top-squark signal, by subtracting the number of events with opposite-flavor ($N^{\mathrm{e}^\pm\mu^\mp}_{\mathrm{in}}$).

In Eq. (6.9), the factor of 0.5 accounts for the additional combinatorial choices in the $e^{\pm}\mu^{\mp}$ sample, while the k-factors in Eq. (6.9) account for the different reconstruction efficiencies of electrons and muons,

$$k_{e^+e^-} = \sqrt{\frac{N_{\text{in}}^{e^+e^-}}{N_{\text{in}}^{\mu^+\mu^-}}} \quad k_{\mu^+\mu^-} = \sqrt{\frac{N_{\text{in}}^{\mu^+\mu^-}}{N_{\text{in}}^{e^+e^-}}} \tag{6.10}$$

Assuming that the simulation accurately models the Z-mass line shape, the expected number of DY events that fall outside of the Z-mass window, $N_{\text{out}}^{\text{obs. DY}}$, can be calculated as a simple scaling of $N_{\text{out}}^{\text{DY MC}}$, the analogous number as estimated by the simulation, i.e.,

$$N_{\text{out}}^{\text{DY obs.}} = \text{SF}_{\text{DD DY}}^{\ell^+\ell^-} N_{\text{out}}^{\text{DY MC}}. \tag{6.11}$$

Figure 6.13 shows the distribution of the dilepton invariant mass for events passing a basic jet selection (2+ jets required). This figure demonstrates that the simulation models the Z-mass line shape with acceptable accuracy. The observed "wiggle" in the ratio of data and simulation around $M_{\ell^+\ell^-} = m_{\text{Z}}$ indicates a slight miscalibration of the lepton energy scales. This discrepancy primarily stems from electrons in the endcap, but we note that the discrepancy is easily covered within the systematic uncertainty on the lepton energy scale calibrations. To further validate the claim made above, a linear function to the ratio returned a value of 8.8×10^{-4} for the slope, implying little systematic bias, if any, in the modeling of the Z-mass line shape.

The aforementioned equations only directly apply for estimating the DY contribution in the same-flavor event samples. While DY processes also contribute to the opposite-flavor event sample, primarily through the leptonic tau decays in the process, $\text{Z} \to \tau\tau$, the relative loss of visible energy to neutrinos in these $\text{Z} \to \tau\tau$

events destroys the resonant peak in the $M_{\ell^+\ell^-}$ distribution.

In the opposite-flavor event sample, we calculate a scale-factor for the opposite-flavor DY MC events using the geometric mean of the two same-flavor scale factors,

$$\mathrm{SF}^{\mathrm{e}^\pm\mu^\mp}_{\mathrm{DD\,DY}} = \sqrt{\mathrm{SF}^{\mu^+\mu^-}_{\mathrm{DD\,DY}} \times \mathrm{SF}^{\mathrm{e}^+\mathrm{e}^-}_{\mathrm{DD\,DY}}}. \tag{6.12}$$

A qualitative justification for this expression is the following. Rewrite the scale factor for one of same flavor channels as the square of a single-lepton scale factor, i.e. $\mathrm{SF}^{\ell^+\ell^-}_{\mathrm{DD\,DY}} = \left(\mathrm{SF}^{\ell}_{\mathrm{DD\,DY}}\right)^2$. Making the association that each application of this single-lepton scale factor is correcting for one lepton at a time, it is natural then to write the DY scale factor for the $\mathrm{e}^\pm\mu^\mp$ channel as the product of one instance each of the single-lepton scale factors for the electron and the muon, which naturally leads to the expression in Eq. (6.12).

6.3.2.1 Data-driven DY estimation results

The calculated normalization scale factors, the measured ratio of events falling inside or outside of the Z-mass window in the simulation, and the estimated number of DY events in the Z-mass window are summarized in Table 6.5, separated by event flavor and for different stages of the selection. These calculations were done using the jet-exclusive $\mathrm{Z} \to \ell^+\ell^-$ simulation samples. We cross-checked these results by calculating the scale factors using the jet-inclusive $\mathrm{Z} \to \ell^+\ell^-$ simulation samples instead; this cross-check yielded similar results for the calculated scale factors.

Table 6.5: Results of the calculation of (top) the data-driven DY normalization scale factor and (bottom) the estimated DY contribution for events outside of the Z-mass window. These numbers were calculated using the jet-exclusive DY simulation samples, c.f. Section 6.2.1, and are shown at various levels of the event selection. The display format is central value ± statistical uncertainty $^{+\text{up syst.}}_{-\text{down syst.}}$, using absolute uncertainty magnitude for the normalization scale factor, and relative (%) uncertainty magnitude for the estimated DY contribution.

Event Flavor	DY DD Scale Factor			
	≥ 2 jets	≥ 2 jets + $\not{E}_{\mathrm{T}} \geq 40\,\mathrm{GeV}$	≥ 2 jets + ≥ 1 b-jets	≥ 2 jets + ≥ 1 b-jets + $\not{E}_{\mathrm{T}} \geq 40\,\mathrm{GeV}$
$\mu^+\mu^-$	$1.08 \pm 0.01\ ^{+0.07}_{-0.07}$	$1.20 \pm 0.01\ ^{+0.26}_{-0.24}$	$1.30 \pm 0.01\ ^{+0.11}_{-0.10}$	$1.43 \pm 0.04\ ^{+0.28}_{-0.26}$
e^+e^-	$1.15 \pm 0.01\ ^{+0.08}_{-0.07}$	$1.29 \pm 0.02\ ^{+0.27}_{-0.25}$	$1.33 \pm 0.02\ ^{+0.10}_{-0.10}$	$1.46 \pm 0.05\ ^{+0.28}_{-0.26}$
$e^\pm\mu^\mp$	$1.12 \pm 0.01\ ^{+0.08}_{-0.07}$	$1.24 \pm 0.01\ ^{+0.26}_{-0.25}$	$1.32 \pm 0.01\ ^{+0.11}_{-0.10}$	$1.44 \pm 0.04\ ^{+0.28}_{-0.26}$

Event Flavor	DY Estimated Yield Outside of Z-mass Window			
	≥ 2 jets	≥ 2 jets + $\not{E}_{\mathrm{T}} \geq 40\,\mathrm{GeV}$	≥ 2 jets + ≥ 1 b-jets	≥ 2 jets + ≥ 1 b-jets + $\not{E}_{\mathrm{T}} \geq 40\,\mathrm{GeV}$
$\mu^+\mu^-$	$30{,}210.02 \pm 0.70\%\ ^{+1.75\%}_{-1.89\%}$	$5{,}242.55 \pm 1.60\%\ ^{+4.15\%}_{-4.10\%}$	$3{,}387.52 \pm 2.32\%\ ^{+1.96\%}_{-2.29\%}$	$597.49 \pm 5.30\%\ ^{+3.44\%}_{-3.49\%}$
e^+e^-	$15{,}727.39 \pm 0.90\%\ ^{+1.15\%}_{-1.15\%}$	$2{,}667.97 \pm 2.11\%\ ^{+2.83\%}_{-3.00\%}$	$1{,}701.29 \pm 2.97\%\ ^{+3.53\%}_{-0.97\%}$	$312.44 \pm 6.86\%\ ^{+8.70\%}_{-3.94\%}$
$e^\pm\mu^\mp$	$2{,}187.84 \pm 0.97\%\ ^{+2.97\%}_{-3.02\%}$	$1{,}481.35 \pm 1.34\%\ ^{+15.61\%}_{-15.05\%}$	$234.86 \pm 2.95\%\ ^{+3.83\%}_{-3.21\%}$	$155.26 \pm 4.48\%\ ^{+14.67\%}_{-12.83\%}$

6.3.2.2 Cross-check of the Prediction in the $e^{\pm}\mu^{\mp}$ Channel

In order to validate our estimate of the DY scale factor in the $e^{\pm}\mu^{\mp}$ channel,we performed a cross-check where we fit the DY component of the observed $M_{\ell^+\ell^-}$ distribution in the $e^{\pm}\mu^{\mp}$ channel.

Figure 6.14 shows the results of this fit. We extracted scale factors by directly comparing the fit's estimated DY yield against the simulation's estimated contribution of DY to the $e^{\pm}\mu^{\mp}$ channel.

Table 6.6: Data-driven Drell-Yan background estimation using a fit to $M_{\ell^+\ell^-}$ in the $e^{\pm}\mu^{\mp}$ channels (c.f. Fig. 6.14) compared with simulation, for several steps of the analysis.

$e^{\pm}\mu^{\mp}$	$\geq$ 2 jets	$\geq$ 2 jets + $\geq$ 1 b-jets
DY MC	2,206.8 $\pm$ 16.4	205.88 $\pm$ 4.83
DY Data-fit Estimate	2455.3 $\pm$ 168.3	506.28 $\pm$ 140.40
SF Data/MC	1.11 $\pm$ 0.08	2.46 $\pm$ 0.68

Table 6.6 shows these calculated scale factors for various stages of the selection. A direct comparison of Table 6.6 with Table 6.5 shows that, at the jet selection level, the results are remarkably consistent. At later stages of the preselection, there start to be deviations. These deviations are about 1.7σ, but it should be noted that no systematic uncertainties were considered on the DY fit in the $e^{\pm}\mu^{\mp}$ channel other than MC statistics.

Referring to Table 6.5, the systematic uncertainties on the DY scale factors tend to be of the same order as the statistical uncertainties. As well, the uncertainties on the DY fit in the $e^{\pm}\mu^{\mp}$ channel completely dominate the relative uncertainty

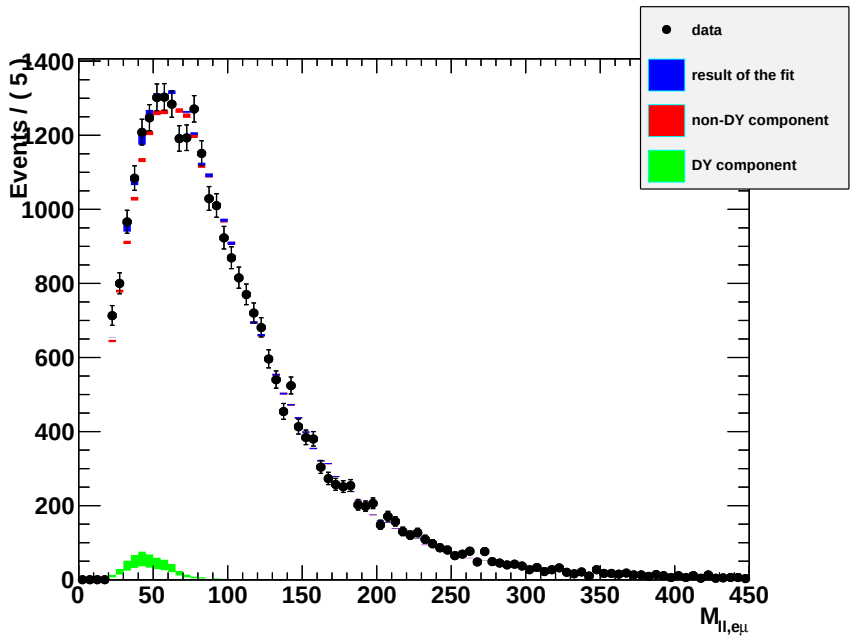

Figure 6.14: Result of the fits used to estimate the $e^{\pm}\mu^{\mp}$ background at the pre-selection level. The green represents the DY component. The red represents the non-DY component. The result of the fit (blue) and the data points show strong agreement.

of the comparison of the scale factors. Thus, it would be reasonable to assume that a more comprehensive treatment of the systematic uncertainties on the DY fit in the $e^{\pm}\mu^{\mp}$ channel would yield large enough overall uncertainties such that the scale factors calculated with the two separate methods would be consistent within 1σ.

6.3.2.3 Cross-check: DY Scale Factors at Large $M_{\mathrm{T2}}(\ell\ell)$

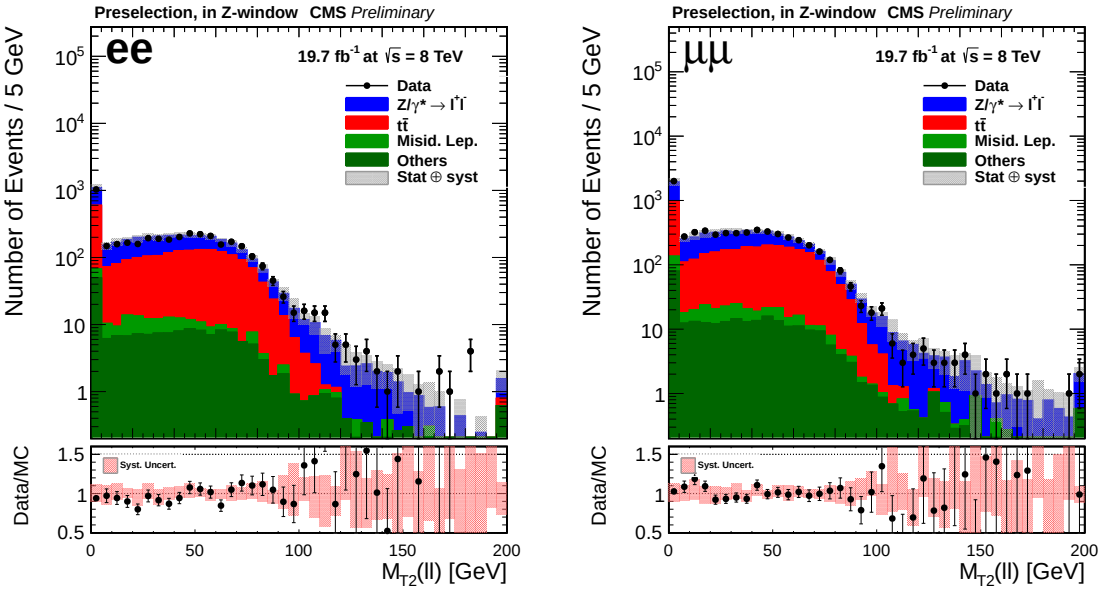

(a) The $M_{\mathrm{T2}}(\ell\ell)$ distribution in the $\mathrm{e}^+\mathrm{e}^-$ channel.

(b) The $M_{\mathrm{T2}}(\ell\ell)$ distribution in the $\mu^+\mu^-$ channel.

Figure 6.15: The $M_{\mathrm{T2}}(\ell\ell)$ distribution in $\mathrm{e}^+\mathrm{e}^-$ and $\mu^+\mu^-$ events passing the basic preselection, but where $M_{\ell^+\ell^-}$ falls *within* the Z-mass window.

Table 6.7: Results of the calculation of the data-driven DY normalization scale factor using the full preselection with additional, varying thresholds on $M_{\mathrm{T2}}(\ell\ell)$. The display format is central value $\pm$ statistical uncertainty $^{+\text{up syst.}}_{-\text{down syst.}}$.

Event Flavor	DY DD Scale Factor			
	$M_{\mathrm{T2}}(\ell\ell)$ Threshold			
	≥ 40 GeV	≥ 60 GeV	≥ 80 GeV	≥ 100 GeV
$\mu^+\mu^-$	$1.54 \pm 0.08\ ^{+0.23}_{-0.22}$	$1.51 \pm 0.12\ ^{+0.21}_{-0.22}$	$1.48 \pm 0.16\ ^{+0.22}_{-0.27}$	$1.51 \pm 0.22\ ^{+0.30}_{-0.22}$
e^+e^-	$1.52 \pm 0.09\ ^{+0.22}_{-0.24}$	$1.52 \pm 0.13\ ^{+0.27}_{-0.24}$	$1.82 \pm 0.19\ ^{+0.32}_{-0.25}$	$2.38 \pm 0.33\ ^{+0.53}_{-0.29}$

Table 6.8: Results of the calculation of the data-driven DY normalization scale factor using the full preselection, except with various thresholds of the $\not{E}_{\mathrm{T}}$ cut, The display format is central value $\pm$ statistical uncertainty $^{+\text{up syst.}}_{-\text{down syst.}}$.

Event Flavor	DY DD Scale Factor				
	Same-flavor dilepton channel $\not{E}_{\mathrm{T}}$ Threshold				
	≥ 0 GeV	≥ 20 GeV	≥ 40 GeV	≥ 60 GeV	≥ 80 GeV
$\mu^+\mu^-$	$1.30 \pm 0.01\ ^{+0.11}_{-0.10}$	$1.34 \pm 0.02\ ^{+0.15}_{-0.13}$	$1.43 \pm 0.04\ ^{+0.28}_{-0.26}$	$1.67 \pm 0.15\ ^{+0.42}_{-0.41}$	$2.13 \pm 0.53\ ^{+0.31}_{-0.40}$
e^+e^-	$1.33 \pm 0.02\ ^{+0.10}_{-0.10}$	$1.37 \pm 0.02\ ^{+0.14}_{-0.13}$	$1.46 \pm 0.05\ ^{+0.28}_{-0.26}$	$1.74 \pm 0.17\ ^{+0.39}_{-0.38}$	$2.03 \pm 0.53\ ^{+0.28}_{-0.48}$

In order to normalize the Drell-Yan background in our signal region, we take the scale factors shown in the last column of Table 6.5. We do not calculate the scale factor in specific $M_{\mathrm{T2}}(\ell\ell)$ regions. We have checked for possible biases introduced by this choice by applying varying $M_{\mathrm{T2}}(\ell\ell)$ thresholds and re-calculating the scale factors. The results of this check are shown in Table 6.7. Except for the last two columns in the $\mathrm{e^+e^-}$ channel, all the calculated scale factors for both channels are consistent within $\mathcal{O}(5\%)$ of the nominal scale factors calculated with no $M_{\mathrm{T2}}(\ell\ell)$ threshold. The last two columns in the $\mathrm{e^+e^-}$ channel show significantly larger deviations from the nominal scale factor for this channel.

We investigated these deviations further. They seem to be due primarily to statistical fluctuations in the data in the $\mathrm{e^+e^-}$ channel. Figure 6.15 shows the $M_{\mathrm{T2}}(\ell\ell)$ distributions for the $\mathrm{e^+e^-}$ and $\mu^+\mu^-$ channels, for events passing the preselection except that the reconstructed $M_{\ell^+\ell^-}$ falls inside the Z-mass window. As can be seen from the distribution in the $\mathrm{e^+e^-}$ channel, there is a clear excess of events starting around $M_{\mathrm{T2}}(\ell\ell) \sim 100\,\mathrm{GeV}$. There is no clear indicator that this excess (which is slightly above 2σ after accounting for systematic uncertainties) is anything more than a statistical fluctuation.

6.3.2.4 Systematic Uncertainties on the DY Estimation

As with the $\mathrm{t\bar{t}}$ estimation (Section 6.3.1), we recalculate the DY normalization scale factors individually for each systematic variation of the simulation. We also derive an additional systematic uncertainty based on the statistical imprecision of

this normalization. We do not apply any additional systematic uncertainties to the DY estimation.

The previous analyses [156, 157] that utilized this method applied an additional systematic uncertainty of 30% to their overall DY contribution. This 30% systematic was added primarily to cover mis-modeling of the $\not{E}_{\mathrm{T}}$ in the simulated DY events, in particular the underestimation of $\not{E}_{\mathrm{T}}$ resolution stemming from jet energy resolution. However, because we developed specific corrections to mitigate this mismodeling (Section 6.2.4.3), we chose not to apply this 30% systematic. We checked the dependence of the estimated DY scale factors on the $\not{E}_{\mathrm{T}}$ cut after applying smearing to the $\not{E}_{\mathrm{T}}$ — the results are shown in Table 6.8. The scale factors with different $\not{E}_{\mathrm{T}}$ thresholds are consistent at the 1σ level or better.

6.3.2.5 Signal Contamination in the DY Estimation

As noted above, the subtraction of opposite flavor events in Eq. (6.11) serves to remove the contamination by non-DY events. Thus, we expect there to be little, if any, signal contamination in our net estimate of the contribution of $\mathrm{Z} \to \ell^+\ell^-$ to our final signal region.

6.3.3 Misidentified Lepton Estimation

The "misidentified lepton" background consists of events in which non-prompt leptons, i.e. leptons not coming directly from the decay of a Z or W boson, pass our tight lepton ID criteria. The largest category of events falling in this group are semi-leptonic $t\bar{t}$ events and leptonically-decaying W events in which a jet, or a lepton within a jet, is misreconstructed as an isolated, prompt lepton. To guard against the possibility that these misidentified leptons are not well-modeled in the simulation, we perform a data-driven estimate of there expected contribution using the "tight-to-loose" method.

In this section, we highlight the key details of this method. More details are provided in Appendix F, or alternatively in Refs. [158, 159].

We start by defining a relaxed set of lepton quality requirements, a "loose" selection, to be contrasted against our "tight" lepton selection defined in Section 6.2.2. This "loose" selection defines a correspondingly named "loose lepton".

For both flavors of lepton candidate, we relax the relative isolation requirement; for muons (electrons), we require the relative isolation to be less than at 0.5 (1.0) instead of the value of 0.15 that is used for tight leptons. There are no other changes to the selection for muon candidates. For electron candidates, we apply a looser set of cuts on the selection variables discussed in Section 6.2.2. The exact values are shown in Table A.3.

Using these two lepton selections, we then measure the following,

- The lepton misidentification rate f (also called the "fake rate"), the probability

for a non-prompt lepton that passes the loose selection to also pass the tight selection.

- The lepton "prompt rate" p, the probability for a prompt lepton that passes the loose selection to also pass the tight selection.

After measuring f and p, we create of data sample of "loose-loose" dilepton events, events selected as per the requirements of Section 6.2.3, except that the 2 chosen lepton candidates only have to pass the "loose" lepton selection requirements mentioned above.

The expected contribution of misidentified leptons in our signal region is estimated by calculating a weighted sum of all events that fall into the signal region in this "loose-loose" sample, where the weights for each event are based upon whether either, neither, or both of the chosen lepton candidates in a given event also passed the "tight" lepton selection requirements, see Eq. (6.14) below.

6.3.3.1 Measurement of the Lepton Prompt Rates

The muon and electron prompt rates are obtained with a standard tag-and-probe technique using $\mathrm{Z} \to \ell^+\ell^-$ events[2]

For muons, p ranges from 0.84 (for $p_\mathrm{T} < 15\,\mathrm{GeV}$) to 0.99 (for $p_\mathrm{T} > 50\,\mathrm{GeV}$), with little to no dependence on the muon η.

For electrons, in the EB, p ranges from 0.83 (for $p_\mathrm{T} < 15\,\mathrm{GeV}$) to 0.98 (for $p_\mathrm{T} > 50\,\mathrm{GeV}$).

[2]Again, see Section 6.1 of Ref. [151] for a concise pedagogical discussion on the tag-and-probe technique.

In the EE, the analogous numbers are 0.75 and 0.94. For both flavors of lepton candidate, the statistical uncertainties on p are $\mathcal{O}(1\%)$ for $p_{\mathrm{T}} < 15\,\mathrm{GeV}$, and completely negligible otherwise.

6.3.3.2 Measurement of the Lepton Fake Rates

The muon and electron fake rates are extracted from a phase space dominated by QCD dijets events. We ensure that there is a reconstructed lepton candidate (to measure the lepton fake rate) by utilizing events that passed a single-lepton trigger. These triggers were discussed in Section 6.2.1.

In this sample of single-lepton triggered events, events with W decays are rejected by requiring both the PF $\not{E}_{\mathrm{T}}$ and the PF M_{T} — the transverse mass constructed with the lepton $\vec{p}_{\mathrm{T}}$ and PF $\vec{\not{E}}_{\mathrm{T}}$ — to each individually be less than $20\,\mathrm{GeV}$. We apply cuts on $M_{\ell^+\ell^-}$ in events with two or more reconstructed lepton candidates with the same-flavor but opposite charge; we reject the majority of events coming from Z decays by requiring $M_{\ell^+\ell^-}$ to fall outside of the Z-mass window, i.e. $M_{\mu^+\mu^-} \notin [76, 106]\,\mathrm{GeV}$ and $M_{\mathrm{e}^+\mathrm{e}^-} \notin [60, 120]\,\mathrm{GeV}$. For these dilepton events, we also reject low-mass dilepton resonances by requiring $M_{\ell^+\ell^-} > 20\,\mathrm{GeV}$, in line with the requirement used in our preselection. We remove the remaining expected contribution of W and Z events using estimates from simulation.

The remaining sample of events is heavily enriched in QCD multijet events. Using this sample, we apply what is effectively a variant of the standard lepton tag-and-probe technique to measure the lepton fake rate. In this QCD multijet sample,

we select events with a good jet (our "tag") — a jet passing all nominal selection requirements except that it has $p_{\mathrm{T}} > 50\,\mathrm{GeV}$ for — and a well separated "loose" lepton with $\Delta R(\mathrm{jet}, \ell) > 1.0$. Relative to this the loose lepton, we also require a jet that is close, $\Delta R(\mathrm{jet}, \ell) < 1.0$. We then measure the fake rate as per its written definition above,

$$f(p_{\mathrm{T}},\, |\eta|) \equiv \frac{N_{\mathrm{tight}}^{\mathrm{QCD}}}{N_{\mathrm{loose}}^{\mathrm{QCD}}}. \tag{6.13}$$

The altered p_{T} requirement for the "tag" jet was motivated by studies in a same-sign dilepton sample enriched in QCD multijet events, see the discussion below in Section 6.3.3.3.

6.3.3.3 Systematic Biases in the Fake Rate Measurement

There are two notable possible systematic biases to this method. First, the energy spectrum of jets misidentified as leptons can be different in events in this sample used to measured the fake rate, as compared with the events producing the misidentified leptons that contaminate our signal region. The relative isolation of a loose lepton, and thus the robustness of our estimate, is sensitive to these differences in jet energy.

In order to study any possible systematic biases stemming from this, we utilized a same-sign dilepton control sample enriched in QCD multijet events. Since the lepton fake rate should be more-or-less independent of lepton charge, this same-sign QCD sample provides a statistically independent measure in data of the robustness of our fake rate measurement.

To select this sample, we required two same-charge lepton candidates, each passing the loose lepton requirements, to be present in each event. In this sample, we tested how the estimated yield of misidentified leptons depended upon the tag jet's p_{T} requirement, and, in particular, how the estimated yield for a given choice of the tag jet's p_{T} requirement depended upon the relative isolation requirement placed on the loose lepton.

We found that requiring the "tag" jet's p_{T} to be within $\sim 10\,\mathrm{GeV}$ of $50\,\mathrm{GeV}$ resulted in misidentified lepton yields that were relatively robust to changes in the loose lepton's relative isolation requirement.

As an additional cross-check on the choice of tag jet p_{T} requirement, we also compared the estimated yield of misidentified leptons in the same-sign dilepton QCD sample against the estimated yield in a same-sign dilepton sample enriched in W + jet events.

We selected the latter sample by requiring two reconstructed, same-sign, tight leptons. This sample was enriched in W + jets events by using $\not{E}_{\mathrm{T}}$ and M_{T} requirements as per the discussion above. From this W-enriched sample, we subtracted, using simulation, the estimated contribution of backgrounds with two genuine, prompt leptons.

Figure 6.16 shows both checks that we performed. As you can see, the choice of tag jet $p_{\mathrm{T}} > 50\,\mathrm{GeV}$ is rather stable against variations in the loose lepton's relative isolation requirement. Moreover, our choice of tag jet p_{T} requirement results in an estimated misidentified lepton yield that is also fairly consistent with the expected yield from the same-sign, W-enriched control region. As a final step, in

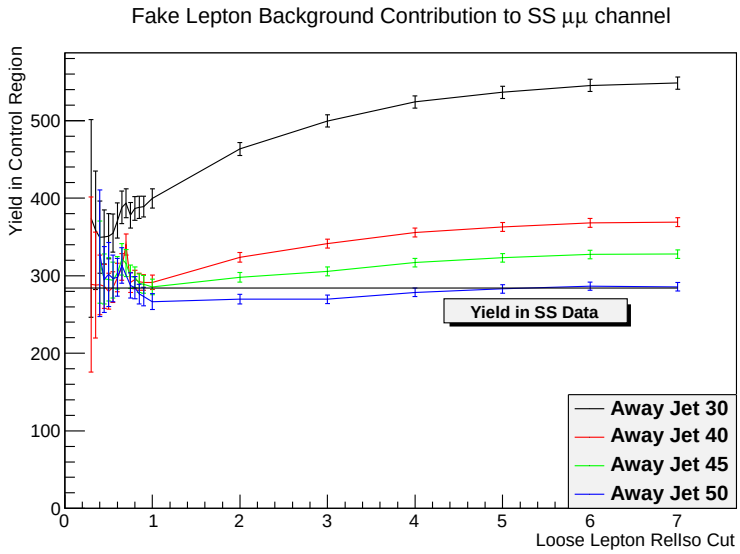

Figure 6.16: The dependence of the estimated yield of misidentified leptons in a loose-loose, same-sign dimuon QCD sample, see Section 6.3.3.3, on the relative isolation requirement placed on the leptons. The different colored lines with error bars represent this estimated yield as calculated using different requirements on the "tag" jet p_{T}. The black horizontal line with error bars refers to the estimated yield of misidentified leptons in a tight-tight same-sign dimuon sample enriched in W + jets events, where the expected contribution from backgrounds with two prompt leptons have been subtracted using simulation samples.

order to ensure a statistically conservative estimate of the fake rate, we calculated a systematic uncertainty on f, by varying the tag jet p_{T} requirement by $\pm 10\,\mathrm{GeV}$ and then recalculating f for each variation.

The second possible systematic bias is the residual contamination from any W or Z events that still remain in the sample. This is particularly relevant at high p_{T}^{ℓ}, where the contribution from QCD multijet events drops out. To mitigate this bias, we assumed that the lepton fake rate flattens out for $p_{\mathrm{T}}^{\ell} > 35\,\mathrm{GeV}$.

For muons, f ranges from 0.27 (for $p_{\mathrm{T}}^{\mu} < 15\,\mathrm{GeV}$) to ~ 0.08 (for $p_{\mathrm{T}}^{\mu} > 30\,\mathrm{GeV}$), with no systematic dependence on muon η. The statistical uncertainties on f are $\sim 40\%$ ($\sim 75\%$) for low (high) p_{T}^{μ}.

For electrons, f ranges from ~ 0.34 (for $p_{\mathrm{T}}^{\mathrm{e}} < 15\,\mathrm{GeV}$) to ~ 0.15 (for $p_{\mathrm{T}}^{\mathrm{e}} > 30\,\mathrm{GeV}$), with no systematic dependence on electron η. The statistical uncertainties are $\sim 25\% - 40\%$ for all $p_{\mathrm{T}}^{\mathrm{e}}$.

For both lepton flavors, the systematic uncertainties on f (from varying the tag jet p_{T}) are, on average, about a factor of 3 to 4 lower than the corresponding statistical uncertainties.

6.3.3.4 Calculating Event Weights

Each event in the loose-loose dilepton sample receives an event weight based upon whether both, either, or neither of the chosen two lepton candidates also passed

tight selection requirements,

$$\begin{aligned}
\text{Pass} - \text{Pass}: \;\; w^{pass\,pass} &= -\frac{\epsilon_1\eta_1 + \epsilon_2\eta_2 - \epsilon_1\epsilon_2\eta_1\eta_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Fail} - \text{Fail}: \;\; w^{fail\,fail} &= -\frac{\epsilon_1\epsilon_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Pass} - \text{Fail}: \;\; w^{pass\,fail} &= \frac{\epsilon_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Fail} - \text{Pass}: \;\; w^{fail\,pass} &= \frac{\epsilon_1}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)}
\end{aligned} \tag{6.14}$$

where the 1 (2) refers to the (sub-)leading lepton and, to simplify the equations, p and f have been rewritten in terms of the quantities, $\eta = \frac{1-p}{p}$ and $\epsilon = \frac{f}{1-f}$.

6.3.3.5 Signal Contamination in the Fake Lepton Estimation

There are three primary components to the misidentified lepton estimation where signal contamination from top-squark events might present an issue,

1. Measuring the lepton prompt rate, p.

2. Measuring the lepton fake rate, f.

3. Applying p and f on the loose-loose data sample to derive the final expected contribution of misidentified leptons.

For component 1, if the signal does exist, it does not bias our measurement of p, as the majority of the leptons coming from top-squark decays that would populate the relevant region used to measure p are "prompt" leptons, as per our definition above, and thus should not bias the measurement of p. An analogous argument applies to component 2, namely, that if the signal does exist, it should not bias our measurement of f.

For the final component, as we noted above, we expect that semi-leptonic $t\bar{t}$ events are one of the largest contributors to the overall misidentified lepton background. By that same logic, if the signal exists, there are likely semi-leptonic top-squark events in the loose-loose sample used as part of component 3.

However, using Table 6.9, we can see that the misidentified lepton background is never the dominant background in any of the signal regions considered. Thus, we expect that signal contamination in the misidentified lepton background estimate should have a negligible impact on our final results of this analysis.

Table 6.9: Table showing the expected contribution, along with total systematic uncertainties, of the separate background sources considered in this analysis for each choice of $M_{\mathrm{T2}}(\ell\ell)$ threshold.

$M_{\mathrm{T2}}(\ell\ell)$ threshold	≥ 80 GeV	≥ 90 GeV	≥ 100 GeV	≥ 110 GeV	≥ 120 GeV
$\mathrm{t\bar{t}}$	$1532.93\ ^{+111.21}_{-92.31}$	$349.45\ ^{+50.42}_{-29.34}$	$71.83\ ^{+13.22}_{-7.82}$	$14.10\ ^{+4.36}_{-1.99}$	$2.60\ ^{+1.94}_{-0.98}$
$\mathrm{Z}/\gamma^* \to \ell^+\ell^-$	$30.25\ ^{+9.07}_{-6.55}$	$20.95\ ^{+7.73}_{-6.24}$	$11.09\ ^{+6.18}_{-3.11}$	$7.48\ ^{+1.81}_{-1.87}$	$5.37\ ^{+1.68}_{-1.39}$
Misid. Leptons	$22.02\ ^{+19.47}_{-19.34}$	$9.42\ ^{+7.45}_{-7.46}$	$4.12\ ^{+3.40}_{-3.46}$	$2.58\ ^{+2.32}_{-2.40}$	$2.08\ ^{+1.88}_{-1.95}$
Others	$92.51\ ^{+9.80}_{-7.44}$	$31.78\ ^{+4.03}_{-4.21}$	$13.39\ ^{+2.72}_{-2.38}$	$7.01\ ^{+1.82}_{-1.77}$	$3.94\ ^{+1.19}_{-0.88}$
Total	$1677.72\ ^{+120.57}_{-99.16}$	$411.60\ ^{+50.42}_{-29.34}$	$100.43\ ^{+19.72}_{-9.02}$	$31.17\ ^{+5.44}_{-3.82}$	$13.99\ ^{+3.53}_{-2.55}$

6.3.4 Summary

The central values of the background predictions for the 5 $M_{\mathrm{T2}}(\ell\ell)$ signal regions considered are shown in Table 6.9. Table 6.9 also shows the estimated impact of systematic uncertainties on each notable background. We now turn to discussion of these systematic uncertainties and how we estimated them.

6.4 Systematic Uncertainties

The sensitivity of this analysis is affected both by uncertainties on the background contribution in the signal region as well as uncertainties on the acceptance and efficiency for the signal models considered. In this section, we provide details on the systematic uncertainties that we accounted for and their impact on the expected contribution to our $M_{\mathrm{T2}}(\ell\ell)$ signal regions for both background and signal events.

6.4.1 Systematic Uncertainties Affecting the Background and Signal

In this section, we provide details on the systematic uncertainties in our analysis that affect both the background and the top-squark signal. Note that, for

both the background and signal estimates, we do include the intrinsic statistical uncertainty of the simulation samples as a systematic uncertainty.

6.4.1.1 $\not{E}_{\mathrm{T}}$ Uncertainties Propagated to $M_{\mathrm{T2}}(\ell\ell)$

The $\vec{\not{E}}_{\mathrm{T}}$ reconstruction in the simulation is affected by any systematic uncertainties on the energy scales and resolutions for all reconstructed visible objects in the event. Of special concern are the uncertainties on the lepton energy scale, jet energy scale, jet energy resolution, and the energy scale of unclustered objects in the event. In order to evaluate the effect of these uncertainties, we utilized a combination of several prescriptions.

For the lepton and jet energy scales, we varied individual object momentum four-vectors within systematic uncertainties taken from dedicated studies of the energy scales [84, 116, 117, 118, 120]. We then propagated the shifted $\vec{p}_{\mathrm{T}}$ back into the $\vec{\not{E}}_{\mathrm{T}}$ calculation. For the leptons, the shifted $\vec{p}_{\mathrm{T}}$ themselves are also used in the calculation of $M_{\mathrm{T2}}(\ell\ell)$.

For the unclustered energy scale and jet energy resolution uncertainties, see Section 6.4.1.2 below.

We propagated all calculated systematic uncertainties on the $\vec{\not{E}}_{\mathrm{T}}$ measurement to the $M_{\mathrm{T2}}(\ell\ell)$ measurement by recalculating $M_{\mathrm{T2}}(\ell\ell)$ for each systematic shift version of the $\vec{\not{E}}_{\mathrm{T}}$ measurement. The individual systematic uncertainties are treated as uncorrelated; the total systematic uncertainty on the $M_{\mathrm{T2}}(\ell\ell)$ measurement is the addition in quadrature of the individual uncertainties. This particular treatment of

the individual systematic uncertainties, namely viewing them as uncorrelated, has precendence within CMS, particularly with heavily $\not{E}_{\mathrm{T}}$-based analyses [71].

6.4.1.2 Uncertainties on the $\not{E}_{\mathrm{T}}$ Resolution and Unclustered Energy Scale

In Section 6.2.4.3, we provided details on corrective "smearing" that we applied to both the $\not{E}_{\mathrm{T}}$ magnitude and direction in our simulated samples in order to improve the agreement with data.

As noted in Section 6.2.4.3, the smearing template samples that we used to correct the modeling of $\vec{\not{E}}_{\mathrm{T}}$ in our simulation were derived from a separate set of simulation samples. In addition to the nominal smeared $\vec{\not{E}}_{\mathrm{T}}$, these separate simulation samples also contain variants of the smeared $\vec{\not{E}}_{\mathrm{T}}$ where the energy scale of unclustered objects (PF candidates with $p_{\mathrm{T}} < 10\,\mathrm{GeV}$) was varied within $\pm 1\sigma$ (10%) and where the jet smearing scale factors (i.e. jet smearing magnitude) were varied within $\pm 1\sigma$ (see Table 5.3 in Section 5.7.1.4). We used these variants of the smeared $\not{E}_{\mathrm{T}}$ to derive additional smearing templates to replicate the systematic variations of the unclustered energy scale and jet energy resolution in our own simulation samples. These two systematic variations of the $\not{E}_{\mathrm{T}}$, and subsequent propagation into $M_{\mathrm{T2}}(\ell\ell)$, were done independently of one another.

6.4.1.3 b-tagging Efficiency

We propagate the systematic uncertainties (Section 6.2.4.2) on the heavy-flavor jet tagging efficiency scale factors and the light-flavor jet mis-tag rate scale factors into our estimate on the number of jets passing our b-tagging requirement.

Since the measurements of the two sets of scale factors are statistically independent, in order to derive a total systematic uncertainty associated with the CSV b-tagging algorithm, we independently vary the two scale factors within their systematic uncertainties and then add in quadrature the resulting systematic shifts in our background and signal estimates.

6.4.1.4 Lepton Efficiency Scale Factors

As a reminder from Section 6.2.4.1, we apply a conservative 1% systematic uncertainty to the calculated lepton efficiency scale factors in order to account for possible systematic biases.

We propagate these systematic uncertainties into our estimates of the simulated backgrounds in the $M_{\mathrm{T2}}(\ell\ell)$ signal region; they primarily affect the normalization of the non-$\mathrm{t\bar{t}}$, non-DY simulated backgrounds, as there is no substantial correlations between the shape of the $M_{\mathrm{T2}}(\ell\ell)$ distribution and the lepton efficiency scale factors. Consequently, the systematic uncertainties on these efficiency scale factors have a minimal impact on our final results.

6.4.2 Systematic Uncertainties Affecting only the Background

In this section, we provide details on the systematic uncertainties in our analysis that affect only the background estimates.

6.4.2.1 Misidentified Lepton Estimate

The estimate of the misidentified leptons (Section 6.3.3) is affected by two systematic uncertainties: the statistical uncertainty on the measured lepton fake and prompt rates and the systematic uncertainty on the measurement of the lepton fake rate. These uncertainties are propagated into the event weights used to estimate the number of misidentified lepton events in our signal region.

This leads to an $\mathcal{O}(80\%)$ uncertainty on the normalization of the misidentified lepton background, with a relatively smaller uncertainty on the shape of this background. This large uncertainty is primarily driven by the statistical uncertainty on the measured lepton fake and prompt rates.

6.4.2.2 Normalization Uncertainties

We account for the statistical uncertainty of the data-driven normalizations of the $t\bar{t}$ and DY backgrounds (Section 6.3.1.1 and 6.3.2.4 respectively) as an additional systematic uncertainty.

As well, we apply a conservative cross section uncertainty of 50% on the sum of the VV and VG electroweak backgrounds. This is done to to guard against inaccuracies in the calculation of their respective cross sections.

6.4.2.3 Generator-level Shape Uncertainties

As noted in Section 6.2.4.5, we apply generator-level reweighting, based upon particle p_{T} values, to both our top-squark signal as well as our background samples that contain $\mathrm{t\bar{t}}$ pairs. Associated with these event weights is a conservative, 100% systematic uncertainty.

6.4.3 Systematic Uncertainties Affecting Only the Signal

In this section, we provide details on the systematic uncertainties in our analysis that affect only the estimates of the signal efficiency.

6.4.3.1 PDF Uncertainties

We utilize the canonical method for calculating the systematic uncertainty stemming from the choice of PDF sets [160, 161]. For processes where the cross-section is well-known, i.e. $\mathrm{Z} \to \ell^+\ell^-$ or $\mathrm{t\bar{t}}$, the eigenvectors of the PDF set with which the simulation samples were generated are varied within their 1σ uncertainties, yielding non-unity event weights. These weights are applied in order to calculate relative changes in the shapes of relevant kinematic variables, such as $M_{\mathrm{T2}}(\ell\ell)$. A total systematic uncertainty is then estimated by adding in quadrature the estimated uncertainties from the shifted event weights for each individual PDF eigenvector.

For unconfirmed processes, such as the top-squark signal, the above calculation is repeated for three different PDF sets, CT10, MSTW, and NNPDF. The total systematic uncertainty stemming from PDF uncertainties is then taken as the

Table 6.10: The relative systematic effect (%) from varying the $M_{\mathrm{T2}}(\ell\ell)$ shape in the simulated $\mathrm{t\bar{t}}$ POWHEG sample according to the individual PDF eigenvector set uncertainties.

$M_{\mathrm{T2}}(\ell\ell)$ cut	fractional yield	CT10	MSTW	NNPDF
$\geq 80\,\mathrm{GeV}$	3.31×10^{-2}	$^{+0.31\%}_{-0.28\%}$	$^{+0.13\%}_{-0.14\%}$	$^{+1.91\%}_{-1.62\%}$
$\geq 90\,\mathrm{GeV}$	7.69×10^{-3}	$^{+1.09\%}_{-0.77\%}$	$^{+0.29\%}_{-0.29\%}$	$^{+2.40\%}_{-2.43\%}$
$\geq 100\,\mathrm{GeV}$	1.58×10^{-3}	$^{+2.22\%}_{-1.64\%}$	$^{+0.53\%}_{-0.50\%}$	$^{+4.63\%}_{-5.02\%}$
$\geq 110\,\mathrm{GeV}$	3.38×10^{-4}	$^{+4.56\%}_{-3.34\%}$	$^{+0.99\%}_{-0.94\%}$	$^{+7.75\%}_{-7.84\%}$
$\geq 120\,\mathrm{GeV}$	$7.43 \times 10^{+-5}$	$^{+12.53\%}_{-\ 8.03\%}$	$^{+2.67\%}_{-2.82\%}$	$^{+22.15\%}_{-21.82\%}$

envelope (e.g. the max uncertainty) of the three individual PDF set uncertainties.

When we performed this calculation with our top-squark signal, we found that the magnitude of the uncertainty depended both upon the top-squark decay mode and also upon the position in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. However, for most mass points considered, the magnitude of the PDF envelope uncertainty was around 5–10%.

We considered applying the PDF uncertainty to our background estimates as well. As a test, we calculated its expected impact on the $\mathrm{t\bar{t}}$ estimate. Table 6.10 shows the relative fractional yield of $\mathrm{t\bar{t}}$, and the overall relative (%) change in this fractional yield when varying the three standard PDF eigenvector sets within their systematic uncertainties. Because the $\mathrm{t\bar{t}}$ process is treated as "known" in this analysis — i.e. we are not trying to measure its cross section — the canonical treatment for PDF uncertainties is to only consider variations in the PDF eigenvector set (CT10) that was used to make the $\mathrm{t\bar{t}}$ sample.

Direct comparisons of Table 6.10 with Table 6.11 (which will be described in more detail later) show that, for every $M_{\mathrm{T2}}(\ell\ell)$ threshold choice except $M_{\mathrm{T2}}(\ell\ell) =$

120 GeV, the impact of the CT10 eigenvector variations is negligible relative to the pre-existing systematic uncertainties. For the threshold choice, $M_{\mathrm{T2}}(\ell\ell) = 120\,\mathrm{GeV}$, the impact is $\mathcal{O}(50\%)$ relative to the other systematic uncertainties, meaning that we could be underestimating the total systematic uncertainty for that $M_{\mathrm{T2}}(\ell\ell)$ threshold choice by a relative 10% at worst. Thus, we chose not to account for any PDF eigenvector set uncertainties in our background estimates.

6.4.4 Other Systematics

We also considered the effect of variations of the W mass on the background yield. However the current world average uncertainty on the W mass is only 15 MeV and the uncertainty on the width is only 42 MeV [162]. Since these uncertainties are much smaller than the $\not{E}_{\mathrm{T}}$ resolution, and subsequently much smaller than the $M_{\mathrm{T2}}(\ell\ell)$ resolution, we do not use this uncertainty in the final result.

6.4.5 Correlations between Systematics

Individual systematics are treated as uncorrelated relative to one another. Each individual source of systematic uncertainty that affects both the background and signal is treated as 100% correlated when expected background and signal yield are compared, e.g. when we calculated the exclusion contours shown in Section 6.8.

Table 6.11: The relative impact (in percentage of the total background yield) for relevant sources of systematic uncertainty on the background estimate for each signal region used in the limit setting. From left to right, the systematic uncertainty sources are: simulation sample statistics, lepton energy scale, jet energy scale, unclustered energy scale, $\not{E}_{\mathrm{T}}$ energy resolution from jets, uncertainty on b-tagging scale factors, lepton selection efficiency, ISR reweighting, the total uncertainty on the misidentified lepton estimate, and the combined normalization uncertainty on the $\mathrm{t\bar{t}}$, DY, and electroweak backgrounds.

	Systematic Uncertainty (%)										
$M_{\mathrm{T2}}(\ell\ell)$ cut	MC stat.	ℓ ES	JES	uncl.	JER	b-tag	ℓ eff.	recoil	Misidentified Lepton	Sample σ	total
≥ 80 GeV	± 1.08	$^{+4.42}_{-4.55}$	$^{+2.04}_{-1.27}$	$^{+2.91}_{-0.78}$	$^{+3.34}_{-3.27}$	$^{+0.69}_{-0.00}$	$^{+0.88}_{-0.83}$	$^{+1.06}_{-1.14}$	$^{+1.14}_{-1.14}$	$^{+0.91}_{-0.91}$	$^{+6.99}_{-6.23}$
≥ 90 GeV	± 2.25	$^{+6.32}_{-6.22}$	$^{+5.38}_{-2.26}$	$^{+7.35}_{-1.42}$	$^{+6.71}_{-3.95}$	$^{+1.97}_{-0.00}$	$^{+1.01}_{-0.95}$	$^{+0.24}_{-0.15}$	$^{+1.77}_{-1.77}$	$^{+1.04}_{-1.04}$	$^{+13.50}_{-8.46}$
≥ 100 GeV	± 4.14	$^{+6.13}_{-5.41}$	$^{+9.24}_{-1.70}$	$^{+9.53}_{-1.16}$	$^{+12.15}_{-2.04}$	$^{+0.92}_{-0.50}$	$^{+0.75}_{-0.69}$	$^{+1.46}_{-1.01}$	$^{+3.23}_{-3.29}$	$^{+2.10}_{-2.10}$	$^{+19.92}_{-8.47}$
≥ 110 GeV	± 7.23	$^{+8.93}_{-5.12}$	$^{+8.75}_{-1.44}$	$^{+4.44}_{-0.00}$	$^{+4.75}_{-0.00}$	$^{+1.14}_{-1.86}$	$^{+0.30}_{-0.26}$	$^{+3.23}_{-2.23}$	$^{+6.92}_{-7.18}$	$^{+4.66}_{-4.66}$	$^{+18.23}_{-12.74}$
≥ 120 GeV	± 9.85	$^{+3.93}_{-4.69}$	$^{+11.90}_{-2.78}$	$^{+2.19}_{-0.27}$	$^{+4.54}_{-0.00}$	$^{+3.15}_{-0.82}$	$^{+0.20}_{-0.17}$	$^{+6.13}_{-4.04}$	$^{+11.89}_{-12.31}$	$^{+5.06}_{-5.06}$	$^{+22.23}_{-17.92}$

The composite effect of all systematic uncertainties is shown, among other plots, in Fig. 6.12 and in Tables 6.9 and 6.11.

6.5 $M_{\mathrm{T2}}(\ell\ell)$ Shape Validation

This analysis depends on a clear and precise understanding of the $M_{\mathrm{T2}}(\ell\ell)$ shape. We therefore performed a number of studies with toy simulations and the full CMS simulation to determine how well this shape is understood. We also compared the shape in control regions in data with the prediction of the simulation to check for unanticipated effects.

6.5.1 Effects Governing the $M_{\mathrm{T2}}(\ell\ell)$ Shape

Because $M_{\mathrm{T2}}(\ell\ell)$ is constructed from $\not{E}_{\mathrm{T}}$ and two leptons, there are numerous factors that affect the distribution of $M_{\mathrm{T2}}(\ell\ell)$ and more importantly, the modeling of $M_{\mathrm{T2}}(\ell\ell)$ in the simulation.

SM $\mathrm{t\bar{t}}$ events that contain leptonically decaying τs nominally do not have a kinematic edge at m_{W} due to the additional neutrinos from the τ decays. However, the calculated $M_{\mathrm{T2}}(\ell\ell)$ values for these events tends to be quite low as the added $\not{E}_{\mathrm{T}}$ from the τ neutrinos tends to be close to the selected leptons (leading to small values of M_{T} for the nominal neutrino-lepton pairings).

The intrinsic resolution and scale on the lepton $\vec{p}_{\mathrm{T}}$ affect $M_{\mathrm{T2}}(\ell\ell)$, not only through the $\vec{\not{E}}_{\mathrm{T}}$ used in the calculation, but also through the lepton $\vec{p}_{\mathrm{T}}$ used. However, from studies using $\mathrm{Z} \to \ell^{+}\ell^{-}$ events, both the resolution and scale on the

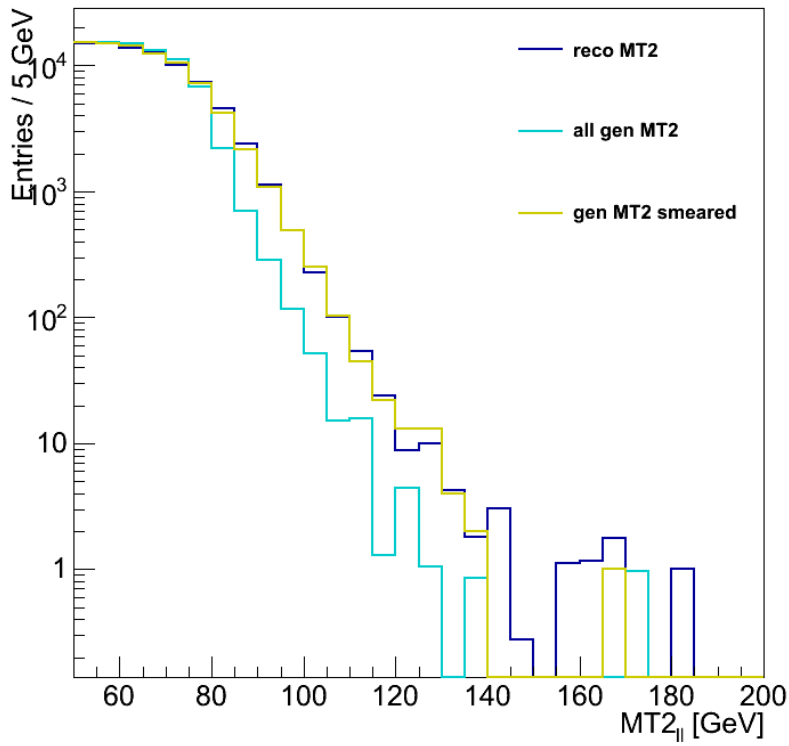

Figure 6.17: $M_{\mathrm{T2}}(\ell\ell)$ distribution for simulated $\mathrm{t\bar{t}}$ in three different cases. In dark blue, the distribution is shown using reconstructed $\not{E}_{\mathrm{T}}$ as would be done for data events. In light blue, the same distribution in shown substituting generator-level $\not{E}_{\mathrm{T}}$ for reconstructed $\not{E}_{\mathrm{T}}$. In gold, we smear the generated $\not{E}_{\mathrm{T}}$ by the method indicated in the text, and recover essentially the same result as for reconstructed $\not{E}_{\mathrm{T}}$.

lepton $\vec{p}_{\mathrm{T}}$ are well-modeled in the simulation.

In SM $\mathrm{t\bar{t}}$, the Gaussian core of the $\not{E}_{\mathrm{T}}$ resolution as well as the intrinsic width of the intermediate W bosons are the two main driving factors that turn the hard cut-off at $M_{\mathrm{T2}}(\ell\ell) = m_{\mathrm{W}}$ into a steeply falling kinematic edge. Figure 6.17 shows a study we performed to confirm this fact.

From this figure, you can see that when one only considers the $M_{\mathrm{T2}}(\ell\ell)$ constructed with the generator-level $\not{E}_{\mathrm{T}}$ (light-blue), the $M_{\mathrm{T2}}(\ell\ell)$ in dileptonic $\mathrm{t\bar{t}}$ events has a strong kinematic edge around $M_{\mathrm{T2}}(\ell\ell) \approx m_{\mathrm{W}}$. This edge is not a sharp cut-off

because the Ws produced in the top-quark decays have a finite Breit-Wigner width to them.

We fit a simple Gaussian distribution to the core of the $\not{E}_{\mathrm{T}}$ resolution for this selected sample and smeared the generator-level $\not{E}_{\mathrm{T}}$ by this resolution before then recalculating $M_{\mathrm{T2}}(\ell\ell)$. This "by-hand" smeared $M_{\mathrm{T2}}(\ell\ell)$ (gold) matches the shape of the nominal reconstructed $M_{\mathrm{T2}}(\ell\ell)$ (dark blue) quite well, corroborating our claims made above.

The final effect driving the $M_{\mathrm{T2}}(\ell\ell)$ shape is unusual or extreme mismeasurements that populate the tails of the $\not{E}_{\mathrm{T}}$ resolution distribution. These mismeasurements, especially if they're not well-modeled by the simulation, can present a notable challenge for both the accuracy and precision of the analysis.

One of the most common causes of extreme $\not{E}_{\mathrm{T}}$ mismeasurement events are $\not{E}_{\mathrm{T}}$ noise events, events where detector-related effects led to a spuriously high calculated $\not{E}_{\mathrm{T}}$ value. As noted in Section 5.7.2, the CMS $\not{E}_{\mathrm{T}}$ group has a dedicated set of event filters that work extremely well (c.f. Fig. 5.11) for removing these $\not{E}_{\mathrm{T}}$ noise events.

Beyond these noisy $\not{E}_{\mathrm{T}}$ events, though, extreme $\not{E}_{\mathrm{T}}$ mismeasurement events are also possible just through normal mismeasurements of visible objects. It is important to confirm the accuracy of the simulation in modeling this particular class of events, as these events are what comprise the majority of our expected contribution to the high $M_{\mathrm{T2}}(\ell\ell)$ tail from backgrounds such as $\mathrm{Z} \to \ell^+\ell^-$. To check that our simulation models this class of events with acceptable accuracy, we performed a comparison between data and simulation in several control regions.

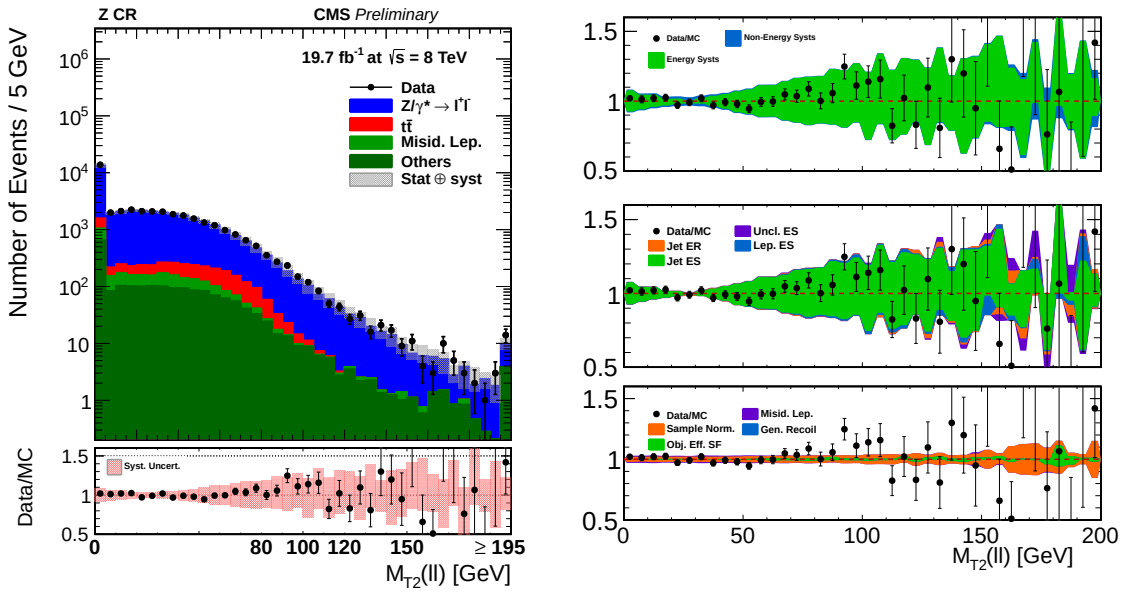

(a) (b)

Figure 6.18: The distribution of $M_{\mathrm{T2}}(\ell\ell)$ in events in a b-vetoed Z-peak control region (e.g. Table 6.3, but with inversions applied on the b-jet selection and Z-mass window cuts). Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

6.5.2 Z-enriched Control Region

We checked for catastrophic failures of $M_{\mathrm{T2}}(\ell\ell)$ reconstruction in a control region similar to our main preselection sample, but enriched in DY events. We created this sample by using our preselection requirements, except we require that there be 0 reconstructed b-jets in each event and that $M_{\ell^+\ell^-}$ fall within 15 GeV of m_{Z}.

The results of this check are shown in Fig. 6.18. Examination of the high $M_{\mathrm{T2}}(\ell\ell)$ tail clearly shows that the simulation does an acceptable job of modeling this important region. There are no major systematic trends in the ratio of data and simulation, and this ratio is consistent with 1 within the systematic uncertainties.

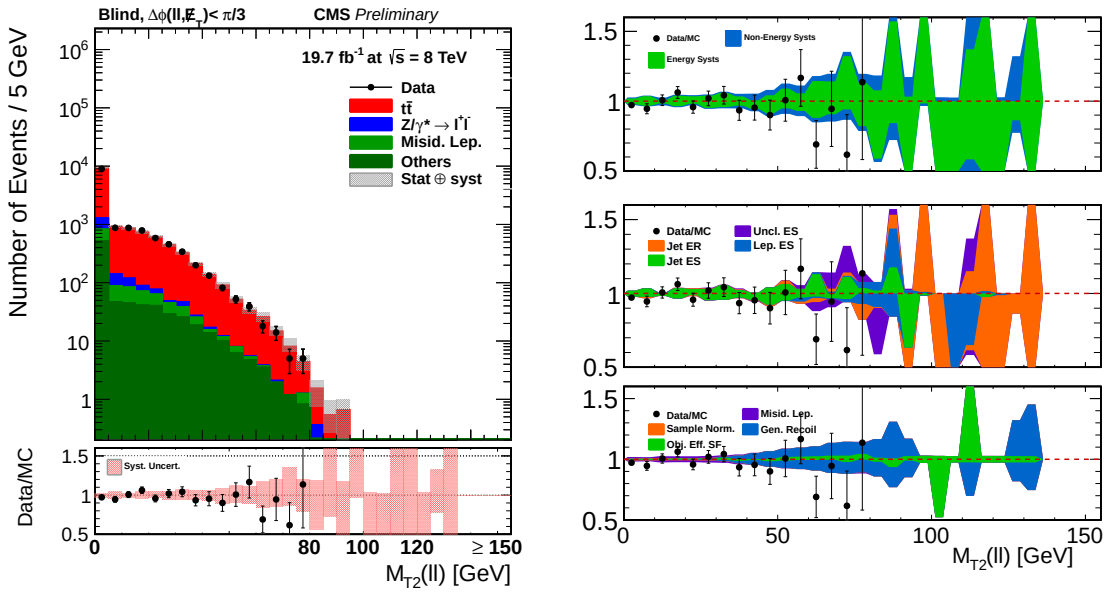

(a) (b)

Figure 6.19: The distribution of $M_{\mathrm{T2}}(\ell\ell)$ at the preselection level for events with $\Delta\phi(\ell^+\ell^-, \not{E}_{\mathrm{T}}) < \pi/3$. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

It is worth noting that, from Fig. 6.18b, the systematic uncertainty on the simulation stemming from the jet energy scale calibration uncertainty completely dominates the overall systematic uncertainties at high $M_{\mathrm{T2}}(\ell\ell)$ in this sample of events.

6.5.3 $\Delta\phi$ Studies

When the $M_{\mathrm{T2}}(\ell\ell)$ cut is low, the number of observed events above the cut is dominated by the shape of the falling edge and not the extreme tail. If the tail fell more or less sharply in data than in simulation, our simulation can predict too many or too few events in the signal region. We check for evidence of any possible mis-modeling in a $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$ control region in order to put an upper limit

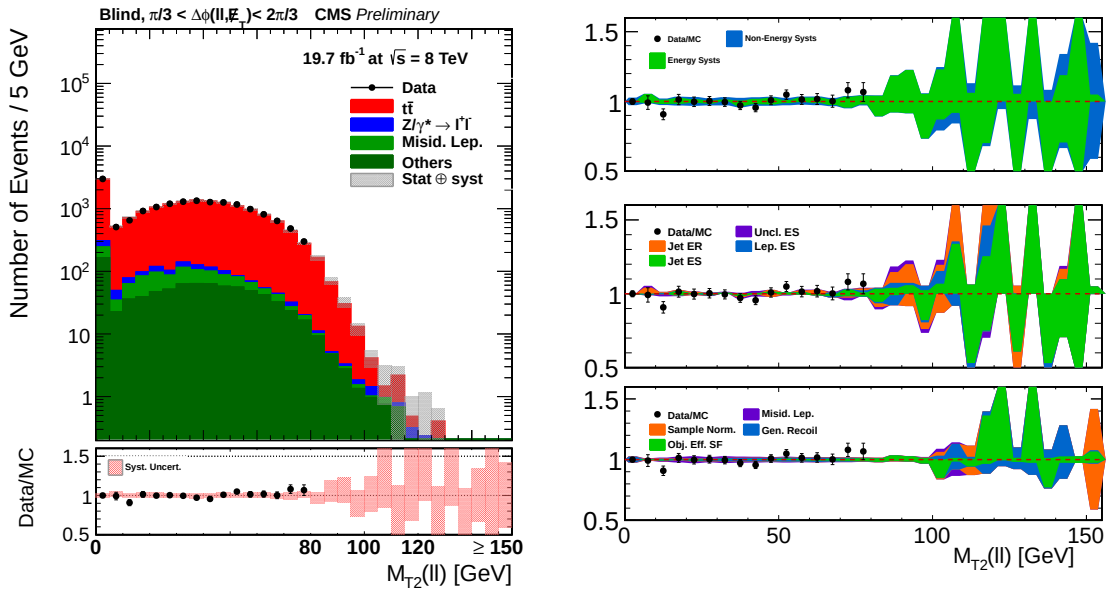

(a) (b)

Figure 6.20: The distribution of $M_{\mathrm{T2}}(\ell\ell)$ at the preselection level for events with $\pi/3 < \Delta\phi(\ell^+\ell^-, \not{E}_{\mathrm{T}}) < 2\pi/3$. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

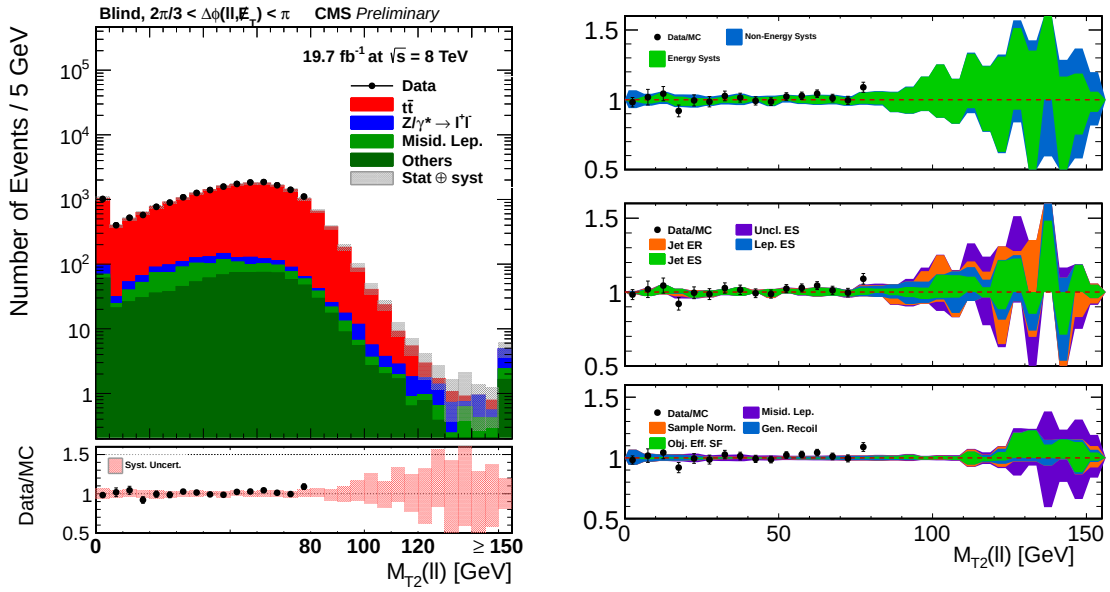

(a) (b)

Figure 6.21: The distribution of $M_{\mathrm{T2}}(\ell\ell)$ at the preselection level for events with $\Delta\phi(\ell^+\ell^-, \not{E}_{\mathrm{T}}) > 2\pi/3$. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

on the size of the effect.

Our strategy revolves around the use of the $\Delta\phi$ between the dilepton system and the $\not{E}_{\mathrm{T}}$. By looking at the smallest angle configurations, we can look at the falling edge of the $M_{\mathrm{T2}}(\ell\ell)$ distribution while remaining blinded to the signal region, as shown in Fig. 6.19. If there are sources of spurious high-$M_{\mathrm{T2}}(\ell\ell)$ events in data (due to additional backgrounds not considered in the analysis, mismeasurement of high $\not{E}_{\mathrm{T}}$ tails, or other unanticipated sources) we should expect to see an excess here. There are no major systematic trends in the ratio of data and simulation and the existing systematics cover the observed values of $M_{\mathrm{T2}}(\ell\ell)$ in data, thus suggesting that no such effects are present in our selected sample.

Further check using other $\Delta\phi$ regions, as shown in Figs. 6.20 and 6.21, show that there is consistently accurate modeling of the $M_{\mathrm{T2}}(\ell\ell)$ shape throughout the entire range of $\Delta\phi$ values, further improving our confidence that the simulation adequately models the $M_{\mathrm{T2}}(\ell\ell)$ variable.

6.5.4 Data-MC Mixed WW Control Sample

We may also inspect the $M_{\mathrm{T2}}(\ell\ell)$ shape in data with another technique. We select, in both the data and MC, a sample dominated with single W-boson events. Then, we mix the lepton and neutrino from a simulated $\mathrm{W} \to \ell\nu$ event into the reconstructed event to produce a WW event where the calculated $M_{\mathrm{T2}}(\ell\ell)$ should display the same kinematic edge at m_{W} as the real $\mathrm{t\bar{t}}$ dominated selection we use for the main analysis. In order to select a relatively clean sample of events with a single

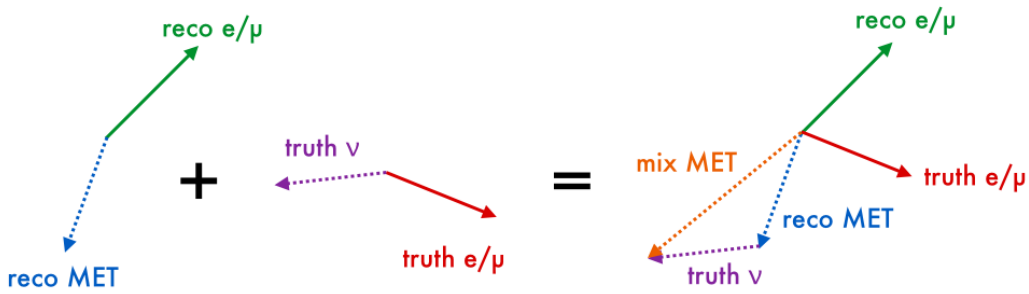

Figure 6.22: Mixing scheme used to overlay MC truth W events on reconstructed W events.

W-boson decaying leptonically, we utilize events from the single-lepton triggered datasets discussed in Section 6.2.1. Using our same object definitions as for the main analysis, we require these events to have exactly one electron or muon with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$. In order to reject QCD multijet events where a jet is misreconstructed as a lepton, we require at least 20 GeV of $\not{E}_{\mathrm{T}}$. For the simulated single W-boson events, we require only that the event contain an electron or muon originating from a W boson.

Having identified a reconstructed and truth-level W event to mix, the event information is combined as depicted in Fig. 6.22. This results in an event with one reconstructed lepton, one MC truth lepton, and a $\vec{\not{E}}_{\mathrm{T}}$ that combines the original, reconstructed one with the MC truth neutrino. The $M_{\mathrm{T2}}(\ell\ell)$ is computed from the mixed $\vec{\not{E}}_{\mathrm{T}}$ and the two leptons. The result is shown in Fig. 6.23. We find agreement between the data and the simulation at the 10% level, which is consistent with the systematic uncertainties on the full result (Table 6.11).

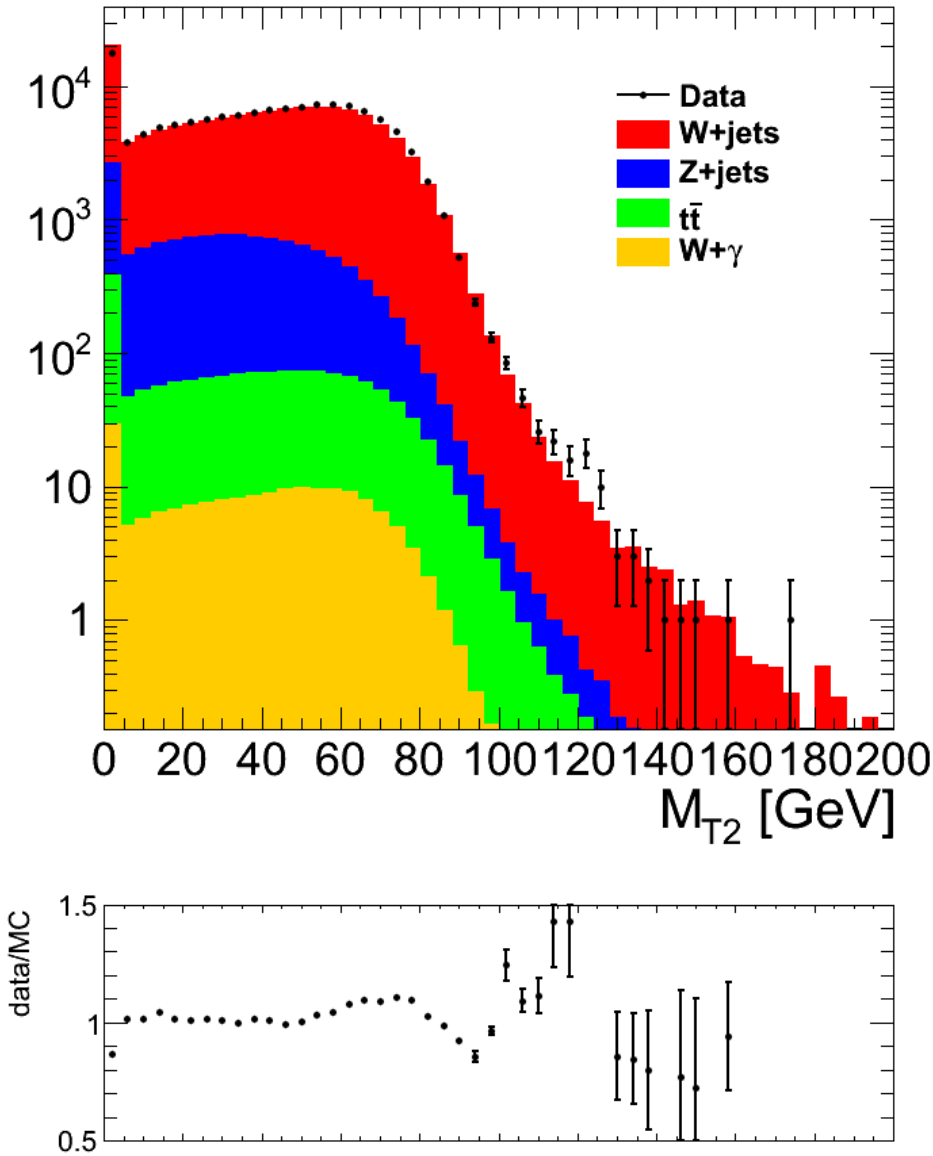

Figure 6.23: $M_{\mathrm{T2}}(\ell\ell)$ shape from the data-MC mixed events.

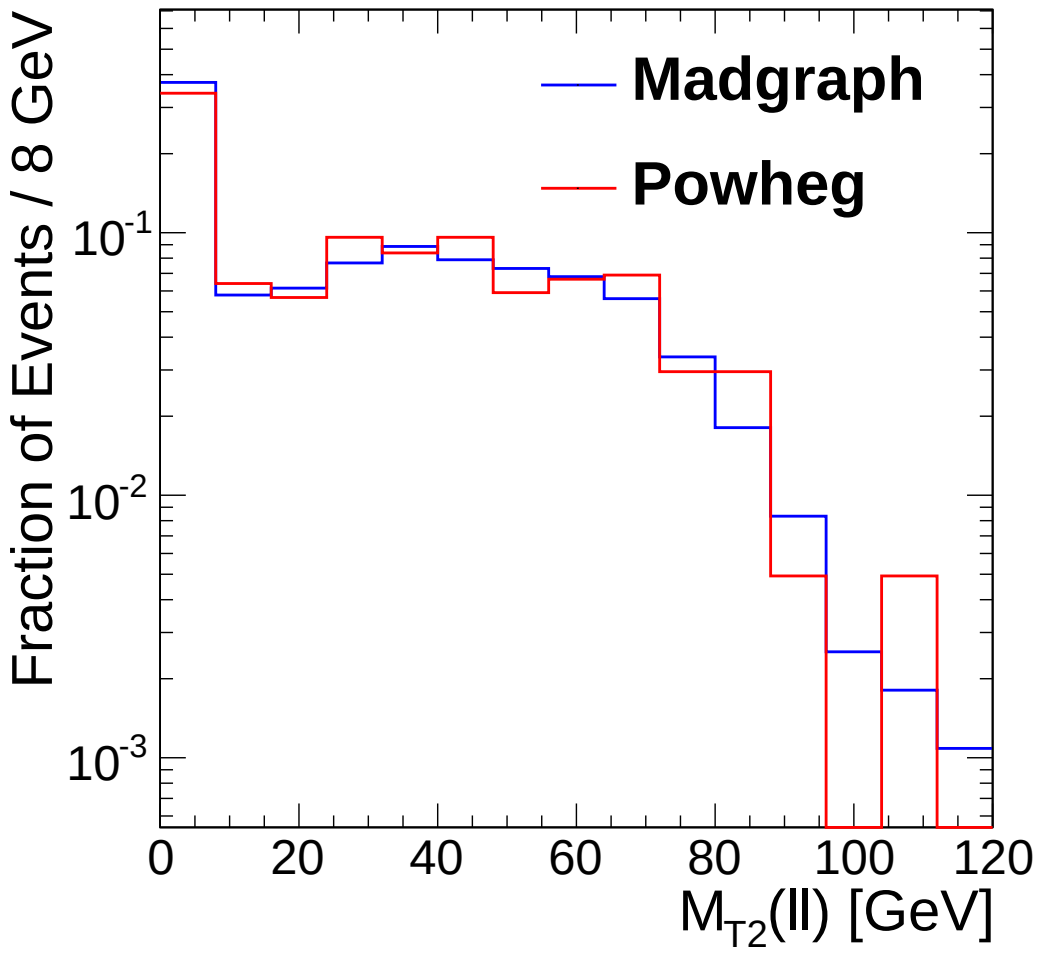

Figure 6.24: $M_{\mathrm{T2}}(\ell\ell)$ shape for WW events in MadGraph and powheg.

Summary of the Control Region Checks

From the aforementioned data/simulation comparison studies in control regions, we conclude that there is no evidence for significant inaccuracies in the simulation modeling of the high $M_{\mathrm{T2}}(\ell\ell)$ tail.

6.5.5 MC WW Shape with Width Cutoff

The default (MadGraph) WW MC has a hard-coded cutoff in the natural width of the W. We checked using a powheg sample that was generated without this cutoff to see if the $M_{\mathrm{T2}}(\ell\ell)$ shape is affected by this cutoff. We applied the full selection and then plotted the shape for each sample. The results are shown in Fig. 6.24. No significant difference is observed: the fraction of events passing the cut $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$ are $(3.2 \pm 0.3)\%$ and $(3.9 \pm 1.0)\%$ for MadGraph and powheg respectively, where the error is the MC statistical error.

6.5.6 Summary

To summarize the results of the $M_{\mathrm{T2}}(\ell\ell)$ shape analysis:

- The shape of the falling edge near m_{W} is primarily motivated by the intrinsic width of the W.
- The smearing of the shape due to $\not{E}_{\mathrm{T}}$ mismeasurements can be well approximated using only the Gaussian core of the $\not{E}_{\mathrm{T}}$ response in simulation, suggesting that unusually large mismeasurements play only a minor role.
- Control samples in data were used to establish that both the shape of the falling edge and the number of events in the far tails are well predicted by simulation when compared to the known systematic uncertainties.

As a result, we do not apply any additional systematic for the $M_{\mathrm{T2}}(\ell\ell)$ shape.

6.6 The Top-squark Signal

In this section, we discuss the aspects of this analysis directly related to our top-squark signal.

6.6.1 Calculating Signal Yields

An accurate and robust estimate of the signal efficiency in the SUSY mass plane (for each considered decay mode) is a key component of this analysis. In this section, we detail the calculations used to estimate these efficiencies.

Calculation of Efficiencies

In the remainder of this section, we will refer several times to calculating the efficiency, with respect to a specific cut, of a given point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, for a specific decay mode. Using S_0 to signify the total dataset of signal events for this chosen point and decay mode, we calculate the efficiency by taking the ratio of weighted events from S_0 passing the specific cut divided by the total number of generated events for S_0, where the weights applied are only to correct for the pileup reweighting (c.f. Section 6.2.4.4) and the ISR reweighting (c.f. Section 6.2.4.5).

Combination of Efficiencies

For both the T2tt and T2bw decay modes, as Tables 6.1 and 6.2 show, certain points in the 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane are covered by multiple simulation samples.

For these mass points, when we calculate the relevant cut efficiencies that we care about, we combine the information from these statistically orthogonal simulation samples in order to maximize the precision of our overall estimate of the signal efficiency efficiency. We do this by taking a variance-weighted average of the two available measurements of the efficiency,

$$\bar{\epsilon} = \frac{\sum_i \left(\epsilon_i \sigma_i^{-2}\right)}{\sum_i \left(\sigma_i^{-2}\right)}, \tag{6.15}$$

where σ_i is the statistical uncertainty on the efficiency measurement from sample i.

6.6.2 Calculation of the "Coarse" Efficiency

For a chosen top-squark decay mode and $M_{\mathrm{T2}}(\ell\ell)$ threshold, we calculated the efficiency for each mass point in the 2D $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane to pass the preselection as well as the chosen threshold.

Systematic uncertainties on this efficiency were calculated by using systematic variations of the signal simulation, as per the discussion in Section 6.4.

Figure 6.25 shows the estimated "coarse" (i.e. taken just from the individual simulation samples) efficiency for signal events to pass the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, for each of the unpolarized top-squark decay modes considered in this analysis. As a reference, the estimated efficiency for the $\mathrm{t\bar{t}}$ background to pass the same cut is $\approx 0.033\%$.

For each of these figures, there are two lines demarcating notable kinematic regions. For both the T2tt and T2bw decay modes, there is the region where the final state particles would be off-shell, $m_{\tilde{\mathrm{t}}_1} < m_{\tilde{\chi}_1^0} + m_{\mathrm{W}} + \mathrm{m_b}$. We did not consider

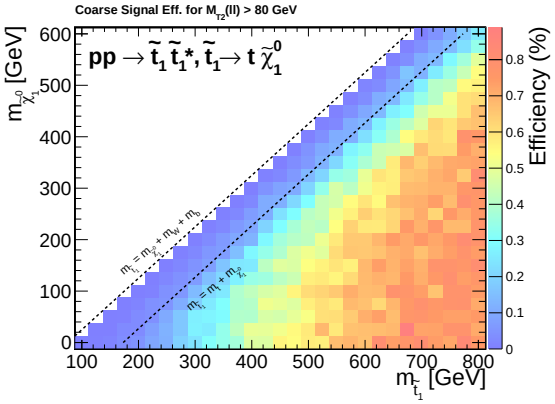

(a) The signal efficiency for the unpolarized T2tt decay mode.

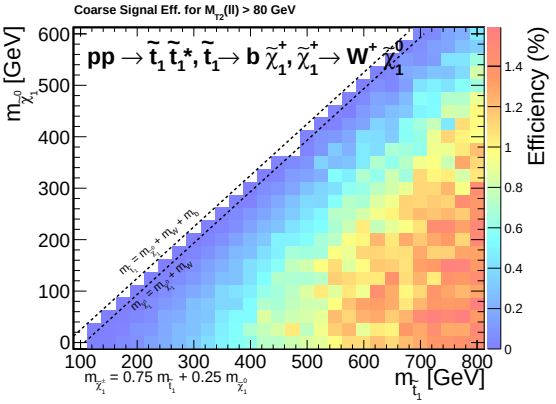

(b) The signal efficiency for the unpolarized T2bw decay mode with $x = 0.75$.

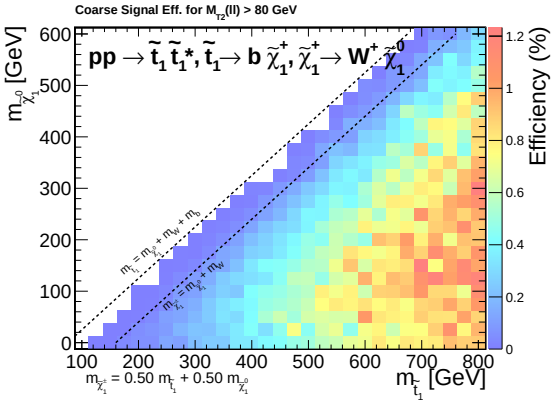

(c) The signal efficiency for the unpolarized T2bw decay mode with $x = 0.50$.

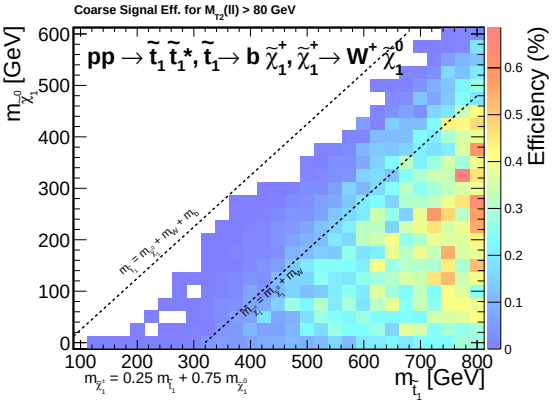

(d) The signal efficiency for the unpolarized T2bw decay mode with $x = 0.25$.

Figure 6.25: The signal efficiency for the cut, $M_{\mathrm{T2}}(\ell\ell) > 80$ GeV, for the four, unpolarized top-squark decay modes considered.

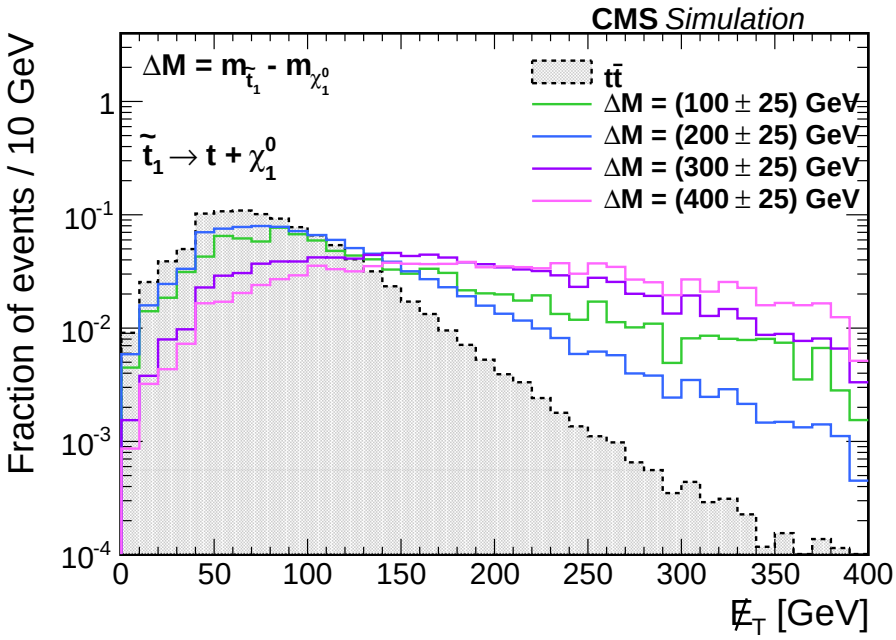

Figure 6.26: The distributions, at the preselection level, of $\not{E}_{\mathrm{T}}$ in the $\mathrm{t\bar{t}}$ background and the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}^0_1$ decay mode for different $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}^0_1}$ regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}^0_1})$ plane; distributions are normalized to the same area.

this region in our analysis. As noted in Section 3.2.6.3, in this region, the $\widetilde{\mathrm{t}}_1$ is expected to decay primarily to the final state, $\tilde{t}_1 \to c\tilde{\chi}^0_1$, so our analysis would have very little sensitivity.

For both the T2tt and T2bw decay modes, there is also the region where the decay products of the final decay involving SUSY particles would be off-shell. For the T2tt decay mode, this is $m_{\widetilde{\mathrm{t}}_1} < m_{\widetilde{\chi}^0_1} + m_{\mathrm{t}}$. For the T2bw decay mode, this is $m_{\widetilde{\chi}^\pm_1} < m_{\widetilde{\chi}^0_1} + m_{\mathrm{W}}$.

We turn now to discussing the primary driving factors for the shape and magnitude of the observed signal efficiencies.

6.6.2.1 Primary Factors that Influence the Signal Efficiency

As we noted in Section 6.2.5, the shape of the $M_{\mathrm{T2}}(\ell\ell)$ distribution depends upon the value of ΔM between the $\widetilde{\chi}_1^0$ and its mother, as this directly influences the momentum of the $\widetilde{\chi}_1^0$ and subsequently, the additional $\vec{\not{E}}_{\mathrm{T}}$ that it adds to the event. This can be seen in Fig. 6.26, which compares the normalized distributions of $\not{E}_{\mathrm{T}}$ in the $\mathrm{t\bar{t}}$ background and the T2tt decay mode for several $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0}$ regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane.

Moreover, the relative mass splitting between the SUSY mother and daughter particles for decay stages involving SUSY particles also directly influences the amount of momentum the visible SM particles receive, subsequently affecting the relative efficiency with which they pass our object selection requirements, and thereby directly influencing the expected number of signal events that pass our preselection. It should be noted that as $m_{\widetilde{\mathrm{t}}_1}$ grows, this relative effect coming from the changes of ΔM is not as strong. This is most readily visible in Fig. 6.25a. The simple explanation for this is that larger values of $m_{\widetilde{\mathrm{t}}_1}$ in general lead to higher momentum decay products overall, which leads to better selection efficiencies for the individual reconstructed objects.

Figure 6.27 compares the normalized p_{T} distributions of the leading lepton and leading b-tagged jet in the $\mathrm{t\bar{t}}$ background and the T2bw decay mode with $x = 0.25$ for several $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0}$ regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane. The full set of object selection requirements have been applied to both for these plots. The shape of the lead-lepton p_{T} in the $\Delta\mathrm{M} = 100$–$200\,\mathrm{GeV}$ regions implies that a significant fraction

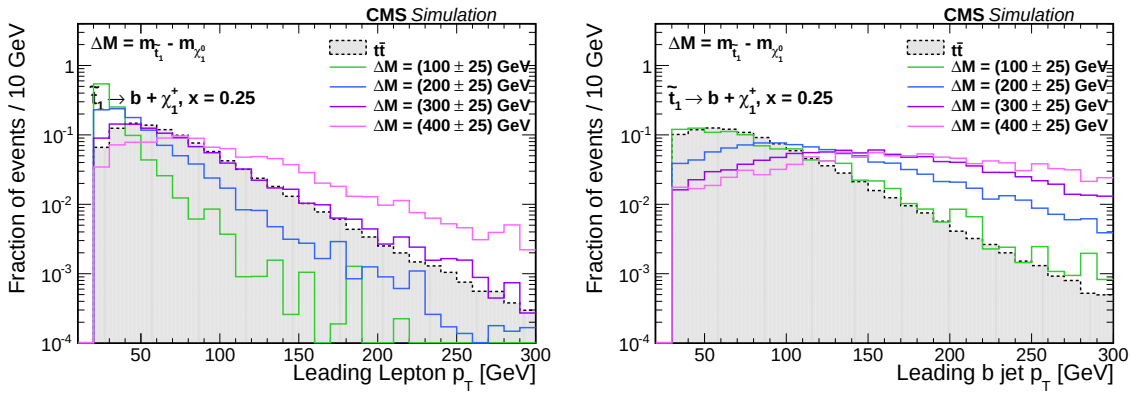

(a) The leading lepton p_{T}

(b) The leading b-tagged jet p_{T}

Figure 6.27: The distributions, at the preselection level, of the p_{T} for the leading lepton and leading b-tagged jet in the $\mathrm{t\bar{t}}$ background and the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode with chargino mass-splitting parameter $x = 0.25$ for different $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0}$ regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane; distributions are normalized to the same area.

of the top-squark events from these regions likely failed to pass the preselection due to low lepton p_{T}, It is important to note that, since they are input objects for the $M_{\mathrm{T2}}(\ell\ell)$ calculation, the shape of the leptons' p_{T} distributions also directly influences the shape of the $M_{\mathrm{T2}}(\ell\ell)$ distribution.

The ΔM value can also influence the relative angular correlations between the input objects for the $M_{\mathrm{T2}}(\ell\ell)$ calculations, subsequently influencing $M_{\mathrm{T2}}(\ell\ell)$ as per the discussion in Section 6.2.5. Figure 6.28, which is set up analogously to Figs. 6.26 and 6.27 except it shows the distribution of $\Delta\phi(\ell^+, \ell^-)$, demonstrates that, of these two top-squark decay modes shown, the ΔM value only significantly influences the $\Delta\phi$ distribution in the T2tt decay mode.

With those discussion points in hand, we can then clarify the major reasons for the relatively strong differences in calculated efficiency for each of the investigated top-squark decay modes.

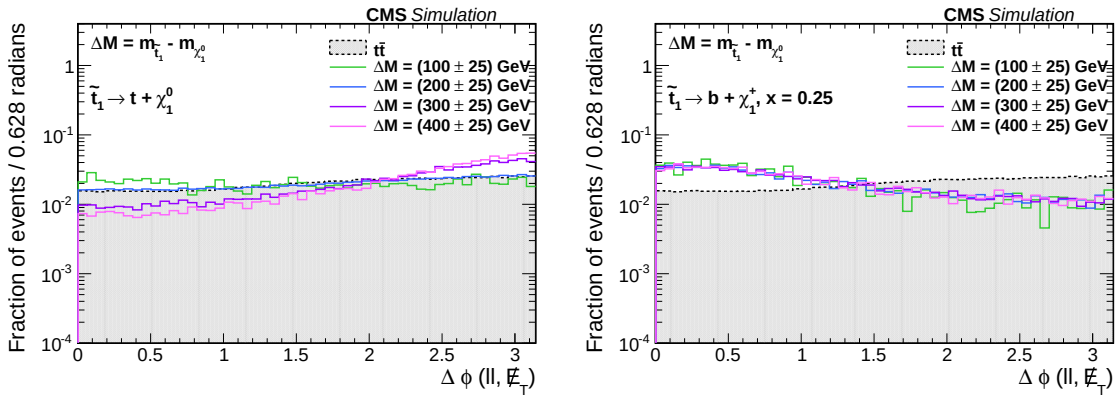

(a) $M_{\mathrm{T2}}(\ell\ell)$ vs. $\Delta\phi(\ell^+, \ell^-)$ for the T2tt decay mode.

(b) $M_{\mathrm{T2}}(\ell\ell)$ vs. $\Delta\phi(\ell^+, \ell^-)$ for the T2bw decay mode with $x = \frac{m_{\tilde{\chi}_1^+} - m_{\tilde{\chi}_1^0}}{m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}} = 0.25$.

Figure 6.28: The distributions, at the preselection level, of $\Delta\phi(\ell^+\ell^-, \not{E}_{\mathrm{T}})$ in the $\mathrm{t\bar{t}}$ background and two of the top-squark decay modes, shown for different $\Delta\mathrm{M} = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$ regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane; distributions are normalized to the same area.

6.6.2.2 Understanding the Relative Differences in Top-squark Decay Mode Efficiencies

Figures 6.30 and 6.29 compare the normalized distributions of several variables in the $\mathrm{t\bar{t}}$ and different top-squark decay modes for $\Delta\mathrm{M} = 300\,\mathrm{GeV}$.

The shapes of these variables' distributions, as per the discussions above, influences the absolute signal efficiency for the considered top-squark decay modes to fall into the region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$. In particular, from investigations of these figures and similar ones for other $\Delta\mathrm{M}$ regions (which we do not show here), we can summarize the salient differences between the different top-squark decay modes:

- The observed relative differences in efficiency for the various top-squark decay modes are not significantly driven by the magnitude of the $\not{E}_{\mathrm{T}}$, as the shape of the $\not{E}_{\mathrm{T}}$ is fairly consistent across the top-squark decay modes.

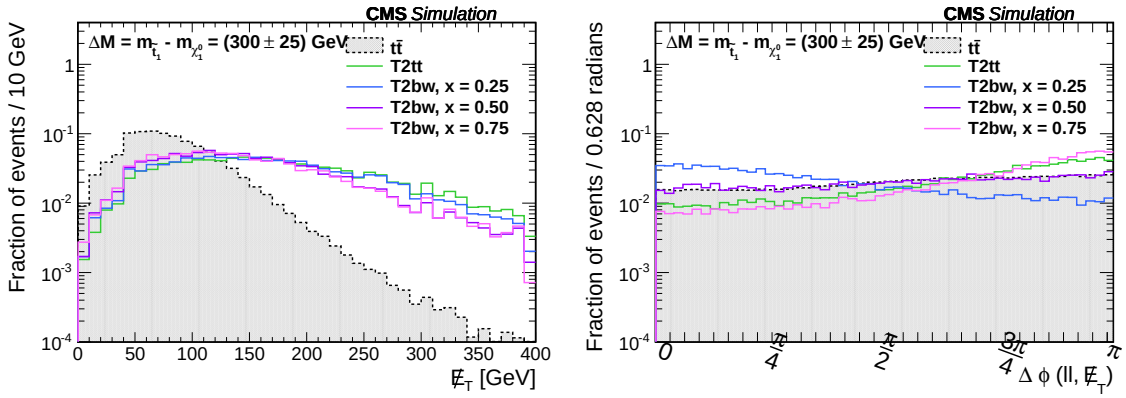

(a) The reconstructed $\not{E}_{\mathrm{T}}$

(b) The opening angle between the dilepton system and the $\vec{\not{E}}_{\mathrm{T}}$, $\Delta\phi\,(\ell^{+}\ell^{-}, \not{E}_{\mathrm{T}})$

Figure 6.29: The distributions at the preselection level, of $\not{E}_{\mathrm{T}}$ and $\Delta\phi\,(\ell^{+}\ell^{-}, \not{E}_{\mathrm{T}})$ in the $\mathrm{t\bar{t}}$ background and different top-squark decay modes for a ΔM value of 300 GeV; distributions are normalized to the same area.

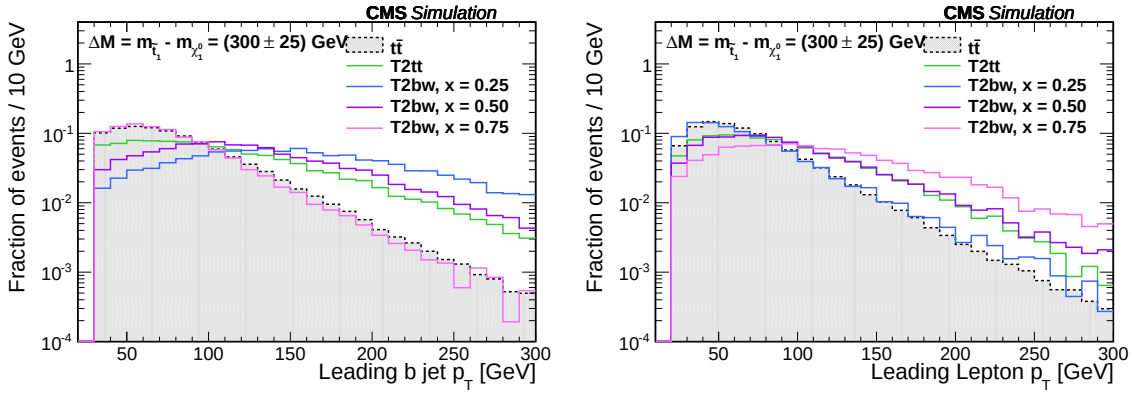

(a) The leading b-tagged jet p_{T}

(b) The leading lepton p_{T}

Figure 6.30: The distributions at the preselection level, of leading selected visible object p_{T} in the $\mathrm{t\bar{t}}$ background and different top-squark decay modes for a ΔM value of 300 GeV; distributions are normalized to the same area.

- The shape of the $\Delta\phi\,(\ell^+,\,\ell^-)$ distribution for the T2tt and T2bw, $x = 0.75$, decay mode favor larger values of $M_{\mathrm{T2}}\,(\ell\ell)$ relative to the $\mathrm{t\bar{t}}$ background, while the analogous shape in the T2bw, $x = 0.25$ decay mode favors smaller values of $M_{\mathrm{T2}}\,(\ell\ell)$. The $x = 0.50$ T2bw decay mode falls in the middle of these two bounds.

- The b-tagged jets have characteristically higher p_{T} for the lower x T2bw decay modes, meaning that less events from these scenarios fail the b-tagged jet requirement of the preselection.

- The shape of the leading b-tagged jet p_{T} for the $x = 0.75$ T2bw decay mode is notably softer than the other top-squark decay modes. However, to contrast this, the leptons for this decay mode have characteristically higher p_{T} relative to all the other top-squark decay modes considered. This not only means that less T2bw, $x = 0.75$ events fail the lepton requirement of the preselection, but furthermore, the $M_{\mathrm{T2}}\,(\ell\ell)$ value is also characteristically higher in these events because of the higher lepton p_{T}.

To summarize this information into one statement, the signal efficiency in the $M_{\mathrm{T2}}\,(\ell\ell) > 80\,\mathrm{GeV}$ region is higher for the $x = 0.75$ T2bw decay mode relative to the other top-squark decay modes because the events from this decay mode pass the preselection requirements more often, the input leptons for the $M_{\mathrm{T2}}\,(\ell\ell)$ calculation are characteristically higher-p_{T}, and the angular configurations between the leptons and the $\not{E}_{\mathrm{T}}$ in these events tend to favor larger $M_{\mathrm{T2}}\,(\ell\ell)$ values.

Similar arguments can be used to explain the relative differences in efficiency

for the other top-squark decay modes.

6.6.3 Calculation of the Signal Yield

For a chosen top-squark decay mode and $M_{\mathrm{T2}}(\ell\ell)$ threshold, we calculated the estimated signal yield at the nominal SUSY cross section for each mass point in the 2D $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane using the following equation,

$$N_s\left(m_{\tilde{\mathrm{t}}_1},\, m_{\tilde{\chi}_1^0},\, M_{\mathrm{T2}}(\ell\ell)\right) = \mathcal{L}\,\sigma_{\tilde{\mathrm{t}}_1\tilde{\mathrm{t}}_1*}\left(m_{\tilde{\mathrm{t}}_1}\right)\,\epsilon\left(m_{\tilde{\mathrm{t}}_1},\, m_{\tilde{\chi}_1^0},\, M_{\mathrm{T2}}(\ell\ell)\right). \tag{6.16}$$

As a reminder, Fig. 6.3 shows the dependence of the 8 TeV top-squark pair production cross section on $m_{\tilde{\mathrm{t}}_1}$; the integrated luminosity of our collision data set is 19.66 ± 0.51 fb^{-1};

Figure 6.31 shows the expected number of signal events passing the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, for the four, unpolarized top-squark decay modes considered in this analysis. For purposes of visualization, a 2D Gaussian kernel-based smoothing, Appendix H.2.3, was applied to the signal efficiencies before calculating signal yields.

Figure 6.32 shows, for various points in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane for the unpolarized T2tt top-squark decay mode, the relative magnitude of the systematic uncertainties on the estimated number of top-squark events that pass the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$. For most of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, the systematic uncertainties are at the level of 5–10%. For smaller ΔM values, however, the magnitude of the systematic uncertainties grows, up to relative values of 50% for the lowest ΔM regions (note that the y-axis maximum of Fig. 6.32 has been set at 1.25). When looking at the other top-

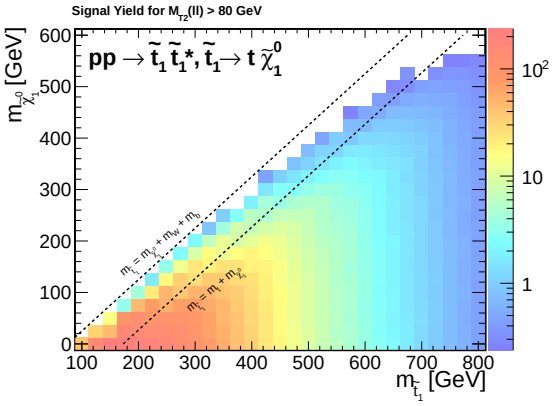

(a) The expected number of signal events for the unpolarized T2tt decay mode.

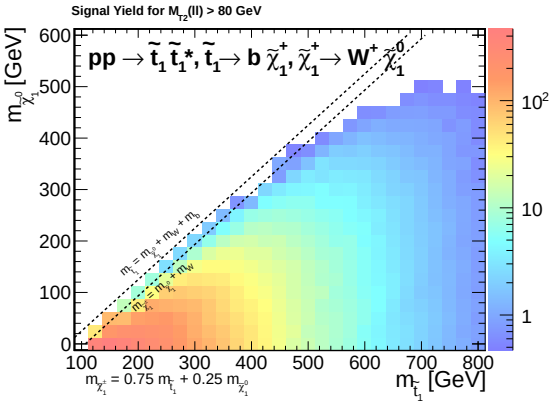

(b) The expected number of signal events for the unpolarized T2bw decay mode with $x = 0.75$.

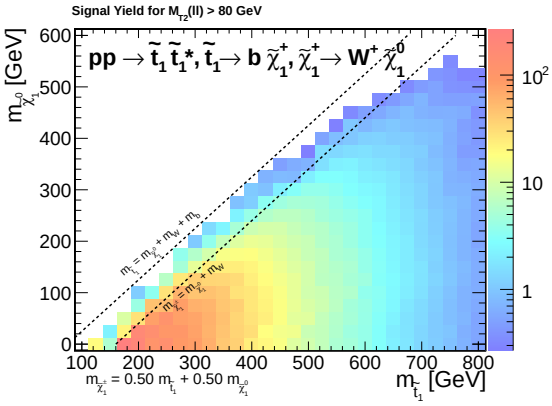

(c) The expected number of signal events for the unpolarized T2bw decay mode with $x = 0.50$.

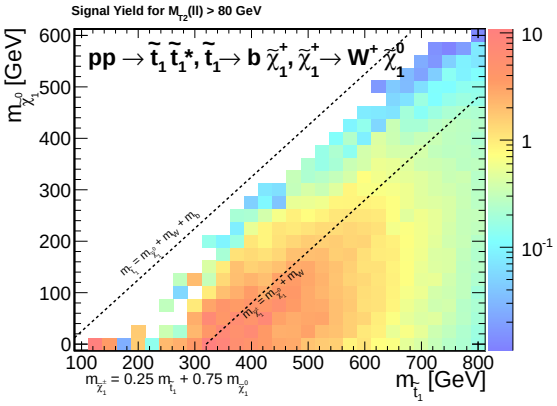

(d) The expected number of signal events for the unpolarized T2bw decay mode with $x = 0.25$.

Figure 6.31: The expected number of signal events passing the cut, $M_{\mathrm{T2}}(\ell\ell) > 80$ GeV, for the four, unpolarized top-squark decay modes considered. For purposes of visualization, a 2D Gaussian kernel-based smoothing, Appendix H.2.3, was applied to the signal efficiencies before calculating signal yields.

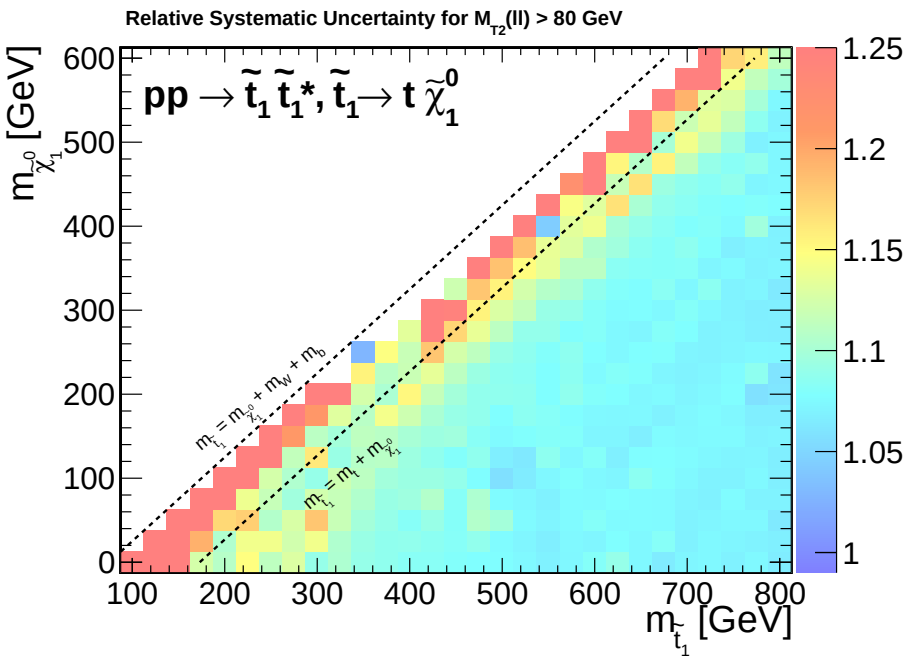

Figure 6.32: The relative change, due to systematic uncertainties on the signal efficiency, Section 6.4, in the expected number of signal events that pass the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, for the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode.

squark decay modes, we found that the systematic uncertainties had similar relative magnitudes, with similar dependencies on ΔM.

6.6.4 Calculating the Magnitude of Signal Contamination

In Sections 6.3.1.2, 6.3.2.5, and 6.3.3.5 we discussed in general terms the expected impact of signal contamination on the data-driven estimations described in each of the respective sections. We noted how the data-driven estimation of the $\mathrm{t}\bar{\mathrm{t}}$ was the only data-driven estimation where signal contamination could notably impact the final results of this analysis.

6.6.5 Calculating the Signal Contamination in the $t\bar{t}$ Estimation

In this section we provide details on how we calculated the impact of the signal contamination on the $t\bar{t}$ estimation.

As a reminder, Eqs. (6.7) and (6.8) in Section 6.3.1.2 show respectively how we calculate the scale factor used to normalize the expected $t\bar{t}$ contribution in our general signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$.

Equation (6.7) contains the implicit assumption,

$$N^{\mathrm{ctrl}}_{\mathrm{Data}} = N^{\mathrm{ctrl}}_{t\bar{t}} + N^{\mathrm{ctrl}}_{\text{Non-}t\bar{t}\text{ Bkg.}}$$

If the signal exists, this assumption is obviously no longer valid, and instead must be altered to,

$$N^{\mathrm{ctrl}}_{\mathrm{Data}} = N^{\mathrm{ctrl}}_{t\bar{t}} + N^{\mathrm{ctrl}}_{\mathrm{signal}} + N^{\mathrm{ctrl}}_{\text{Non-}t\bar{t}\text{ Bkg.}}$$

Subsequently, we can change Eq. (6.7),

$$\begin{aligned} \mathrm{SF}'_{\mathrm{DD}\,t\bar{t}} &= \frac{N^{\mathrm{ctrl}}_{\mathrm{Data}} - N^{\mathrm{ctrl}}_{\text{Non-}t\bar{t}\text{ Bkg.}} - N^{\mathrm{ctrl}}_{\mathrm{signal}}}{N^{\mathrm{ctrl}}_{t\bar{t}\ \mathrm{MC}}} \\ &= \mathrm{SF}_{\mathrm{DD}\,t\bar{t}} - \frac{N^{\mathrm{ctrl}}_{\mathrm{signal}}}{N^{\mathrm{ctrl}}_{t\bar{t}\ \mathrm{MC}}} \end{aligned} \tag{6.17}$$

which leads to a change in Eq. (6.8),

$$N'^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{DD}} = \mathrm{SF}'_{\mathrm{DD}\,t\bar{t}} N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{MC}} \quad = N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{DD}} - \frac{N^{\mathrm{ctrl}}_{\mathrm{signal}} N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{MC}}}{N^{\mathrm{ctrl}}_{t\bar{t}\ \mathrm{MC}}} \tag{6.18}$$

From Eq. (6.18), the overall effect of the signal contamination is to reduce the estimate of the $t\bar{t}$ background by,

$$\frac{N^{\mathrm{ctrl}}_{\mathrm{signal}} N^{\mathrm{sig.}}_{t\bar{t}\ \mathrm{MC}}}{N^{\mathrm{ctrl}}_{t\bar{t}\ \mathrm{MC}}}.$$

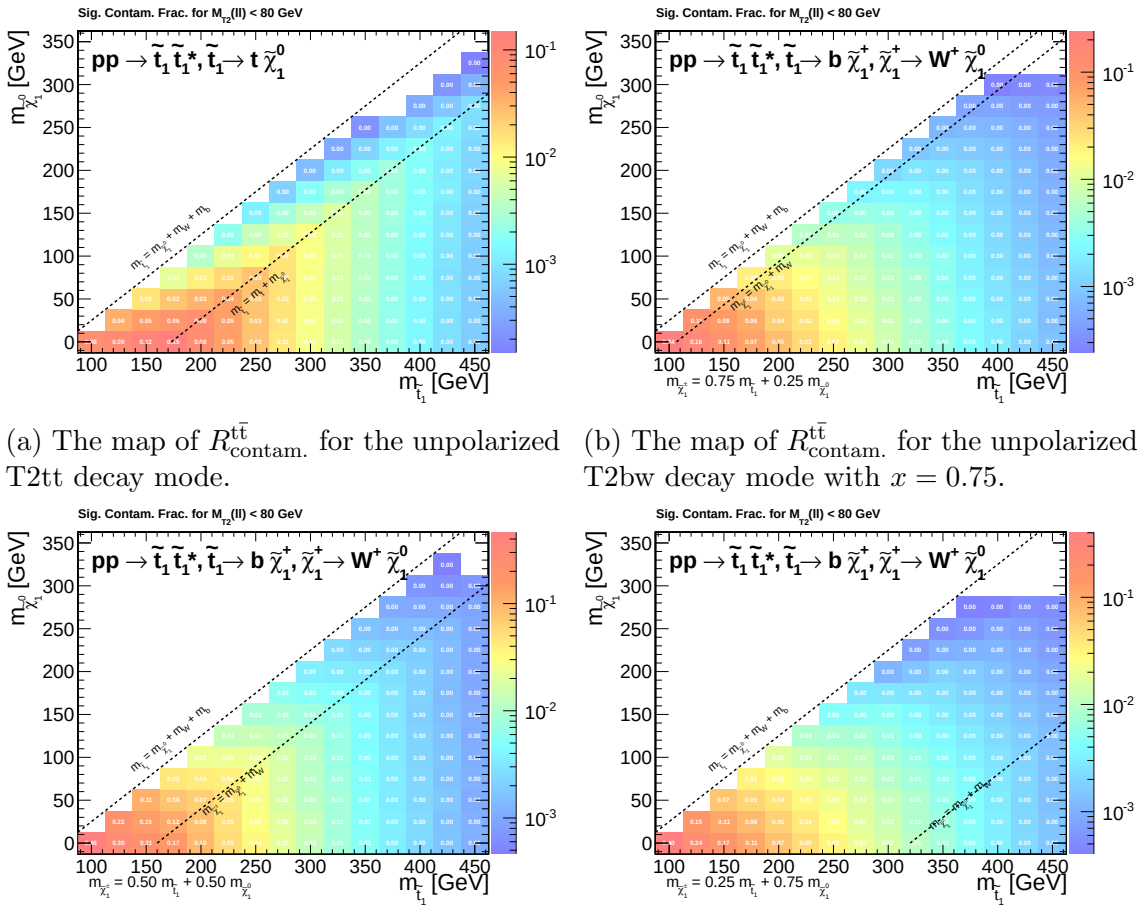

(a) The map of $R^{t\bar{t}}_{\text{contam.}}$ for the unpolarized T2tt decay mode.

(b) The map of $R^{t\bar{t}}_{\text{contam.}}$ for the unpolarized T2bw decay mode with $x = 0.75$.

(c) The map of $R^{t\bar{t}}_{\text{contam.}}$ for the unpolarized T2bw decay mode with $x = 0.50$.

(d) The map of $R^{t\bar{t}}_{\text{contam.}}$ for the unpolarized T2bw decay mode with $x = 0.25$.

Figure 6.33: The relative signal contamination in the control region, $M_{\text{T2}}(\ell\ell) < 80\,\text{GeV}$, for the four, unpolarized top-squark decay modes considered. The relative signal contamination is written in terms of the quantity, $R^{t\bar{t}}_{\text{contam.}}$, c.f. Eq. (6.19).

Calculations of the impact of signal contamination per point of the 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass plane

Figure 6.33 shows the relative level of signal contamination in the $t\bar{t}$ control region, $M_{\text{T2}}(\ell\ell) < 80\,\text{GeV}$, for the four, unpolarized top-squark decay modes considered in this analysis. In each sub-figure, the value for each point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$

plane is the ratio, $R^{\mathrm{t\bar{t}}}_{\mathrm{contam.}}$,

$$R^{\mathrm{t\bar{t}}}_{\mathrm{contam.}} = \frac{N^{\mathrm{ctrl}}_{\mathrm{signal}}}{N^{\mathrm{ctrl}}_{\mathrm{t\bar{t}\ \ MC}}}. \tag{6.19}$$

As can be seen, the relative level of signal contamination is quite small for a large part of the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}^0_1})$ mass plane, but becomes nontrivial — $\mathcal{O}\,(10\%)$ — around low $m_{\widetilde{\mathrm{t}}_1}$.

We properly account for this changed estimate of the $\mathrm{t\bar{t}}$ background contribution when we set upper limits on the top-squark pair production cross section in Section 6.8.

6.6.6 Dealing with the Handedness of SUSY Particles and Couplings in $\mathrm{pp} \to \widetilde{\mathrm{t}}_1\widetilde{\mathrm{t}}^*_1$ Production

The relative couplings of top-squarks to left- and right-handed daughter particles are model dependent, c.f. Section 3.2.6.3 or Refs. [38, 39]; the signal simulation samples we used assumed unpolarized couplings between the $\widetilde{\mathrm{t}}_1$ and its daughter particles (as well as unpolarized couplings between the $\widetilde{\chi}^{\pm}_1$ and its daughter particles in the T2bw decay mode).

However, as we have discussed before, the $M_{\mathrm{T2}}\,(\ell\ell)$ value in dilepton events has a strong dependence upon the angular configuration of the two leptons in the transverse plane. Subsequently, the relative efficiency for signal events to pass $M_{\mathrm{T2}}\,(\ell\ell)$ cuts has a strong dependence on the handedness of the couplings between the top-squark and its daughter particles (and the $\widetilde{\chi}^{\pm}_1$ and its daughter particles in the T2bw decay mode).

We investigated the impact of this effect on our analysis's final results (see Section 6.8.3.1) by reweighting the unpolarized samples to match several different scenarios. The technical details on the actual weight calculations are discussed in Appendix D.

For the T2tt decay mode, we investigated fully left- and fully right-handed couplings between the $\widetilde{\mathrm{t}}_1$ and its daughter t.

For the T2bw decay mode, we investigated fully left- and fully right-handed couplings between the $\widetilde{\mathrm{t}}_1$ and its daughter $\widetilde{\chi}_1^{\pm}$, with the $\widetilde{\chi}_1^{\pm}$ then either decaying to a left- or right-handed W boson.

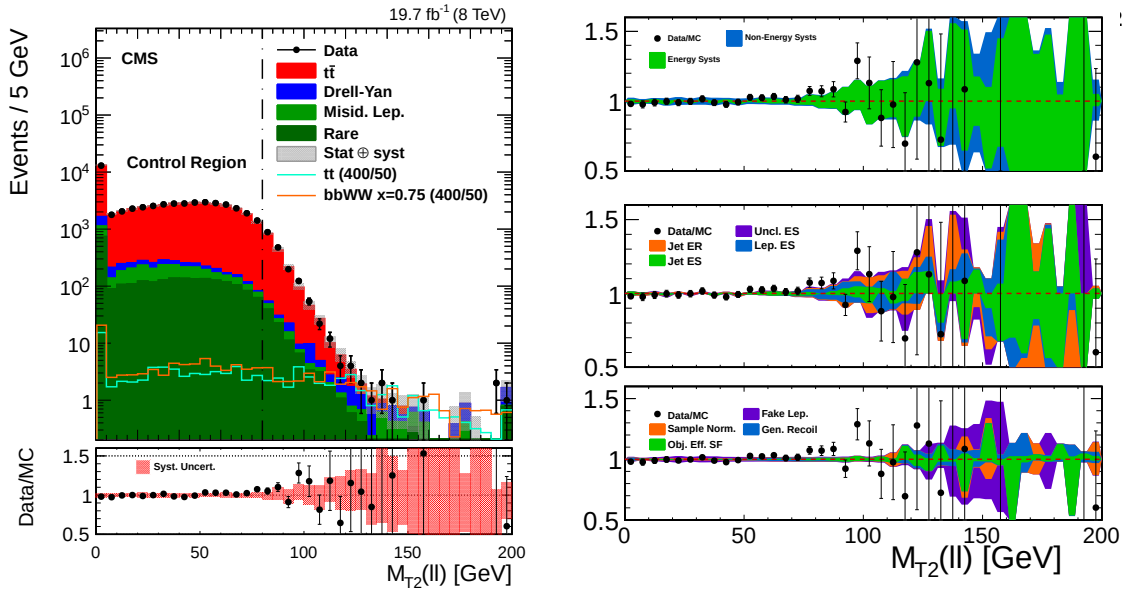

(a) (b)

Figure 6.34: The full unblinded $M_{\mathrm{T2}}(\ell\ell)$ distribution. In 6.34a, "tt" and "bbww x = 0.75" refer to example $M_{\mathrm{T2}}(\ell\ell)$ distributions from the point (400,50) GeV in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, where "tt" represents the unpolarized T2tt decay mode and "bbww x = 0.75" represents the unpolarized, $x = 0.75$ T2bw decay mode. Shown on the right are itemized breakdowns of the contributions from individual sources of systematic uncertainty to the overall level of data and simulation agreement.

6.7 Results

Figure 6.34 shows a comparison between the full, unblinded $M_{\mathrm{T2}}(\ell\ell)$ distribution in the collision data and the expected composite distribution from background contributions. Also shown in Fig. 6.34 are example distributions from the point $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (400\,\mathrm{GeV}, 50\,\mathrm{GeV})$ for two possible decay modes of the top-squark: "tt" represents the expected $M_{\mathrm{T2}}(\ell\ell)$ distribution for this mass point from the unpolarized T2tt decay mode, while "bbww x = 0.75" represents the expected $M_{\mathrm{T2}}(\ell\ell)$ distribution for this mass point from the unpolarized $x = 0.75$ T2bw decay mode.

Table 6.12 shows the background predictions of Table 6.9 along with the ob-

Table 6.12: Data yields and background expectation for the five different $M_{\mathrm{T2}}(\ell\ell)$ cut values used in this analysis. The asymmetric uncertainties quoted for the background indicate the total systematic uncertainty, including the statistical uncertainty on the background expectation.

$M_{\mathrm{T2}}(\ell\ell)$ threshold	80 GeV	90 GeV	100 GeV	110 GeV	120 GeV
Data	1785	427	106	30	14
Expected background	$1677.7^{+120.6}_{-99.2}$	$411.6^{+52.6}_{-34.0}$	$100.4^{+19.7}_{-9.0}$	$31.2^{+5.4}_{-3.8}$	$14.0^{+3.5}_{-2.5}$

served number of events in the unblinded signal region in data.

As you can see, there is no clear excess of events in the collision data; furthermore, the agreement between the data and the background prediction is excellent throughout the signal region. Thus, we can proceed to set limits to exclude regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass plane. First, though, we provide details on some thorough examinations we performed on the few events with high $M_{\mathrm{T2}}(\ell\ell)$ that are observed in the collision data.

6.7.1 Examination of Data Events with High $M_{\mathrm{T2}}(\ell\ell)$

We chose three events in data with $M_{\mathrm{T2}}(\ell\ell) > 190\,\mathrm{GeV}$ to examine in detail, as a cross-check in case the large observed value was due to unexpected detector effects or reconstruction failures. We note that the number of high $M_{\mathrm{T2}}(\ell\ell)$ events is predicted with good accuracy by the simulation, so we expect to find mostly genuine physics events where well-understood detector acceptance or resolution effects have introduced spuriously high $\not{E}_{\mathrm{T}}$ values.

For all events, we checked the following:

- Consistency between PF $\not{E}_{\mathrm{T}}$ and Calo $\not{E}_{\mathrm{T}}$ (tracking failures can introduce high

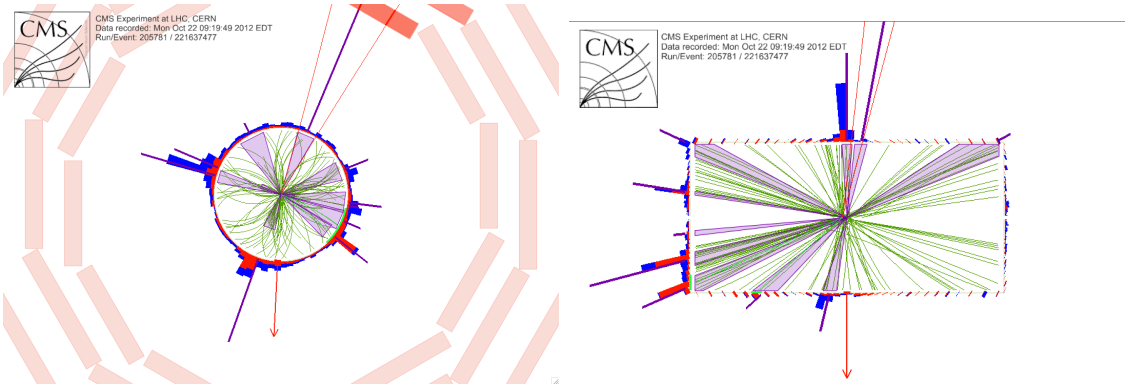

Figure 6.35: High-$M_{\mathrm{T2}}(\ell\ell)$ event number 1. This is a $\mu^+\mu^-$ event with $M_{\mathrm{T2}}(\ell\ell) =$ 190 GeV. The opening angle between the $\not{E}_{\mathrm{T}}$ in this event (180 GeV), and the dilepton system ($M_{\ell^+\ell^-} = 43$ GeV) is 2.7 radians. There are six jets with $p_{\mathrm{T}} >$ 50 GeV.

momentum PF candidates that do not affect Calo $\not{E}_{\mathrm{T}}$).

- Consistency of jet constituents with the primary vertex of the event (reconstruction failures can lead to tracks with unphysical impact parameters).

- The presence of jets with unusually high charged-energy fractions (very close to one), as these can be the product of known issues with the reconstruction.

All the events inspected in data passed these checks.

We show these events in Figs. 6.35 through 6.37.

In these pictures, ECAL tower cells are represented by red blocks; HCAL tower cells are represented by blue blocks; reconstructed tracks are green; the $\vec{\not{E}}_{\mathrm{T}}$ is the red arrow; clustered jets are represented by purple lines with associated purple triangles; muons are long red lines; finally, electrons are the bright green lines.

The first event is shown in Figure 6.35. The $M_{\mathrm{T2}}(\ell\ell)$ for this event is 190 GeV. It is a $\mu^+\mu^-$ event containing six jets that have $p_{\mathrm{T}} > 50$ GeV. The large value of $M_{\mathrm{T2}}(\ell\ell)$ comes from the $\not{E}_{\mathrm{T}}$ pointing opposite to the high-p_{T} $\mu^+\mu^-$ system. The

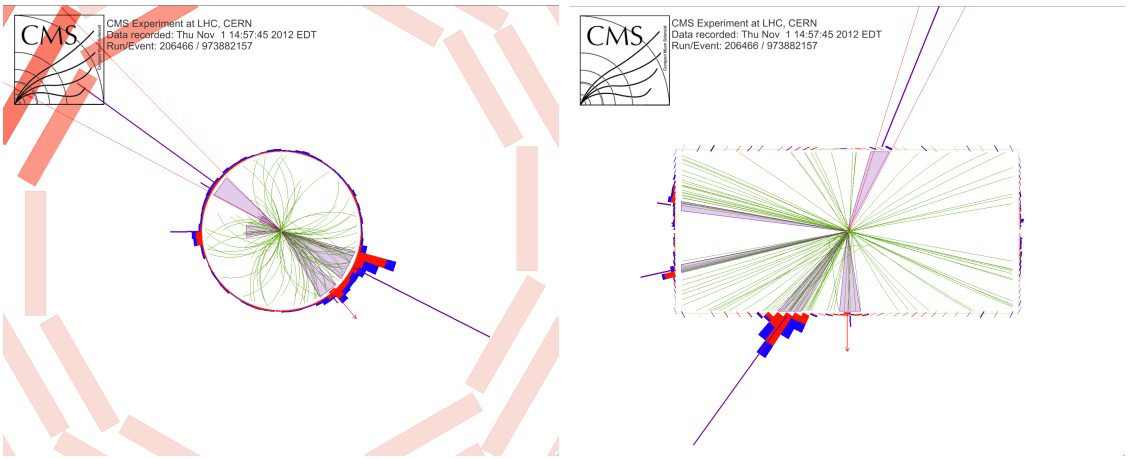

Figure 6.36: High-$M_{\mathrm{T2}}(\ell\ell)$ event number 2. This is a $\mu^+\mu^-$ event with $M_{\mathrm{T2}}(\ell\ell) = 190$ GeV. The opening angle between the $\not{E}_{\mathrm{T}}$ in this event (100 GeV), and the dilepton system ($M_{\ell^+\ell^-} = 75$ GeV) is 2.9 radians. There are three jets with $p_{\mathrm{T}} > 50$ GeV.

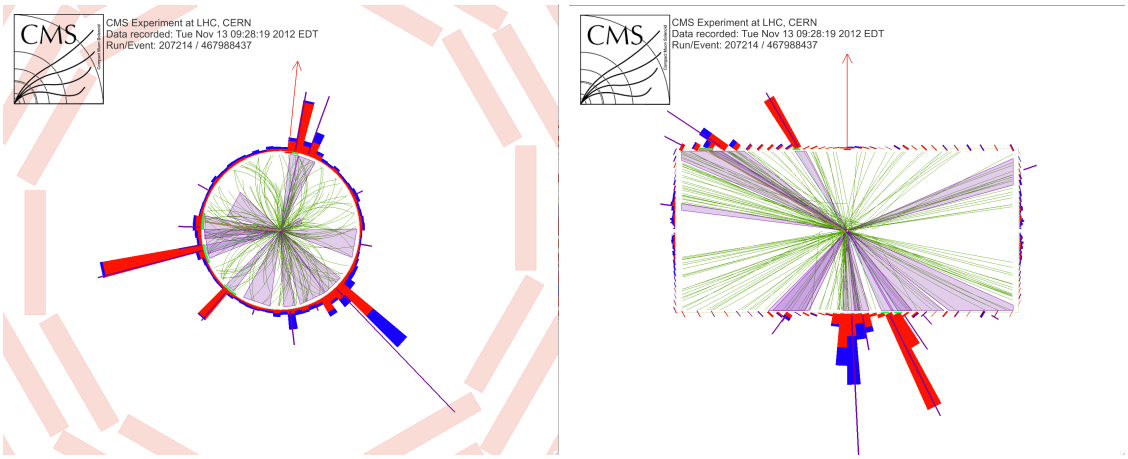

Figure 6.37: High-$M_{\mathrm{T2}}(\ell\ell)$ event number 3. This is an $\mathrm{e}^+\mathrm{e}^-$ event with $M_{\mathrm{T2}}(\ell\ell) = 270$ GeV. There are seven jets with $p_{\mathrm{T}} > 50$ GeV.

mass of the dilepton system is 43 GeV, the $\not{E}_{\mathrm{T}}$ is 140 GeV, and the angle between them is 2.7 radians. The $\not{E}_{\mathrm{T}}$ points near a high (180 GeV) p_{T} jet that is almost perfectly back-to-back with the dilepton system. One possible explanation for this event is that this recoiling jet is badly mismeasured. Close inspection of the individual objects in the event did not reveal any irregularities.

The second event is shown in Figure 6.36. The $M_{\mathrm{T2}}(\ell\ell)$ for this event is also 190 GeV. It is a $\mu^+\mu^-$ event with three jets with p_{T} above 50 GeV. Again, the $\not{E}_{\mathrm{T}}$ points opposite to the high-p_{T} $\mu^+\mu^-$ system. The mass of the dilepton system is 75 GeV, falling just outside the Z veto window, which starts at 76 GeV. The $\not{E}_{\mathrm{T}}$ is 100 GeV and the angle between the leptons and the $\not{E}_{\mathrm{T}}$ is 2.9 radians. Likely this is a $\mathrm{Z} \to \mu^+\mu^-$ event where the hadronic recoil is mismeasured.

The third event is shown in Figure 6.37. It is an $\mathrm{e}^+\mathrm{e}^-$ event, but the electrons (in light green) are not easily seen due to the large multiplicity of high p_{T} particles in this event. The $M_{\mathrm{T2}}(\ell\ell)$ for this event is about 270 GeV, a remarkable value. The event had an extremely large amount of activity with seven jets above 50 GeV. The invariant mass of the electron pair is 106.3 GeV, falling just above the Z-mass veto window, which ends at 106 GeV. The $\not{E}_{\mathrm{T}}$ is aligned near two of the high p_{T} recoiling jets, so it is likely that this is a $\mathrm{Z} \to \mathrm{e}^+\mathrm{e}^-$ event where the hadronic recoil is mismeasured.

All three events are same-flavor as anticipated by the simulation, as high jet-multiplicity $\mathrm{Z} \to \ell^+\ell^- + \mathrm{X}$ events contribute to the tails much more in the same-flavor channels.

6.8 Limit Setting

In this section we interpret the results of our search in the context of several models for top-squark pair production.

As a reminder, we considered two possible decay modes of the top-squark: T2tt, i.e. $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$, and T2bw, i.e. $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$. For the T2bw decay mode, the mass of the intermediate chargino, $m_{\widetilde{\chi}_1^+}$, is defined by the chargino mass-splitting parameter, $x = \frac{m_{\widetilde{\chi}_1^+} - m_{\widetilde{\chi}_1^0}}{m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0}}$. We considered three values for x, $x = 0.25$, $x = 0.50$, and $x = 0.75$.

In this section we provide an overview on how, for each signal point in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane, we set statistical limits on the top-squark pair-production cross section for a hypothetical signal at that point. By comparing these cross section limits against the theory-calculated cross section, we subsequently exclude regions of the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ mass plane.

The basic notion is that we are performing a likelihood comparison between two different hypotheses: the *background*-only hypothesis, represented by B and the *signal + background* hypothesis, represented by $\mu S + B$, where μ is the signal strength parameter: the ratio of the observed pp $\to \widetilde{\mathrm{t}}_1\widetilde{\mathrm{t}}_1^*$ cross section divided by the theory-calculated cross section,

$$\mu = \frac{\sigma^{\mathrm{obs.}}_{\widetilde{\mathrm{t}}_1\widetilde{\mathrm{t}}_1*}}{\sigma^{\mathrm{calc.}}_{\widetilde{\mathrm{t}}_1\widetilde{\mathrm{t}}_1*}},$$

6.8.1 Accounting for the Signal Contamination in the $t\bar{t}$ Estimation

In Section 6.6.5, we provided details on how we calculated the impact of the signal contamination on the data-driven normalization of the $t\bar{t}$ background. In particular, the effect is vanishingly small for most of the 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass plane, but grows to be as large as 30% at very low $m_{\tilde{t}_1}$.

In the B-only hypothesis, i.e. signal $\mu = 0$, there's no change needed ($B \to B$). In the $\mu S + B$ hypothesis, the signal contamination changes the background estimate $B \to B'$. This, alternatively can be represented in terms of changing the signal estimate in accordance with Eq. (6.18),

$$\begin{aligned}
\mu S_{\mathrm{s}} + B_{\mathrm{s}} \to \mu S_{\mathrm{s}} + B_{\mathrm{s}}^{'} &= \mu S_{\mathrm{s}} + \left(B_{\mathrm{s}}^{'} - B_{\mathrm{s}}\right) + B_{\mathrm{s}} \\
&= \mu S_{\mathrm{s}} + \left(B_{\mathrm{s}} - B_{\mathrm{s}}\frac{\mu S_{\mathrm{c}}}{B_{\mathrm{c}}} - B_{\mathrm{s}}\right) + B_{\mathrm{s}} \\
&= \mu\left(S_{\mathrm{s}} - B_{\mathrm{s}}\frac{S_{\mathrm{c}}}{B_{\mathrm{c}}}\right) + B_{\mathrm{s}} \\
&= \mu S_{\mathrm{s}}^{'} + B_{\mathrm{s}},
\end{aligned} \tag{6.20}$$

where S_{c} and B_{c} represent respectively the yield of signal and background events in the $t\bar{t}$ control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, and S_{s} and B_{s} are the analogous quantities for the signal region under consideration. The magnitude of this correction is primarily driven by the relative fraction of signal in the $M_{\mathrm{T2}}(\ell\ell)$ control region. However, there is also an $M_{\mathrm{T2}}(\ell\ell)$ dependence on the correction as well, due to the B_{s} in the correction term. This means that, for a specific $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ point for a chosen top-squark decay scenario, the magnitude of the correction relative to the original signal yield, $S_{\mathrm{s}}^{'}/S_{\mathrm{s}}$, will change with $M_{\mathrm{T2}}(\ell\ell)$. Notably, for most of the top-

squark decay modes considered, the relative loss in signal efficiency from cutting harder on $M_{\mathrm{T2}}(\ell\ell)$ is lower than the corresponding relative decrease in $\mathrm{t\bar{t}}$. Consequently, for most signal points, the relative magnitude of the signal contamination correction goes down with increasing $M_{\mathrm{T2}}(\ell\ell)$.

6.8.2 Optimizing the $M_{\mathrm{T2}}(\ell\ell)$ Cut

Using the estimated signal yield maps shown in Fig. 6.31, along with the analogous maps for each of the relevant signal systematic uncertainties, we calculate the "data-blind" median expected upper limit (Appendix G) on the signal strength, μ for each point in the 2D $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ mass plane and for each of the four main decay modes considered in this analysis.

We perform this calculation for all of the 5 $M_{\mathrm{T2}}(\ell\ell)$ cut values considered.

We then calculate the CMS standard observed and expected limits (median, $\pm 1\sigma$) on the signal strength μ using whichever of the 5 $M_{\mathrm{T2}}(\ell\ell)$ cuts yields the optimal "data-blind" median expected upper limit[3].

We found that, for a large part of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, the optimal $M_{\mathrm{T2}}(\ell\ell)$ cut was $M_{\mathrm{T2}}(\ell\ell) \gtrsim 110\,\mathrm{GeV}$. Notably, however, in regions where the signal efficiency is relatively low, such as most of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane for the $x = 0.25$ T2bw decay mode, or $\Delta\mathrm{M} = m_{\mathrm{t}}$ for the T2tt decay mode, the optimal $M_{\mathrm{T2}}(\ell\ell)$ cut is close to the edge of the signal region $M_{\mathrm{T2}}(\ell\ell) \gtrsim 80\,\mathrm{GeV}$.

[3] we considered up to 140 GeV cuts on $M_{\mathrm{T2}}(\ell\ell)$, but found that there was no significant gain in expected exclusion power – there are certainly points in the 2D $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ mass plane where $M_{\mathrm{T2}}(\ell\ell)$ cuts higher than 120 GeV yield stronger limits on μ, but, for these points, there are not enough expected signal events to actually exclude the points in question

6.8.3 Exclusion Limits

Figure 6.38 shows the expected and observed mass exclusions for the four top-squark decay modes considered. For each sub-figure, the color map shows the 95% confidence level (CL) upper limit on the top-squark pair-production cross section, derived by calculating an upper limit on μ and multiplying this number by the corresponding theory-calculated cross section. Overlaid on these color maps are contours that denote regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL (i.e. the regions where the calculated upper limit on μ is $\mu < 1$).

Note that in these figures, points that either look abnormally purple relative to their neighbors (occurring primarily at low $m_{\tilde{t}_1}$) or are blank entirely are points where either the signal efficiency was estimated as 0 or the limit setting could not set an upper limit on μ due to the expected signal efficiency being too low.

As well, in order to improve the visual interpretation of the results, the 2D maps of the signal strength were smoothed using a 2D Gaussian kernel with dynamic, independent widths along the $m_{\tilde{t}_1}$ and $\Delta\mathrm{M} = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$ axes. More details on this Gaussian kernel-based smoothing can be found in Appendix H.2.3.

For the T2tt decay mode, depending upon the top-squark mass and the mass splitting between the top-squark and the LSP, this analysis can probe top-squark masses up to 415 GeV and LSP masses up to 120 GeV.

The final results of this analysis for the T2bw decay mode depend strongly upon the value of x. In the $x = 0.75$ scenario, our analysis can probe top-squark

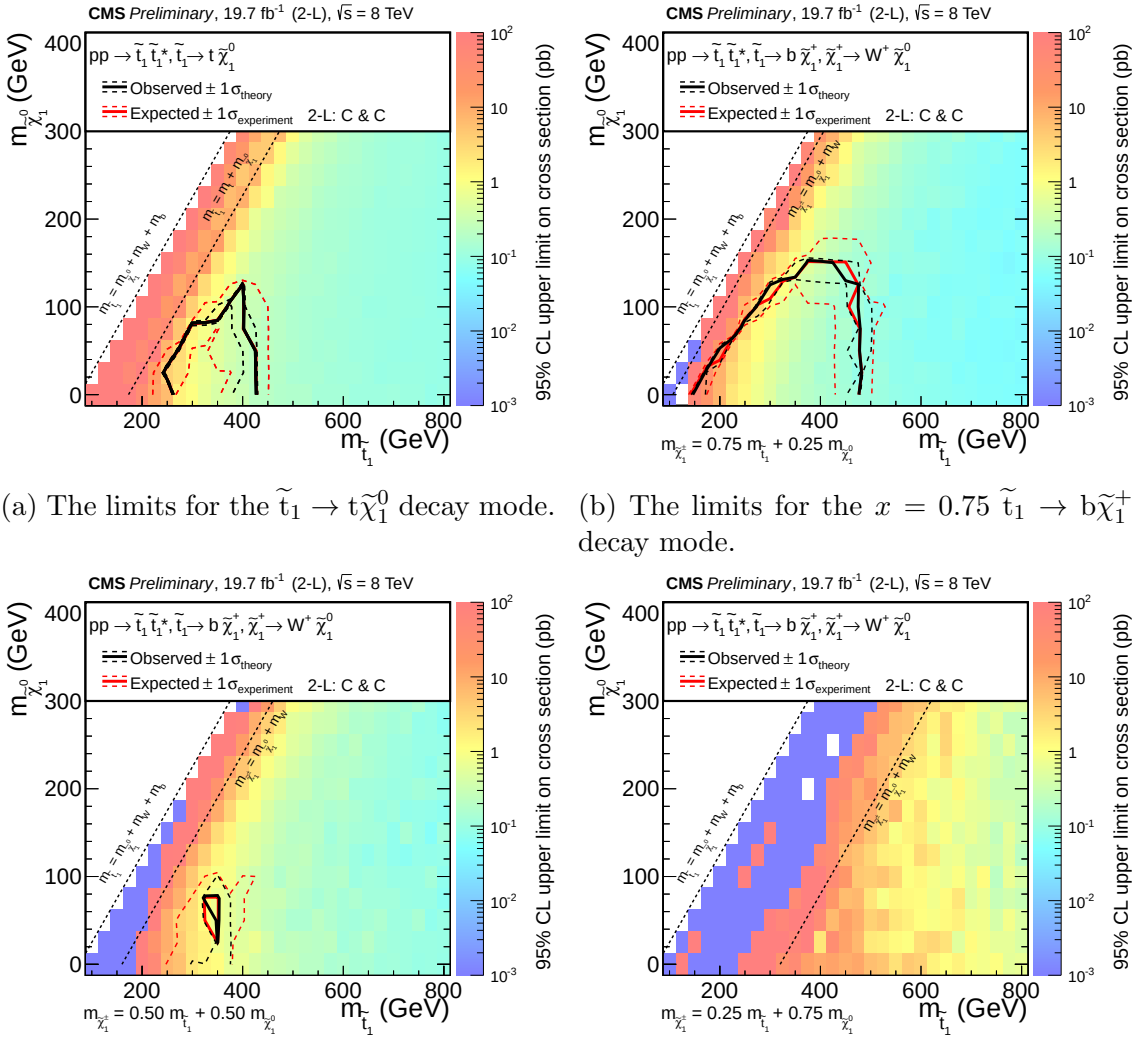

(a) The limits for the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode.

(b) The limits for the $x = 0.75$ $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode.

(c) The limits for the $x = 0.50$ $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode.

(d) The limits for the $x = 0.25$ $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode.

Figure 6.38: 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the cut-and-count dilepton analysis. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The two solid contour lines were made using the median expected cross-section limit (red) and the observed cross-section limit (black). The dashed red lines represent the exclusion region when using the 0.16 and 0.84 percentile expected limits, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

masses up to 475 GeV and $\widetilde{\chi}_1^0$ masses up to 150 GeV. For the $x = 0.50$ scenario, the analysis only probes a small region around $m_{\widetilde{\mathrm{t}}_1} = 350\,\mathrm{GeV}$, $m_{\widetilde{\chi}_1^0} = 75\,\mathrm{GeV}$. Our analysis has no observed sensitivity to the $x = 0.25$ scenario.

There are several effects that drive this large difference in analysis sensitivity for the respective x scenarios. When this relative SUSY particle mass-splitting (i.e. x) is larger, the decay products of the $\widetilde{\chi}_1^{\pm}$ have higher momentum. This not only increases the additional $\vec{\not{E}}_{\mathrm{T}}$ coming from the $\widetilde{\chi}_1^0$ and W, but also improves the relative efficiency of the visible objects to pass selection requirements. This relative mass splitting also influences the relative location and size of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ region where the daughter W boson is produced off-shell, $m_{\mathrm{W}^*} < m_{\mathrm{W}}$. In these regions, the relative signal efficiency in the region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, tends to be characteristically quite low. The reason for this is rather straightforward. In these events, if you were ignore the additional $\vec{\not{E}}_{\mathrm{T}}$ coming from the $\widetilde{\chi}_1^0$, you would subsequently recover the topology of a SM $\mathrm{t\bar{t}}$ event. however, when the W bosons are produced off-shell, the upper bound of $M_{\mathrm{T2}}(\ell\ell)$, for this hypothetical situation where the $\widetilde{\chi}_1^0$ add no additional $\vec{\not{E}}_{\mathrm{T}}$, is no longer m_{W}, but is instead $m_{\mathrm{W}^*} < m_{\mathrm{W}}$.

6.8.3.1 Sensitivity of the Dilepton Analysis to Varying Top-squark Daughter Chiralities

As discussed in Section 6.6.6, the $M_{\mathrm{T2}}(\ell\ell)$ value for dilepton events has a strong dependence upon the angular configuration of the two leptons in the transverse plane. Subsequently, the relative acceptance of signal events has a strong dependence on the polarization of the top-squark daughter particles. In order to quantify the impact of this effect, we reweighted our simulation samples to match several different top-squark daughter handedness-chirality scenarios.

For the T2tt decay mode, we considered scenarios where the coupling between the $\widetilde{\mathrm{t}}_1$ and the t is either fully left- or fully right-handed; for the regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane where the t is produced on-shell, these two scenarios correspond to left- or right-polarized top quarks respectively.

Figure 6.39 shows the same cross section upper limit map as Fig. 6.38a. Overlaid on this cross section map are contours demarcating the regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane where we can exclude top-squark pair-production at the 95% CL using our observed cross section limits. The black line denotes the unpolarized T2tt decay mode limits, while the red line denotes the T2tt decay mode limits in the fully right-handed, i.e. $\widetilde{\mathrm{t}}_1 \to \mathrm{t}_R + \widetilde{\chi}_1^0$, coupling scenario. Nominally, there would also be a blue line for the T2tt decay mode limits in the fully left-handed, $\widetilde{\mathrm{t}}_1 \to \mathrm{t}_L + \widetilde{\chi}_1^0$, coupling scenario. However, our analysis does not have any sensitivity to this scenario. The gain in sensitivity for the fully right-handed coupling scenario is $\mathcal{O}(25\,\mathrm{GeV})$ in all directions.

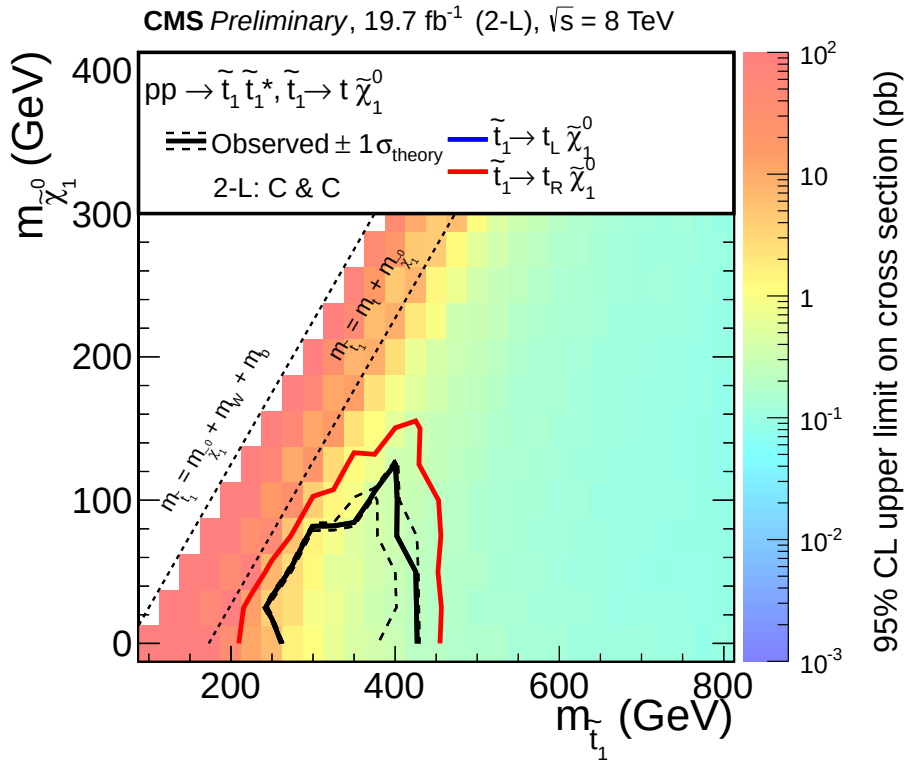

Figure 6.39: A 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ map of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode, as calculated using the cut-and-count dilepton analysis. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The black contour denotes the observed exclusion region, where the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. The red (blue) solid line denotes the observed exclusion region when the $\widetilde{t}_1$ only has a (left-) right-handed coupling to the t. Note that the absence of the blue line indicates that the cut-and-count analysis has no sensitivity to this decay mode. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

For the T2bw decay mode, we considered scenarios where the $\widetilde{\chi}_1^{\pm}$ produced in the $\widetilde{\mathrm{t}}_1$ decay have either left- or right-handed chirality. We further separate this into scenarios where the $\widetilde{\chi}_1^{\pm}$ decays into either a left- or right-handed W boson. We observed that these variations do not significantly change the observed exclusion regions in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane. Thus, we did not include the figures here.

6.8.3.2 Contextualizing the Dilepton Analysis Results

The dileptonic top-squark search that we have detailed in this chapter represents the only 8 TeV CMS search for direct top-squark pair production in the dileptonic final state. It is natural to ask where this result falls in the context of other searches for top-squark pair-production.

In Figs. 3.3 and 3.4, we showed the most current summary plots of top-squark search results in the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode, as produced by the CMS and ATLAS experiments. Directly comparing Fig. 6.38a against these results, most notably Fig 3.4, which actually contains a few ATLAS search results in the dilepton final state, shows that the reach of the cut-and-count analysis described in this chapter is not as large as other analyses. We turn now, to discuss why this is the case.

Comparison against Other Top-squark Final States

The biggest difference between top-squark searches in the dileptonic final state against those in the semileptonic or fully-hadronic final states is, in short, a much lower expected signal yield. This difference in expected signal yield primarily stems

from the difference in the branching ratios for these final states. Ignoring the production of τ leptons, the dileptonic final state BR is 4%, to be compared against 30% and 46% for the semileptonic and fully-hadronic final states.

These BRs represent the maximal signal efficiency for direct top-squark pair-production; the actual efficiencies are often, of course, much lower than that due to analysis cuts. This is particularly relevant at high $m_{\tilde{t}_1}$, as the cross section for top-squark pair production drops by many orders of magnitude from low $m_{\tilde{t}_1}$ to high $m_{\tilde{t}_1}$ (c.f. Fig. 6.3).

One can see how much additional signal statistics can improve things by comparing the various analyses using a single final state against those using statistical combinations of multiple final states, for example in Fig. 3.3.

Comparison against ATLAS's Dilepton Results

The primary reason why ATLAS's dileptonic top-squark searches have stronger reaches relative to this analysis stems from differences in analysis strategies.

For example, the ATLAS search for high mass $\tilde{t}_1$ signatures [163] utilized a multivariate approach based on training BDTs (Boosted Decision Trees [164]). BDTs are very powerful tools that can provide much more powerful discrimination between signal and background relative to a cut-and-count approach.

The advantage of the cut-and-count approach is that it is relatively straightforward to ensure the robustness of background estimates and the modeling of signal efficiencies.

ATLAS used two approaches to derive the exclusion limits for dileptonic top-squark signatures in the $\Delta\mathrm{M} \approx m_{\mathrm{t}}$ region. In this region, top-squark pair events look extremely similar to SM $\mathrm{t\bar{t}}$ events, as the $\widetilde{\chi}_1^0$ tend to be produced at rest. The main difference is that the produced tops in top-squark decays have different angular correlations relative to SM $\mathrm{t\bar{t}}$ production, leading to differing angular correlations of the t daughter particles.

The first ATLAS analysis exploited these differences between $\ell^+\ell^-$ spin-correlations in the top-squark pair events and the SM $\mathrm{t\bar{t}}$ [165]. The second ATLAS analysis measured the dileptonic $\mathrm{t\bar{t}}$ cross section and compared it against the predicted value from theory calculations [166].

6.9 Summary of the Cut-based Dilepton Top-squark Search

In the chapter, we provided details on a "cut-and-count" search for signatures of top-squark pair-production in the dileptonic final state.

This search utilized the stransverse mass variable constructed with two leptons, $M_{\mathrm{T2}}(\ell\ell)$, as the primary means of demarcating our top-squark signal from our main background of dileptonic events coming from SM $\mathrm{t\bar{t}}$ pair-production.

We utilized robust data-driven methods to estimate the contributions of three of our notable backgrounds, including the aforementioned dileptonic $\mathrm{t\bar{t}}$, as well as $\mathrm{Z} \to \ell^+\ell^-$ and "fake-lepton" events.

We performed extensive studies both to validate our understanding of the $M_{\mathrm{T2}}(\ell\ell)$ variable as well as to confirm the accuracy of our simulation's modeling of

the $M_{\mathrm{T2}}(\ell\ell)$ shape.

Looking within our predefined signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, we found no clear signs of an excess and thus proceeded to set frequentist-based upper limits on the top-squark pair-production cross section.

In order to maximize the statistical power of our analysis for each expected decay mode and mass point in the 2D $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, we optimized the choice of $M_{\mathrm{T2}}(\ell\ell)$ threshold that would be used to set our final set of limits, based upon the data-blind median expected upper limits on the top-squark pair-production cross section.

Utilizing these optimized thresholds, we found that, for an unpolarized choice of the handedness of the coupling between the top-squark and its daughter particles, we could exclude top-squark pair-production for $m_{\tilde{\mathrm{t}}_1}$ up to $435\,\mathrm{GeV}$ and for $m_{\tilde{\chi}_1^0}$ up to $160\,\mathrm{GeV}$.

In the T2tt decay mode, these exclusion limits strongly depend upon the nature of the coupling between the top-squark and its daughter top quark. For scenarios where the top-squark only couples to right-handed tops, the observed exclusion regions improve by approximately $25\,\mathrm{GeV}$ in all directions. For scenarios where the top-squark only couples to left-handed tops, unfortunately, the cut-and-count dileptonic top-squark search has no sensitivity.

Unlike the T2tt decay mode, we found that, in the T2bw decay mode, the observed exclusion limits did not depend strongly upon the chirality structure of the SUSY couplings.

Chapter 7: Extensions to the Dilepton Top-squark Search

In the previous chapter, we provided details on a "cut-and-count" search for signatures of top-squark pair-production in the dileptonic final state.

Comparisons of these results, along with the results of other searches for top-squark pair production at CMS and ATLAS, show that no *direct* search for top-squark pair-production has sensitivity in the interesting region, $\Delta\mathrm{M} = m_{\tilde{\mathrm{t}}_1} - m_{\tilde{\chi}_1^0} \approx m_{\mathrm{t}}$.

As we have repeatedly stressed throughout the previous chapter, this region is experimentally difficult as the top-squark signal looks quite similar to the $\mathrm{t}\bar{\mathrm{t}}$ background. Thus, further improvements to this region either require refined analysis techniques[1] or additional data. This can be contrasted against the high $m_{\tilde{\mathrm{t}}_1}$ region, where the limits are driven entirely by the available data (related to the kinematic limits of the LHC).

In this chapter we provide details on a number of extensions to the basic cut-and-count dilepton top-squark search. These extensions were designed to both improve the statistical power of the dilepton analysis by itself, as well as to improve the overall current experimental reach into the aforementioned interesting region,

[1] For example, two ATLAS analyses that do have sensitivity in this region, c.f. Fig. 3.4, utilized comparisons of visible object correlations and SM precision measurements to place constraints on top-squark pair-production.

$\Delta \mathrm{M} \approx m_{\mathrm{t}}$.

The layout of the rest of this chapter is as follows.

In Section 7.1 we provide details on a statistical combination we performed between this analysis and another CMS analysis [167] looking for top-squark pairs in the 1-lepton final state.

In Section 7.2 we describe a set of improvements to the dilepton analysis that can significantly improve its statistical sensitivity to top-squark pair-production. Of these improvements, the most notable one is the construction and utilization of two variants of the M_{T2} variable. These variants provide additional signal discriminating power in scenarios where the $M_{\mathrm{T2}}(\ell\ell)$ variable cannot. The information from all three M_{T2} variants is combined in a multidimensional, binned, shape analysis. This multidimensional, binned, shape analysis has significantly stronger statistical discrimination power compared with the cut-and-count dilepton analysis.

7.1 Statistical Combination of the Cut-based Dilepton Top-squark Search with a Semi-leptonic Top-squark Search

Beyond the search for top-squark pair-production in the dileptonic final state, there have been analyses at CMS that have looked for top-squark pair-production in the fully hadronic and semi-leptonic final states [60, 62], as well as another analysis re-evaluating and further optimizing the search in the semi-leptonic final state [167].

In order to further improve the current experimental limits on top-squark pair-production, we developed a statistical combination of the dilepton top-squark

analysis described in Chapter 6 and the second of the two semileptonic top-squark searches mentioned above.

In the remainder of this section, we first provide a quick overview of the semi-leptonic top-squark search. We then discuss the particulars of the combination.

7.1.1 Overview of the Semi-leptonic Top-squark Search

The semi-leptonic top-squark search utilizes M_{T} as a primary signal discrimination variable.

The preselection for the semi-leptonic top-squark search is similar to the preselection used in the dileptonic top-squark search. Crucially, in order to remove the majority of dileptonic $\mathrm{t\bar{t}}$ events (where the additional $\vec{\not{E}}_{\mathrm{T}}$ from the extra neutrino pushes the M_{T} tail past m_{W}), additional vetos are applied on reconstructed leptons or lepton-like signatures (e.g. isolated tracks in the tracker) beyond the first.

After applying the preselection, the major backgrounds for the semi-leptonic top-squark search are W + jet events, where the M_{T} extends past m_{W} either due to additional $\not{E}_{\mathrm{T}}$ coming from object mismeasurements or high-mass, off-shell Ws; semi-leptonic $\mathrm{t\bar{t}}$ events, for similar reasons as the W + jet events (although there are less high-mass off-shell Ws produced in semi-leptonic $\mathrm{t\bar{t}}$ events); and dileptonic $\mathrm{t\bar{t}}$ events where one of the leptons is "lost" (i.e. either because it is not reconstructed or falls outside of the detector acceptance).

In order to optimize the discrimination between signal and background, the semi-leptonic top-squark search utilizes a multivariate approach based on boosted

decision trees (BDTs) [164]. Separate BDTs are trained for each top-squark decay mode and for various regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, primarily demarcated by $\Delta M = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$. The full list of variables fed into each separate BDT's training can be found in Ref. [37].

The output of these BDTs, when applied to a collision event (simulated or real) is a BDT discriminant, a variable that represents how "signal-like" a given event is. The distribution of this variable tends to be *approximately* monotonically decreasing for both signal and background but, relative to the background, the signal tends to have at least a stronger tail, if not a peak at higher BDT discriminator values.

Once these BDT discriminant distributions are produced for both signal and background, a cut-and-count experiment is performed in the tail of the BDT discriminant distribution. The exact BDT discriminant cut values used depend upon the BDT, and are optimized based on minimizing the data-blind median expected upper limit (see Appendix G.2) calculated using an asymptotic method [168].

Appendix E provides more details on the semi-leptonic top-squark search, including overviews of the background estimation strategies, systematic uncertainties, and calculated upper limits on the top-squark pair-production cross section.

7.1.2 Inter-Analysis Correlations

For the remainder of this dissertation, when discussing the two lepton-based top-squark searches, we will refer to the semi-leptonic top-squark search as the 1ℓ

search, and the dileptonic top-squark search as the 2ℓ search.

We executed comprehensive studies to understand the relative correlations between the 1ℓ and 2ℓ searches. In particular, we checked for correlations in the collision data samples, the background estimates, the background systematic uncertainties, the expected signal efficiencies, and the systematic uncertainties on the expected signal efficiencies. We turn now to discuss each of the individual categories.

7.1.2.1 Orthogonality Check on the Datasets Used

Table E.1 contains the relevant details on the preselection utilized by the 1ℓ search. The applied veto of a second lepton in the 1ℓ search, should, in principle, yield statistically orthogonal selections between the 1ℓ and 2ℓ searches.

Checking this orthogonality explicitly in the collision data showed that there was a statistical overlap at the level of 1 part in 1.2×10^{-5}. This overlap is thus completely negligible in terms of its expected effects on the overall result of the combination.

Orthogonality of the Background and Signal Estimates

It immediately follows from the prior discussion that, by construction, for the two analyses, the central value estimates for both the expected background contributions and the expected signal efficiencies should be effectively orthogonal from one another.

There is a small caveat to this statement; The 1ℓ analysis utilizes a dilepton

control region to derive a data-driven background estimate of their dileptonic $t\bar{t}$ "lost-lepton" background.

For this control region, the veto on the second lepton is relaxed, and subsequently there is the possibility of overlap between the sample of events in this region and those used in the main sample of the 2ℓ analysis.

We explicitly calculated the magnitude of this correlation. For the following discussion, we use $A_c^{1\ell}$ to denote the 1ℓ's $t\bar{t}$ control region, and $B_c^{2\ell}$ and $B_s^{2\ell}$ to denote the 2ℓ's $t\bar{t}$ control region and signal region respectively.

Both analyses use $A_c^{1\ell}$ and $B_c^{2\ell}$ to set the normalization of the dileptonic $t\bar{t}$ background in the respective analyses. In both analyses, systematic uncertainties are applied to the dileptonic $t\bar{t}$ estimate based upon the statistical uncertainties in these control regions. Calculating the overlap, we found that $\sim 55\%$ of $A_c^{1\ell}$ overlapped with $\sim 32\%$ of $B_c^{2\ell}$. Ostensibly, these overlaps would be accounted for in a relative correlation; however, the expected impact of this correlation is negligible: The systematic uncertainty that is derived from $A_c^{1\ell}$ represents typically a ~ 3–4% uncertainty, out of a total background uncertainty of $\mathcal{O}\,(20-40\%)$. The analogous numbers for $B_c^{2\ell}$ and the 2ℓ search are 1% and $\mathcal{O}\,(8-20\%)$.

Comparing $A_c^{1\ell}$ with $B_s^{2\ell}$, we found that $\sim 3.3\%$ of $A_c^{1\ell}$ overlapped with $\sim 52\%$ of $B_s^{2\ell}$. Assuming that $B_s^{2\ell}$ was comprised entirely of signal events, this would represent a definite signal contamination in $A_c^{1\ell}$ of 3.3%. From Table 6.9, however, over 90% of the events in $B_s^{2\ell}$ are expected to be background events, implying that this signal contamination is $\mathcal{O}\,(0.1\%)$ at worst, which is negligible compared to the statistical uncertainty in $A_c^{1\ell}$.

7.1.2.2 Correlations of Background Systematic Uncertainties

The majority of the systematic uncertainties for the 1ℓ search, especially the highest impact ones, are derived as part of their data-driven background estimations. Barring the dileptonic $t\bar{t}$ control region, which we have already discussed, the data-driven background estimates for the 1ℓ search are completely orthogonal from the control regions used in the 2ℓ search, thereby rendering the resulting systematic uncertainties on the background estimates for both analyses statistically orthogonal.

7.1.2.3 Correlations of Signal Systematic Uncertainties

The estimates of signal efficiency for both searches are derived from simulation. Thus, in contrast to the background estimates, there are a number of correlated systematic uncertainties between them.

In particular, among other systematic uncertainties, the systematic uncertainties on the b-tagging scale factors; the lepton ID, isolation, and trigger efficiency scale factors; the integrated luminosity normalization; the energy scale calibration of jets; the generator-level ISR reweighting; and the PDF shape uncertainties are accounted for in both analyses.

Each of these individual systematic uncertainties is treated as 100% correlated between the two analyses[2].

[2] It is worth noting that both analyses found the CT10 PDF uncertainty set yielded the largest (while still sensible) uncertainties on signal efficiency; if this had not been the case, for the purposes of the combination, we would have had to account for partial correlations in the PDF uncertainty.

7.1.3 Performing the Combination

We utilize the methods described in Appendix G to perform the statistical combination of the two analyses. To derive the upper limit on the signal strength for a given point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass plane, for each analysis we take the central value yields (and associated uncertainties) for signal and background as determined by each analysis's respective optimization procedures.

As a reminder, the 2ℓ optimization procedure for a chosen SUSY mass point is to choose the $M_{\mathrm{T2}}(\ell\ell)$ threshold that yields the best data-blind median expected upper limit. Similarly, the 1ℓ search optimizes a cut on a BDT discriminant value that yields the best median expected upper limit[3].

These yields are combined into a 2-bin counting experiment (one bin for each analysis) with correlations set as per the discussion above.

7.1.3.1 Results of the Combination

Figure 7.1 shows, for the unpolarized scenarios for the four top-squark decay modes considered, the expected and observed mass exclusions for the statistical combination of the 1ℓ and 2ℓ analyses.

Depending upon the top-squark mass and the mass-splitting between the top-squark and the $\tilde{\chi}_1^0$, and, for the T2bw decay modes, the value of x, the combination of the 1ℓ and 2ℓ analyses can probe top-squark masses up to 700 GeV and LSP masses up to 260 GeV.

[3] Note that we did not consider re-optimizing the relevant cuts, $M_{\mathrm{T2}}(\ell\ell)$ for dilepton, BDT discriminant for single lepton, based on the combined information of the two analyses.

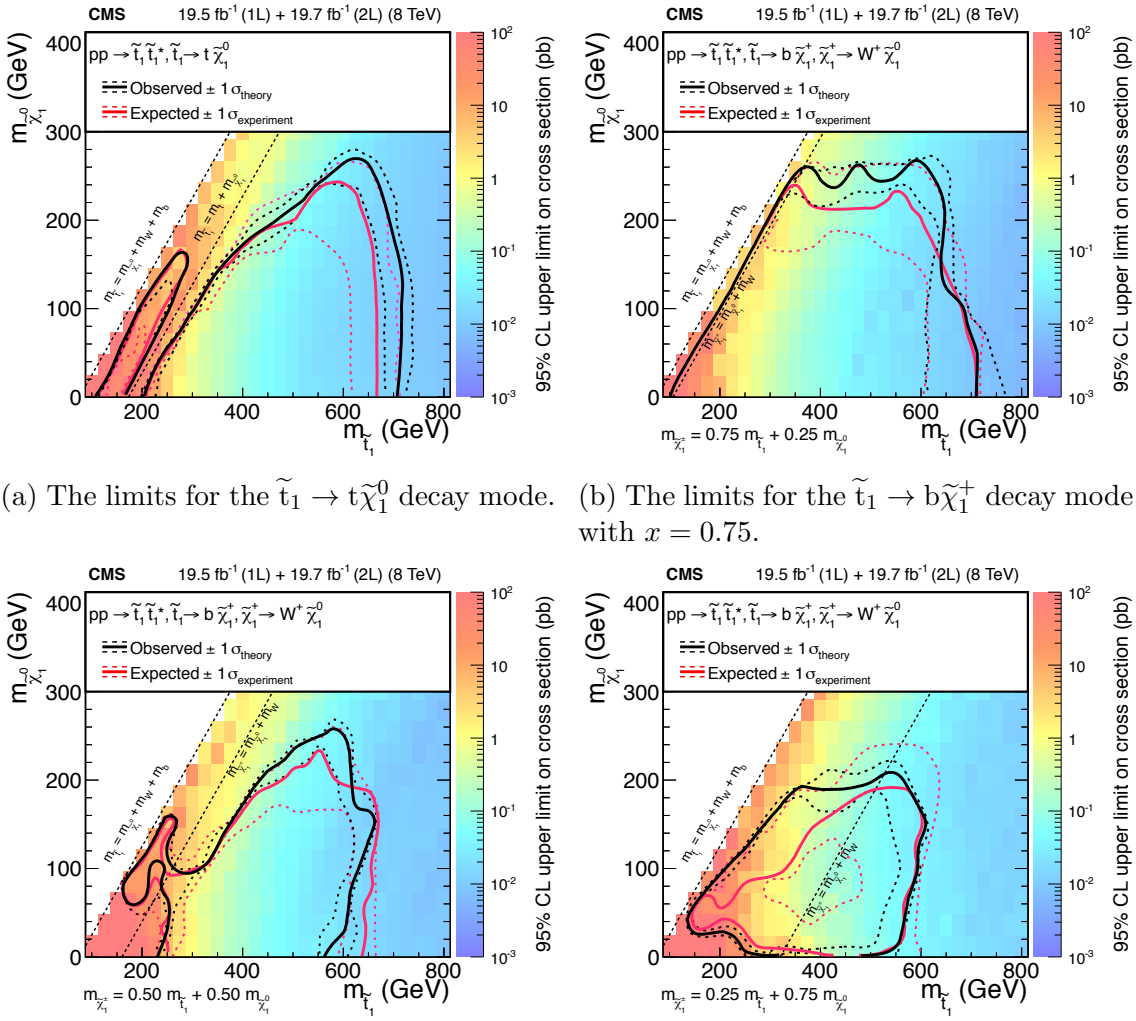

(a) The limits for the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode.

(b) The limits for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with $x = 0.75$.

(c) The limits for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with $x = 0.50$.

(d) The limits for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with $x = 0.25$.

Figure 7.1: 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the statistical combination of the 1L and 2L analyses. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The two solid contour lines were made using the median expected cross-section limit (red) and the observed cross-section limit (black). The dashed red lines represent the exclusion region when using the 0.16 and 0.84 percentile expected limits, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. A small amount of by-hand smoothing has been applied in order to aid the visual interpretation of the results.

Relative to the 2ℓ analysis, the 1ℓ analysis has stronger statistical sensitivity to signatures of top-squark pair-production, c.f. Fig. E.4. Thus, to understand how the lepton analysis combination improves the overall sensitivity in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, relative comparisons should be made between Figs. 7.1 and E.4.

Comparisons of the relative mass exclusions in the T2tt decay mode, Figs. 7.1a and E.4a, show that the 2ℓ analysis brings additional sensitivity into the region, $170\,\text{GeV} \lesssim \Delta\text{M} \lesssim 225\,\text{GeV}$, particularly for low $m_{\tilde{\chi}_1^0}$. As noted before, this is a theoretically interesting region due to considerations from natural SUSY. Moreover, compared with Fig. 3.3, this lepton combination analysis presents the strongest (to be) published limits in this region.

For the unpolarized T2bw decay mode with $x = 0.75$, Figs. 7.1b and E.4b, the 2ℓ analysis slightly extends the edges of the mass-sensitivity for the upper middle and upper right corners of the exclusion contours. Moreover, the overall strength of the exclusion (based upon the relative coloration of the exclusion map), in the central bulk of the excluded region is strengthened as well.

For the unpolarized T2bw decay mode with $x = 0.50$, Figs. 7.1c and E.4c, the 2ℓ analysis slightly extends all of the edges of the mass-sensitivity. Most notably, there is a small exclusion "bridge" in the observed exclusions, around $m_{\tilde{\chi}_1^0} \approx 80\,\text{GeV}$, that links the regions $\Delta\text{M} \approx 150\,\text{GeV}$ and $\Delta\text{M} \approx 200\,\text{GeV}$. This observed exclusion bridge is widened by the 2ℓ analysis. As well, the 2ℓ analysis also serves to open this bridge in the expected exclusion contours.

For the unpolarized T2bw decay mode with $x = 0.25$, Figs. 7.1d and E.4d, the 2ℓ analysis does not bring much additional sensitivity (which is expected given the

relatively low signal efficiency for this decay scenario).

7.1.3.2 Sensitivity of the Lepton Combination to Different Top-squark Coupling Chirality Scenarios

As in Section 6.8.3.1, we also calculated the dependence of the observed and expected exclusion regions on the polarizations and chiralities of the top-squark daughter particles.

It is important first to note that, due to time constraints on the BDT training, the 1ℓ analysis's BDTs were not optimized for different coupling chirality scenarios. That is, when calculating the expected background and signal yields for a given decay mode and mass point in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, the BDTs based upon the *unpolarized* couplings were used regardless of which coupling scenario was considered. This should not in principle strongly impact the results for most regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. However, in regions, such as those where ΔM is small, where the kinematics of top-squark events are expected to change rapidly from one mass point to another, this can strongly impact the overall effectiveness of the BDTs.

To contrast this, the 2ℓ analysis's $M_{\mathrm{T2}}(\ell\ell)$ cuts *were* optimized based upon the coupling scenario being considered. With that caveat, Fig. 7.2 shows the effects of coupling variations on the observed exclusion limits for the four top-squark decay modes considered.

For the T2tt decay mode, in the region $\Delta\mathrm{M} < m_{\mathrm{t}}$, where the daughter top quark is being produced off-shell in the top-squark decay, we observe that the com-

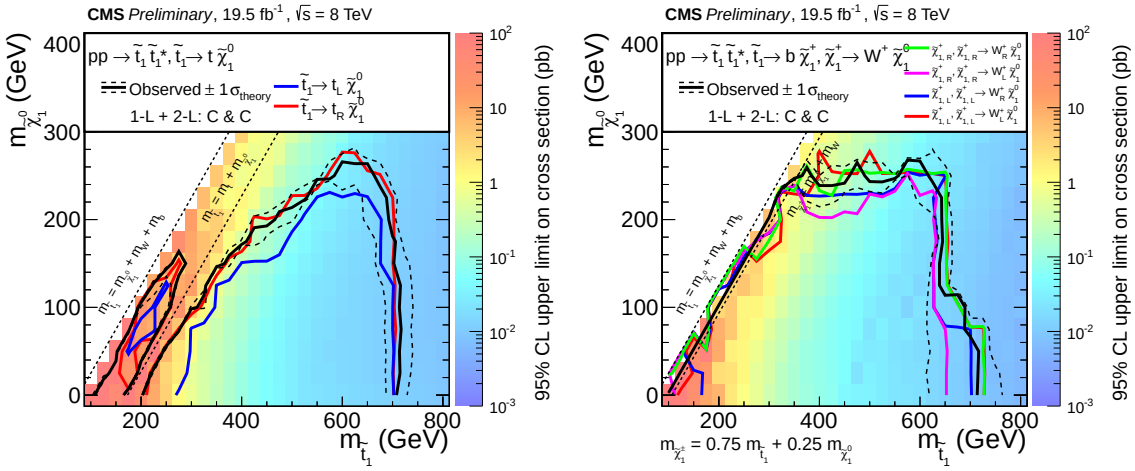

(a) The effects of varying the handedness of the $\widetilde{t}_1$-t coupling for $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode.

(b) The effects of varying the handedness of the $\widetilde{t}_1$-$\widetilde{\chi}_1^\pm$ and $\widetilde{\chi}_1^\pm$-W couplings for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$, $x = 0.75$ decay mode.

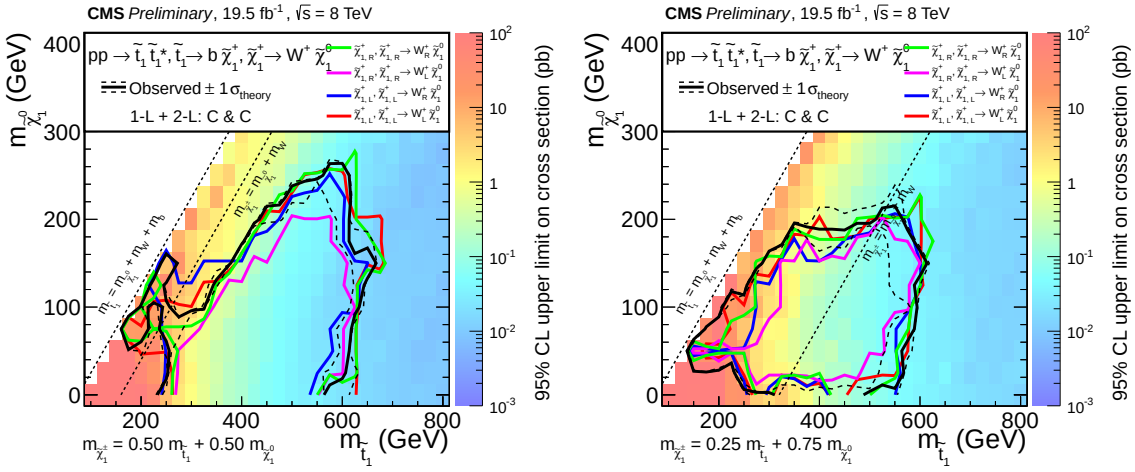

(c) The effects of varying the handedness of the $\widetilde{t}_1$-$\widetilde{\chi}_1^\pm$ and $\widetilde{\chi}_1^\pm$-W couplings for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$, $x = 0.50$ decay mode.

(d) The effects of varying the handedness of the $\widetilde{t}_1$-$\widetilde{\chi}_1^\pm$ and $\widetilde{\chi}_1^\pm$-W couplings for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$, $x = 0.25$ decay mode.

Figure 7.2: 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the statistical combination of the 1L and 2L analyses. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The solid black line was made using the observed cross-section limit, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. For the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode [upper left], the red (blue) solid line denotes the observed exclusion region when the $\widetilde{t}_1$ only has a (left-) right-handed coupling to the t. For the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay modes [all others], the green (pink) solid line denotes the observed exclusion region when the $\widetilde{\chi}_1^\pm$ is right-handed and decays to a (left-) right-handed W boson. The blue (red) solid line denotes the observed exclusion region when the $\widetilde{\chi}_1^\pm$ is left-handed and decays to a (left-) right-handed W boson. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

bination of the two analyses seems to have less sensitivity to both fully left- and fully right-polarized $\widetilde{\mathrm{t}}_1 - \mathrm{t}$ couplings. This is a counterintuitive result, and one that ostensibly should be investigated further. However, it is important to note that, for the T2tt decay mode, in this region of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane the kinematics of the top-squark events change very rapidly as one moves from one ΔM diagonal to another. As noted above, this is expected to strongly impact the statistical power of the BDTs in this region, as they were trained assuming unpolarized top-squark decays.

For the T2tt decay mode, In the ΔM regions where the top quark is produced on-shell in the top-squark decays, we observe that the combined lepton analysis is significantly more sensitive to the scenario where the top-squark decays to right-handed top quarks. However, the overall sensitivity to the left-handed polarization scenario is notably worse, being reduced overall by $\sim$ 25–50 GeV throughout this section of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane.

The sensitivity to top-squark coupling chiralities is even more notable for the low x T2bw decay modes. Notably, for the scenario where the top-squark decays to a right-handed $\widetilde{\chi}_1^+$, which then decays to a left-handed W^+, for both the $x = 0.25$ and $x = 0.50$, the observed sensitivity to top-squark pair-production is worse by as much as ~ 100 GeV for $\Delta\mathrm{M} \lesssim 200$ GeV. When $x = 0.50$, if the top-squark decays to a left-handed chargino, the observed sensitivity improves by as much as ~ 50 GeV in this region. Given the prior observation that the dileptonic cut-and-count analysis is not significantly sensitive to differences in coupling chirality for the T2bw decay mode, these results indicate that it is the semi-leptonic top-squark search that is

strongly sensitive to top-squark coupling-chiralities.

7.1.3.3 Conclusions and Summary of the Lepton Combination

In this section we have presented a statistical combination of the dileptonic cut-and-count top-squark search with a BDT-based semi-leptonic top-squark search. These two analyses are almost completely orthogonal in terms of their event selection and background systematic uncertainties. The two analyses share a number of systematic uncertainty sources for the signal efficiencies; each of these systematic uncertainties are treated as 100% correlated between the two analyses.

The two analyses are combined into a multi-bin counting experiment and exclusion limits on top-squark pair-production are calculated. The semi-leptonic top-squark search by itself is significantly more sensitive to top-squark production through most of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. Nevertheless, the additional data from the dileptonic analysis extend the combined sensitivity in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. Most notably, the 2ℓ analysis brings additional sensitivity into the region, $170\,\mathrm{GeV} \lesssim \Delta\mathrm{M} \lesssim 225\,\mathrm{GeV}$, particularly for low $m_{\tilde{\chi}_1^0}$. This lepton combination consequently presents the strongest published limits from the CMS experiment on direct top-squark pair-production in this theoretically interesting region.

The lepton combination analysis has some degree of sensitivity to different top-squark coupling chirality scenarios. For the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0$ decay mode, the overall sensitivity when the top-squark decays to left-handed top quarks is reduced by ~ 25 to $50\,\mathrm{GeV}$ throughout the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. This sensitivity to different top-squark

coupling chirality scenarios is particularly notable for the T2bw decay mode when the chargino mass-splitting parameter is low. As an example, for the $x = 0.50$ T2bw decay mode, in the region, $\Delta\mathrm{M} \lesssim 200\,\mathrm{GeV}$, the observed sensitivity can vary by as much as $\sim 100\,\mathrm{GeV}$ depending upon the chirality of the $\widetilde{\chi}_1^{\pm}$ and its daughter W boson.

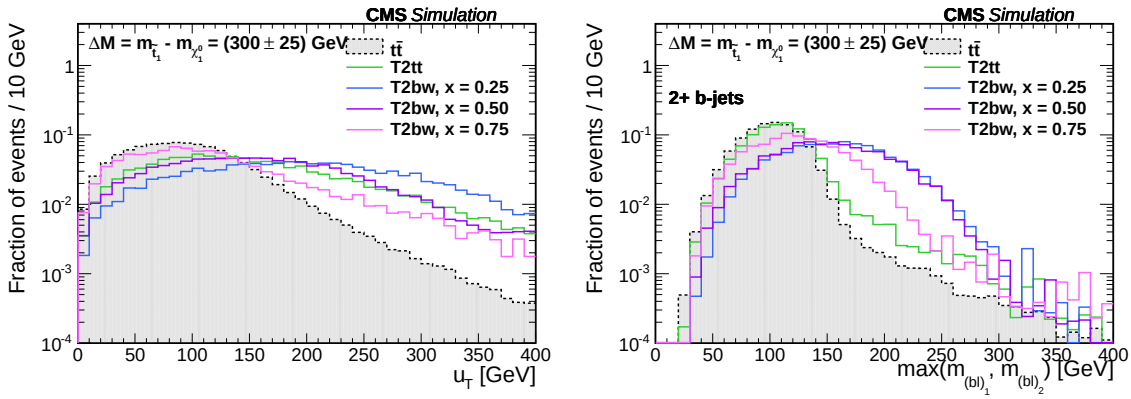

(a) The "hadronic recoil" u_{T}

(b) The maximum invariant mass of the two individual ℓb pairs, for events where two or more reconstructed b jets are required.

Figure 7.3: The distributions at the preselection level, of the maximum $m_{\ell\mathrm{b}}$ and the "hadronic recoil" u_{T} in the $\mathrm{t\bar{t}}$ background and different top-squark decay modes for a ΔM value of 300 GeV; distributions are normalized to the same area. Note that the events shown in Fig. 7.3b were also required to have 2 or more b jets.

7.2 Improving the Dilepton Top-squark Search: Additional M_{T2} Variants

One of the interesting findings from Chapter 6 is the stark differences in signal efficiency in the region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV}$, between the various decay modes considered in that analysis, c.f. Fig. 6.25.

Specifically, in the low-x T2bw decay modes, the lepton have characteristically softer momenta, which, along with the relative angular configurations of the leptons with the $\vec{\not{E}}_{\mathrm{T}}$, directly correlates to lower $M_{\mathrm{T2}}(\ell\ell)$ values and subsequently lower signal efficiencies. Therefore, the $M_{\mathrm{T2}}(\ell\ell)$ variable by itself is not an optimal discriminating variable for the low-x T2bw top-squark decay modes.

Although the leptons in low-x T2bw events are typically softer, these events tend to have harder b jets (Fig. 6.30a). Also, in these low-x T2bw events, the

fact that the dilepton system tends to be aligned with the $\vec{\not{E}}_{\mathrm{T}}$ means that the "hadronic recoil" $u_{\mathrm{T}} = |\vec{q}_{\mathrm{T}} + \vec{\not{E}}_{\mathrm{T}}|$ tends to be large (see Fig. 7.3a, which compares the distributions of u_{T} in the $\mathrm{t\bar{t}}$ background and the various top-squark decay modes). Finally, because there are not real top quarks being produced in the T2bw decay modes, the individual invariant masses of the ℓb pairs do not have the typical upper bound of m_{t}, in contrast to the $\mathrm{t\bar{t}}$ background and T2tt decay mode (see Fig. 7.3b, which compares the distributions of $\max(m_{\ell\mathrm{b}})$ in the $\mathrm{t\bar{t}}$ background and the various top-squark decay modes).

This naturally motivates the consideration of event variables that utilize this information. In particular, we propose two additional variants of M_{T2} that capitalize on the b jet, hadronic recoil, and $m_{\ell\mathrm{b}}$ information.

7.2.1 M_{T2} with b Jets and Invisible W Bosons: $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$

The standard M_{T2} calculation involves scanning through viable partitions of the $\vec{\not{E}}_{\mathrm{T}}$ into two hypothetical neutrinos. If, instead of $\vec{\not{E}}_{\mathrm{T}}$, one adds the leptons' $\vec{p}_{\mathrm{T}}$ to the $\vec{\not{E}}_{\mathrm{T}}$, then these partititions of the modified $\vec{\not{E}}_{\mathrm{T}}$ will effectively be hypothetical W bosons. This, in essence, creates an M_{T2} calculation in which the W bosons ($\ell+\nu$) are treated as invisible. We construct one of the M_{T2} variants, $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb};\, m_{\widehat{\mathrm{W}}}\right)$, where the $\not{\mathrm{W}}$ signifies these "invisible W bosons", by pairing the reconstructed b jets in the event with the net $\vec{p}_{\mathrm{T}}$ of these invisible W bosons,

$$M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb};\, m_{\widehat{\mathrm{W}}}\right) \equiv \min_{\vec{p}_{\mathrm{T}}^{\hat{\mathrm{W}}_1} + \vec{p}_{\mathrm{T}}^{\hat{\mathrm{W}}_2} = -\vec{u}_{\mathrm{T}}} \left(\max \left[M_{\mathrm{T}}\left(\vec{p}_{\mathrm{T}}^{\mathrm{b}_1}, \vec{p}_{\mathrm{T}}^{\hat{\mathrm{W}}_1};\, m_{\hat{\mathrm{W}}}\right),\, M_{\mathrm{T}}\left(\vec{p}_{\mathrm{T}}^{\mathrm{b}_2}, \vec{p}_{\mathrm{T}}^{\hat{\mathrm{W}}_2};\, m_{\hat{\mathrm{W}}}\right) \right] \right), \tag{7.1}$$

where we denote the sum of the leptons' $\vec{p}_{\mathrm{T}}$ with the $\vec{\not{E}}_{\mathrm{T}}$ as the negative of $\vec{u}_{\mathrm{T}}$, the "hadronic recoil"[4]. As a result of the differences in its construction, $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb};\, m_{\widehat{\mathrm{W}}}\right)$ behaves differently from $M_{\mathrm{T2}}\left(\ell\ell\right)$. As an example, the particular form of the M_{T} calculation in Eq. (7.1),

$$M_{\mathrm{T}}^{2}\left(\vec{p}_{\mathrm{T}}^{\,\mathrm{b}},\, \vec{p}_{\mathrm{T}}^{\,\widehat{\mathrm{W}}};\, m_{\widehat{\mathrm{W}}}\right) = m_{\mathrm{b}}^{2} + m_{\widehat{\mathrm{W}}}^{2} + 2\left(E_{\mathrm{T}}^{\mathrm{b}} E_{\mathrm{T}}^{\widehat{\mathrm{W}}} - \vec{p}_{\mathrm{T}}^{\,\mathrm{b}} \cdot \vec{p}_{\mathrm{T}}^{\,\widehat{\mathrm{W}}}\right), \tag{7.2}$$

means that the lower bound of $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb};\, m_{\widehat{\mathrm{W}}}\right)$ is the mass of the b quark, m_{b}, plus the input mass for the "invisible" W, $m_{\widehat{\mathrm{W}}}$. For future discussions and calculations, we set $m_{\widehat{\mathrm{W}}}$ to be the actual W-boson mass, $m_{\widehat{\mathrm{W}}} = m_{\mathrm{W}}$, and hereafter denote $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb};\, m_{\widehat{\mathrm{W}}}\right)$ with the shorter $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb}\right)$. Just like $M_{\mathrm{T2}}\left(\ell\ell\right)$, the $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb}\right)$ calculation will always return its lower bound value $(m_{\mathrm{b}} + m_{\mathrm{W}})$ for events where the "invisible" $\vec{p}_{\mathrm{T}}$ falls in the smaller of the two opening angles defined by the two visible $\vec{p}_{\mathrm{T}}$.

Furthermore, Eq. (7.2) also implies that, in dileptonic $\mathrm{t\bar{t}}$ events, the upper bound of $M_{\mathrm{T2}}^{\not{\mathrm{W}}}\left(\mathrm{bb}\right)$, if it is constructed with the correct input objects, will be m_{t}, the mass of the top quark. This can be seen by considering the reconstruction of the invariant mass of a top quark that decayed to the leptonic final state,

$$\begin{aligned} m_{\mathrm{t}} &= \left(p_{\mathrm{b}} + p_{\ell} + p_{\bar{\nu}}\right)^{\mu}\left(p_{\mathrm{b}} + p_{\ell} + p_{\bar{\nu}}\right)_{\mu} \\ &= m_{\mathrm{b}}^{2} + m_{\mathrm{W}}^{2} + 2\left(E_{\mathrm{T}}^{\mathrm{b}} E_{\mathrm{T}}^{\mathrm{W}} \cosh\left(\Delta y\right) - \vec{p}_{\mathrm{T}}^{\,\mathrm{b}} \cdot \vec{p}_{\mathrm{T}}^{\,\mathrm{W}}\right) \\ &= m_{\ell\mathrm{b}}^{2} + m_{\bar{\nu}}^{2} + 2\left(E_{\mathrm{T}}^{\ell\mathrm{b}} E_{\mathrm{T}}^{\bar{\nu}} \cosh\left(\Delta y\right) - \vec{p}_{\mathrm{T}}^{\,\ell\mathrm{b}} \cdot \vec{p}_{\mathrm{T}}^{\,\bar{\nu}}\right), \end{aligned} \tag{7.3}$$

where the bottom two lines of Eq. (7.3) signify that the final expression depends upon whether one pairs the four-momentum vectors of the b quark and lepton (ℓb) or the lepton and neutrino (W), and $\cosh\left(\Delta y\right)$ signifies the difference in rapidity between

[4]The term "hadronic recoil" is nomenclature that is more suited for γ + jets or Z + jet events — see Ref. [71] — and should not be taken literally here.

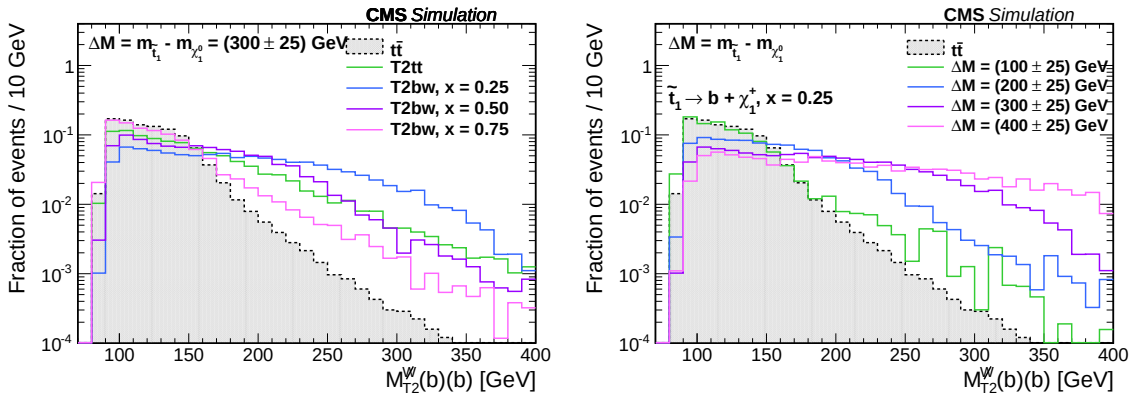

(a) The $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distributions for several different top-squark decay modes for a $\Delta\mathrm{M}$ value of 300 ± 25 GeV.

(b) The $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution for several different $\Delta\mathrm{M}$ regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure 7.4: The preselection level distributions of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in the $\mathrm{t\bar{t}}$ background and the top-squark signal; distributions are normalized to the same area.

the paired system and the unpaired particle. Remembering that $\cosh(x) \geq 1$, direct comparisons of Eq. (7.2) and the second-to-last line of Eq. (7.3) corroborate our claim made above about the upper bound of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$.

Figure 7.4a compares the normalized distributions of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in the $\mathrm{t\bar{t}}$ background and different top-squark decay modes for $\Delta\mathrm{M} = 300\,\mathrm{GeV}$. For the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ calculation, if there are at least two reconstructed b jets available in the event, the two leading b jets are chosen for the input b jet $\vec{p}_{\mathrm{T}}$. For events where only one reconstructed b jet is available, this single b jet is used as the first $\vec{p}_{\mathrm{T}}$ and the leading non b-tagged jet is selected for the second $\vec{p}_{\mathrm{T}}$.

The distribution of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in the $\mathrm{t\bar{t}}$ background is a slightly-sloped plateau between 80 GeV and 150 GeV that rapidly drops off after that. Compared with $M_{\mathrm{T2}}(\ell\ell)$ — i.e. Fig. 6.10 — the relative fraction of events in the $\mathrm{t\bar{t}}$ that survive past the respective nominal kinematic endpoints — m_{W} for $M_{\mathrm{T2}}(\ell\ell)$, m_{t} for $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$—

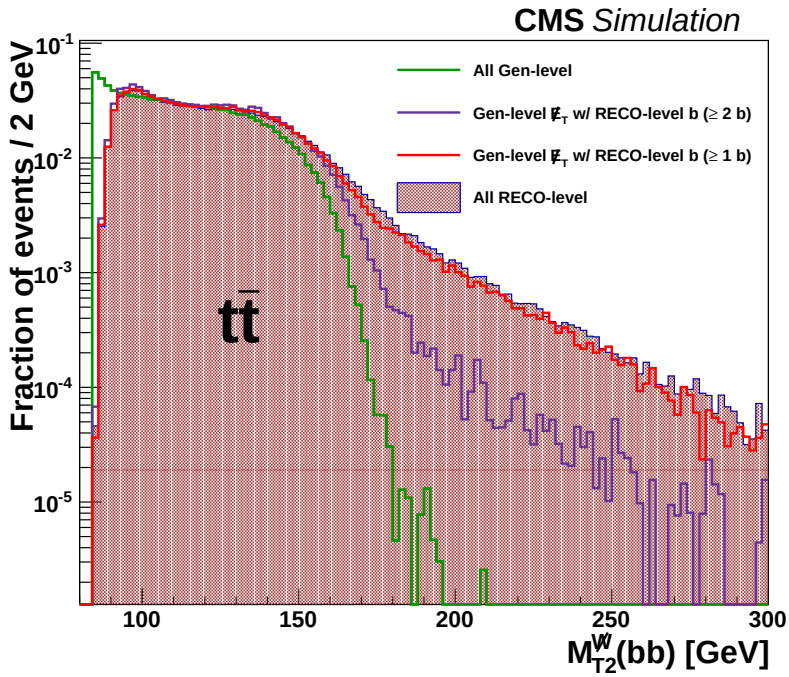

Figure 7.5: $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution for simulated $\mathrm{t\bar{t}}$ in four different cases. In dark green, the distribution is shown using the generator-level $\vec{\not{E}}_{\mathrm{T}}$ as well as the generator-level b quarks from the $\mathrm{t\bar{t}}$ decay. In purple, the distribution is shown using the generator-level $\vec{\not{E}}_{\mathrm{T}}$ and with reconstructed b jets. The red distribution is analogous to the purple, except that events with only one reconstructed b jet are allowed (these events utilize the leading non b-tagged jet as the second input "b jet") Finally, the filled dark red histogram is analogous to the red distribution, except it was constructed using reconstruction-level objects, as would be done for data events.

is much larger for $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ primarily because the energy resolution on the b jets is worse than the leptons.

As well, however, there is a subset of events where the input "b jets" selected for the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ calculation do not coincide with the "true" b jets in the event; this notably occurs more often for events where there was only one reconstructed b jet. The distribution of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in this class of events will not have a kinematic edge at m_{t}, and a (large) portion of the long tail of events in the $\mathrm{t\bar{t}}$ background is comprised of these events.

Figure 7.5 provides some intuition for the relative importance of some of the above mentioned effects on the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ reconstruction in $\mathrm{t\bar{t}}$ events. This figure shows the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution for simulated $\mathrm{t\bar{t}}$ in four different cases.

- In dark green, the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution is shown using the generator-level $\vec{\not{E}}_{\mathrm{T}}$ as well as the generator-level b quarks from the $\mathrm{t\bar{t}}$ decay. The kinematic edge, $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < m_{\mathrm{t}}$, is quite visible here.

- In purple, the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution is shown using the generator-level $\vec{\not{E}}_{\mathrm{T}}$, but using the two highest momentum reconstructed b jets. As can be seen, the kinematic edge gets notably smeared by the imprecise energy resolution of the input b jets.

- The light red distribution is analogous to the purple, except that events with only one reconstructed b jet are allowed (these events utilize the leading non b-tagged jet as the second input "b jet"). The relaxation of the b jet requirement introduces a class of events where one of the input jets for the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ calculation is not a "true" b jet coming from a top quark decay. As can clearly be seen, these kinds of events tend to have much larger $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ values.

- Finally, the filled dark red histogram is analogous to the red distribution, except it was constructed using all reconstruction-level objects, as would be done for data events. The salient difference of this distribution and the light red distribution is the replacement of the generator-level $\vec{\not{E}}_{\mathrm{T}}$ with the reconstructed $\vec{\not{E}}_{\mathrm{T}}$. As can be seen by comparing the two red distributions. This replacement does not severely impact the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ resolution relative to other effects.

Bringing the discussion back to Fig. 7.4a, the $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$ distributions in the top-squark signal, particularly for the low-x T2bw decay modes, have much stronger tails in comparison to the $\mathrm{t\bar{t}}$ background. Furthermore, it is worth noting that as with $M_{\mathrm{T2}}(\ell\ell)$, the ΔM value strongly influences the overall shape of $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$, as can be seen in Fig. 7.4b, which shows the $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$ distribution for the $x = 0.25$ top-squark decay mode for different values of ΔM.

7.2.2 M_{T2} with ℓb Pairs: $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$

We can also construct an M_{T2} variant by pairing b jets with leptons to make the input $\vec{p}_{\mathrm{T}}$,

$$M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) \equiv \min_{\vec{p}_{\mathrm{T}}^{\widehat{\nu}_1}+\vec{p}_{\mathrm{T}}^{\widehat{\nu}_2}=\vec{\not{E}}_{\mathrm{T}}} \left(\max\left[M_{\mathrm{T}}(\vec{p}_{\mathrm{T}}^{(\ell\mathrm{b})_1}, \vec{p}_{\mathrm{T}}^{\widehat{\nu}_1}),\ M_{\mathrm{T}}(\vec{p}_{\mathrm{T}}^{(\ell\mathrm{b})_2}, \vec{p}_{\mathrm{T}}^{\widehat{\nu}_2})\right]\right). \tag{7.4}$$

In analogy with $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$, the differences in the construction of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}(\ell\ell)$ cause $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ to behave in a notably different fashion from $M_{\mathrm{T2}}(\ell\ell)$. Unlike $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$ and $M_{\mathrm{T2}}(\ell\ell)$, the lower bound of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, from the form of the M_{T} calculation in Eq. (7.4),

$$M_{\mathrm{T}}^2\left(\vec{p}_{\mathrm{T}}^{\ell\mathrm{b}}, \vec{p}_{\mathrm{T}}^{\widehat{\nu}}\right) = m_{\ell\mathrm{b}}^2 + m_{\widehat{\nu}}^2 + 2\left(E_{\mathrm{T}}^{\ell\mathrm{b}} E_{\mathrm{T}}^{\widehat{\nu}} - \vec{p}_{\mathrm{T}}^{\ell\mathrm{b}} \cdot \vec{p}_{\mathrm{T}}^{\widehat{\nu}}\right) \tag{7.5}$$

is the larger of the two invariant masses of the b-lepton pairs, $\max(m_{\ell\mathrm{b}})$. Similar to the other M_{T2} variants, the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation will always return this lower bound value for angular configurations where the $\vec{\not{E}}_{\mathrm{T}}$ falls in between the two $\vec{p}_{\mathrm{T}}^{\ell\mathrm{b}}$. Although its lower bound differs from $M_{\mathrm{T2}}^{\not{\mathrm{W}}}(\mathrm{bb})$, direct comparisons of Eq. (7.5) and the last line of Eq. (7.3) show that upper bound of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ will also be m_{t},

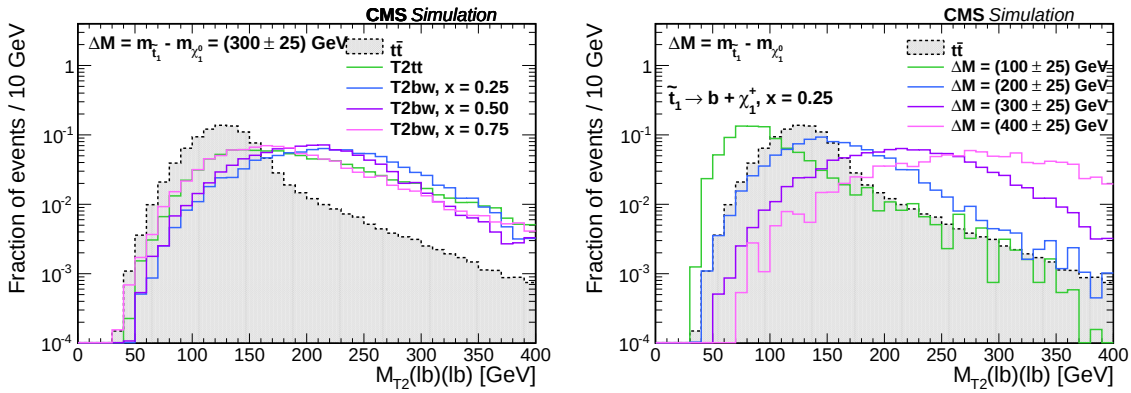

(a) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distributions for several different top-squark decay modes for a ΔM value of 300 ± 25 GeV.

(b) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure 7.6: The preselection level distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the $\mathrm{t\bar{t}}$ background and the top-squark signal; distributions are normalized to the same area.

assuming that it has been constructed with the correct input objects.

Figure 7.6a compares the normalized distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the $\mathrm{t\bar{t}}$ background and different top-squark decay modes for $\Delta\mathrm{M} = 300\,\mathrm{GeV}$. For the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation we utilize the same procedure to select the input b jets as was described above for the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ calculation.

There are 2 possible sets of pairings for the input b jets and leptons. We calculate $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for each pairing set and choose the pairing set that yields the smaller $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ value. One quick justification for this choice is that it mimics the construction of the other M_{T2} variants. As an example, during the construction of $M_{\mathrm{T2}}(\ell\ell)$, every partition of the $\vec{\not{E}}_{\mathrm{T}}$ has an ambiguity, specifically as to which hypothetical neutrino gets paired with which lepton. The construction of $M_{\mathrm{T2}}(\ell\ell)$ intrinsically checks both possible pairings and favors the pairing that yields a smaller maximum M_{T}, thereby resolving this ambiguity.

Another, more subtle justification is that this choice will naturally lead to a $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution that has a smaller tail, nominally improving the separation between the top-squark signal and the SM backgrounds.

The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in $\mathrm{t\bar{t}}$ background has a very prominent peak located at $\approx 120\,\mathrm{GeV}$, before dropping somewhat rapidly up to $\approx 170\,\mathrm{GeV}$ and then decreasing more slowly past that point. This is notably different behavior from $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$, where the relative fraction of $\mathrm{t\bar{t}}$ events per $10\,\mathrm{GeV}$ past the point $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) \gtrsim 250\,\mathrm{GeV}$ is at the per-mille level. The main component of this large tail in the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution are the subset of events where the pairings between the b jets and leptons are incorrect, as these events tend to have large values for the invariant masses of the ℓb pairs. We discuss this large tail in more detail in the next section.

As with $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$, the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution in the top-squark signal can peak at larger values relative to the $\mathrm{t\bar{t}}$. However, as shown in Fig. 7.6b, the distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ is much more strongly dependent upon the value of $\Delta\mathrm{M}$.

7.2.2.1 Tempering the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ Tail

We claimed in Section 7.2.2 the main component of the large tail in the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution are events where the the pairings between the b jets and leptons are incorrect.

Figure 7.7a shows how the choice of ℓb pairing affects the generator-level distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in $\mathrm{t\bar{t}}$ events that pass the preselection. The two choices

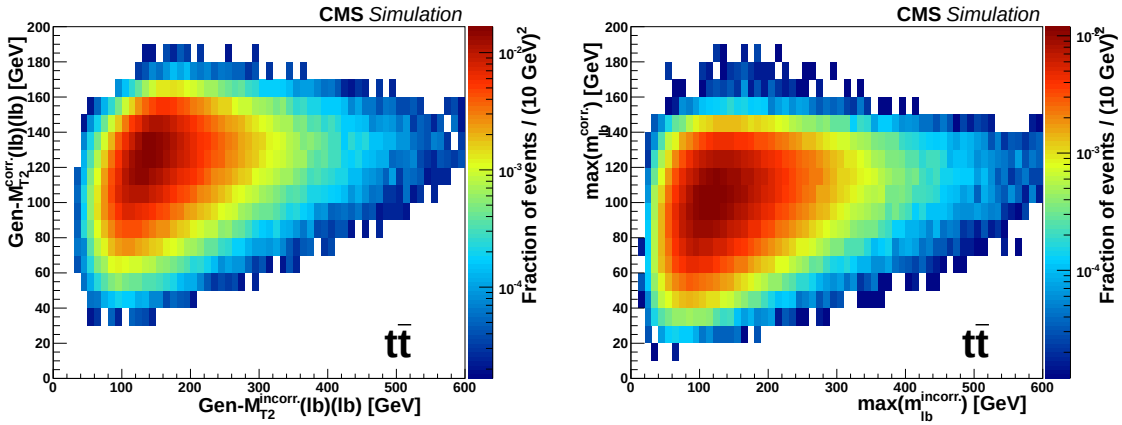

(a) The distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ with either correct or incorrect b-ℓ pairings.

(b) The distributions of $\max(m_{\ell\mathrm{b}})$ with either correct or incorrect b-ℓ pairings.

Figure 7.7: Normalized distributions for the $\mathrm{t\bar{t}}$ background comparing the dependence of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $\max(m_{\ell\mathrm{b}})$ on the choice of b-ℓ pairing (correct of incorrect). Note the difference in x- and y-axis ranges.

of pairing considered are the correct or incorrect pairings, where this distinction is based on the generator-level information for the b quarks and leptons. As can clearly be seen, the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ tail is much more significant when using the incorrect pairing relative to the correct pairing.

This effect is almost completely driven by the contribution to $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ from the invariant masses of the paired ℓb systems, as can be seen in Fig. 7.7b, which is analogous to Fig 7.7a, but shows the generator-level distribution of $\max(m_{\ell\mathrm{b}})$.

It is thus clearly important to utilize a pairing algorithm that has a strong efficiency for choosing the correct ℓb pairings for the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation. However, even after choosing a pairing algorithm, additional cuts or requirements can further reduce the contribution of events with incorrect pairings.

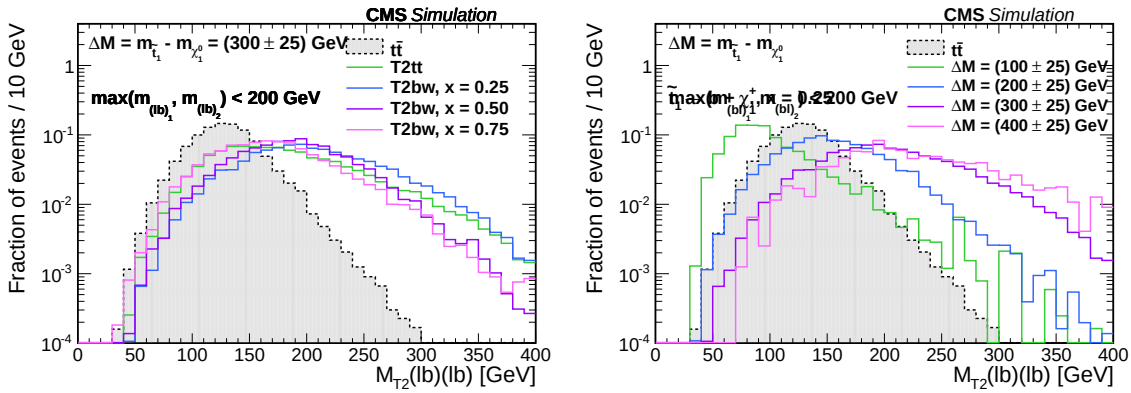

(a) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distributions for several different top-squark decay modes for a ΔM value of 300 ± 25 GeV.

(b) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure 7.8: The distributions, using the preselection but also requiring $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$, of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the $\mathrm{t\bar{t}}$ background and the top-squark signal; distributions are normalized to the same area.

Requiring Low ℓb Masses

One intuitive way to reduce the contribution of mis-paired events is to require that the chosen ℓb pairs have invariant masses below a threshold. Figure 7.8 shows the same distributions as Fig. 7.6, except that each event is required to have $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$. Assuming perfect object reconstructions, the threshold would nominally be m_{t}, as shown in Fig. 7.7b. However, the threshold of 200 GeV accommodates the imperfect object energy resolutions, most notably for the b jets.

As can be seen by comparing Figs. 7.6 and 7.8, this cut does not strongly impact the top-squark signal, most notably for top-squark events with large values of ΔM.

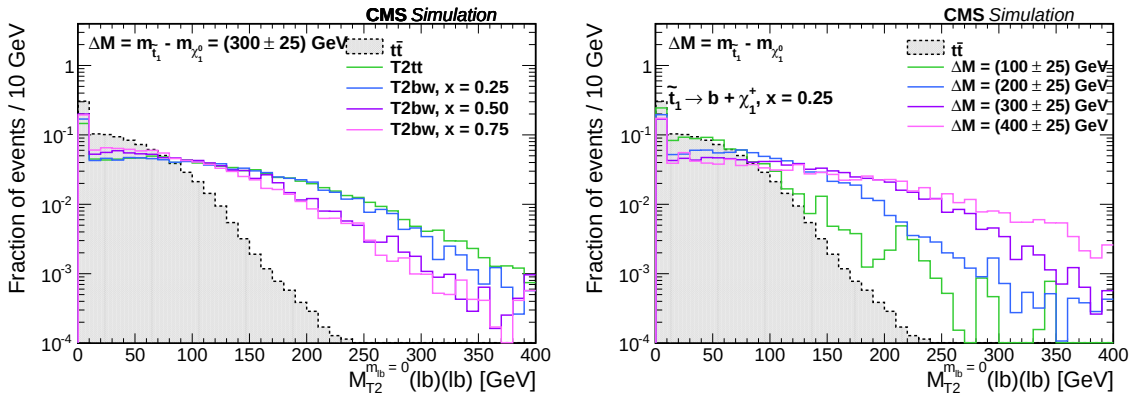

(a) The $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ distributions for several different top-squark decay modes for a ΔM value of 300 GeV.

(b) The $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure 7.9: The preselection-level distributions of $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ in the $\mathrm{t\bar{t}}$ background and the top-squark signal; distributions are normalized to the same area.

Removing the ℓb Masses from the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ Calculation

It should be clear now from the above discussions that the events in the high-$M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ tail are primarily comprised of events with large invariant masses for the ℓb pairs.

One method to reduce the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ tail is to assume that, in the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation, the input ℓb pairs are massless — i.e. $m_{\ell\mathrm{b}} = 0$ in Eq. (7.5). We denote this slight variant of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ by the symbol $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$.

The properties of $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ naturally differ from those of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$. Specifically, its lower bound is now $m_{\widehat{\nu}} = 0$, and its upper bound is now strictly less than m_{t}.

Figure 7.9 shows the analogous distributions as Fig. 7.6, but for $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$.

As can be seen, there is no clear kinematic edges in the $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution in $\mathrm{t\bar{t}}$ events; however, the $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ tail is much softer relative

to $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$. As with $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, the top-squark events have a much harder $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution relative to the $\mathrm{t\bar{t}}$.

7.2.2.2 Maximizing the Contribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ to the Top-squark Search

From the above discussions on $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, it is clear that the particular methods of constructing $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, i.e. whether you assume $m_{\ell\mathrm{b}} = 0$ or not, as well as any additional requirements, e.g. $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$, can strongly affect the overall shape of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, most notably the high-$M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ tail. What is not immediately obvious, however, is which of these particular methods for constructing $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, along with any additional requirements, yields the best overall variable in terms of the variable's statistical power to discriminate between the top-squark signal and the SM backgrounds.

We utilized the Punzi parameter, Appendix I, to rank these different $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ flavors based on how well they would perform in a hypothetical cut-and-count experiment where the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ variable is used to separate the top-squark signal from the SM background.

We found that $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ constructed with the requirement that $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$ had the optimal performance in terms of discriminating between the top-squark signal and the SM background. More details on this can be found in Appendix I.3.

7.2.3 Correlations between the M_{T2} Variants

It is clear from their defining equations, Eqs. (6.6), (7.1), and (7.4), that the three variants of M_{T2} that we have described in this dissertation will have some degree of correlation with one another. That does not preclude combining the information from each, however, so long as the correlations between each variant are relatively different in the top-squark signal and SM backgrounds.

Figure 7.10 compares the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathbb{W}}(\mathrm{bb})$ in the full data-driven background estimate against the T2tt and the $x = 0.25$ T2bw top-squark decay modes, both with $\Delta\mathrm{M} = 300\,\mathrm{GeV}$. Events are required to pass the "3D M_{T2} Shape Selection", which is described in more detail in Section 7.2.4.1.

The majority of the background events ($> 90\%$) that pass the preselection are expected to be $\mathrm{t\bar{t}}$ events. Thus, the overall distribution of the background is predominantly clustered in the region $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}, M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) \lesssim m_{\mathrm{t}}$. For the top-squark signal events, the overall correlation of $M_{\mathrm{T2}}(\ell\ell)$ and $M_{\mathrm{T2}}^{\mathbb{W}}(\mathrm{bb})$ is notably different between the two decay modes considered, but regardless of decay mode, a significant fraction of top-squark events are expected to be located at high values of either $M_{\mathrm{T2}}(\ell\ell)$ or $M_{\mathrm{T2}}^{\mathbb{W}}(\mathrm{bb})$.

For the T2bw decay mode, the shape of this correlation between $M_{\mathrm{T2}}(\ell\ell)$ and $M_{\mathrm{T2}}^{\mathbb{W}}(\mathrm{bb})$ is particularly dependent upon the chargino mass-splitting parameter, x. Figure 7.11 compares the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathbb{W}}(\mathrm{bb})$ for the $x = 0.50$ and $x = 0.75$ T2bw decay modes, both with $\Delta\mathrm{M} = 300\,\mathrm{GeV}$. As can be seen from comparing Figs. 7.10c , 7.11a, and 7.11b, the relative mass-splitting between the $\widetilde{\mathrm{t}}_1$,

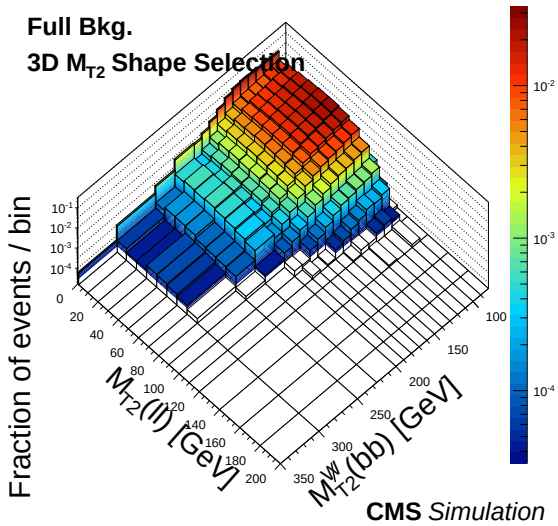

(a) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for the full data-driven background estimate.

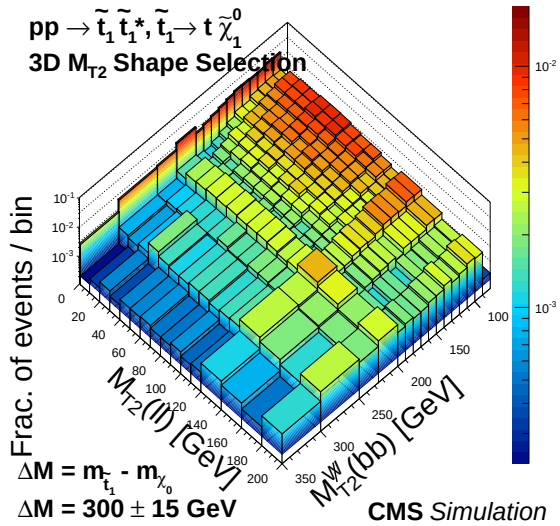

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0$ decay mode.

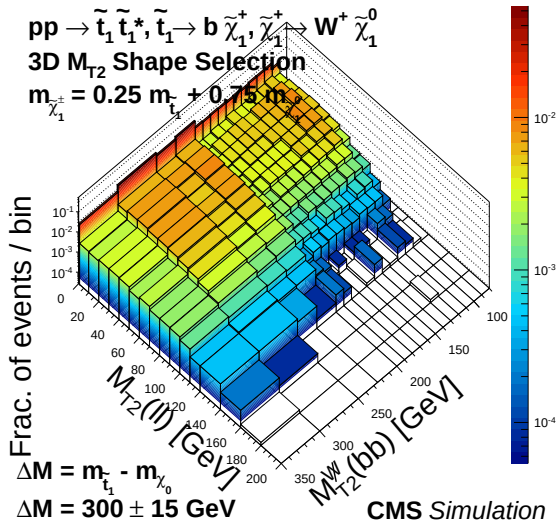

(c) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the $x = 0.25$ $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode.

Figure 7.10: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in the full data-driven background estimate and two top-squark decay modes with $\Delta\mathrm{M} = (300 \pm 15)\,\mathrm{GeV}$. Note the small slivers located near $M_{\mathrm{T2}}(\ell\ell) \approx 0\,\mathrm{GeV}$ are each individual bins.

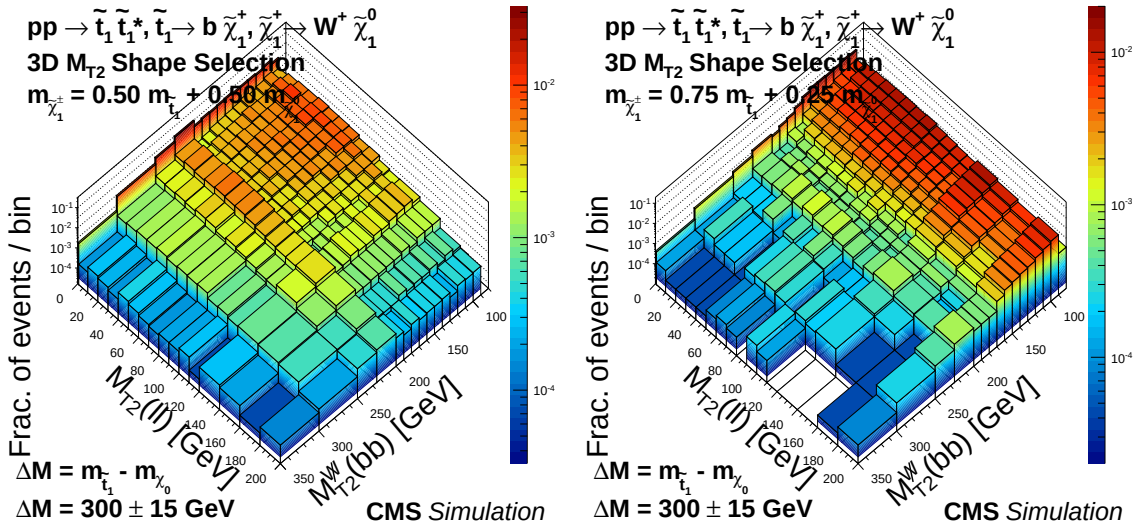

(a) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the T2bw, $x = 0.50$ decay mode.

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the T2bw, $x = 0.75$ decay mode.

Figure 7.11: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for the T2bw top-squark decay mode with $\Delta\mathrm{M} = 300\,\mathrm{GeV}$, for two different values of the chargino mass-splitting parameter, x, Note the small slivers located near $M_{\mathrm{T2}}(\ell\ell) \approx 0\,\mathrm{GeV}$ are each individual bins.

$\widetilde{\chi}_1^{\pm}$, and $\widetilde{\chi}_1^0$ strongly influences where the top-squark events tend to populate in this 2D parameter space.

In Appendix A.2.1, we have included figures showing the other permutations of the comparisons between the 3 M_{T2} variants as well as a comparison of the correlation among all 3 M_{T2} variants. In general, we observe similar behavior in these other correlation comparisons. The background, which is dominated by dileptonic $\mathrm{t\bar{t}}$ events, strongly populates the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$, while the signal tends to strongly populate the general signal region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV} \cup M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) > 170\,\mathrm{GeV} \cup M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) > 170\,\mathrm{GeV}$, although the specific regions where the signal is largest depend upon the top-squark decay mode considered.

7.2.4 Building the 3D M_{T2} Shape Analysis

From the discussions in the prior sections, it is clear that, for the purposes of distinguishing the top-squark signal from the SM backgrounds, each of the 3 M_{T2} variants provide correlated, but nevertheless complementary, information. In order to combine this information, we created a 3D shape where each axis is one of the M_{T2} variants.

We also considered training advanced machine learning algorithms such as a BDT or neural network. For this analysis, there are several key advantages of the 3D shape over BDTs (which are a representative proxy for advanced machine learning algorithms in general):

- In a general sense, the construction and development of a BDT often requires a significant amount of time dedicated to understanding systematic uncertainties beyond those of the basic simulated samples; for example, properly characterizing and controlling possible overtraining of the BDT, as well as ensuring that each input variable adds worthwhile information to the BDT, are both tasks that do not have definitive success metrics nor guaranteed completion times.

- By comparison, the modeling of the individual M_{T2} variants, most notably their correlation with one another, which is one of the key pillars of the 3D shape analysis, is fairly well understood in the context of the pre-existing systematic uncertainties.

- The behavior of the M_{T2} variants in top-squark events relative to the SM

background is fairly uniform in a general sense: the SM background favors small M_{T2} values, no matter the M_{T2} variant, while the top-squark signal predominantly favors large M_{T2} values for at least one of the M_{T2} variants. This behavior aids the interpretation and understanding of the final results.

- It is true that there are phenomenological reasons to believe that we might expect to find signatures of R-parity conserving top-squark pair=production in our data. However, the high-M_{T2} signal behavior is more general than just R-parity conserving top-squark pair-production. Thus, the 3D M_{T2} shape should be fairly sensitive to any BSM model that produces a $\mathrm{t\bar{t}}$-like final state with extra $\vec{\not{E}}_{\mathrm{T}}$. This is to be contrasted against a BDT approach, where the training of a BDT is always in the context of a specific signal, so every separate signal model that one wishes to test requires a separate BDT training in order to achieve the optimal results. A notable example in this dissertation is Section 7.1.3.2 where we showed the sensitivity of the lepton combination to different chirality scenarios for the T2tt decay mode. The analysis's sensitivity was worse for both the left- and right-handed chirality scenarios relative to the unpolarized T2tt decay mode. This was due to using BDTs that were trained for the unpolarized T2tt top-squark decay mode.

7.2.4.1 Object and Event Selection for the Shape Analysis

We utilize the same object and event selection requirements that were described in Section 6.2.3, with one additional requirement: based on the findings of

Section 7.2.2.2, we require $\max(m_{\ell\mathrm{b}}) < 200$ GeV, where ℓb refers to the ℓb pairings produced during the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation. Hereafter, we call this selection the "3D M_{T2} Shape Selection".

7.2.4.2 Background Estimation in the Shape Analysis

We utilize the same data-driven techniques to estimate our DY and misidentified lepton background contributions (Sections 6.3.2 and 6.3.3 respectively). Our technique for the $\mathrm{t\bar{t}}$ estimation is similar to Section 6.3.1, except instead of normalizing in the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, we instead utilize the region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$, to normalize our $\mathrm{t\bar{t}}$ contribution.

The treatment for possible signal contamination in the $\mathrm{t\bar{t}}$ estimate is the exact same as Section 6.6.5, but we update it to account for the changed M_{T2} control region. We found that the magnitude of the signal contamination remains the same with $\sim 10\%$.

7.2.4.3 Systematic Uncertainties in the Shape Analysis

The 3D M_{T2} shape analysis utilizes the same set of systematic uncertainties and the same overall procedure for estimating their impact as described in Section 6.4.

As a reminder, the majority of the individual sources of systematic uncertainty are, for the most part, treated on an object- and event-basis. This subsequently

enables the construction of 3D M_{T2} shapes that represent the expected shape variations from individual sources of systematic uncertainty, thus yielding effective shape uncertainties. When we compare the 3D M_{T2} shapes in the collision data against the expected shape of our background and signal, these shape uncertainties allow for morphing of the expected background and signal shape to better match the observed data for the hypothesis being tested, c.f. Appendix H.3. It is worth noting that we utilize the same procedure to propagate the systematic uncertainties on the $\not{E}_{\mathrm{T}}$ calculation into the calculation of $M_{\mathrm{T2}}\left(\ell\mathrm{b}\right)\left(\ell\mathrm{b}\right)$ and $M_{\mathrm{T2}}^{\mathrm{W}}\left(\mathrm{bb}\right)$.

It is important to validate this treatment of the systematic uncertainties; we devote additional discussion to this in the next section.

For the calculation of $M_{\mathrm{T2}}\left(\ell\mathrm{b}\right)\left(\ell\mathrm{b}\right)$ and $M_{\mathrm{T2}}^{\mathrm{W}}\left(\mathrm{bb}\right)$, as with the $M_{\mathrm{T2}}\left(\ell\ell\right)$ calculation, we do not force the systematic variations of these two M_{T2} variants to use the same input objects as the nominal simulation. That is, if there is a change in the specific choice of input "b jets" or the particular pairings of b jets and leptons, e.g. if a given input b jet is no longer considered b-tagged due to a variation in the b-tagging efficiency scale factor, we keep the change along with the subsequently different $M_{\mathrm{T2}}\left(\ell\mathrm{b}\right)\left(\ell\mathrm{b}\right)$ or $M_{\mathrm{T2}}^{\mathrm{W}}\left(\mathrm{bb}\right)$ values.

7.2.4.4 The Shapes of $M_{\mathrm{T2}}\left(\ell\mathrm{b}\right)\left(\ell\mathrm{b}\right)$ and $M_{\mathrm{T2}}^{\mathrm{W}}\left(\mathrm{bb}\right)$

As an addendum to the discussion in Section 6.5, this expanded version of the dilepton top-squark search also relies on a clear and precise understanding of the shape of the $M_{\mathrm{T2}}\left(\ell\mathrm{b}\right)\left(\ell\mathrm{b}\right)$ and $M_{\mathrm{T2}}^{\mathrm{W}}\left(\mathrm{bb}\right)$ variables, as well as the correlations

between all three M_{T2} variants considered.

Shapes of the Additional M_{T2} Variants

In SM $\mathrm{t\bar{t}}$, as with $M_{\mathrm{T2}}(\ell\ell)$, the Gaussian core of the $\not{E}_{\mathrm{T}}$ resolution is one of the main driving factors that turns a nominally hard cut-off for $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ at m_{t} into a falling kinematic edge. Because both $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ rely on b jet information as well, the Gaussian core of the jet energy resolution conflates into this resolution component of the falling kinematic edge for both variables. For $M_{\mathrm{T2}}(\ell\ell)$, the intrinsic width of the W boson was the other, main driving component in its shape. Similarly, for $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$, the intrinsic width of the top quark also affects the shape of the falling kinematic edge.

As noted above in our initial discussions of these variables, the correct choice of input b jets as well as the correct pairing of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ also contribute notably to the shape of these two variables.

We already discussed the relative importance of these effects on the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ reconstruction, as shown in Fig. 7.5.

For $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, as demonstrated in Fig. 7.7, the shape of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$, specifically the relative fraction of high-$M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ events, has a particularly strong dependence upon the correct choice of pairings between the input leptons and b jets. At a more fundamental level, this strong dependence on the correct pairing stems from the overall dependence of $\max(m_{\ell\mathrm{b}})$ on the correct choice of pairings, as shown in Fig. 7.7b.

Control Region Checks

As with Section 6.5.2, we utilized a Z-enriched control region to check the modeling of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$. We used the same selection as described in Section 6.5.2. Remember that as part of this selection, we require there be **no** reconstructed b jets in the event. Thus, the input "b jets" for the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ calculation are likely not actually "true" b jets.

Figure 7.12 shows the distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ and $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for events in this Z-enriched control region. There seems to be a slight systematic trend in the ratio of data and simulation, wherein the simulation slightly underestimates the higher values of the two respective M_{T2} variant, but we note that the modeling of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ is acceptable within the systematic uncertainties. The modeling of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ is not as good however; notably, It is clear that there is an overestimate of low $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ values in the SM background.

One possible explanation that would nominally explain both the slight systematic trend as well as the overestimate of low $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ values is that the resolution of either the $\not{E}_{\mathrm{T}}$ or the jet energies, or both, is underestimated in the simulation; however, as discussed in Section 6.2.4.3, the underestimation of $\not{E}_{\mathrm{T}}$ resolution in the simulation is controlled by smearing the $\vec{\not{E}}_{\mathrm{T}}$ in simulation. Although we did not perform an analogous smearing on the jet-energies, careful checks of the modeling of variables that would be directly affected by underestimated jet energy resolutions in the simulation found no signs that jet energies are being significantly underestimated by the simulation.

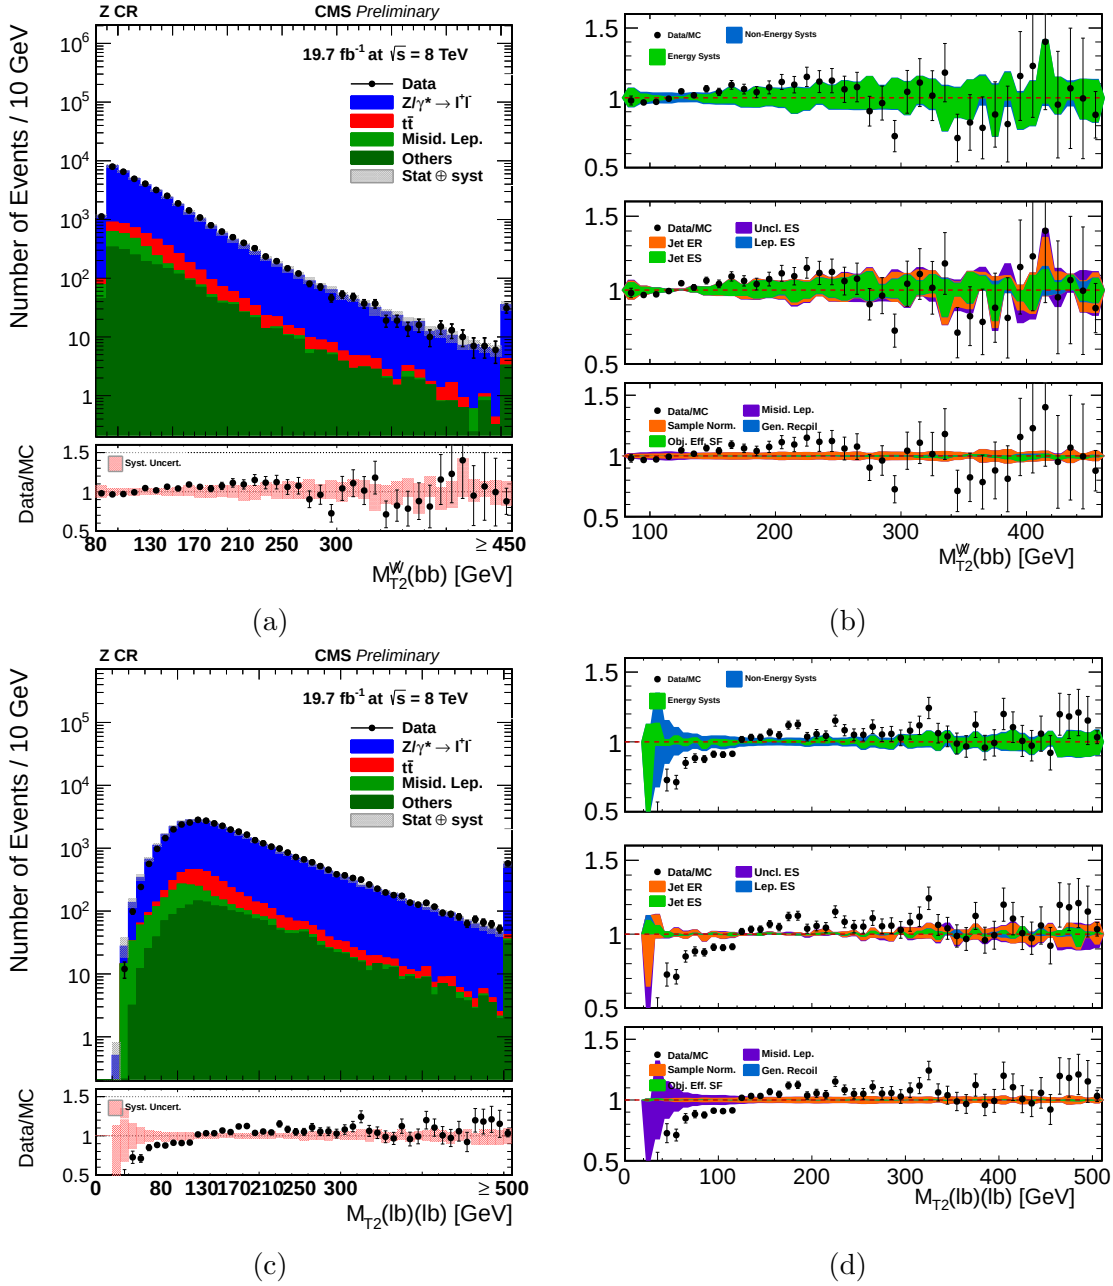

(a) (b)

(c) (d)

Figure 7.12: The distributions of M_{T2}^{W} (bb) and M_{T2} (ℓb) (ℓb) in the sample of events in the Z-enriched control region. Shown on the right of each distribution is an itemized breakdown of the contributions from individual sources of systematic uncertainty

Instead, more careful investigations have shown that the overestimate of low $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ values by the background is due to an overestimate of the misidentified lepton contribution in this region. Although there is an overestimate of the misidentified lepton background at low $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in this control region, the applied systematic uncertainties – i.e. the purple in the bottom panel of Fig. 7.12d – on the misidentified lepton estimate almost cover the observed discrepancy. The overall systematic uncertainty on just the misidentified lepton background estimate in this low $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ region is $\mathcal{O}(65\%)$ and is completely dominated ($\sim 99.9\%$) by the statistical uncertainties on the measured fake rate f.

The systematic uncertainties on the misidentified lepton estimate are treated as a shape uncertainty. Thus, discrepancies between the data and simulation can be (and are) mitigated by morphing the background M_{T2} shapes within their $\pm 1\sigma$ shape systematic uncertainties to best match the observed data, c.f. Appendix H.3.

We have performed other cross-checks on the modeling of the new M_{T2} variants and found no notable issues. Details on some of these cross-checks can be found in Appendices C.1 and C.3.

7.2.4.5 M_{T2} Binning Choices

Although the overall shape of the 3 M_{T2} variants is acceptably modeled by the simulation, we improve the robustness of the shape analysis by binning each of the M_{T2} variants into relatively coarse ranges.

Table 7.1 displays the binning choices we made for the individual M_{T2} vari-

Table 7.1: The M_{T2} ranges used for each of the M_{T2} variants.

M_{T2} Variant	M_{T2} value range for bin [GeV]			
	Bin 1	Bin 2	Bin 3	Bin 4
$M_{\mathrm{T2}}\,(\ell\ell)$	< 80	80:100	100:120	> 120
$M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$	< 170	170:200	200:250	> 250
$M_{\mathrm{T2}}^{\mathrm{W}}\,(\mathrm{bb})$	< 170	170:200	200:250	> 250

ants. Each M_{T2} variant was partitioned into 4 ranges. This partitioning is applied coherently across the whole 3D M_{T2} shape, leading to a total of 64 bins in the 3D M_{T2} phase space. The bin corresponding to the range, $M_{\mathrm{T2}}\,(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}\,(\mathrm{bb}) < 170\,\mathrm{GeV}$, serves as a control bin for normalizing the expected dileptonic $\mathrm{t\bar{t}}$ contributions to the other 63 bins — c.f. Sections 6.3.1 and 7.2.4.2.

Dynamically Rebinning Bins with 0 Background Events

The choice of binning shown in Table 7.1 leads to bins where there is no expected contribution from background events — e.g. the region $M_{\mathrm{T2}}\,(\ell\ell) > 120\,\mathrm{GeV} \cap M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b}) > 250\,\mathrm{GeV}$ in Fig. A.1a. When comparing the 3D M_{T2} shapes of data, background and our top-squark signal, the bins where there is no expected contribution from background events are added into neighboring bins in order to improve the robustness of our statistical methods (Appendix G) used to calculate the statistical limits and observed significances in Section 7.2.5.

For the calculation of statistical limits, this rebinning is performed dynamically depending upon the top-squark decay mode hypothesis being tested and the specific

ΔM region being considered.

Iterating backwards from the highest M_{T2} values for the 3 M_{T2} variants, every bin B_i where there are less than a threshold b_0 events in the full background estimate is rebinned with one of the neighboring bins that has an equal or lesser position along each M_{T2} axis, denoted by B_j. The choice of which bin B_j to rebin B_i with is made utilizing the Punzi parameter to compare each of the possible bin recombinations, Appendix I.2.2.

In order to minimize the relative loss of independent sources of information, we utilized a threshold $b_0 = 0.25$ events. With this threshold, we rebin ~ 30 bins in the 3D M_{T2} shape, leaving ~ 33 signal region bins.

Note that, when calculating observed significances for a given top-squark decay mode, we fixed the binning to one of five ΔM-based binning choices. This ensures the relative consistency of the data and background shapes across the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane. As a cross-check, we compared the statistical results when using other binning choices and found they did not significantly vary.

7.2.5 Results of the Top-squark Search using the 3D M_{T2} Shape

Figure 7.13 shows the projected 1D distributions of the three M_{T2} variants used for the 3D M_{T2} shape analysis. The modeling of all three variables is excellent throughout most of the individual M_{T2} phase spaces. There is a $\sim 15\%$ excess in the data in the first bin of the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ distribution, corresponding to the range, $80\,\mathrm{GeV} < M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 90\,\mathrm{GeV}$. However, this is almost completely consistent

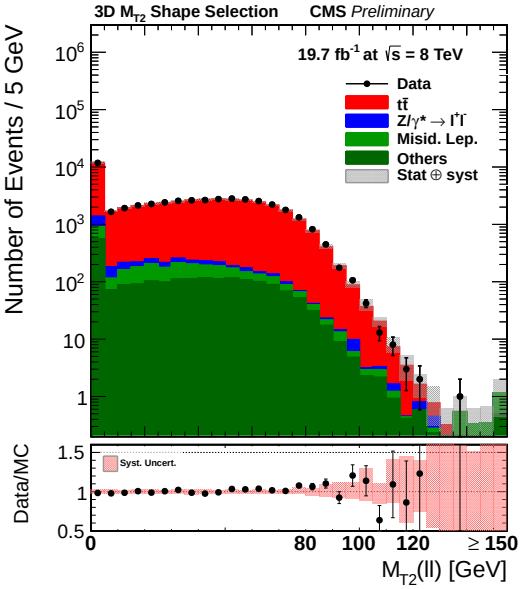

(a) The distribution of $M_{\mathrm{T2}}(\ell\ell)$.

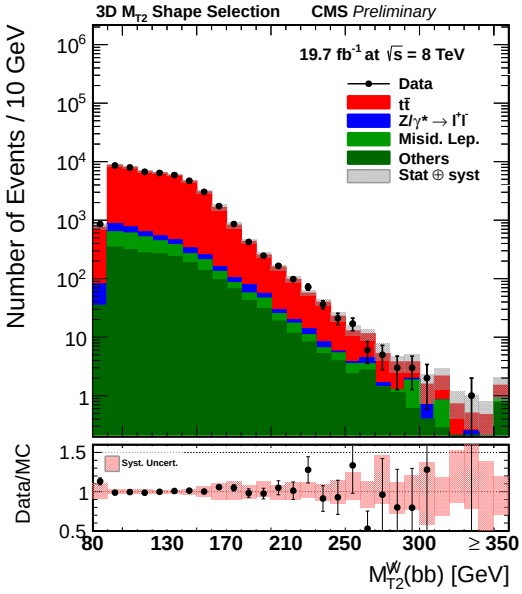

(b) The distribution of $M_{\mathrm{T2}}^{W/W}(\mathrm{bb})$.

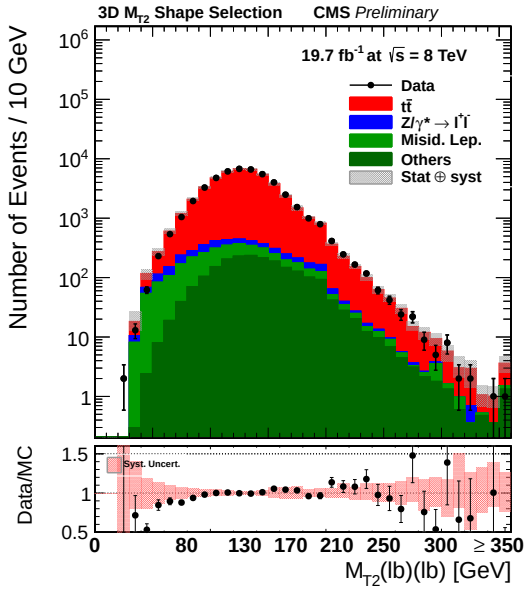

(c) The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$.

Figure 7.13: The distributions of the three M_{T2} variables used in the basic version of the 3D M_{T2} shape analysis; each distribution was made using the sample of events passing the 3D M_{T2} shape selection, defined in Section 7.2.4.1.

within the expected systematic uncertainties for this range; as a side note, for this bin, more than 90% of the total systematic uncertainty stems from the systematic uncertainty on the JES. The discrepancy between the data and background expectations for $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 100\,\mathrm{GeV}$ is almost entirely due to the modeling of the $\max(m_{\mathrm{b}\ell})$ variable in the misidentified lepton estimate range, similar to our findings in Section 7.2.4.4. As with Fig. 7.12d, the systematic uncertainty on the background estimate in this region is completely dominated by the statistical uncertainty on the measured fake rate f.

7.2.5.1 "De-stacking" the 3D M_{T2} Shape

Before we performed the statistical limit and significance calculations, after performing the dynamic rebinning described in Section 7.2.4.5, we first "de-stacked" the 3D M_{T2} shape – every bin in the 3D shape was assigned a global bin number, B_G, based upon its values for B_x, B_y, and B_z, the bin numbers for the x $[M_{\mathrm{T2}}(\ell\ell)]$, y $[M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})]$, and z $[M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})]$ axes respectively. This mapping was done via the following equation,

$$B_G = 16\,(B_x - 1) + 4\,(B_y - 1) + B_z. \tag{7.6}$$

For future reference, to "re-stack", i.e. extract the original B_x, B_y, and B_z values for a given B_G, we utilized the following equation,

$$
\begin{aligned}
B_x &= 1 + \text{floor}\left(\frac{B_G}{16}\right) \\
B_y &= 1 + \text{floor}\left(\frac{B_G - 16\,(B_x - 1)}{4}\right) \\
B_z &= 1 + \text{floor}\left(\frac{B_G - 16\,(B_x - 1) - 4\,(B_y - 1)}{1}\right),
\end{aligned}
\tag{7.7}
$$

where "floor" refers to the result of integer division.

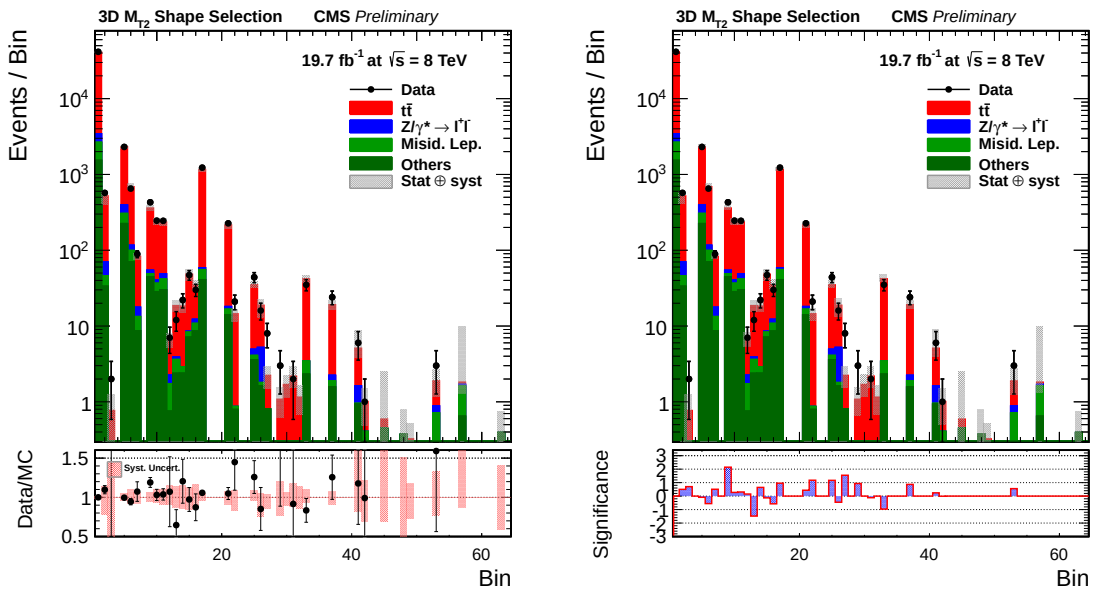

(a) (b)

Figure 7.14: A 1D "de-stacking" of the 3D M_{T2} distribution used in the basic version of the 3D M_{T2} shape analysis. Every bin in this 1D distribution corresponds to an individual bin from the 3D M_{T2} shape, based on the mapping in Eq. (7.6). The top panel is the same in each sub-figure. In the left sub-figure, the data points in the bottom panel represent the ratio of data and simulation, while the red band represents the total systematic uncertainty on the ratio. In the right sub-figure, the bottom panel represents the statistical significance of any deviations of the aforementioned ratio from unity, where the significance, for each bin, has been calculated using the average systematic uncertainty added in quadrature with the statistical uncertainty on the ratio.

The results of this, using the rebinning optimized for the unpolarized T2tt decay mode with $\Delta\mathrm{M} = 500\,\mathrm{GeV}$, are shown in Fig. 7.14. There is agreement

within systematic uncertainties for most of the bins. After accounting for these systematic uncertainties, there is a notable observed local excess of $\sim 2\sigma$ in the 9th bin, corresponding to the region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV} \cap 200\,\mathrm{GeV} \leq M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 250\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$. There are also a number of observed deficits, but these primarily fall in bins with very little available data.

Note that, although the 1st bin, corresponding to the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$, is shown in Fig. 7.14, it is removed for the limit setting calculation because it was the control bin used to normalize the $\mathrm{t\bar{t}}$ contribution.

7.2.5.2 Limits on Top-squark Pair-Production

Similar to Section 6.8.3 we calculated the observed and expected upper limits on the top-squark pair production cross-section. These upper limits were then subsequently transformed into excluded regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ 2D plane.

Figure 7.15 shows the expected and observed mass exclusions for the unpolarized-coupling scenarios for the four top-squark decay modes considered.

Direct comparisons between Fig. 7.15 and the dileptonic cut-and-count results, Fig. 6.38, shows a number of interesting results.

For the T2tt decay mode, the sensitivity at high $m_{\tilde{\mathrm{t}}_1}$ is extended by $\sim$ 25–50 GeV, depending upon $m_{\tilde{\chi}_1^0}$. Moreover, for this decay mode, the 3D M_{T2} shape analysis is expected to have sensitivity in the experimentally challenging region, $\Delta\mathrm{M} = m_{\mathrm{t}}$, for values of $m_{\tilde{\chi}_1^0} \lesssim 50\,\mathrm{GeV}$. The observed exclusions match the expec-

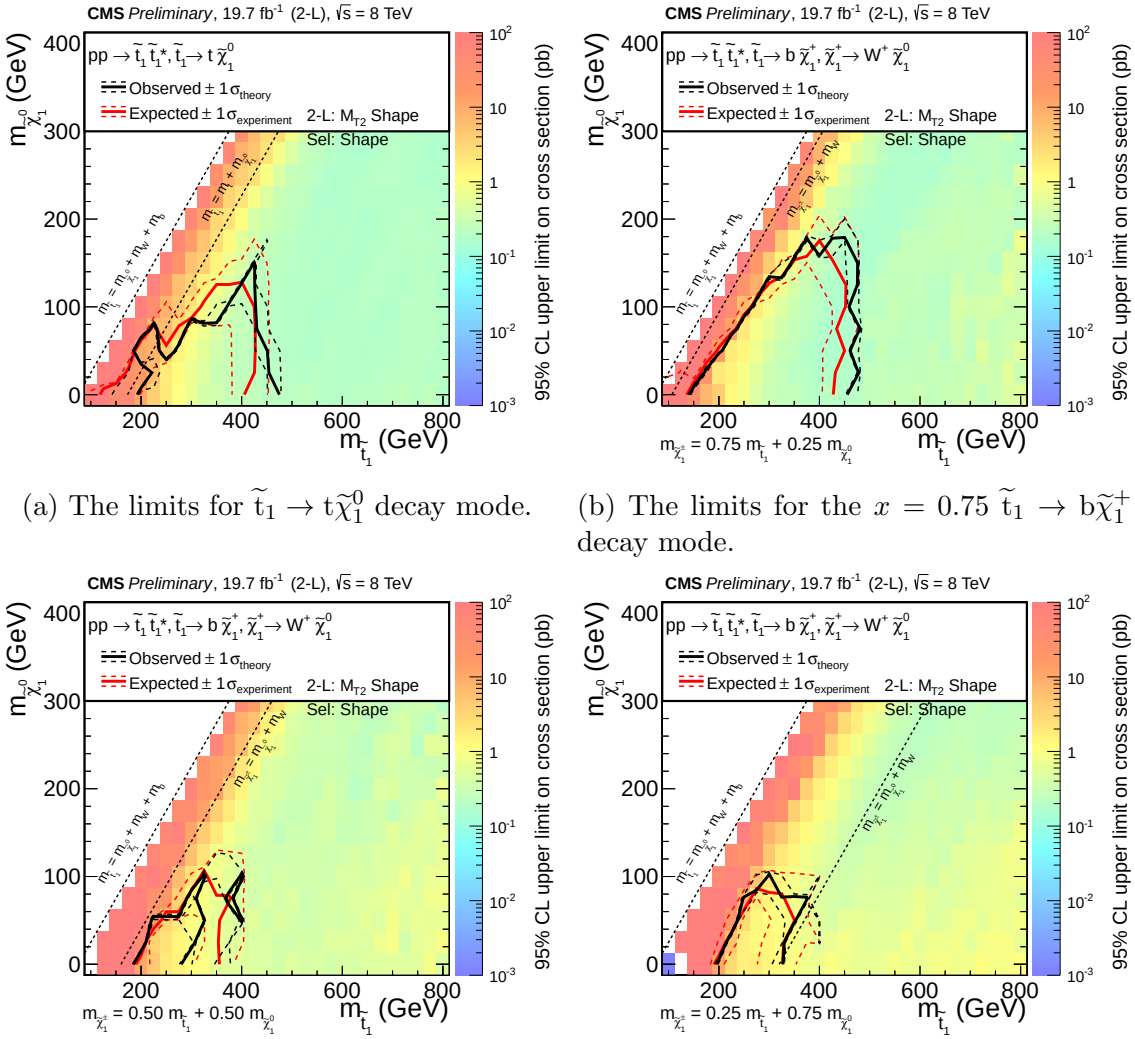

(a) The limits for $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode.

(b) The limits for the $x = 0.75$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode.

(c) The limits for the $x = 0.50$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode.

(d) The limits for the $x = 0.25$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode.

Figure 7.15: 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the 3D M_{T2} shape analysis. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The two solid contour lines were made using the median expected cross-section limit (red) and the observed cross-section limit (black). The dashed red lines represent the exclusion region when using the 0.16 and 0.84 percentile expected limits, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

tations for larger $m_{\widetilde{\chi}_1^0}$, but notably, do not for $m_{\widetilde{\mathrm{t}}_1} \approx m_{\mathrm{t}}$, $m_{\widetilde{\chi}_1^0} \approx 0\,\mathrm{GeV}$.

For the T2bw decay mode, when $x = 0.75$, we do not observe any significant differences between the sensitivity of the 3D M_{T2} shape analysis and the dileptonic cut-and-count analysis. However, there are notable increases in sensitivity for the $x = 0.25$ and $x = 0.50$ scenarios. This can be seen both from the observed and expected exclusion contours as well as the relative lack of missing parts (i.e. no purplish-blue squares) in the cross-section limit map for the 3D M_{T2} shape analysis.

7.2.5.3 Sensitivity of the 3D M_{T2} Shape Analysis to Varying Top-squark Daughter Chiralities

As in Section 6.8.3.1, we also calculated the dependence of the observed and expected exclusion regions on the polarizations and chiralities of the top-squark daughter particles.

Figure 7.16a is analogous to Fig. 6.39, except it shows the sensitivity of the 3D M_{T2} shape analysis to different coupling chiralities between the $\widetilde{\mathrm{t}}_1$ and t. The (left-handed) right-handed coupling is shown in (blue) red. This figure demonstrates that the 3D M_{T2} shape analysis is much more sensitive, relative to the cut-and-count dilepton analysis, to the coupling chirality between the $\widetilde{\mathrm{t}}_1$ and t. For example, in the right-handed scenario, there is a complete observed exclusion of the region $\Delta\mathrm{M} = m_{\mathrm{t}}$, $m_{\widetilde{\chi}_1^0} \lesssim 50\,\mathrm{GeV}$. As well, the observed sensitivity is extended by $\sim 50\,\mathrm{GeV}$ for the relatively large $m_{\widetilde{\mathrm{t}}_1}$ when $m_{\widetilde{\chi}_1^0}$ is relatively large as well. To constrast this, in the left-handed scenario, there is no observed sensitivity in the region, $\Delta\mathrm{M} = m_{\mathrm{t}}$,

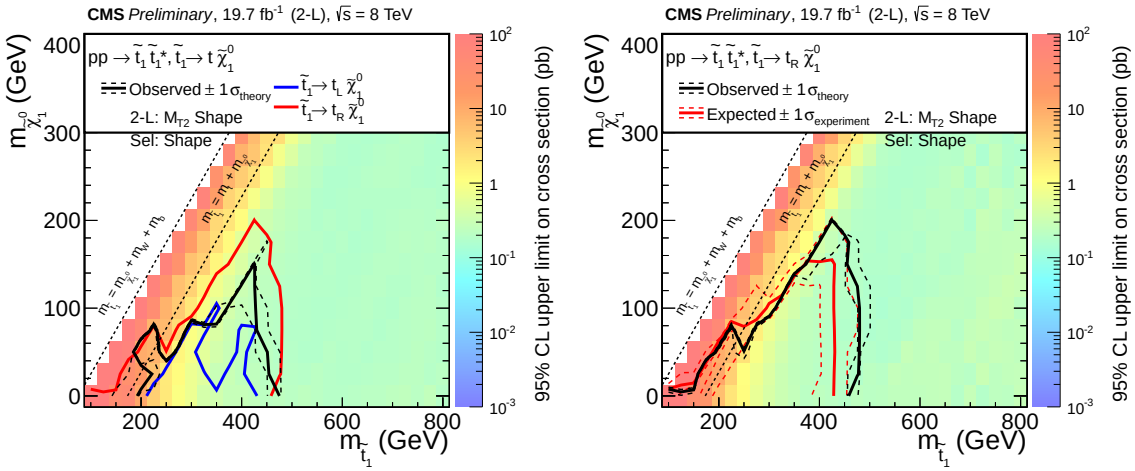

(a) The observed limits for all three variations of the $\widetilde{t}_1$-t coupling.

(b) The observed limits for the $\widetilde{t}_1 \to t_R\widetilde{\chi}_1^0$ decay mode.

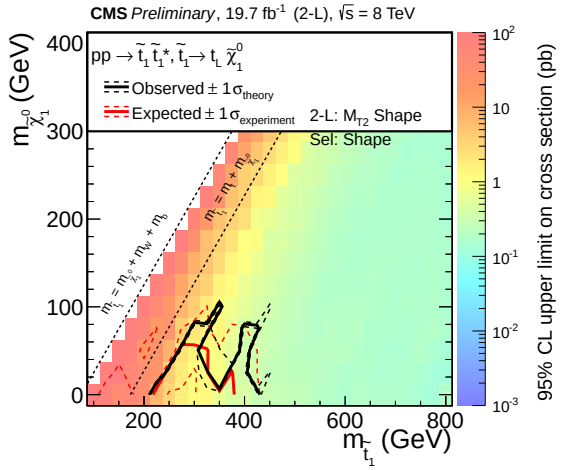

(c) The observed limits for the $\widetilde{t}_1 \to t_L\widetilde{\chi}_1^0$ decay mode.

Figure 7.16: 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the variants of the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode. For the top left, the black line denotes the observed exclusion region for the unpolarized, i.e. $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$, decay mode. The red (blue) solid line denotes the observed exclusion region when the $\widetilde{t}_1$ only has a (left-) right-handed coupling to the t. The upper right [bottom middle] sub-figure shows the observed (black) and expected (red) exclusion regions for the $\widetilde{t}_1 \to t_R\widetilde{\chi}_1^0$ [$\widetilde{t}_1 \to t_L\widetilde{\chi}_1^0$] decay mode. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

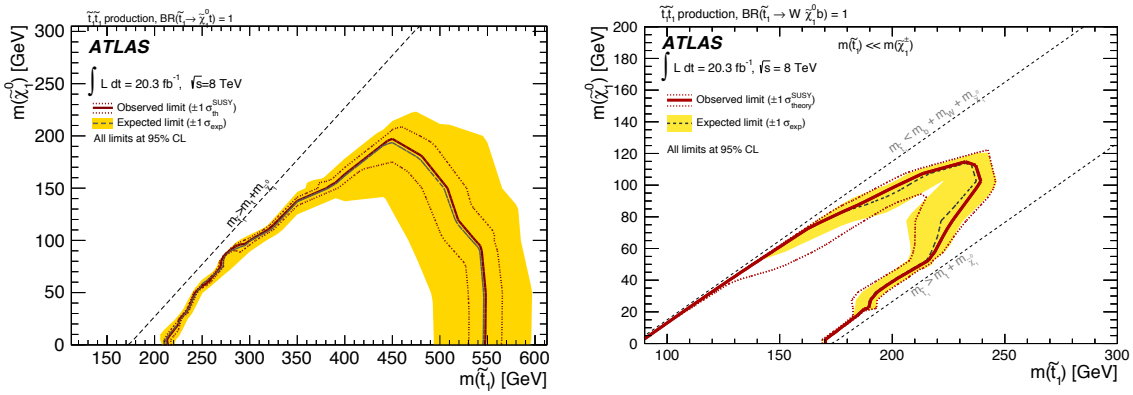

(a) The limits for the region, $\Delta\mathrm{M} \geq m_{\mathrm{t}}$, where the top-quark is produced on-shell in the top-squark decay. Reprinted from Fig. 18 of [59].

(b) The limits for the region, $\Delta\mathrm{M} < m_{\mathrm{t}}$, where the top-quark is produced off-shell in the top-squark decay. Reprinted from Fig. 19 of [59].

Figure 7.17: Exclusion limits from an ATLAS analysis [59] searching for direct top-squark pair-production in the dileptonic final state. Matching with the terminology of this dissertation, the assumed decay mode is the right-handed T2tt decay mode.

and in general the observed sensitivity is worse, notably for the region, $\Delta\mathrm{M} \approx 300\,\mathrm{GeV}$, $m_{\tilde{\chi}_1^0} \gtrsim 50\,\mathrm{GeV}$.

For comparison, Figs. 7.16b and 7.16c provide comparisons of the observed and expected sensitivity for the right- and left-handed polarization scenarios.

These results provide a notably strong endorsement of the 3D M_{T2} shape analysis. To the author's knowledge, the 3D M_{T2} shape analysis is the **only** direct search for top-squark pair-production, across all top-squark pair final states, that, using data from Run 1 of the LHC, has observed exclusions in the region, $\Delta\mathrm{M} = m_{\mathrm{t}}$, for the T2tt decay mode. Moreover, the 3D M_{T2} shape analysis can achieve relatively comparable results in other regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane as more complicated approaches.

For example, the ATLAS experiment performed an analogous dileptonic top-squark search [59]. Their approach was a multivariate one based on a BDT con-

structed with $M_{\mathrm{T2}}(\ell\ell)$, a slight variant of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$, and several other discriminating variables. Figure 7.17 shows their observed and expected exclusion contours for what is effectively the right-handed T2tt decay mode.

Direct comparisons of Figs. 7.16b and 7.17 show that ATLAS's observed exclusion is better by about 10% for high $m_{\tilde{\mathrm{t}}_1}$. This result implies that the usage of BDTs for this decay mode does not bring significantly more sensitivity for this region of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane.

As with the cut-and-count dileptonic top-squark search, we also investigated how the variations of $\tilde{\chi}_1^{\pm}$ and W-boson chiralities affect the observed and expected $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ exclusion regions for the T2bw decay mode with different x values. We observed no significant change in the calculated limits for any of the considered chirality scenarios.

7.2.5.4 Observed Significance of Excesses in the Data

There are several bins in Fig. 7.14 that show an observed excess in the data. These observed excesses lead to slightly worse observed exclusion limits compared to what's expected. For example, Fig. 7.15a shows that there is a slight discrepancy between the expected and observed exclusion of top-squark pair production in the region, $m_{\tilde{t}_1} = m_t$, $m_{\tilde{\chi}_1^0} = 0\,\text{GeV}$.

We quantified these observed excesses in terms of local p-values for the *null* hypothesis. These null p-values are calculated using likelihood ratios based on comparisons of the *background-only* hypothesis against an unconstrained *signal + background* hypothesis. The unconstrained qualifier signifies that we allow our signal normalization to float to best match the observed data. Null p-values calculated in this fashion thus quantify how *signal-like* an observed excess is, see Appendix G.3. We then directly relate these null p values into observed local significances (the infamous $n\sigma$) through use of a 1-sided Gaussian tail integral, see Eq. (G.10).

It is important to note that there were a few slight changes between the statistical limit calculations and the o As part of these significance calculations, we removed the signal contamination corrections to the $t\bar{t}$ estimation we included the information from

A number of points in the 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane yielded strong local significances. However, the overall $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ map of the local significances did not have any particularly striking structure. Thus, in order to better quantify these excesses, we performed a number of extensions to the 3D M_{T2} analysis. The first of these

was splitting the event samples by dilepton flavor in each event (i.e. e^+e^-, $e^\pm\mu^\mp$, $\mu^+\mu^-$).

7.2.6 Extending the 3D M_{T2} Shape Analysis: Splitting by Dilepton Event Flavor

The relevant backgrounds for this analysis do not contribute equally to each dilepton channel. For example, the $\mathrm{Z} \to \ell^+\ell^-$ background contributes primarily in the same-flavor channels,

Our methods for estimating the background contributions, including the impact of systematic uncertainties, accounted for these differences between the different dilepton channels. Up to this point, however, when comparing expected background and signal yields against the observed data, we had considered all three dilepton channels as a single unit. This intrinsically ignores the complementary information each individual dilepton channel can provide.

In an attempt to capitalize on this additional information, we reiterated through the 3D M_{T2} shape analysis, splitting the composite 3D M_{T2} shape into separate shapes for each of the individual dilepton channels.

In order to ensure a robust calculation of relevant test statistics for each dilepton shape, we recalculated the dynamic rebinning maps [Section 7.2.4.5], requiring every signal region in each of the three dilepton shapes to have at least b_0 events, where we utilized the same threshold, $b_0 = 0.25$ events.

7.2.7 Results when Splitting into Individual Dilepton Channels

Figures 7.18 through 7.20 show the same information as Fig. 7.20, except the distributions have been split by dilepton event flavor. The modeling by the simulation of these projected distributions is excellent throughout the bulk of the distributions. Cross-checks in other control regions corroborated this finding.

As with the inclusive shape analysis, before calculating statistical limits and observed significances, we "de-stacked" the overall 3D shape, using the same equation to perform the mapping.

The results of this, using the rebinning optimized for the unpolarized T2tt decay mode with $\Delta\mathrm{M} = 500\,\mathrm{GeV}$, are shown in Figs. 7.21 through 7.23. There is agreement within systematic uncertainties for most of the bins. After accounting for these systematic uncertainties, there are still some notable observed excesses.

For example, there is an $\sim 2.1\sigma$ excess observed in the 9th bin of the $\mathrm{e}^{\pm}\mu^{\mp}$ distribution, corresponding to the region $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, $200\,\mathrm{GeV} \leq M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 250\,\mathrm{GeV}$, $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$.

In the $\mathrm{e}^{+}\mathrm{e}^{-}$ distribution, there is an $\sim 1.7\sigma$ excess observed in the 17th bin, corresponding to the region $80\,\mathrm{GeV} \leq M_{\mathrm{T2}}(\ell\ell) < 100\,\mathrm{GeV}$, $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 170\,\mathrm{GeV}$, $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$.

Finally, in the $\mu^{+}\mu^{-}$ distribution, there is an $\sim 2.1\sigma$ excess observed in the 21st bin, corresponding to the region $80\,\mathrm{GeV} \leq M_{\mathrm{T2}}(\ell\ell) < 100\,\mathrm{GeV}$, $170\,\mathrm{GeV} \leq M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 200\,\mathrm{GeV}$, $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$.

There are a number of observed deficits as well, but these are primarily in bins

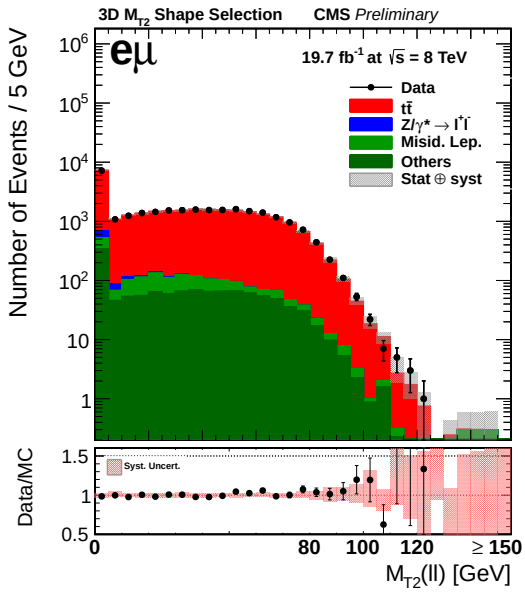

(a) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ for $e^{\pm}\mu^{\mp}$ events.

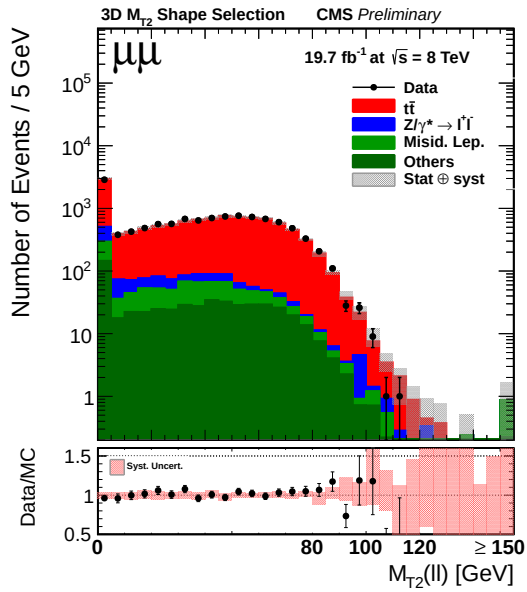

(b) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ for $\mu^+\mu^-$ events.

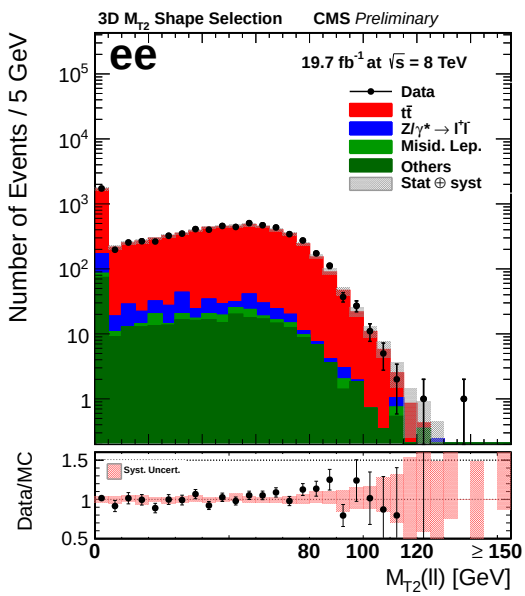

(c) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ for e^+e^- events.

Figure 7.18: The distribution of the $M_{\mathrm{T2}}(\ell\ell)$ variable, split by dilepton event flavor, in the sample of events passing the 3D M_{T2} shape selection, defined in Section 7.2.4.1.

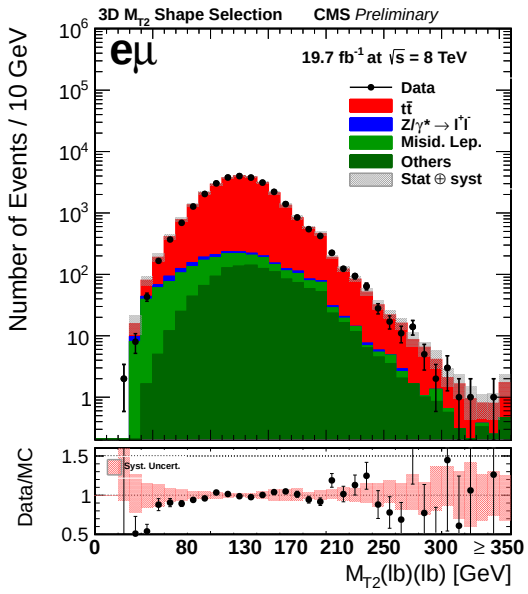

(a) The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for $\mathrm{e}^{\pm}\mu^{\mp}$ events.

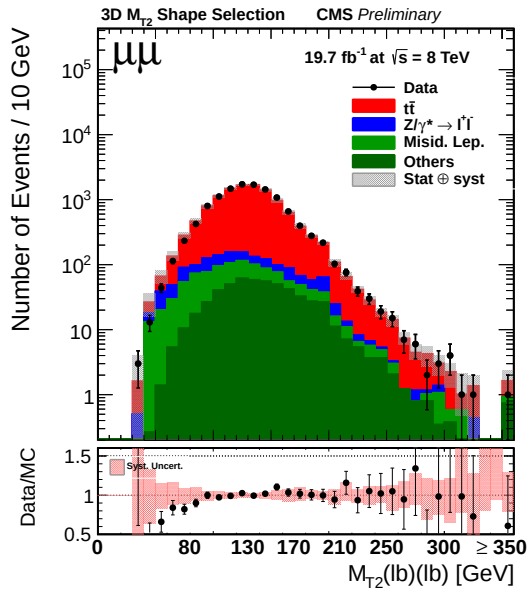

(b) The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for $\mu^+\mu^-$ events.

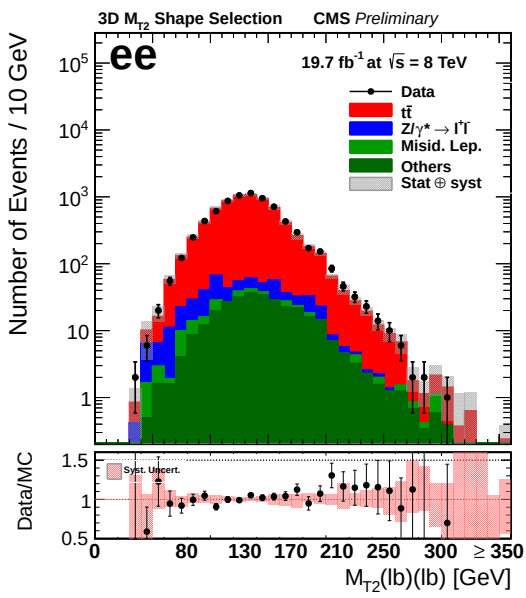

(c) The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for $\mathrm{e}^+\mathrm{e}^-$ events.

Figure 7.19: The distribution of the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ variable, split by dilepton event flavor, in the sample of events passing the 3D M_{T2} shape selection, defined in Section 7.2.4.1.

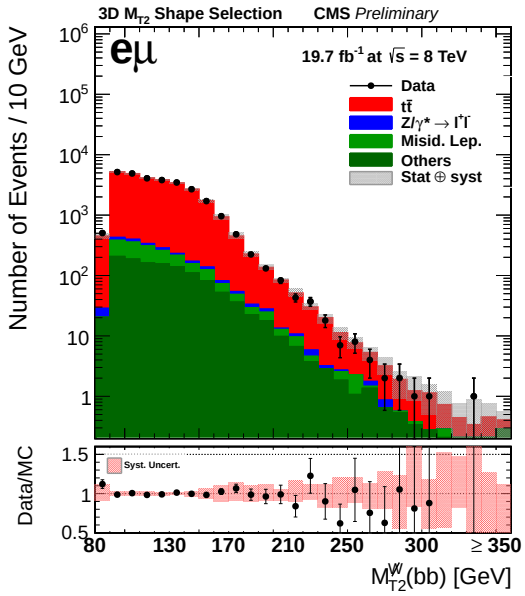

(a) The distribution of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for $\mathrm{e}^{\pm}\mu^{\mp}$ events.

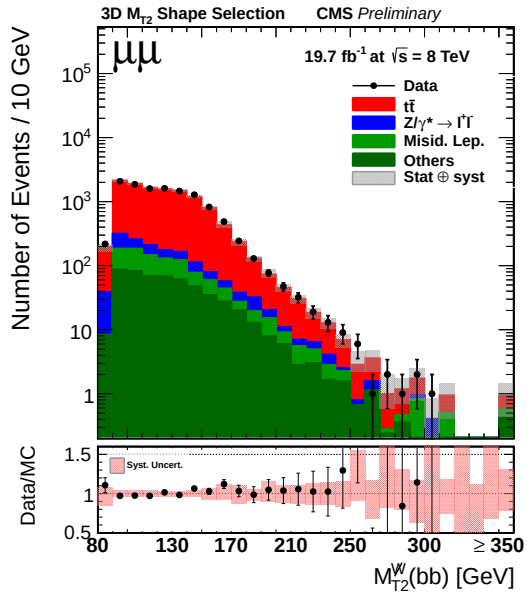

(b) The distribution of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for $\mu^{+}\mu^{-}$ events.

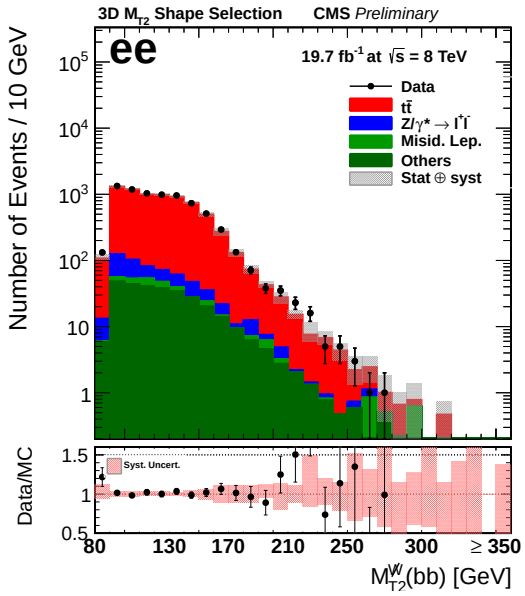

(c) The distribution of $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for $\mathrm{e}^{+}\mathrm{e}^{-}$ events.

Figure 7.20: The distribution of the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ variable, split by dilepton event flavor, in the sample of events passing the 3D M_{T2} shape selection, defined in Section 7.2.4.1.

with relatively low expected backgrounds, where the data seems to have fluctuated downward relative to these expectations.

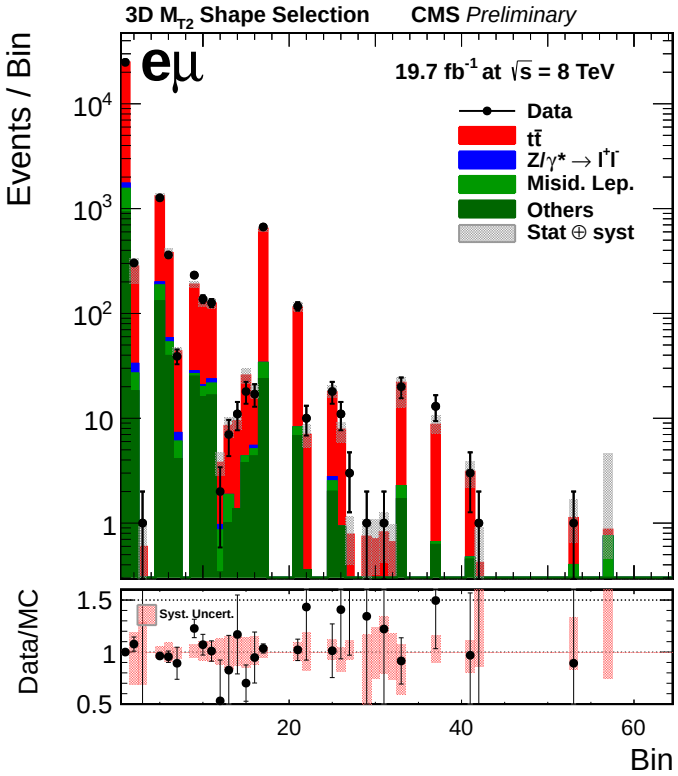

(a)

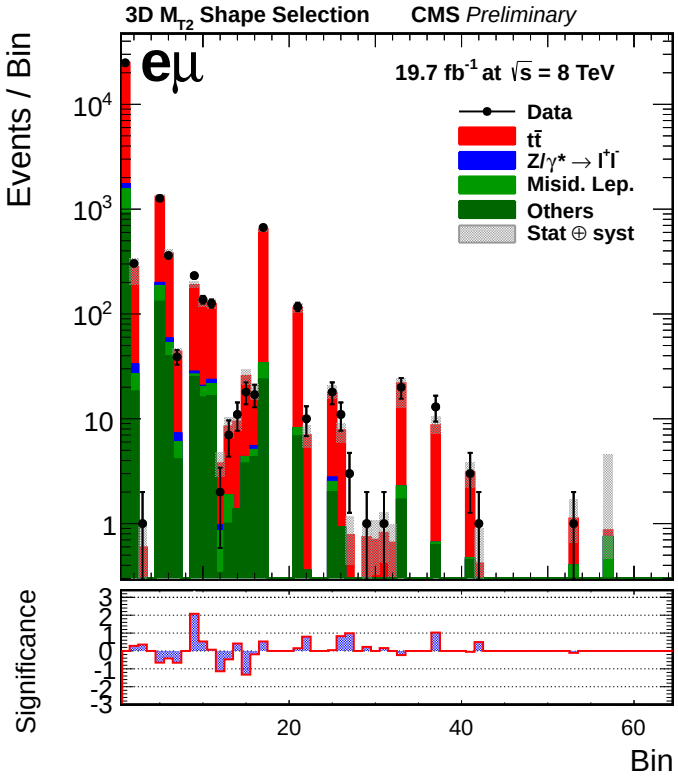

(b)

Figure 7.21: The 1D "de-stacking," Fig. 7.14, but shown for only e^+e^- events.

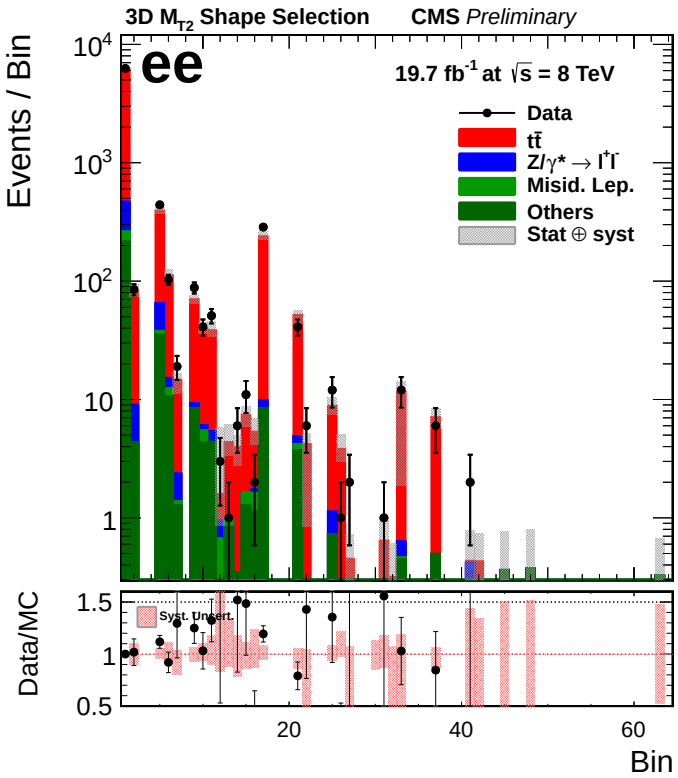

(a)

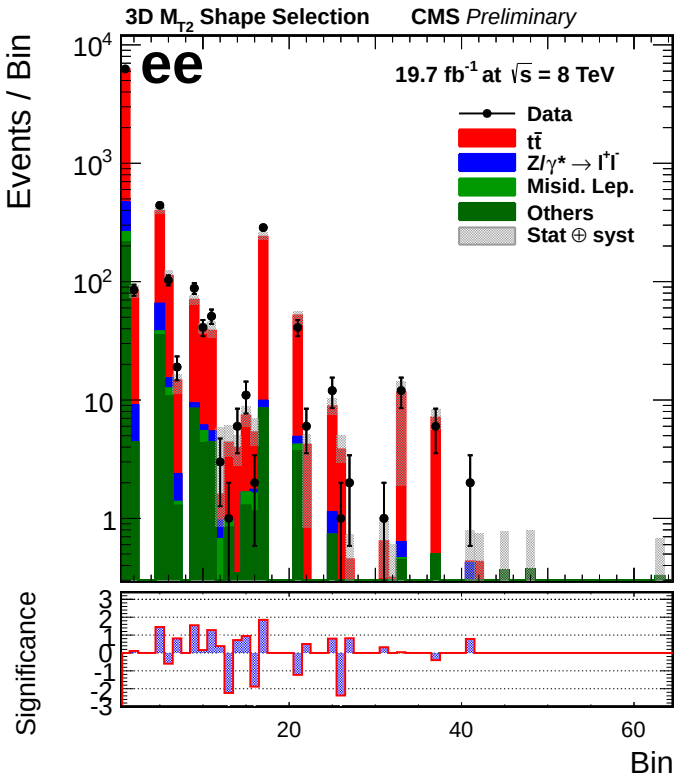

(b)

Figure 7.22: The 1D "de-stacking," Fig. 7.14, but shown for only e^+e^- events.

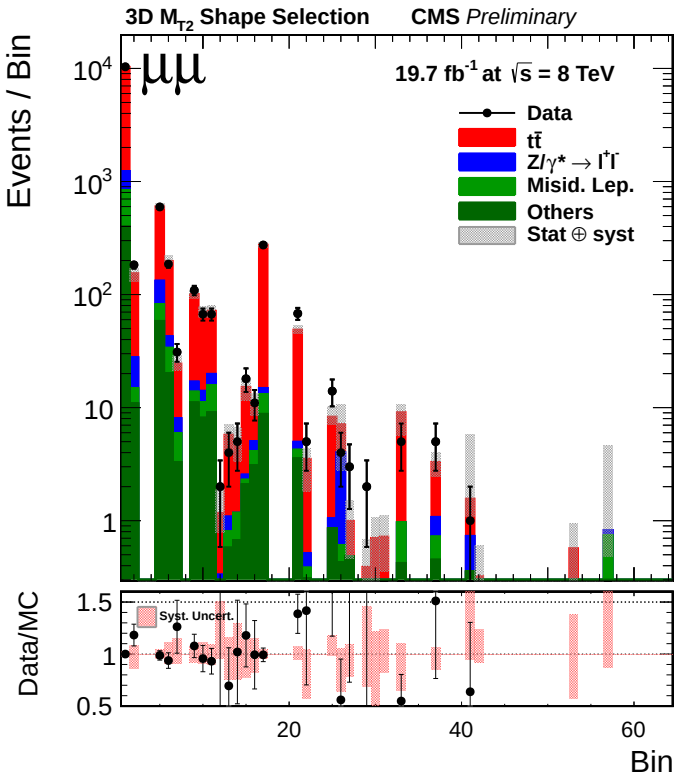

(a)

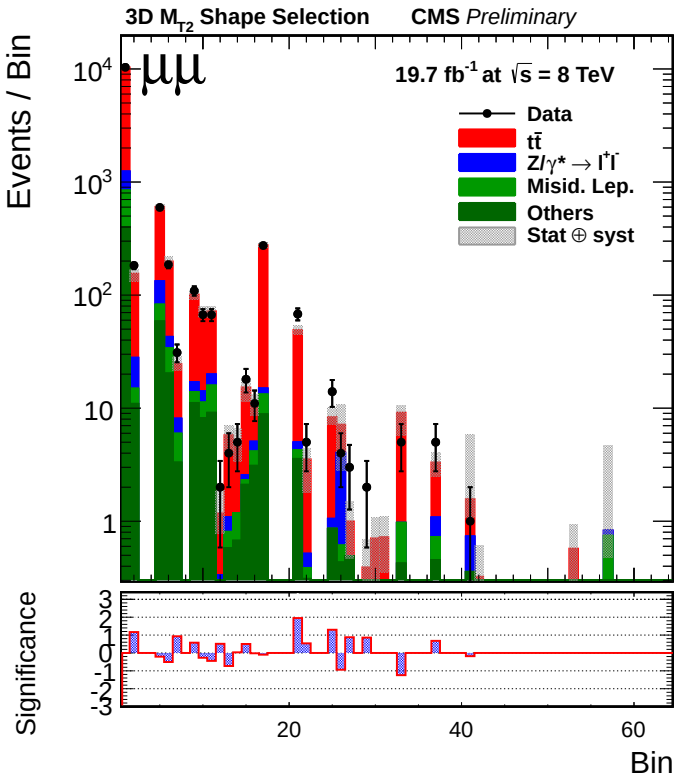

(b)

Figure 7.23: The 1D "de-stacking," Fig. 7.14, but shown for only $\mu^+\mu^-$ events.

7.2.7.1 Observed Significance of Excesses in the Data

Similar to Section 7.2.5.4, we calculated the local significance of observed excesses in the data for the 3D M_{T2} shape split by dilepton flavor. Figure 7.24 shows a map of these observed excess in the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane when testing them against the alternative hypothesis of top-squark production in the T2tt decay mode. Overlaid on top of these maps are contours denoting the regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane where, if top-squark pair-production was occurring with the nominal signal strength $\mu = 1$, it is expected that excesses at the 2σ (red) or 4σ (blue) level would be observed in the data. The solid line represents the median expected significance, while the dashed lines represent, when calculating the significances in a fully-frequentist fashion, the 0.16 and 0.84 percentiles on the expected significance[5].

The points in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane where the observed and expected significances have approximately equivalent values, e.g. the point $m_{\widetilde{\mathrm{t}}_1} = 200\,\mathrm{GeV}$, $m_{\widetilde{\chi}_1^0} = 50\,\mathrm{GeV}$ in Fig. 7.24c, are points for which the best-fit signal strength is close to 1.

The different sub-figures show the effects of varying the chirality of the coupling between the $\widetilde{\mathrm{t}}_1$ and t. All three sub-figures have a string of local significances, with values in the range, 3–4σ, for $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ points with relatively low $m_{\widetilde{\chi}_1^0}$ along the diagonal $\Delta\mathrm{M} = 150\,\mathrm{GeV}$. Moreover, observed local significances for the immediately neighboring $\Delta\mathrm{M}$ diagonals, particularly for low $m_{\widetilde{\chi}_1^0}$, tend to be at least 1σ lower.

This behavior of the observed excess implies that, using the 3D M_{T2} shape analysis, we can possibly achieve an experimental resolution on the $\Delta\mathrm{M}$ variable of

[5]For these particular significance maps, the significances were calculated using an asymptotic formula, and thus do not have associated statistical quantiles.

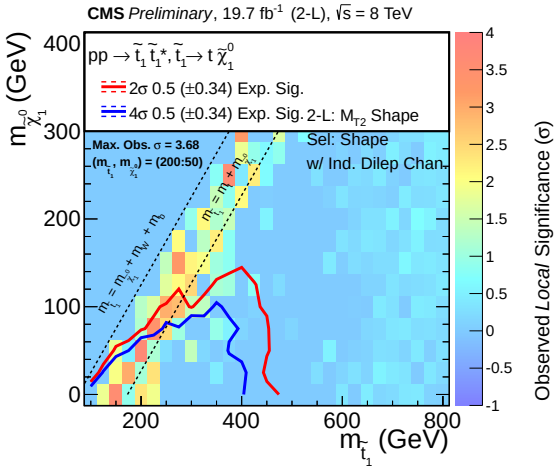

(a) The observed local significance for the unpolarized T2tt decay mode.

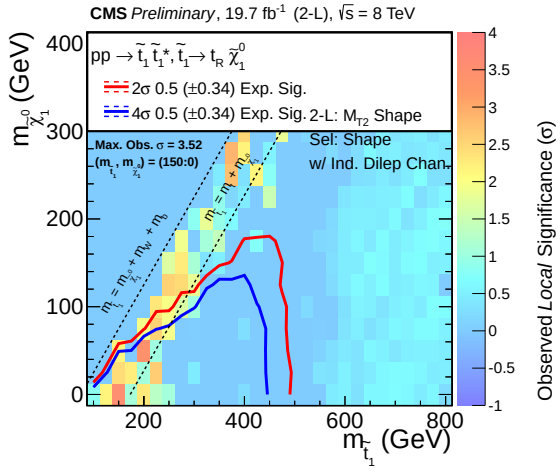

(b) The observed local significance for the right-polarized T2tt decay mode.

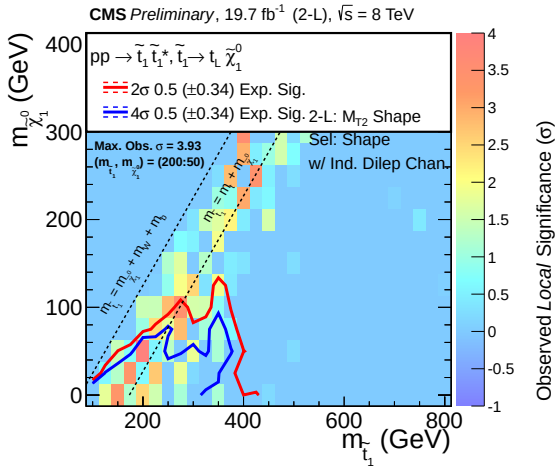

(c) The observed local significance for the left-polarized T2tt decay mode.

Figure 7.24: 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ maps of the *observed local* significance of excesses in the data, using the alternative hypothesis of top-squark pair-production in the different coupling-scenarios of the T2tt decay mode. Overlaid on top of theses maps are contours denoting the regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane where, if top-squark pair-production was occurring with the nominal signal strength $\mu = 1$, it is expected that excesses at the 2σ (red) or 4σ (blue) level would be observed in the data. The solid line represents the median expected significance, while the dashed lines represent the 0.16 and 0.84 percentiles on the expected significance.

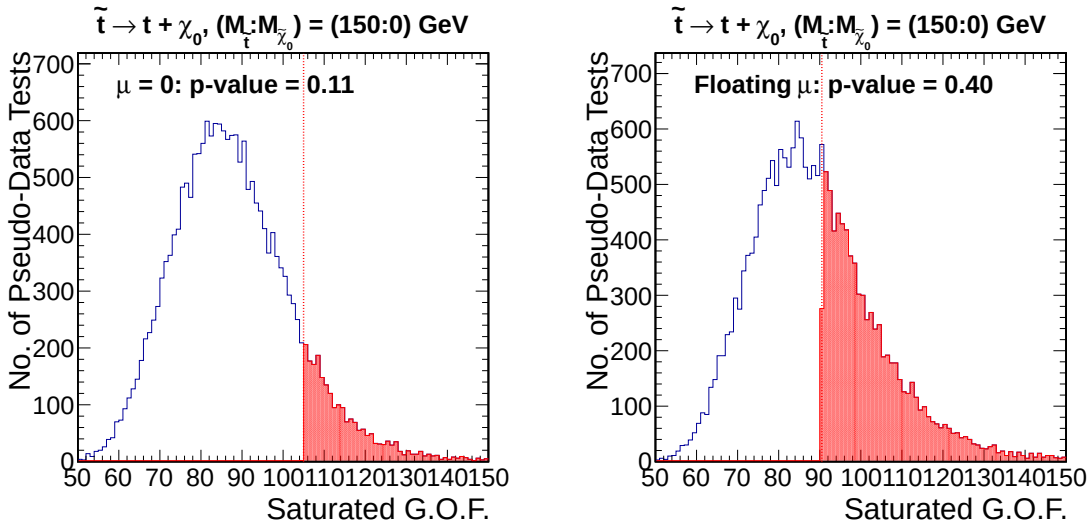

(a) The saturated GOF for the background-only hypothesis.

(b) The saturated GOF for the signal + background hypothesis.

Figure 7.25: A comparison of the goodness-of-fit (GOF) between the background-only and unconstrained signal + background hypotheses. These comparisons were made for the point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\ 0)$ GeV, in the unpolarized T2tt decay mode. The GOF metric used is the "saturated model" [169]. The distributions represent pseudo-data generated from the respective hypotheses. The integrated area represents the observed null p-value for the respective hypotheses.

$\sim$ 25–50 GeV. The experimental resolution on $m_{\tilde{t}_1}$, however, is notably worse. We discuss this further in the next section.

Validation Tests

We have already discussed a number of validations and cross-checks we performed for the 3D M_{T2} shape analysis, notably in the modeling of the additional M_{T2} variants. However, in light of the strength of these observed excesses, and the subsequent implications thereof, we performed a copious amount of additional cross-checks to confirm the robustness of this result. Details on some of these cross-checks can be found in Appendix C.3.

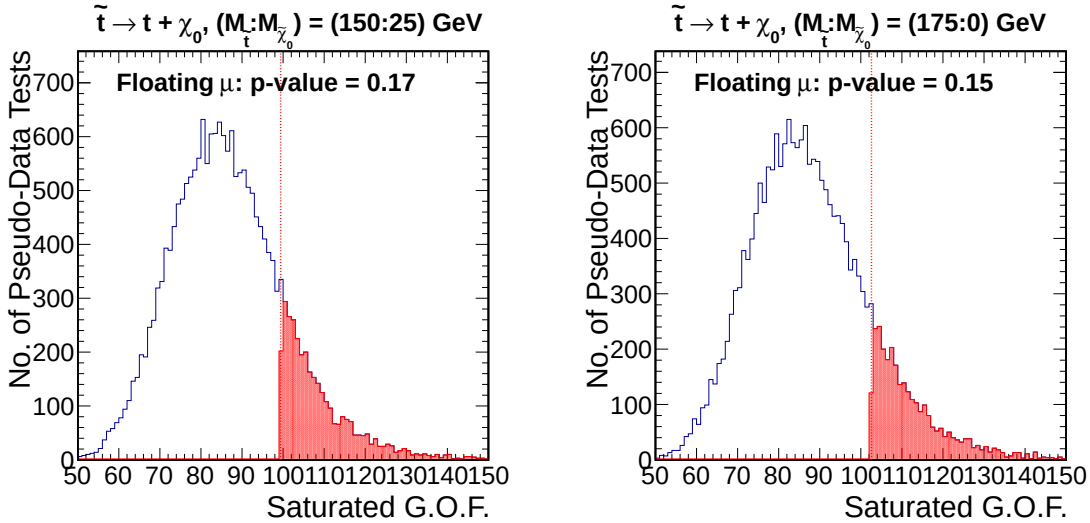

(a) The saturated GOF of the signal + background hypothesis for the point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\ 25)$ GeV.

(b) The saturated GOF of the signal + background hypothesis for the point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (175,\ 0)$ GeV.

Figure 7.26: The observed goodness-of-fit (GOF) for two unconstrained signal + background hypotheses, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\ 25)$ GeV and $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (175,\ 0)$ GeV. The GOF metric used is the "saturated model" [169]. The distributions represent pseudo-data generated from the respective hypotheses. The integrated area represents the observed null p-value for the respective hypotheses.

In this section, we discuss one of these cross-checks that can help develop the reader's intuition for understanding why the particular chain of signal points with $\Delta\mathrm{M} = 150\,\mathrm{GeV}$ yield such strong significances when compared against the null hypothesis.

First, we address why the chain of high significance $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ points are so localized in the $\Delta\mathrm{M}$ plane. The key aspect of these significance calculations is that they compare the relative likelihoods of the background-only hypothesis against the *unconstrained* signal + background hypothesis. This means that points in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane that have similar 3D M_{T2} shapes will yield similar observed local signifances, regardless of their respective theory-calculated top-squark pair-production cross sections, as their normalizations will be floated to best match the observed data.

In other parts of this dissertation, we have emphasized how the overall shapes of the M_{T2} variants depends strongly upon the $\Delta\mathrm{M}$ value of the top-squark signal point, particularly for the T2tt decay mode[6]. This translates into very similar 3D M_{T2} shapes for a given top-squark decay mode (particularly the T2tt decay mode) when $\Delta\mathrm{M}$ is constant, and thus, as per the discussion in the prior paragraph, similar observed significances for signal points with equal $\Delta\mathrm{M}$ values.

In order to better understand why *this* particular set of $\Delta\mathrm{M}$ values ($\Delta\mathrm{M} = 150\,\mathrm{GeV}$) yields such large observed local significances, we performed goodness-of-fit checks on the significance calculation.

[6]The $\Delta\mathrm{M}$ value is correlated with how much momentum the $\tilde{\chi}_1^0$ receives in the top-squark decay, particularly for the T2tt decay mode.

These goodness-of-fit checks are based on likelihood comparisons of the background-only (B, i.e. $\mu = 0$) and unconstrained signal + background ($\hat{\mu}S + B$) hypotheses against the "saturated model" [169], an alternative hypothesis where the observed data is treated as the exact, "true" theory. The variable $\bar{\chi}^2 = -2\ln\lambda$, where $\lambda = \frac{\mathcal{L}(\text{data}|\text{null})}{\mathcal{L}(\text{data}|\text{saturated})}$, follows (asymptotically) a χ^2 distribution.

We calculate a null p-value for the observed value of $\bar{\chi}^2$ by comparing this observed value against the expected distribution of $\bar{\chi}^2$, generated from an ensemble of pseudo-data. Figure 7.25 shows the results of this calculation for the point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\, 0)\,\text{GeV}$. One can see that the unconstrained signal + background hypothesis has a much stronger fit to the data relative to the background-only hypothesis. For comparison, Fig. 7.26 shows the results of this calculation for two other points, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\, 25)\,\text{GeV}$ and $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (175,\, 0)\,\text{GeV}$, each with ΔM values $\pm 25\,\text{GeV}$ relative to the point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\, 0)\,\text{GeV}$. As can be seen, these two points have notably worse fits to the observed excesses in the data compared with the unconstrained signal + background fit for the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150,\, 0)\,\text{GeV}$ point. Notably, the null p-values for the signal fit for these two points is only slightly better than the background-only null p-value.

Observed Significances in the T2bw Decay Mode

As one might expect, these observed excesses in the 3D M_{T2} shape also yield nontrivial observed local significances when tested against 3D M_{T2} shapes from the T2bw decay mode. Figure 7.27 shows the analogous maps for the *unpolarized* T2bw

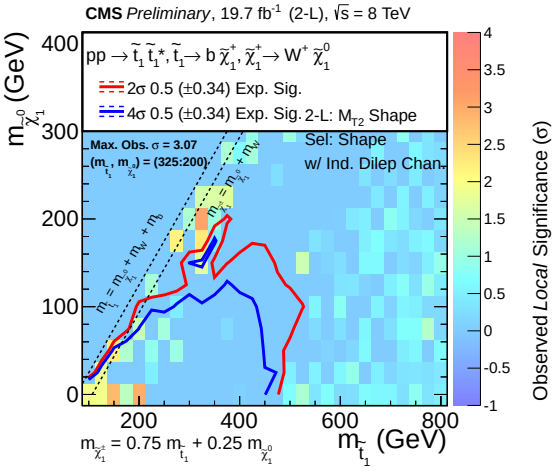

(a) The observed local significance for the unpolarized T2bw, $x = 0.75$ decay mode.

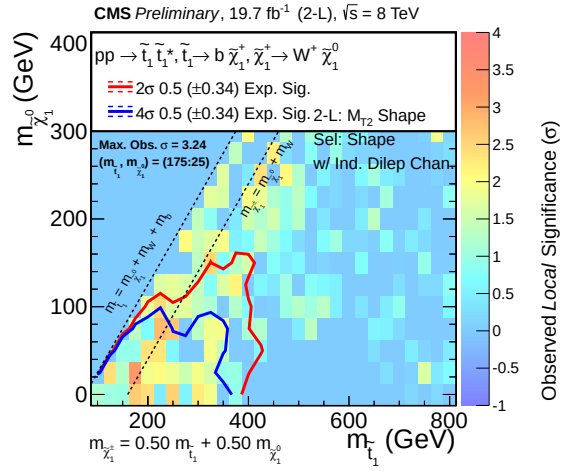

(b) The observed local significance for the unpolarized T2bw, $x = 0.50$ decay mode.

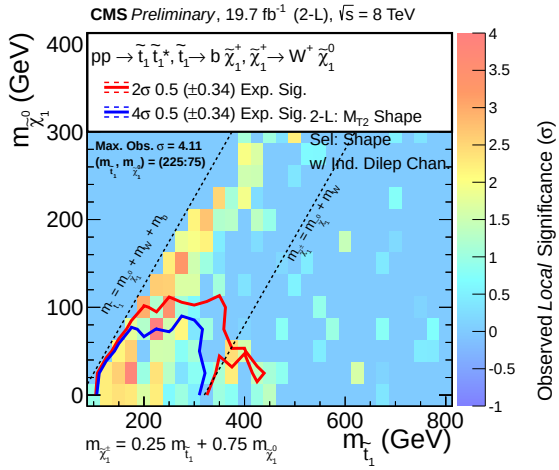

(c) The observed local significance for the unpolarized T2bw, $x = 0.25$ decay mode.

Figure 7.27: 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ maps of the *observed local* significance of excesses in the data, using the alternative hypothesis of top-squark pair-production in the unpolarized T2bw decay mode with different values of x. Overlaid on top of theses maps are contours denoting the regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane where, if top-squark pair-production was occurring with the nominal signal strength $\mu = 1$, it is expected that excesses at the 2σ (red) or 4σ (blue) level would be observed in the data. The solid line represents the median expected significance, while the dashed lines represent the 0.16 and 0.84 percentiles on the expected significance.

decay mode with the three considered values of x.

As can be seen, these excesses also fit with relatively strong local significances to mass points in the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with relatively low $m_{\widetilde{\chi}_1^\pm}$. As with the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode, the strongest fits are for relatively low-mass top-squarks. Relative to the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode, the local significances are $\sim 0.5\sigma$ to 1σ lower in magnitude and spread out over a wider range in the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane, implying a weaker experimental resolution on the ΔM parameter if the top-squark decays through this mode.

We note that in Fig. 7.27a, there is an odd, slight flange in the expected significance contours around $m_{\widetilde{t}_1} = 325\,\text{GeV}$, $m_{\widetilde{\chi}_1^0} = 150\,\text{GeV}$. Closer investigations found that, for the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ points near this region, the signal efficiency seems to be spuriously high. Given the large number of signal points considered, along with the fact that the signal efficiency for each signal point is estimated using independent signal simulation samples, we believe that these spuriously high signal efficiencies are just the result of upward statistical fluctuations in the signal efficiency.

Estimation of the "Look-elsewhere" Effect

Up to this point, we have only discussed the local significance of the observed excesses in the data. In order to quantify the *global* significance of the observed excesses in the data, we performed several calculations to estimate the magnitude of the "look-elsewhere" effect.

The "look-elsewhere" effect is designed to account for the possibility that the

3D M_{T2} shape for the background could coherently fluctuate in such a fashion that one of the various "independent" signal regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane fits the background fluctuation with a large local significance.

The first method we use to estimate the look-elsewhere effect is the following. We begin by generating $\mathcal{O}(10^5)$ pseudo-data 3D M_{T2} shapes using the background estimate. For each of these pseudo-data, background-only 3D M_{T2} shapes, we iterate through points in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, calculate the "observed" local significance at each mass point, and then record the maximum observed local significance across all of the scanned mass points. After performing these maximum observed local significance calculations for each pseudo-data shape, the distribution of these maximum significances is constructed and the *global* significance of the true observed local significance is calculated based upon the fraction of pseudo-data shapes where the maximum local significance from fits of top-squark shapes to the background-only shapes was stronger.

To calculate a preliminary, but nominally statistically conservative estimate of the look-elsewhere effect we scanned in 150 GeV steps in both $m_{\tilde{\mathrm{t}}_1}$ and $m_{\tilde{\chi}_1^0}$. Figure 7.28 shows the calculated distribution, for two of the T2tt decay mode scenarios, of the maximum observed local significance for the set of generated pseudo-data. The shaded red area on these distributions represents the fraction of the toy experiments where, for the scenarios in question, a signal region, other than the region with the largest observed local significance, fit the background-only pseudo data with larger local significance. These fractions directly translate into global p-values and observed significances.

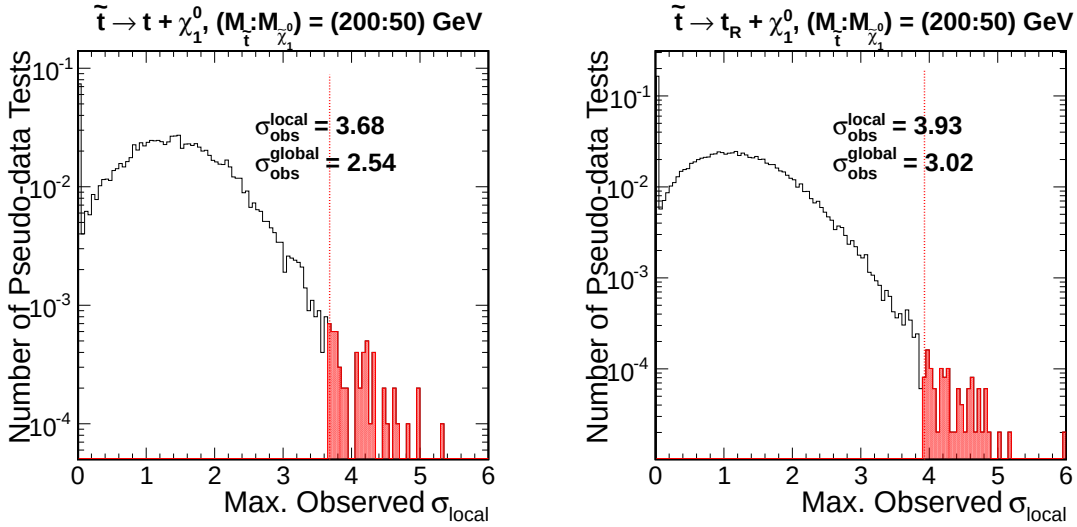

(a) (b)

Figure 7.28: Normalized distributions of the maximum observed *local* significance from a set of scanned points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, where the distribution was constructed using pseudo-data constructed from the background-only hypothesis of the 3D M_{T2} shape analysis, split by dilepton channel. Overlaid on each distribution is a shaded red area representing the global p-value of the background-only hypothesis for the strongest observed local significances, both at the point $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (200{:}50)$ GeV for two T2tt decay mode scenarios: (left) $(3.68 \to 2.54)\sigma$ from the $\tilde{t}_1 \to t\tilde{\chi}_1^0$ decay mode, Fig. 7.24a; (right) $(3.93 \to 3.02)\sigma$ for the $\tilde{t}_1 \to t_L\tilde{\chi}_1^0$ decay mode, Fig. 7.24c.

For the mass point, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (200, 50)\,\text{GeV}$, in the T2tt decay mode, we observe a global significance of 2.54σ (3.02σ) for the unpolarized (left-handed) scenario. The pseudo-data distributions of the maximum observed local significances look fairly similar to one another for these two T2tt scenarios. Thus, we estimate a general "trials" factor, the ratio of the local null p-value divided by the global null p-value for both scenarios, to be $\mathcal{O}\,(20)$.

The Euler Characteristic and Global Significance

An alternative method for estimating the global significance of the observed excesses, discussed in Section G.3.1 or Ref. [170], utilizes the Euler characteristic.

The first steps of this method are very similar to the previously mentioned method. Namely, we begin by generating a number (500 in this case) of pseudo-data 3D M_{T2} shapes using the background estimate. These pseudo-data 3D M_{T2} shapes are used to calculate "observed" local significance for various signal points. Instead of only iterating through a subset of the points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, however, we calculate the observed local significance for each pseudo-data shape for each point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane. Thus, for each pseudo-data shape we can make a complete map of the observed local significance in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane similar to Fig. 7.24c or any one of the other previously mentioned local significance maps.

For each of these local significance maps, we calculate $\phi(A_\sigma)$, the observed Euler characteristic when thresholding the map at various values of the significance (e.g. 0.5σ, 1σ, etc.). We then calculate the average value of the Euler characteristic

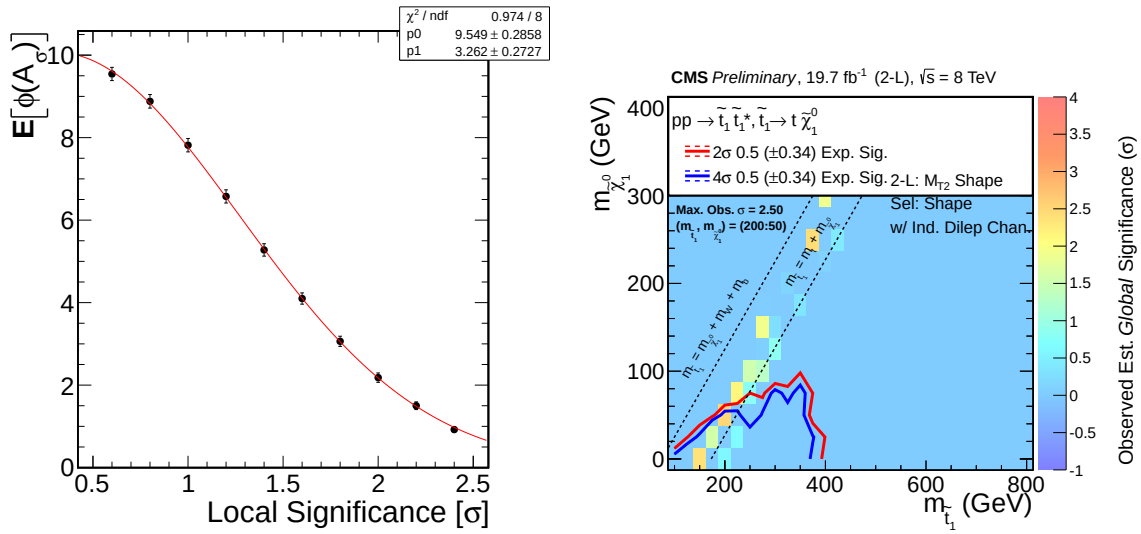

(a) The dependence of the average value of the Euler Characteristic on the local significance threshold.

(b) The observed *global* significance for the unpolarized T2tt decay mode, estimated using the fit shown in Fig. 7.29a.

Figure 7.29: Results from a method for calculating the *global* significance using the average value of the Euler characteristic for background-only pseudo-data.

for each threshold value, $E\left[\phi(A_\sigma)\right]$, averaging across all of the background-only pseudo-data significance maps.

The resulting set of values is fit with the function,

$$E\left[\phi(A_\sigma)\right] = P\left[\sigma\right] + e^{-\sigma^2/2}\left(c_0 + c_1\sigma\right), \tag{7.8}$$

where $P\left[\sigma\right]$ is the null p-value associated with a σ-strength deviation and the c_i are the fit coefficients.

For large significance values, the average value of the Euler characteristic asymptotically equals the global null p-value for the background-only model in question, $E\left[\phi(A_\sigma)\right] \approx p^{\text{global}}$, and thus, the function in Eq. (7.8) provides a way to map large observed *local* significances onto their corresponding *global* significance values.

Figure 7.29a shows the results of the fit of Eq. (7.8) to the background-only pseudo-data shapes. The resulting fit coefficients are, $p_0 = 9.55 \pm 0.29$, $p_1 =$

3.26 ± 0.27. We then applied Eq. (7.8) to our observed local significance maps to calculate observed global significance maps. We assigned a conservative uncertainty of $\mathcal{O}(10\%)$ on our calculated global significances, designed to cover the observed statistical fluctuations in the fit as well as any residual systematic biases in this estimation method[7]. For points where the estimated global null p-value was calculated to be greater than 1, i.e. points with low-values of the observed local significance, we assigned the global significance to be 0.

Figure 7.29b shows an example global significance map for the unpolarized T2tt decay mode. As expected, the overall significance for every point is lower, with most mapped to 0 as per the discussion above. The chain of strong observed significances for points with $\Delta\mathrm{M} = 150\,\mathrm{GeV}$ persists, and the strongest observed local significance point from Fig. 7.24a, which had a value of 3.68σ, maps to $(2.50 \pm 0.25)\,\sigma$. As another example, the strongest observed local significance from Fig. 7.24c, which had a value of 3.93σ, maps to $(2.81 \pm 0.28)\,\sigma$. It is interesting to note that we find consistent results between the two estimation methods for the calculated global significances, i.e. comparing the above numbers against Fig. 7.28.

[7] For example, the fact that the average Euler characteristic only asymptotically maps to the global null p-value.

7.3 Conclusions and Summary of the Extensions to the Dileptonic Top-squark Search

In this chapter, we have presented a number of extensions to the cut-and-count dileptonic top-squark search. The first was a statistical combination of the cut-and-count dilepton search with a top-squark search looking in the semi-leptonic final state. A detailed summary of this extension is provided in Section 7.1.3.3.

The second extension was a powerful, novel approach to the dileptonic top-squark search. In this approach, three M_{T2} variants are constructed in each event, utilizng the two b quarks, two leptons, and $\not{E}_{\mathrm{T}}$ as input variables. These M_{T2} variants are correlated with one another but provide complementary information. A 3D M_{T2} shape is constructed where each axis is one of the unique M_{T2} variants. In this 3D M_{T2} parameter space, the expected background tends to be relatively localized in a single control region. This facilitates a shape-based comparison between the expected background and the nominal top-squark signal, as the top-squark signal populates different tail regions depending upon the top-squark decay mode and the masses of the relevant SUSY particles.

Compared with the cut-and-count dilepton analysis, the 3D M_{T2} shape analysis is much more statistically powerful. Notably, the 3D M_{T2} shape analysis is sensitive to the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode where the $\widetilde{\chi}_1^\pm$ is relatively light. To compare, the cut-and-count dilepton analysis has little to no sensitivity to this mode.

A remarkable result is the observed exclusion of top-squark pair-production

in the $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{t_R}\widetilde{\chi}_1^0$ decay mode for the region $\Delta\mathrm{M} = m_\mathrm{t}$, $m_{\widetilde{\chi}_1^0} < 50\,\mathrm{GeV}$. This result represents the only direct search for top-squark pair-production, across all top-squark pair final states, from the first run of the LHC, that is sensitive to top-squark pair-production in this region.

Another notable aspect of the 3D M_{T2} shape analysis is the presence, in the observed 3D M_{T2} shape, of a number of excesses relative to the background expectation. Models of top-squark pair-production are fit across the entire 3D M_{T2} shape and the strength of these fits are translated into observed local significances in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ 2D plane.

The strongest fits are for the $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{t}\widetilde{\chi}_1^0$ decay mode, where observed local significances of ~ 3.5–4σ are found for a number of mass points with relativealy low $m_{\widetilde{\chi}_1^0}$ and with a relative $\widetilde{\mathrm{t}}_1$, $\widetilde{\chi}_1^0$ mass splitting $\Delta\mathrm{M} = 150\,\mathrm{GeV}$. These observed significances are quite localized in $\Delta\mathrm{M}$, indicating a possible experimental resolution of $\sim 25\,\mathrm{GeV}$ on this parameter. Two separate methods are used to calculate preliminary, conservative estimates of the global significance for these mass points. Both methods yield results that are consistent with one another: The estimated global significance for these mass points is ~ 2.5–3σ, which corresponds to a trials factor of $\mathcal{O}\,(20)$.

These excesses also fit with relatively strong local significances to mass points in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane for the $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{b}\widetilde{\chi}_1^+$ decay mode, notably for scenarios with relatively low $m_{\widetilde{\chi}_1^\pm}$. As with the $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{t}\widetilde{\chi}_1^0$ decay mode, the strongest fits are for relatively low-mass top-squarks, but, relative to the $\widetilde{\mathrm{t}}_1 \rightarrow \mathrm{t}\widetilde{\chi}_1^0$ decay mode, the local significances are ~ 0.5–1σ lower in magnitude and spread out over a wider

range in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, implying a weaker experimental resolution on the ΔM parameter if the top-squark decays through this mode.

In light of these observed excesses, a number of additional cross-checks, were performed to further validate the results of the 3D M_{T2} shape analysis. Currently, none of these cross-checks found any cause for concern about the robustness of the result.

Chapter 8: Summary, Conclusions and Outlook

This dissertation has presented several analyses searching for signatures of pair-production of top-squarks at the LHC.

The first of these analyses was in the dileptonic final state, consisting of two b quarks, two leptons, and missing transverse energy. This analysis used 19.66 fb^{-1} of proton-proton collision data collected with the CMS experiment during the 2012 run of the LHC. The analysis applied a cut-based approach, relying on the stransverse mass variable, $M_{\mathrm{T2}}(\ell\ell)$, to separate the top-squark signal from the large SM $\mathrm{t\bar{t}}$ background. The estimations of background contributions and systematic uncertainties were handled using robust, statistically conservative techniques. No significant excesses above the nominal background expectation were observed. Depending upon the chirality of the coupling between the top-squark and its daughter particles, as well as the specific decay mode considered, this analysis excludes top-squark pair-production for $m_{\tilde{\mathrm{t}}_1}$ in the range, 150 GeV to 435 GeV, for $m_{\tilde{\chi}_1^0}$ up to 160 GeV.

The second analysis presented in this dissertation was a statistical combination of the cut-based dilepton analysis with another statistically orthogonal analysis looking for top-squark pair-production in the semi-leptonic final state, consisting of

two b quarks, one lepton, at least two additional hadronic jets and missing transverse energy. The semi-leptonic top-squark search utilized a multivariate approach based on boosted decision trees to separate the top-squark signal from the $\mathrm{t\bar{t}}$ background. As with the cut-based dileptonic search, it also did not observe any significant excesses above the nominal background expectations.

For the combination of the two aforementioned analyses, correlations between the two analyses were investigated thoroughly and handled in a robust fashion. The two analyses were combined into a two-bin counting experiment. The statistical power of this combination is stronger than either analysis alone. Notably, relative to the semi-leptonic analysis, the lepton combination achieves increased sensitivity to the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode close to a theoretically interesting, but experimentally difficult region, $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0} \approx m_{\mathrm{t}}$. This observed sensitivity is at least comparable, if not better than other published direct searches for top-squark pair-production in this region of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane [30, 31].

The third, and final analysis presented in this dissertation was an extension of the cut-based dilepton search. Two additional M_{T2} variants are constructed that utilize information from reconstructed b quarks. Although these two M_{T2} variants are both correlated with one another and with $M_{\mathrm{T2}}(\ell\ell)$, they provide complementary information to aid in distinguishing between the top-squark signal and the $\mathrm{t\bar{t}}$ background. The information from the three M_{T2} variants is combined into a binned 3D distribution, where each axis is one of the variants. Shape-based comparisons are performed between the observed data and the expected background and signal shapes. The shapes of the background and signal are allowed to shift within

their systematic uncertainties to best fit the observed data. Extensive studies were performed to validate and quantify the underlying systematic uncertainties on the shape of this distribution in both the background and expected signal.

Compared with the cut-based dilepton analysis, the 3D M_{T2} shape analysis is much more statistically powerful. Notably, the 3D M_{T2} shape analysis is sensitive to the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode where the $\widetilde{\chi}_1^\pm$ is relatively light; the cut-and-count dilepton analysis has little to no sensitivity to this mode.

A remarkable result of this 3D M_{T2} shape analysis is the observed exclusion of top-squark pair-production in the $\widetilde{\mathrm{t}}_1 \to \mathrm{t_R}\widetilde{\chi}_1^0$ decay mode for the region $\Delta\mathrm{M} = m_{\mathrm{t}}$, $m_{\widetilde{\chi}_1^0} < 50\,\mathrm{GeV}$. This result represents the only current direct search for top-squark pair-production, across all top-squark pair final states, from the first run of the LHC, that is sensitive to top-squark pair-production in this region.

Another notable aspect of the 3D M_{T2} shape analysis is the observation of a number of excesses, relative to the background expectation, in the observed 3D M_{T2} shape. Models of top-squark pair-production are fit to these excesses and the strengths of these fits are translated into observed local significances in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ 2D plane. In the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode, observed local significances of ~ 3.5–4σ are found for a number of mass points with relatively low $m_{\widetilde{\chi}_1^0}$ and with a mass-splitting $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0} = 150\,\mathrm{GeV}$. These observed significances are quite localized in $\Delta\mathrm{M}$, indicating a nominal experimental resolution of $\sim 25\,\mathrm{GeV}$ on this parameter.

Two separate methods were used to calculate preliminary, conservative estimates of the global significance for these mass points. Both methods yield results

that are consistent with one another: The estimated global significance for these mass points is ~ 2.5–3σ, which corresponds to a trials factor of $\mathcal{O}(20)$.

These excesses also fit with relatively strong local significances to mass points in the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode, for scenarios with relatively low $m_{\widetilde{\chi}_1^\pm}$. As with the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode, the strongest fits are for relatively low-mass top-squarks, $m_{\widetilde{t}_1} \lesssim 300\,\text{GeV}$. Relative to the $\widetilde{t}_1 \to t\widetilde{\chi}_1^0$ decay mode, the local significances are ~ 0.5–1σ lower in magnitude and spread out over a wider range in the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane, implying a weaker experimental resolution on the ΔM parameter if the top-squark decays through this mode.

In light of these observed excesses, a number of additional cross-checks, beyond the initial set, were performed to further validate the results of the 3D M_{T2} analysis. Currently, none of these cross-checks found any cause for concern about the robustness of the result.

8.1 The Dileptonic Top-squark Search: Looking Forward to 13 TeV

The 8 TeV dileptonic top-squark search presented in this dissertation achieved some notable results. It is a natural question to ask what the prospects are for the top-squark search in this final state moving forward.

Run 2 of the LHC is currently ongoing, with proton-proton collisions occurring at a center-of-mass energy, $\sqrt{s} = 13\,\text{TeV}$. This is a particularly exciting time in the search for high mass SUSY partners, as the increased $\sqrt{s}$ leads to larger production cross sections. As an example, Fig. 8.1 shows the predicted

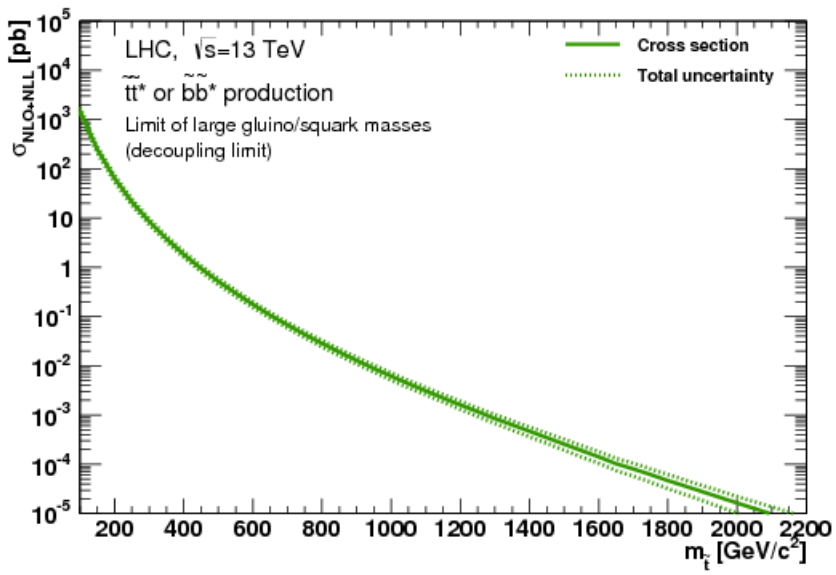

Figure 8.1: The NLO cross-section for direct top-squark pair-production at the LHC at $\sqrt{s} = 13$ TeV. The green lines represent the total cross section and its total uncertainty. Reprinted from Fig. 13 of [171].

top-squark pair-production cross section in proton-proton collisions at the LHC at $\sqrt{s} = 13\,\text{TeV}$ [171].

Compared with the 8 TeV cross section, Fig. 6.3, it is clear that the expected sensitivity to high-mass top-squarks is significantly larger for Run 2 of the LHC; for example, the production cross section for $m_{\tilde{t}_1} = 1\,\text{TeV}$ jumps from 4.35×10^{-4} pb to 6.15×10^{-3} pb, an order-of-magnitude increase. This will be particularly relevant for top-squark searches in the hadronic or semi-leptonic final states, where the relative branching ratios are quite large.

On the other hand, for relatively low values of $m_{\tilde{t}_1}$, the increase in cross section is not nearly as large. For example, the relative increase in production cross section for $m_{\tilde{t}_1} = 150\,\text{GeV}$ is only a factor of 3. This is comparable to the expected increase in the production cross section for SM $t\bar{t}$ pair when moving from 8 TeV to 13 TeV, so

the relative signal to noise will change very little in searches for low-mass top-squarks due to the increased $\sqrt{s}$.

In addition, over the course of Run 2 of the LHC, the instantaneous luminosity of the LHC is expected to increase notably. This will be correlated with an increase in the average number of pileup interactions per bunch crossing. Pileup interactions adversely affect the reconstruction of $\not{E}_{\mathrm{T}}$, see Fig. 5.12, and so, barring novel approaches to $\not{E}_{\mathrm{T}}$ reconstruction, it will become increasingly difficult to distinguish genuine-$\not{E}_{\mathrm{T}}$ signals from spurious-$\not{E}_{\mathrm{T}}$ backgrounds. CMS has been actively working on new $\not{E}_{\mathrm{T}}$ algorithms to mitigate the adverse effects of pileup on $\not{E}_{\mathrm{T}}$ reconstruction, and current results look promising in that regard, see Fig. 5.13 or Ref. [71].

From a statistical precision standpoint, the dileptonic top-squark search is a rather challenging one, as the low branching ratio (4%) to the dilepton final state means that large datasets are needed to ensure large statistical sensitivity to top-squark production. In order to better confirm the possible low-mass top-squark observed in the 3D M_{T2} analysis, it is probably most prudent, then, that future analysts work on controlling the systematic uncertainties associated with the dileptonic top-squark search, rather than wait for additional LHC data.

As an example, the systematic uncertainties on the fake lepton estimate can reach relative magnitudes as large as 80%. Granted, there are statistical components to these large systematic uncertainties, but additional investigations into new data-driven techniques for background estimation could certainly help the precision of low-mass dileptonic top-squark searches.

Appendix A: Summary Information from the Analyses

In this appendix, we provide summary tables and figures for information from Chapters 6 and 7, including things such as selection requirements and signal efficiencies

A.1 Tables of Analysis Object Selection Requirements

Table A.1 contains the list of requirements for our reconstructed muon candidates.

Table A.2 shows the selection requirements applied to the electron candidates in our event. Table A.3 displays the details of the working points for CMS's official electron selection requirements. For the dileptonic analyses described in this dis-

Table A.1: Muon object selection.

Quality Variable	Cut	Cut Value
Muon Type	=	Global, prompt, tight muon
p_{T}	>	20 (10) GeV harder (softer)
$\lvert p_{\mathrm{T}}^{\mathrm{PF}} - p_{\mathrm{T}}^{\mathrm{basic}} \rvert$	<	5 GeV
$\lvert\eta\rvert$	<	2.4
d_0	<	0.2 cm
d_z	<	0.5 cm
$\mathrm{Iso}_\mu/p_{\mathrm{T}}$	<	0.15
Global track fit χ^2/n_{dof}	<	10
n_{hits}(pixel)	>	0
n_{stations}(muon)	>	1
n_{layers}(tracker)	>	5

Table A.2: Electron object selection.

Quality Variable	Cut	Cut Value
p_{T}	$>$	20 (10) GeV harder (softer)
$\lvert p_{\mathrm{T}}^{\mathrm{PF}} - p_{\mathrm{T}}^{\mathrm{basic}} \rvert$	$<$	10 GeV
$\lvert\eta\rvert$	$<$	2.5
		veto EB/EE overlap region ($1.442 < \lvert\eta\rvert < 1.566$)
$\mathrm{Iso_e}/p_{\mathrm{T}}$	$<$	0.15
conversion veto		applied
ID working point		VBTF WP80 (Medium WP in Table A.3)

Table A.3: The requirements placed on electron candidate quality variables that define the working points for CMS's official electron selection requirements.

Quality Variable	Cut	Cut Value		
		Loose	Medium	Tight
$\lvert\Delta\eta\,(\mathrm{SC,\ trk})\rvert$	$<$	0.007 barrel 0.009 endcap	0.004 barrel 0.007 endcap	0.004 barrel 0.005 endcap
$\lvert\Delta\phi\,(\mathrm{SC,\ trk})\rvert$ [radians]	$<$	0.015 barrel 0.010 endcap	0.06 barrel 0.03 endcap	0.03 barrel 0.02 endcap
$\sigma_{i\eta\, i\eta}$	$<$	0.01 barrel 0.03 endcap		
H/E	$<$	0.12 barrel 0.10 endcap		
$\lvert d_0\rvert$ [mm]	$<$	0.02		
$\lvert d_z\rvert$ [mm]	$<$	0.2	0.1	0.1
$\lvert E^{-1} - p^{-1}\rvert[\,\mathrm{GeV}^{-1}]$	$<$	0.05		
n_{miss}	$\leq$	1	1	0
$\mathrm{Iso_{ECAL}}/p_{\mathrm{T}}$	$<$	0.2	-	-
$\mathrm{Iso_{HCAL}}/p_{\mathrm{T}}$	$<$	0.2	-	-
$\mathrm{Iso_{Trk.}}/p_{\mathrm{T}}$	$<$	0.2	-	-

Table A.4: Jet object selection.

Quality Variable	Cut	Cut Value
clustering algorithm		anti-k_{T} ($R = 0.5$)
p_{T}	$>$	30 GeV
$\lvert\eta\rvert$	$<$	2.4
b-tag		CSV medium
additional requirements		loose PF jet ID (Table A.5)

sertation, the "medium" working point is used for the main electron selection. The "loose" working point is used when selecting the "loose" electrons that are utilized in the estimate of the misidentified lepton contribution, Section 6.3.3. For the "loose" working point, the three additional requirements on detector isolation ensure that the offline selection is tighter than the online electron trigger requirement. Specifically, the detector isolation for the electron candidate as calculated by the ECAL, HCAL, and tracker must each independently be less than 20% of the electron's p_{T}.

Table A.5: The quality cuts for the loose working point of the jet identification.

Quality Variable	Cut	Cut Value
f_{CH}	$>$	0.0
f_{NH}	$<$	0.99
f_{γ}	$<$	0.99
f_{EM}	$<$	0.99
n_{charged}	$>$	0
$n_{\mathrm{constituents}}$	$>$	1

Table A.4 shows the hadronic jet selection. Table A.5 shows the selection variables with associated requirements for the loose working point of the PF Jet ID.

The variables used in this ID include the fraction of the jet energy contained in various classes of PF candidates (neutral hadrons, f_{NH}; neutral EM particles, f_{γ}; charged hadrons, f_{CH}; and charged EM particles, f_{EM}); the number of jet constituents, $n_{\mathrm{constituents}}$; and the charged particle multiplicity in the jet, n_{charged}.

A.2 Additional Plots for the Dilepton Shape Analysis

In this appendix, we include additional plots related to the Dilepton Shape analysis.

A.2.1 Correlations between the M_{T2} Variants

In Section 7.2.3, we showed how the correlations of $M_{\mathrm{T2}}(\ell\ell)$ and $M_{\mathrm{T2}}^{\not{W}}(\mathrm{bb})$ differ between the SM background and the top-squark signal. Figures A.1 and A.2 provide the comparisons between the expected background and the top-squark signal for the other two permutations of the 3 M_{T2} variants. Similar to Fig. 7.10, for these two 2D correlations, the background predominantly populates low values for the M_{T2} variants, while the distribution of events in the top-squark signal favors large values for at least one of the considered M_{T2} variants.

Visualizing the M_{T2} correlations in this fashion ignores the possible additional relative correlations that can come about when all three M_{T2} variants are considered together. Although it is somewhat difficult to visualize using only 2D distributions, Fig. A.3 provides some measure for this 3D correlation by showing the dependence of the distribution of $\kappa_{\mathrm{T2}} = M_{\mathrm{T2}}^{\not{W}}(\mathrm{bb}) - M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ on the value of $M_{\mathrm{T2}}(\ell\ell)$. The sub-figures show this 2D correlation for the expected background and the $\Delta\mathrm{M} = (300 \pm 15)\,\mathrm{GeV}$ T2tt and $x = 0.25$ T2bw top-squark decay modes.

As might be expected, the relative correlations between the three M_{T2} variants depends upon the mass-splitting $\Delta\mathrm{M}$. Figure A.4 shows the correlation of $M_{\mathrm{T2}}(\ell\ell)$ with κ_{T2} for the $x = 0.50$ T2bw decay mode for four separate values of the mass-

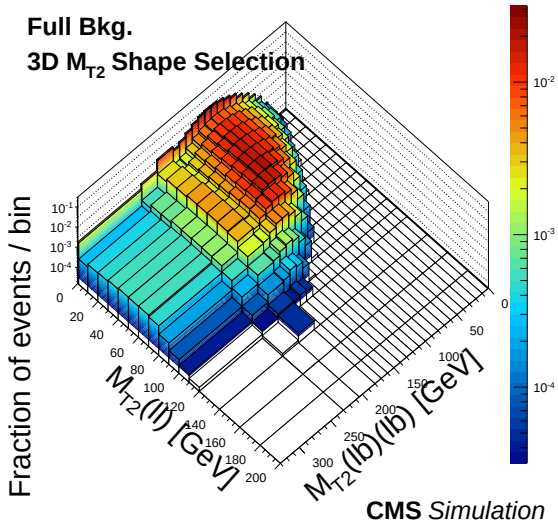

(a) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for the full data-driven background estimate.

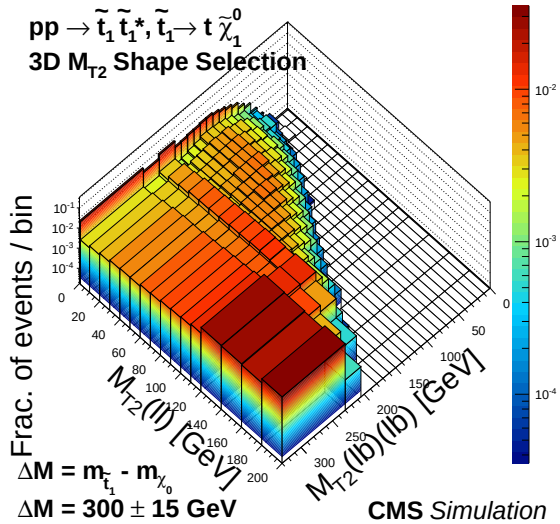

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for top-squark events from the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0$ decay mode.

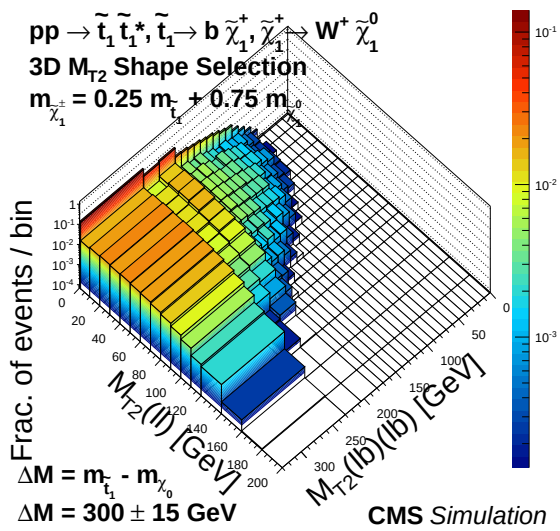

(c) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for top-squark events from the $x = 0.25$ $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode.

Figure A.1: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the full data-driven background estimate and two top-squark decay modes with $\Delta\mathrm{M} = (300 \pm 15)\,\mathrm{GeV}$. Note the small slivers located near $M_{\mathrm{T2}}(\ell\ell) \approx 0\,\mathrm{GeV}$ are each individual bins.

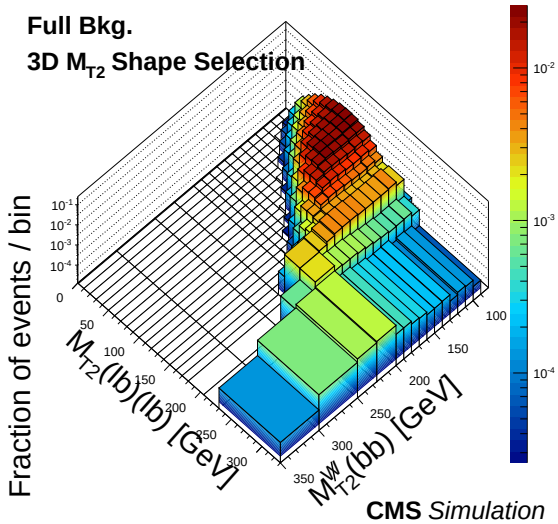

(a) The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for the full data-driven background estimate.

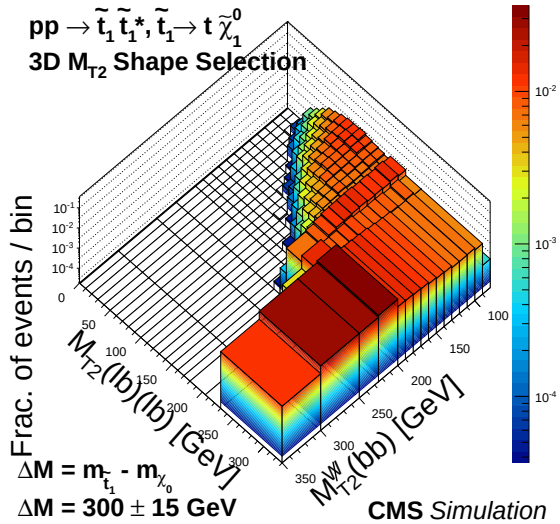

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0$ decay mode.

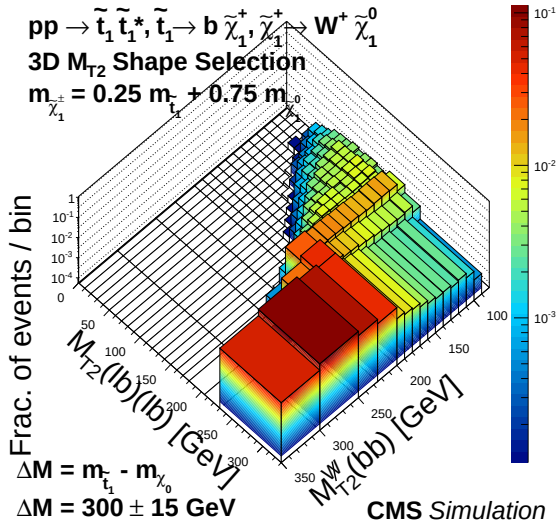

(c) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ for top-squark events from the $x = 0.25$ $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode.

Figure A.2: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ vs. $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$ in the full data-driven background estimate and two top-squark decay modes with $\Delta\mathrm{M} = (300 \pm 15)$ GeV.

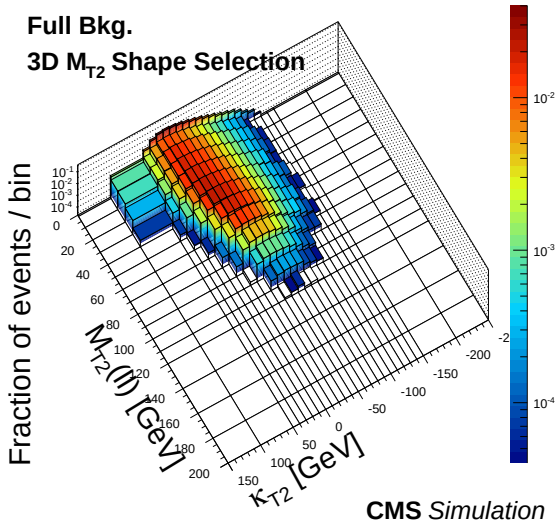

(a) The distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for the full data-driven background estimate.

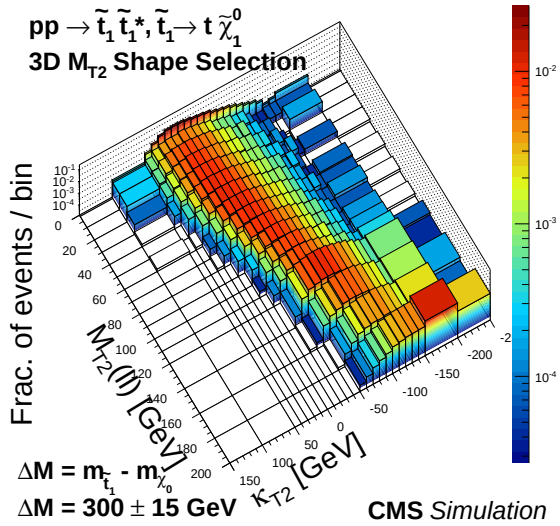

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events from the $\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0$ decay mode.

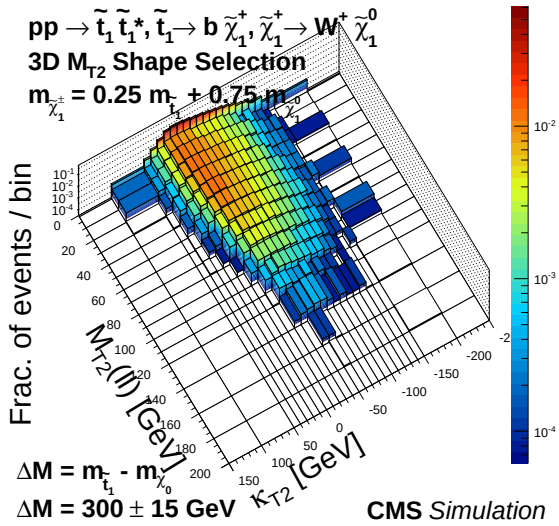

(c) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events from the $x = 0.25$ $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode.

Figure A.3: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $\kappa_{\mathrm{T2}} = M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) - M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the full data-driven background estimate and two top-squark decay modes with $\Delta\mathrm{M} = (300 \pm 15)\,\mathrm{GeV}$. Note the small slivers located near $M_{\mathrm{T2}}(\ell\ell) \approx 0\,\mathrm{GeV}$ are each individual bins.

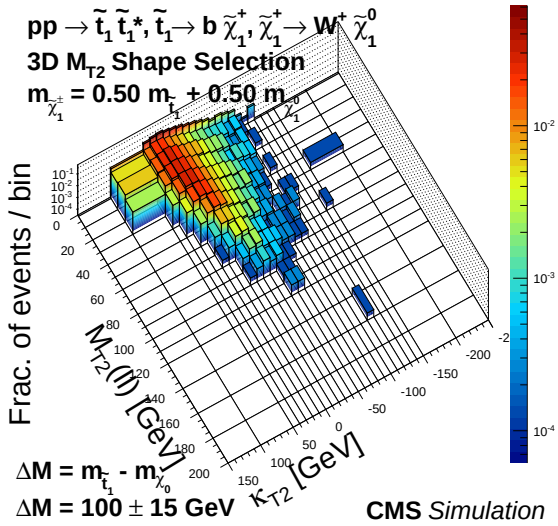

(a) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events with $\Delta\mathrm{M} = 100\,\mathrm{GeV}$

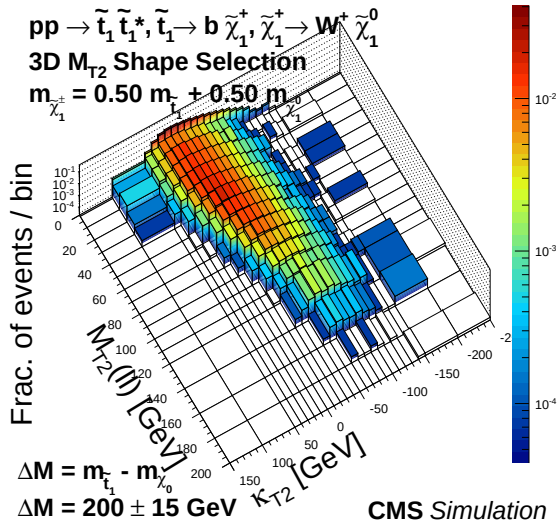

(b) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events with $\Delta\mathrm{M} = 200\,\mathrm{GeV}$

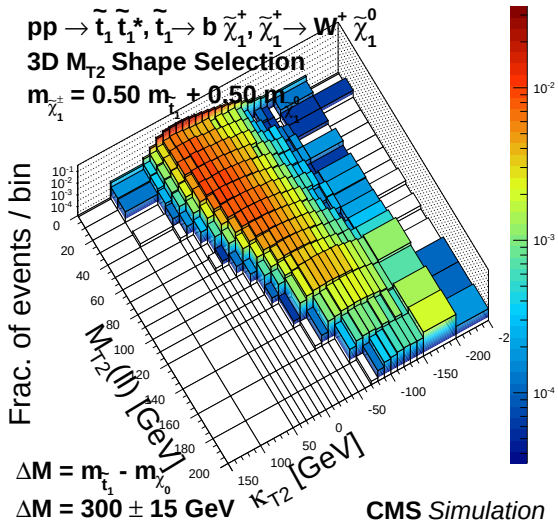

(c) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events with $\Delta\mathrm{M} = 300\,\mathrm{GeV}$

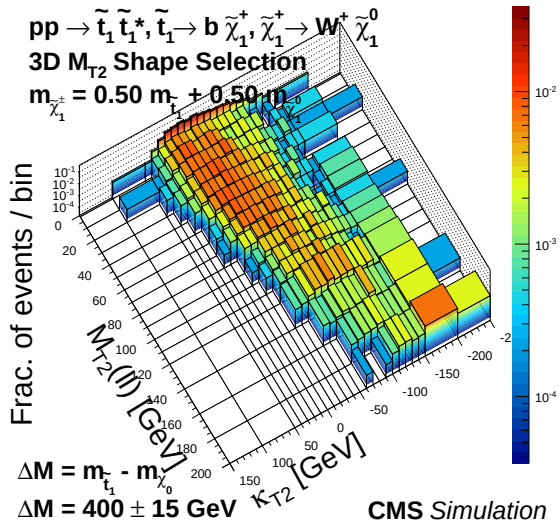

(d) The normalized fractional distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. κ_{T2} for top-squark events with $\Delta\mathrm{M} = 400\,\mathrm{GeV}$

Figure A.4: Comparisons of the 2D distribution of $M_{\mathrm{T2}}(\ell\ell)$ vs. $\kappa_{\mathrm{T2}} = M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) - M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for the $x = 0.50\,\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ top-squark decay mode with varying values of $\Delta\mathrm{M}$. Note the small slivers located near $M_{\mathrm{T2}}(\ell\ell) \approx 0\,\mathrm{GeV}$ are each individual bins.

splitting ΔM.

Appendix B: Additional Analysis Results

In this appendix, we provide additional results that we elected not to show in the main body of this dissertation.

B.1 Sensitivity of the 3D M_{T2} Shape Analysis to Asymmetric Decays of Top-squarks

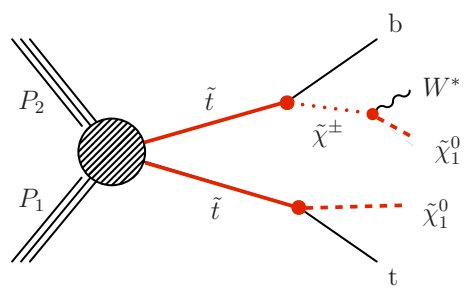

Figure B.1: Feynman diagram representation of the asymmetric decay mode of top-squark pair-production at the LHC. Propagators and vertices for supersymmetric particles are colored in red.

Up to this point, when discussing the decay modes of top-squark pairs, we have only considered symmetric decay modes, i.e. where each $\widetilde{\mathrm{t}}_1$ undergoes the same decays to SUSY daughter particles. However, depending upon the masses of the intermediate particles, asymmetric decay modes, à la the one represented in Fig. B.1, should nominally be feasible.

In this section we provide details on investigations into the sensitivity of the

dilepton analysis to the asymmetric decay mode of top-squark pairs. This decay mode is hereafter referred to, for brevity's sake, by the acronym T2tb.

B.1.1 The T2tb Simulation Samples

As with the T2tt and T2bw simulation samples described in Section 6.2.1, the MadGraph generator was utilized to generate the simulation samples that were used to estimate the signal efficiency of the T2tb decay mode. For these samples, the scan in the 2D $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane did not extend above $m_{\widetilde{\mathrm{t}}_1} = 1$ TeV and did not extend below $\Delta\mathrm{M} < 100$ GeV. For each point of the scan, approximately 125,000 events were generated and then run through the CMS FastSim.

For each individual top-squark, the branching ratios for both allowed decays, $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ and $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$, were both set to 50%. When either top-squark underwent the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay, the mass of the intermediate chargino was set to be slightly higher than the daughter neutralino, $m_{\widetilde{\chi}_1^\pm} = m_{\widetilde{\chi}_1^0} + 5\,\mathrm{GeV}$.

B.1.2 Upper-Limits on Top-squark Pair-production in the T2tb Decay Mode

Figure B.2 shows the observed mass exclusions for the unpolarized T2tb decay mode. As can be seen, the 3D M_{T2} analysis has no observed nor expected sensitivity to this decay mode.

This is notably in contrast to the the T2tt and T2bw decay modes, where at least some sensitivity is observed for all decay modes considered. The simple

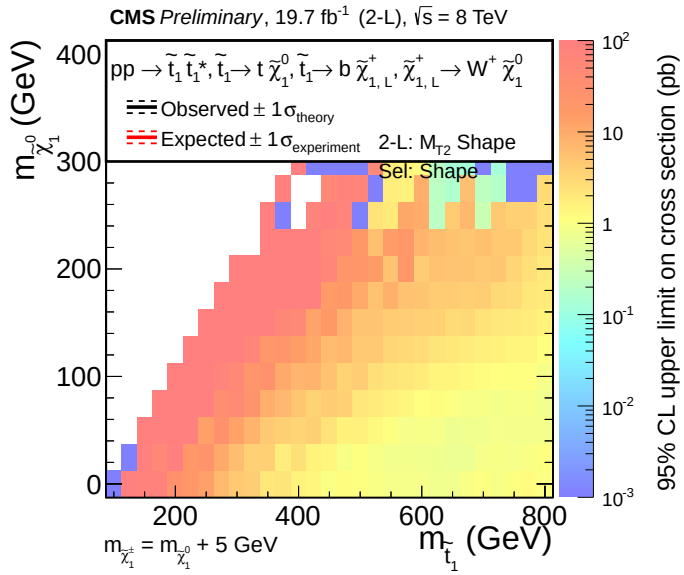

Figure B.2: 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the unpolarized T2tb decay mode. Due the small assumed mass-splitting between the $\tilde{\chi}_1^{\pm}$ and the $\tilde{\chi}_1^0$, the 3D M_{T2} shape analysis has no sensitivity to the top-squark signal.

explanation for this is the fact that the $\tilde{\chi}_1^{\pm}$ is assumed to be nearly degenerate with the $\tilde{\chi}_1^0$, i.e. $m_{\tilde{\chi}_1^{\pm}} = m_{\tilde{\chi}_1^0} + 5$ GeV. This means that the W produced in the decay of the $\tilde{\chi}_1^{\pm}$ is extremely off-shell. This adversely affects both the efficiency with which the W's daughter lepton passes the object selection as well as the efficiency with which the T2tb events fall into the region, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV} \cup M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) > 170\,\mathrm{GeV} \cup M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) > 170\,\mathrm{GeV}$.

B.2 Sensitivity to Top-squark Pair-production when Requiring 2+ b Jets

As shown in Section 7.2.2.1, the relatively significant tail in the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ is primarily driven by events where, for the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation, an incorrect

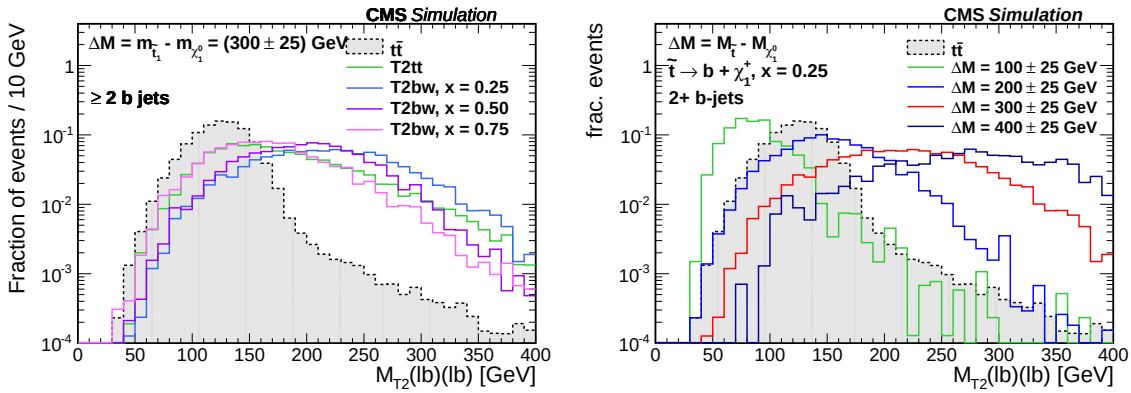

(a) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distributions for several different top-squark decay modes for a $\Delta\mathrm{M}$ value of (300 ± 25) GeV.

(b) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution for several different $\Delta\mathrm{M}$ regions for the $x = 0.25$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode.

Figure B.3: The distributions, using the preselection but also requiring 2+ reconstructed b-jets, of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in the $\mathrm{t\bar{t}}$ background and the top-squark signal; distributions are normalized to the same area.

pairing was made between the b-quarks and leptons.

The contribution of these kind of events can be greatly reduced by requiring $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$.

Another intuitive way to reduce the contribution of this subset of events is to require two or more reconstructed b jets in the event, as this will reduce the contribution from events where either one or both of the input "b-jets" for the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ calculation is not one of the b jets produced from the $\mathrm{t\bar{t}}$ decay.

Figure B.3 shows the same distributions as Fig. 7.6, except that each event is required to have at least 2 reconstructed b-jets.

As can clearly be seen, compared to Fig. 7.6, the high-$M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ tail in the $\mathrm{t\bar{t}}$ events has been greatly reduced. It is also worth noting that there is a much more pronounced kinematic edge around $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) = m_{\mathrm{t}}$.

As an added benefit, this requirement naturally reduces the tail of the $M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb})$

distribution.

Figure B.4 shows the analogous exclusion plots as Fig. 7.15, except that each event is required to have two or more reconstructed b jets; note that this new b jet requirement is in addition to the $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$ requirement.

Directly comparing Figs. B.4 and 7.15, one can see that the additional requirement of two or more reconstructed b jets improves the sensitivity to the $x = 0.25$ T2bw decay mode by ~ 25–$50\,\mathrm{GeV}$, but all other top-squark decay modes lose sensitivity when this additional requirement is placed. Closer investigations found that, for these top-squark decay modes, the dual requirement of $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$ and two or more b jets was excessive and hurt the signal efficiency for these decay modes.

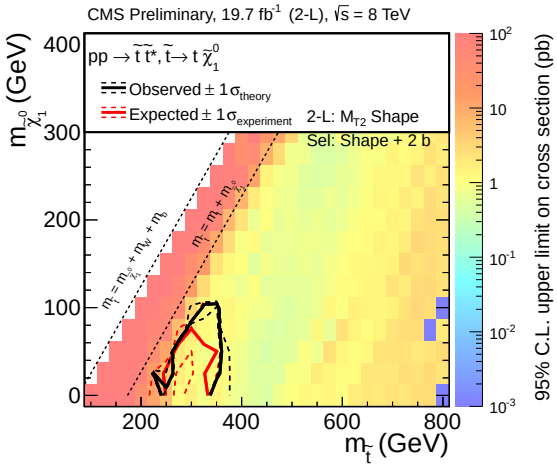

(a) The limits for the unpolarized T2tt decay mode.

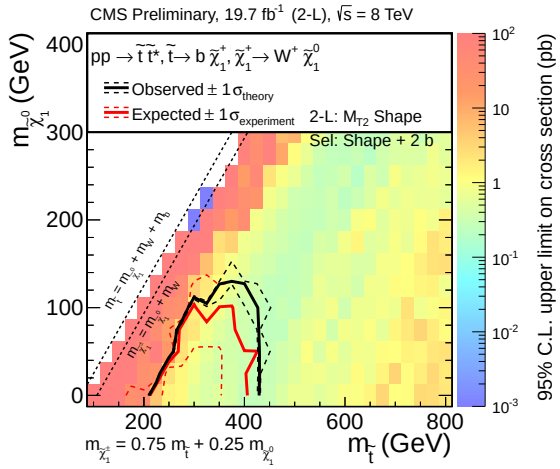

(b) The limits for the unpolarized T2bw decay mode with $x = 0.75$.

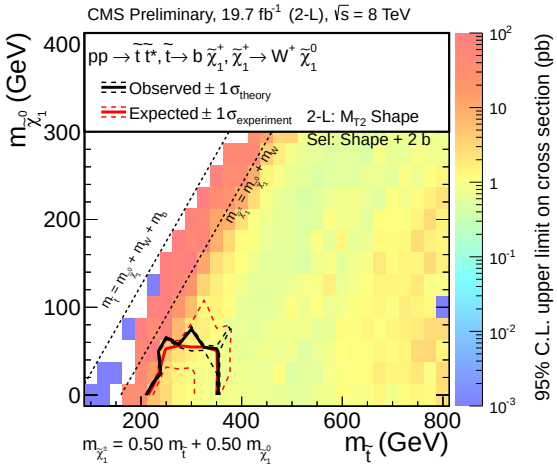

(c) The limits for the unpolarized T2bw decay mode with $x = 0.50$.

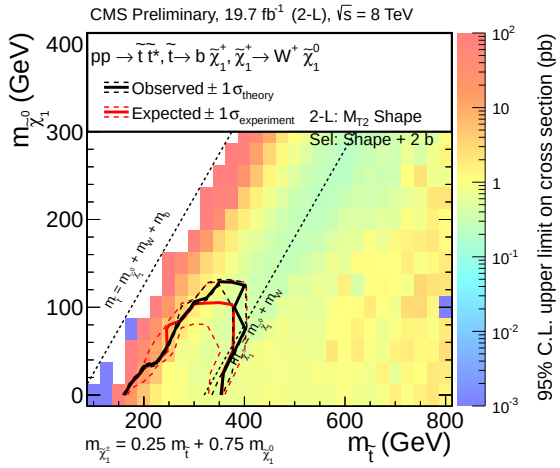

(d) The limits for the unpolarized T2bw decay mode with $x = 0.25$.

Figure B.4: 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the 3D M_{T2} shape analysis, with the additional requirement of two or more reconstructed b jets. Overlaid on top of these maps are contours denoting the regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The two solid contour lines were made using the median expected cross-section limit (red) and the observed cross-section limit (black). The dashed red lines represent the exclusion region when using the 0.16 and 0.84 percentile expected limits, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

Appendix C: Analysis Appendix: Additional Cross-checks

In this section, we provide details on additional cross-checks that we performed to validate the results of the analyses.

C.1 Cross-checks on Modeling of Object Multiplicities and Event Kinematics

This section presents a number of checks on the modeling of object multiplicities and kinematic variables in our simulation samples used to estimate backgrounds for the dilepton analyses described in this dissertation. Overall, when comparing the simulation with the data, we find good to excellent agreement between the two.

Figure C.1 shows the modeling, at the preselection level, of the leading and sub-leading jet p_{T}, the leading b jet p_{T}, and the scalar sum H_{T} of all reconstructed jet p_{T}. All of the jet p_{T} distributions are well-modeled by the simulation. This serves to validate the modeling of individual jet kinematics. The H_{T} distribution is also fairly well-modeled. Although there is a slight systematic trend in the ratio of data and simulation, this is covered within the existing systematic uncertainties.

We did not copy it again here, but an additional distribution that's worth noting is the number of reconstructed jets. This is shown in Fig. C.4d. There is

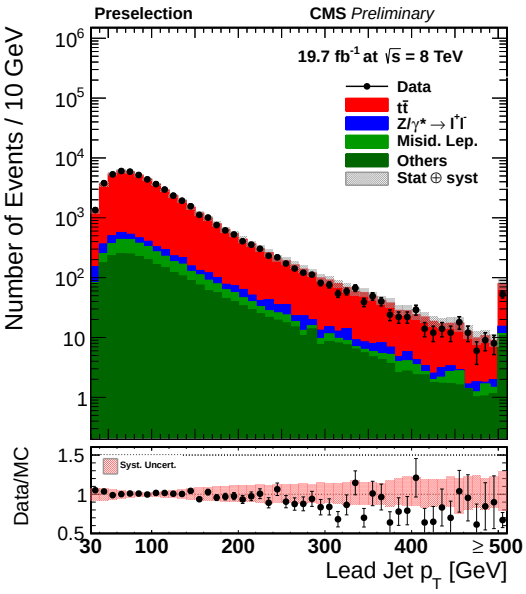

(a) The leading jet p_{T}.

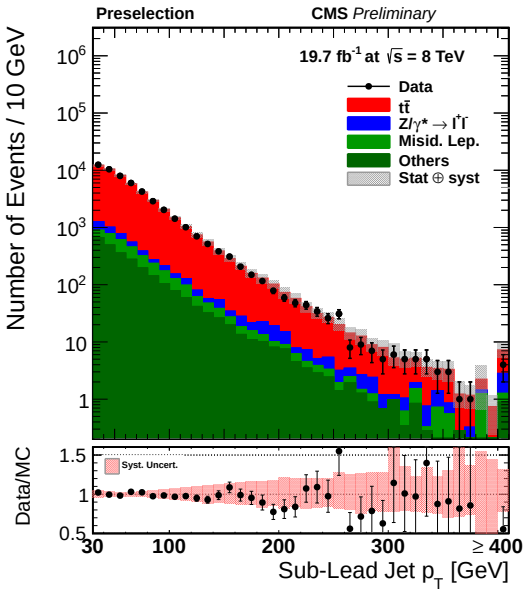

(b) The sub-leading jet p_{T}.

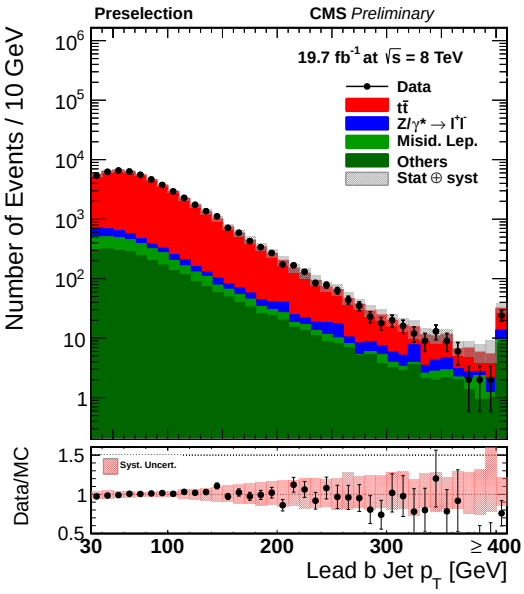

(c) The leading b jet p_{T}.

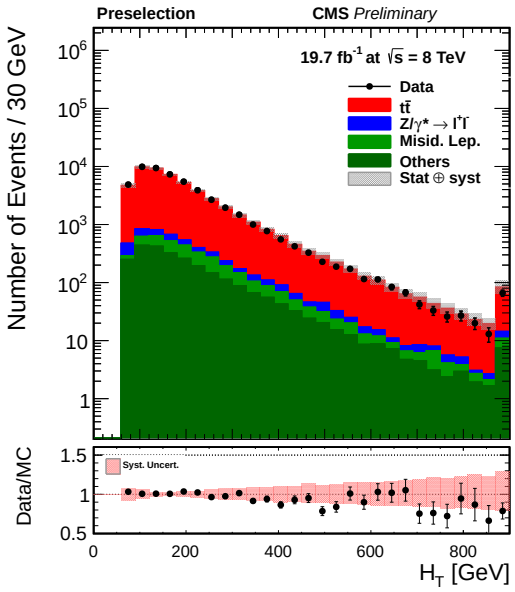

(d) The scalar sum H_{T} of all reconstructed jet p_{T}.

Figure C.1: The distributions, in the preselection event sample, of the leading and sub-leading jet p_{T}, the leading b jet p_{T}, and the scalar sum H_{T} of all reconstructed jet p_{T}.

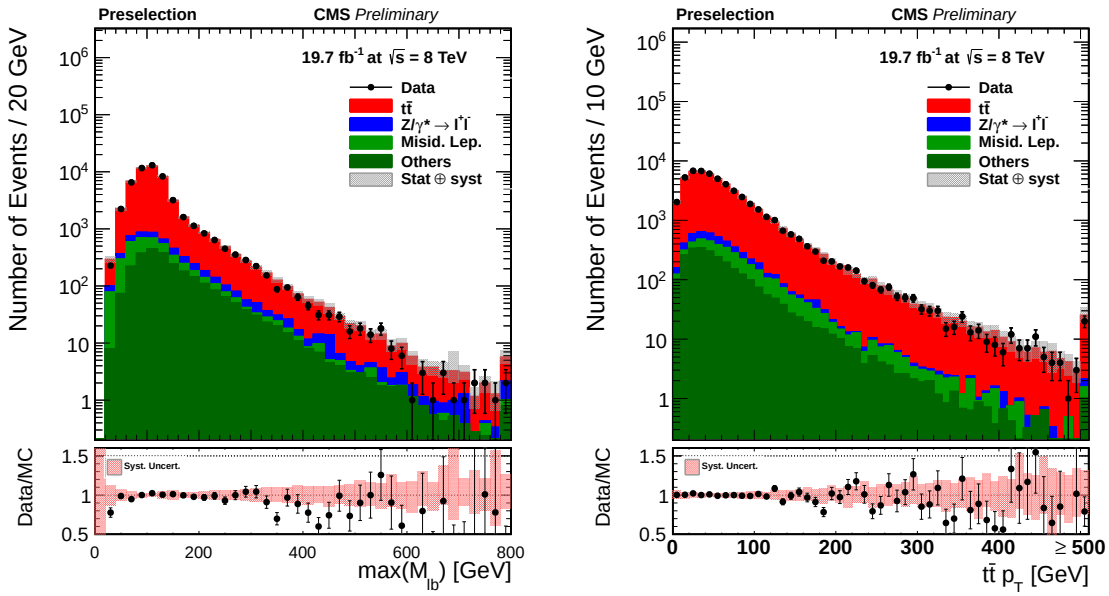

(a) The max($m_{\ell b}$) of the two ℓb systems.

(b) The p_T of the $t\bar{t}$ system.

Figure C.2: The distributions of $t\bar{t}$ related quantities in the preselection event sample.

perhaps a slight systematic trend in the modeling of the number of reconstructed jets, which serves as a metric for how well energetic ISR and FSR are modeled within the simulation. It should be noted that this trend is covered at the $\sim 1\sigma$ level of the systematic uncertainties. In particular, for this distribution, the two leading uncertainties are the jet-energy scale calibration and the generator-level $t\bar{t}$ p_T reweighting. Related to this, the number of reconstructed b jets is shown in Fig. C.5d. There is perhaps a slight systematic trend in the modeling of the number of reconstructed b jets, but this trend is covered at the $\sim 1\sigma$ level of the systematic uncertainties. For this distribution, the uncertainty on the b-tagging efficiency scale factors represents one of the more notable systematics.

Figure C.2 shows two additional validations of the modeling of the $t\bar{t}$ system in the simulation. Figure C.2a shows the modeling of the largest invariant mass of

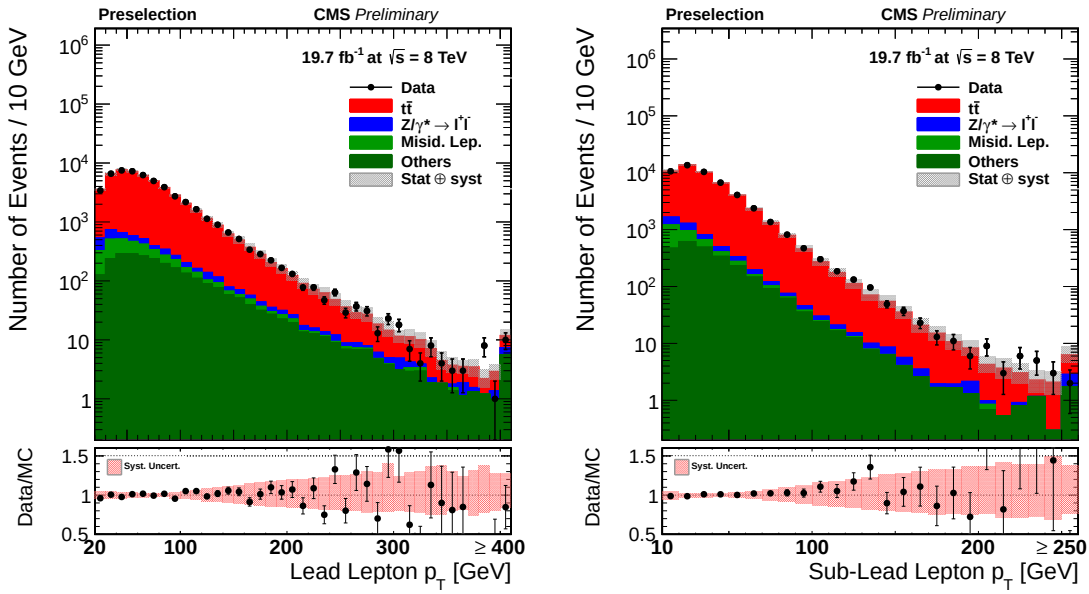

(a) The leading lepton p_{T}.

(b) The sub-leading lepton p_{T}.

Figure C.3: The distributions of the leading and sub-leading lepton p_{T} in the preselection event sample.

the two ℓb systems in each $\mathrm{t\bar{t}}$ event. The discrepancies at low $\max(m_{\ell\mathrm{b}})$ are due to an overestimate of the misidentified lepton contribution in this region. Figure C.2b shows the modeling of the $\mathrm{t\bar{t}}$ system p_{T} (2 b jets' $\vec{p}_{\mathrm{T}}$ + 2 leptons' $\vec{p}_{\mathrm{T}}$ + $\vec{\not{E}}_{\mathrm{T}}$). These two distributions are key checks for the modeling by the simulation of relevant distribution shapes for $\mathrm{t\bar{t}}$ events. These validations are especially crucial for the 3D M_{T2} shape analysis described in Section 7.2.

Figure C.3, which shows the p_{T} distributions for the two leading leptons, validates the modeling of individual lepton kinematics. The relative angular correlations between the leptons, which is shown in Fig. C.7d, is also well modeled within our $\mathrm{t\bar{t}}$ simulation. These three variables, along with the $\not{E}_{\mathrm{T}}$, Fig. 6.4, are key variables for the construction of $M_{\mathrm{T2}}(\ell\ell)$. The fact that all of them are well-modeled within the simulation is important for the robustness of the dilepton analyses described in

this dissertation.

C.2 Cross-check of the Choice of $t\bar{t}$ Generator

In the dileptonic top-squark search described in Chapters 6 and 7, we utilized the POWHEG generator to provide the event samples used to estimate the contribution of $t\bar{t}$ events to our M_{T2} signal regions.

This choice was driven by the results of cross-checks where we found that the other generators we considered (MADGRAPH and MC@NLO) failed to adequately model relevant kinematic variables. In the remainder this section, we provide details on these cross-checks of this choice of MC generator.

It is important to note that for each generated sample of $t\bar{t}$ events, we separately calculated the normalization scale factors utilizing the method described in Section 6.3.1 of Chapter 6. We found that characteristically, relative to the others, the POWHEG generator required larger normalization scale factors.

Using these individual scale factors for each generated $t\bar{t}$ sample, we then checked the modeling of relevant event variables. For the most part, the modeling of relevant event variables was acceptable for all four generated samples; however, for a few of the event variables, most notably $M_{T2}(\ell\ell)$, we found that the modeling of the event variable shapes in one or more of the non-POWHEG samples was quite poor.

Figure C.4 shows the distributions for the number of reconstructed jets in our sample of events passing the preselection. Figure C.5 shows the analogous

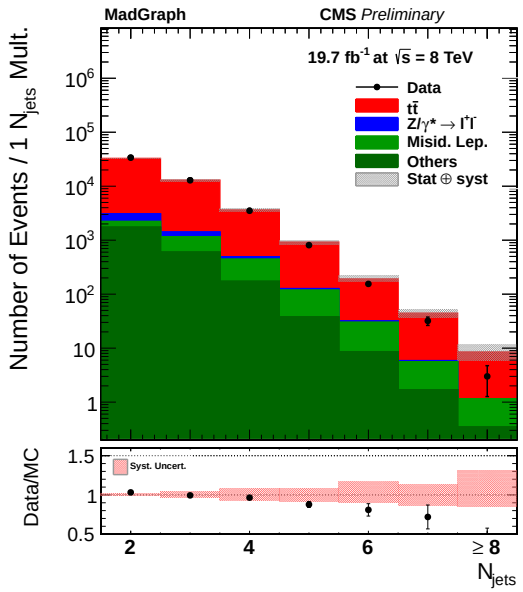

(a) The distribution using the final-state inclusive MadGraph $t\bar{t}$ sample.

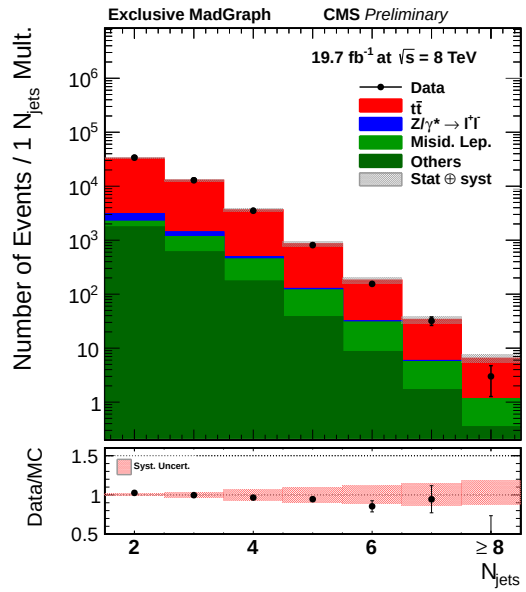

(b) The distribution using the final-state exclusive MadGraph $t\bar{t}$ sample.

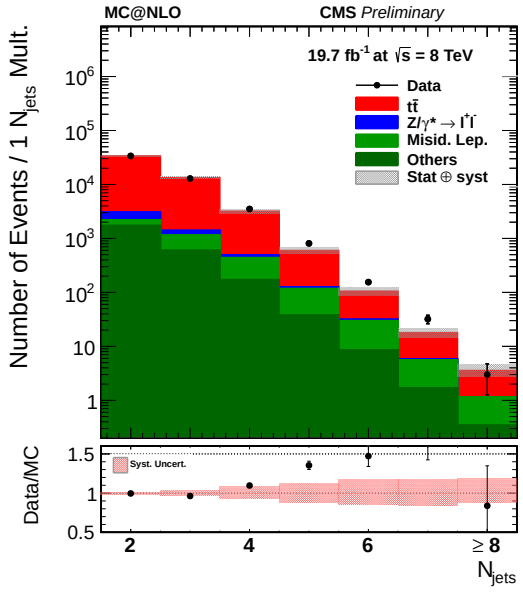

(c) The distribution using the MC@NLO $t\bar{t}$ sample.

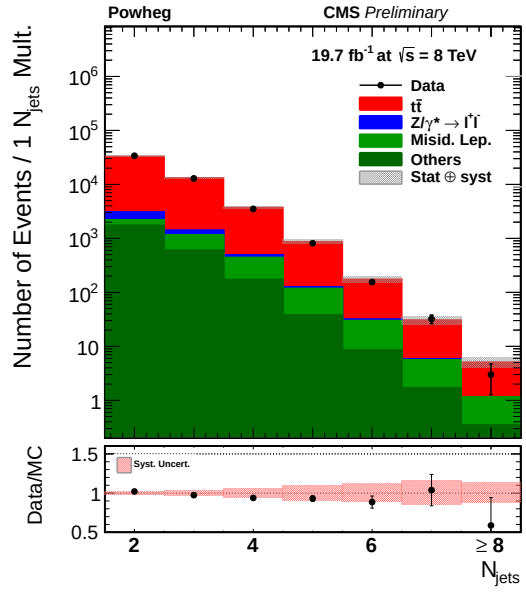

(d) The distribution using the POWHEG $t\bar{t}$ sample.

Figure C.4: The distributions, in the preselection event sample, of the number of jets passing the basic jet identification requirements (Table A.4), where, in each sub-figure, one of four different generated samples of $t\bar{t}$ events has been used to estimate the total contribution of $t\bar{t}$ production to the selected sample.

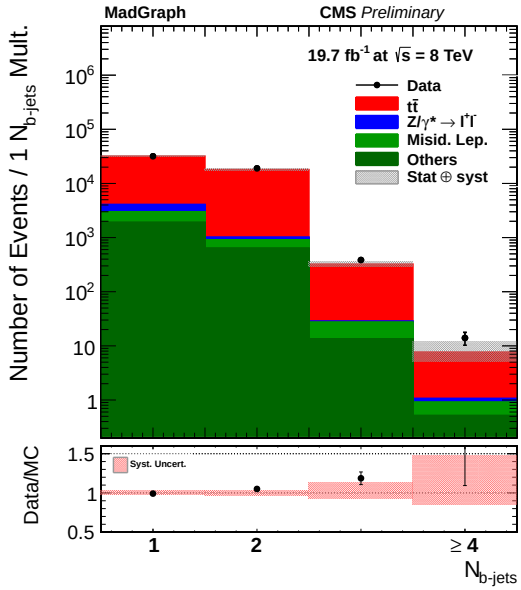

(a) The distribution using the final-state inclusive MadGraph $t\bar{t}$ sample.

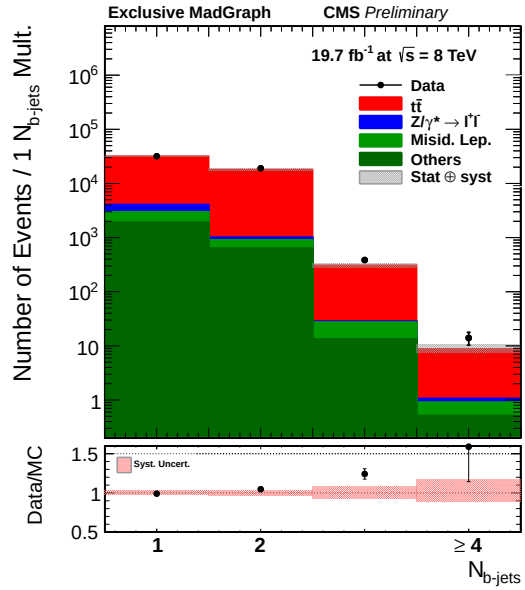

(b) The distribution using the final-state exclusive MadGraph $t\bar{t}$ sample.

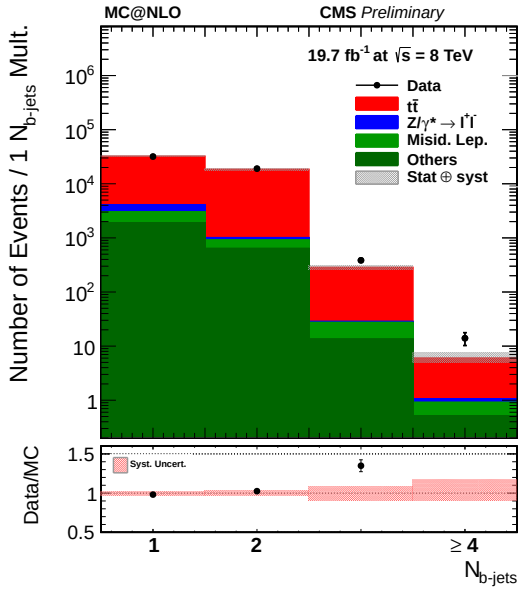

(c) The distribution using the MC@NLO $t\bar{t}$ sample.

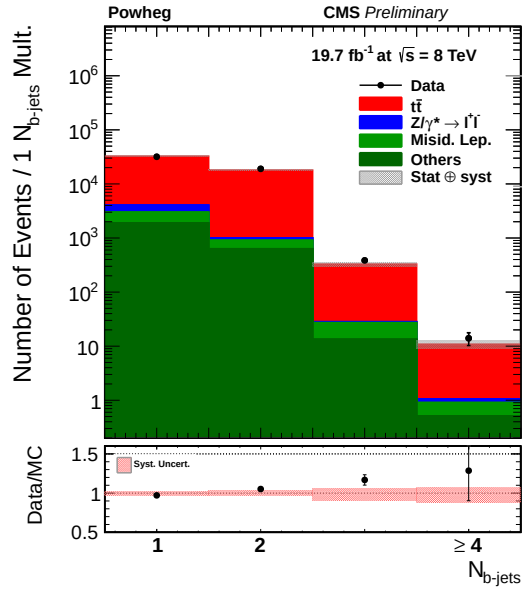

(d) The distribution using the POWHEG $t\bar{t}$ sample.

Figure C.5: The distributions, in the preselection event sample, for the number of b-jets (reconstructed jets passing a b-jet discriminator selection), where, in each sub-figure, one of four different generated samples of $t\bar{t}$ events has been used to estimate the total contribution of $t\bar{t}$ production to the selected sample.

distributions, but for the number of reconstructed b jets. It is clear from these figures that, of the four generated samples, the POWHEG $t\bar{t}$ sample is the only sample with consistently satisfactory modeling of the jet multiplicities. We also note that modeling of jet multiplicities, specifically the expected number of high-jet multiplicity events, is particularly poor in the MC@NLO $t\bar{t}$ sample – Figs. C.4c and C.5c.

Figure C.6 shows the expected distribution of $M_{\mathrm{T2}}(\ell\ell)$ in the four $t\bar{t}$ samples. Comparing the data and simulation in the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV}$, it is immediately clear that in the inclusive MADGRAPH sample, the modeling of this key analysis variable is clearly systematically biases. Of the remaining three samples, the comparisons between data and simulation show that there is acceptable modeling of the $M_{\mathrm{T2}}(\ell\ell)$ variable in all three samples.

In order to better understand why the modeling of the $M_{\mathrm{T2}}(\ell\ell)$ variable is relatively poor in the inclusive MADGRAPH sample, we checked object correlation variables such as the relative opening angles between selected objects.

Figure C.7 shows the distribution of the opening angle between the two selected lepton candidates. It is quite clear that in the inclusive-MADGRAPH $t\bar{t}$ sample, angular configurations where the two leptons are far apart are favored over those where they are close together. The value of the $M_{\mathrm{T2}}(\ell\ell)$ variable is strongly correlated with this dilepton opening angle, however, as small-angle dilepton configurations are more likely to have large $M_{\mathrm{T2}}(\ell\ell)$ values[1].

[1]Remember that $M_{\mathrm{T2}}(\ell\ell)$ is automatically 0 if the $\vec{\not{E}}_{\mathrm{T}}$ falls within the smaller of the two opening angles defined by the two leptons

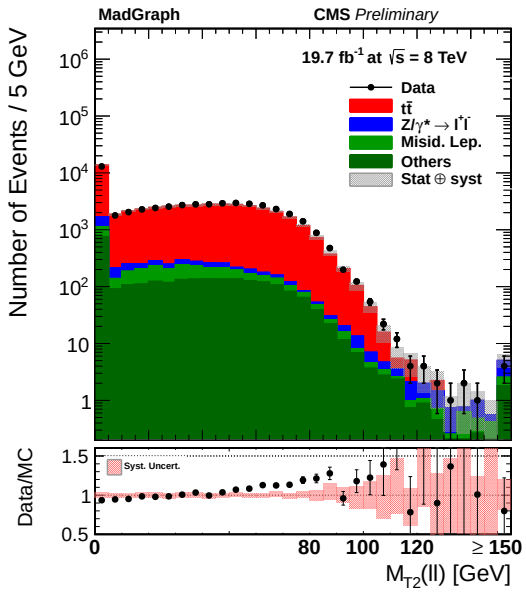

(a) The distribution using the final-state inclusive MadGraph $t\bar{t}$ sample.

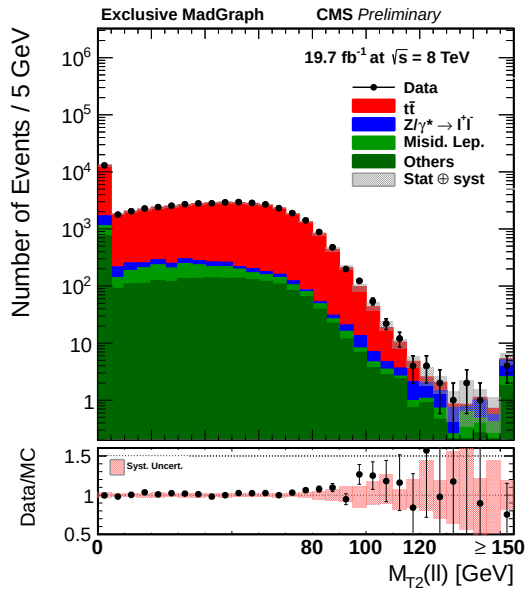

(b) The distribution using the final-state exclusive MadGraph $t\bar{t}$ sample.

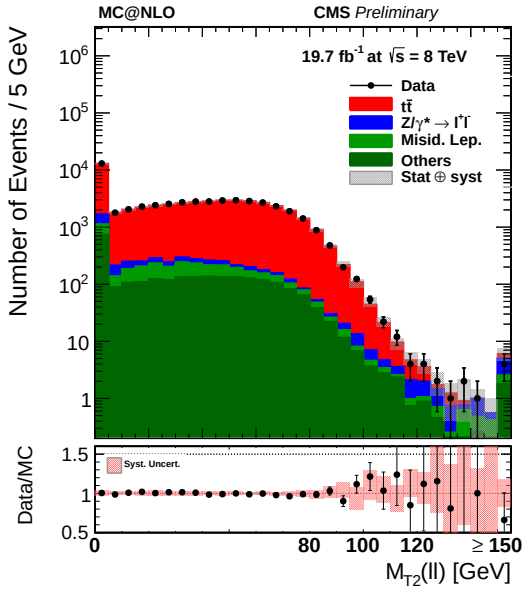

(c) The distribution using the MC@NLO $t\bar{t}$ sample.

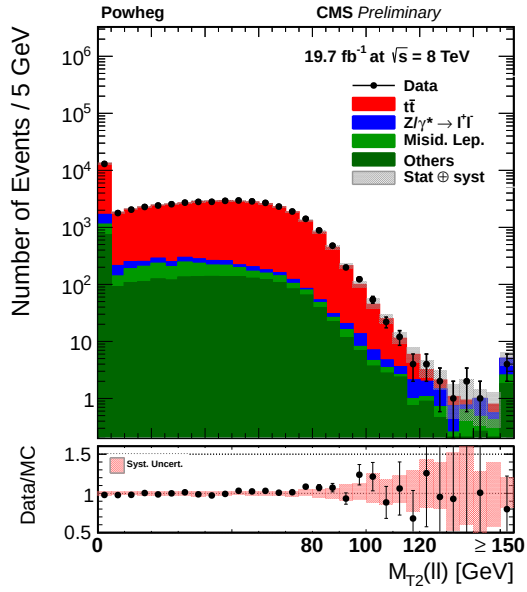

(d) The distribution using the POWHEG $t\bar{t}$ sample.

Figure C.6: The distributions, in the preselection event sample, of $M_{\mathrm{T2}}(\ell\ell)$ where, in each sub-figure, one of four different generated samples of $t\bar{t}$ events has been used to estimate the total contribution of $t\bar{t}$ production to the selected sample.

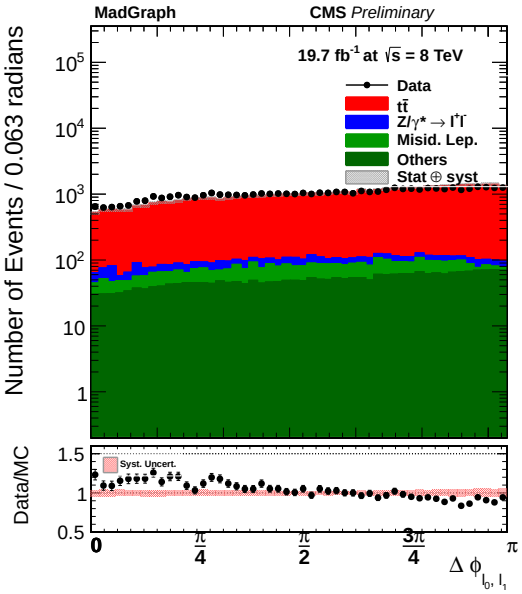

(a) The distribution using the final-state inclusive MadGraph $t\bar{t}$ sample.

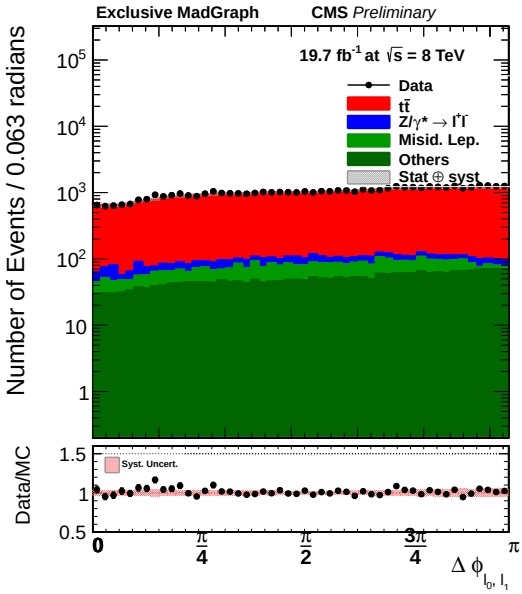

(b) The distribution using the final-state exclusive MadGraph $t\bar{t}$ sample.

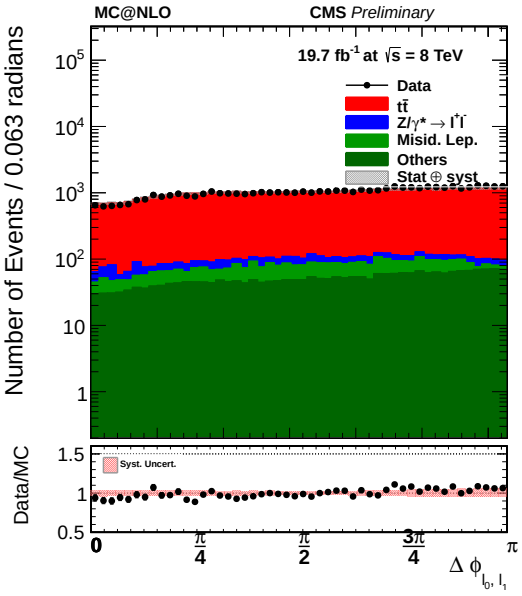

(c) The distribution using the MC@NLO $t\bar{t}$ sample.

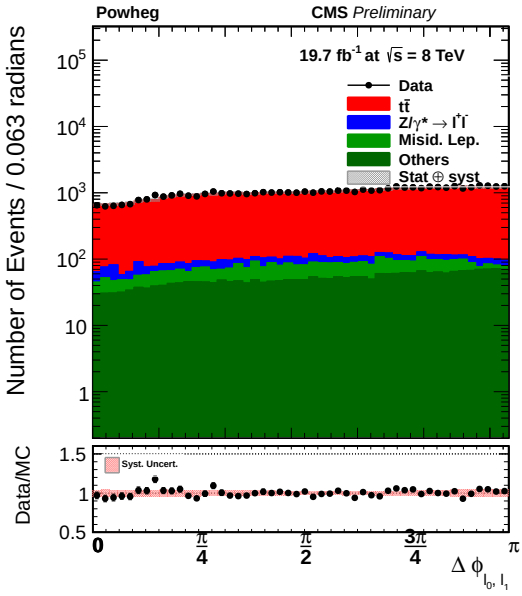

(d) The distribution using the POWHEG $t\bar{t}$ sample.

Figure C.7: The distributions, in the preselection event sample, of $\Delta\phi(\ell_0, \ell_1)$ [the index 0 (1) refers to (sub-)leading selected lepton]. In each sub-figure, one of four different generated samples of $t\bar{t}$ events has been used to estimate the total contribution of $t\bar{t}$ production to the selected sample.

C.2.1 Summary

From the cross-checks described above and other similar ones we found that, in general, the POWHEG generator seems to model the $t\bar{t}$ background with the most accuracy. This was the reason that we chose to use the POWHEG-generated $t\bar{t}$ samples to estimate our $t\bar{t}$ background. It is worth noting though that, for the significance calculations described in Section 7.2.7.1, we did investigate the impact of adding in an additional systematic uncertainty to account for observed differences in the $t\bar{t}$ modeling between the MC@NLO and POWHEG generators, as ostensibly these two generators are both utilizing NLO matrix elements as part of their event generation (in contrast the MADGRAPH, which uses LO matrix elements). Additional discussion on this can be found in the next section.

C.2.2 Accounting for Possible Systematic Biases from the $t\bar{t}$ Generator Choice

In this section, we provide details on an additional systematic uncertainty we investigated that accounts for observed differences in the $t\bar{t}$ modeling between the MC@NLO and POWHEG generators.

Figure C.8 shows a comparison of the 3D M_{T2} shape as it is modeled in the MC@NLO and POWHEG generators. Specifically, Fig. C.8a shows the ratio of the relative distribution of $t\bar{t}$ events in each bin of the 3D M_{T2} shape for the two generators. The statistical significances $\mathcal{S}$ of any deviations between the two shapes

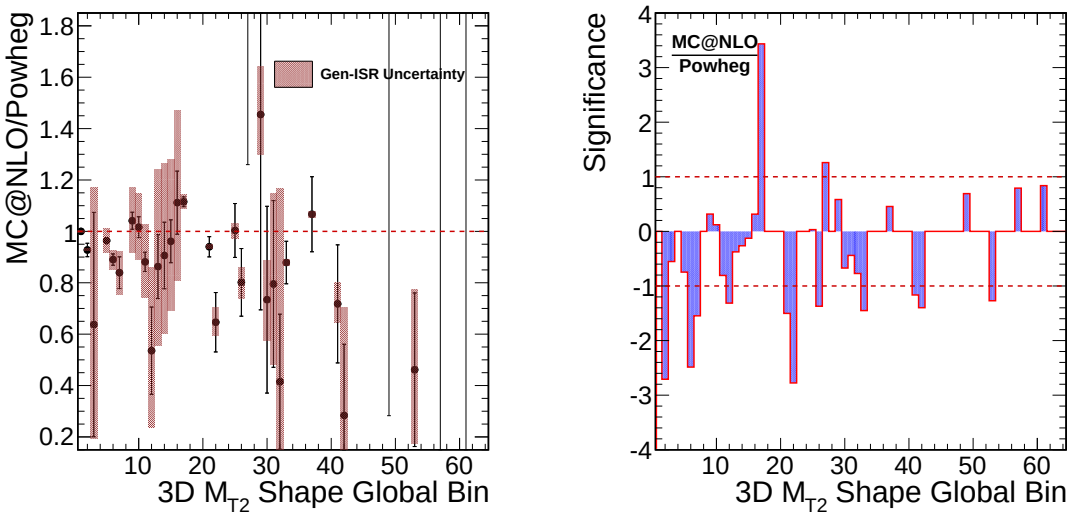

(a) The ratio between the MC@NLO and POWHEG generators, for the relative distribution of $t\bar{t}$ events in each bin of the 3D M_{T2} shape. The error bars represent the statistical uncertainty on this ratio, while the dark red error band represents the systematic uncertainty for the $t\bar{t}$-p_T (ISR) reweighting (Section 6.2.4.5) applied to each sample.

(b) The statistical significance of the deviations shown in Fig. C.8a between the MC@NLO and POWHEG generators.

Figure C.8: Comparisons of the 3D M_{T2} shape between the MC@NLO and POWHEG generators. The 3D M_{T2} Shape Global Bin refers to the de-stacked global bin number, c.f. Eq. (7.6).

are quantified,

$$\mathcal{S} = \frac{\left(\frac{\mathrm{Gen}_a}{\mathrm{Gen}_0} - 1\right)}{\sqrt{\sigma_{\mathrm{stat.}} \oplus \sigma_{\mathrm{syst.}}}}, \tag{C.1}$$

where Gen_a refers to MC@NLO, Gen_0 refers to POWHEG, and $\sigma_{\mathrm{stat.}}$ represents the statistical uncertainty, while $\sigma_{\mathrm{syst.}}$ represents the systematic uncertainty for the $\mathrm{t\bar{t}}$-p_{T} (ISR) reweighting applied to each sample, Section 6.2.4.5. Figure C.8b shows the magnitude of $\mathcal{S}$ for the various bins of the 3D M_{T2} shape.

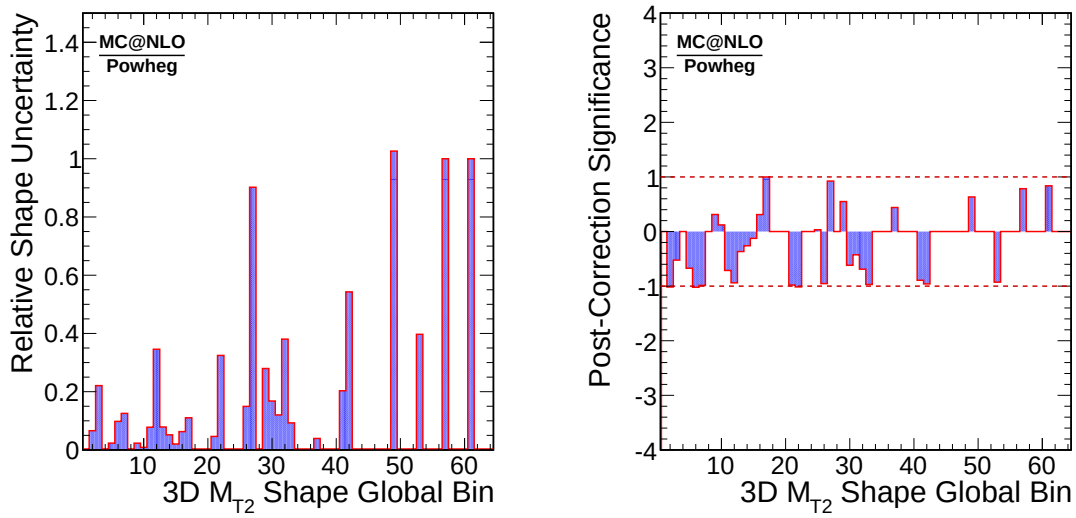

(a) The magnitude of a relative shape-uncertainty between the MC@NLO and POWHEG generators, calculated so as to reduce all inter-generator deviations to have a 1σ statistical significance at most.

(b) The statistical significances of the deviations shown in Fig. C.8a, after including the shape uncertainty from Fig. C.9a.

Figure C.9: The magnitude of a relative shape-uncertainty between the MC@NLO and POWHEG generators, calculated so as to reduce all inter-generator deviations to have a 1σ statistical significance at most. The 3D M_{T2} Shape Global Bin refers to the de-stacked global bin number, c.f. Eq. (7.6).

From Fig. C.8b we derive a relative shape uncertainty,

$$\frac{\sigma_{\mathrm{Gen}}}{\mathrm{Gen}} = \left(\frac{\mathrm{Gen}_a}{\mathrm{Gen}_0} - 1\right)\left(1 - \Phi\left(\frac{\mathcal{S}}{2}\right)\right), \tag{C.2}$$

where $\Phi(x)$ is the Gaussian error function. The magnitude of this uncertainty —

in particular, the inclusion of the $\Phi\left(\frac{\mathcal{S}}{2}\right)$ term — is chosen so that the statistical significance, post-correction, of any deviations between the MC@NLO and POWHEG generators is, at worst, 1σ.

Figure C.9a shows the magnitude of this relative shape uncertainty, while Fig. C.9b shows the results of applying the additional uncertainty and recalculating $\mathcal{S}$ (compare with Fig. C.8b).

C.2.2.1 Including the $t\bar{t}$ Generator Choice Uncertainty in the Significance Calculations

After deriving the relative shape uncertainty, Fig. C.9a, we then included it into our calculation of the observed local significances of the excesses in the 3D M_{T2} shape. Figure C.10 shows the results of including this additional uncertainty. Comparing to the nominal case, where we did not include this uncertainty, Fig. 7.24, we can see that the calculated significances do get smaller. However, it is important to note that the largest observed significances, i.e. the points with $\Delta\mathrm{M} = 150\,\mathrm{GeV}$, are reduced by $\sim 3\%$ at worst. This signifies that the observed excesses in the 3D M_{T2} shape are very likely not due to some systematic bias in the $t\bar{t}$ generator.

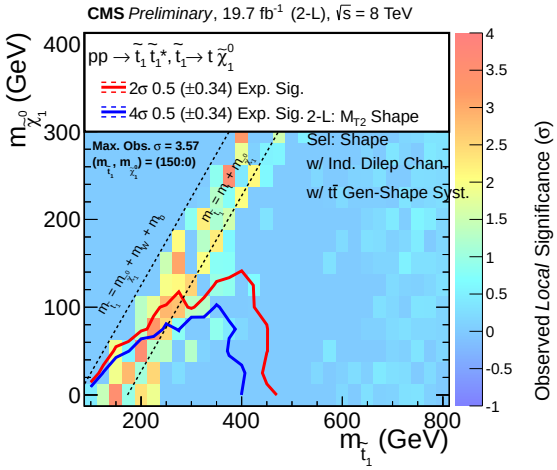

(a) The observed local significance for the $\widetilde{t}_1 \to t\widetilde{\chi}^0_1$ decay mode.

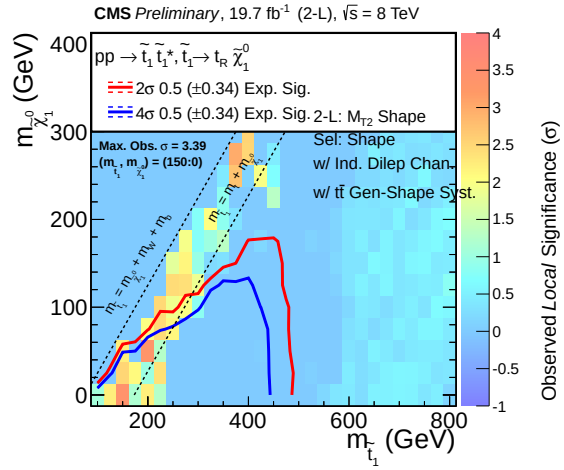

(b) The observed local significance for the $\widetilde{t}_1 \to t_R\widetilde{\chi}^0_1$ decay mode.

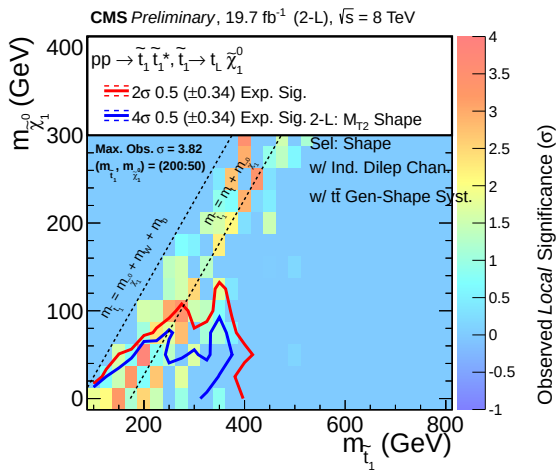

(c) The observed local significance for the $\widetilde{t}_1 \to t_L\widetilde{\chi}^0_1$ decay mode.

Figure C.10: 2D $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}^0_1})$ maps of the *observed local* significance of excesses in the data, using the alternative hypothesis of top-squark pair-production in the different coupling-scenarios of the $\widetilde{t}_1 \to t\widetilde{\chi}^0_1$ decay mode, where an additional shape uncertainty, accounting for $t\bar{t}$ generator differences, see Section C.2.2, has been included in the significance calculation. Overlaid on top of theses maps are contours denoting the regions of the $(m_{\widetilde{t}_1}, m_{\widetilde{\chi}^0_1})$ plane where, if top-squark pair-production was occurring with the nominal signal strength $\mu = 1$, it is expected that excesses at the 2σ (red) or 4σ (blue) level would be observed in the data. The solid line represents the median expected significance, while the dashed lines represent the 0.16 and 0.84 percentiles on the expected significance.

C.3 Cross-checks for the 3D M_{T2} Shape Analysis

In this section we provide details on additional checks we performed to test the robustness of the 3D M_{T2} shape analysis. These are separated into two categories: discussions on the analysis itself, and discussions of the statistical results derived from the analysis.

C.3.1 Validating the 3D M_{T2} Shape Analysis

C.3.1.1 Checking Correlations: $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ with $\mathrm{Max}(m_{\ell\mathrm{b}})$

From the discussion of Section 7.2.2, it is clear that $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ is strongly correlated with the masses $m_{\ell\mathrm{b}}$ of the ℓb pairs used to construct it. As part of the 3D M_{T2} shape analysis we require $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$. In this section we provide a few plots confirming that this cut is not introducing systematic biases into the $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ shape.

Figure C.11 shows the $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ distribution for events passing the basic preselection. The separate sub-figures show this distribution when different requirements are placed upon $\max(m_{\ell\mathrm{b}})$.

It is obvious to see that the shape of the $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ distribution strongly depends upon the specific requirement placed upon $\max(m_{\ell\mathrm{b}})$. The relevant quantity of interest is how well this correlation is modeled within the simulation. We can examine this by comparing the dependence of the ratio of data and simulation, in particular, the shape of this ratio, on the $\max(m_{\ell\mathrm{b}})$ cut.

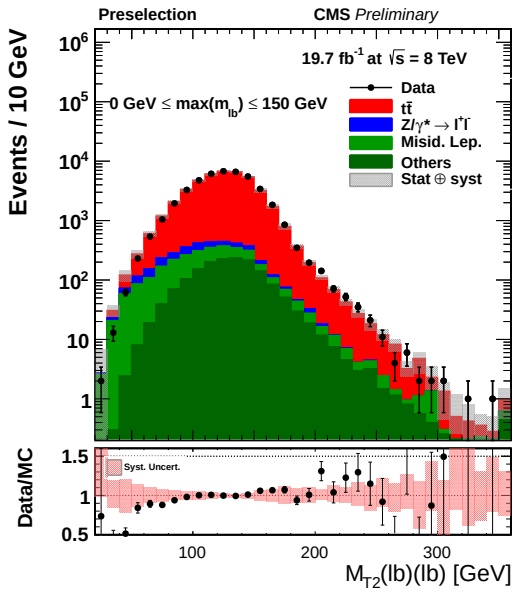

(a) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 150\,\mathrm{GeV}$.

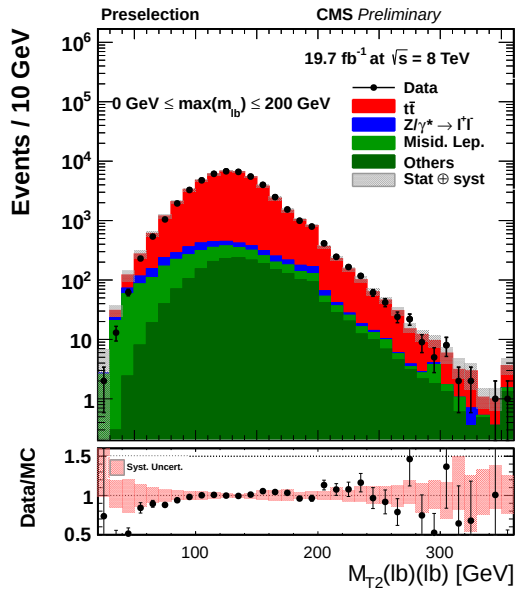

(b) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$.

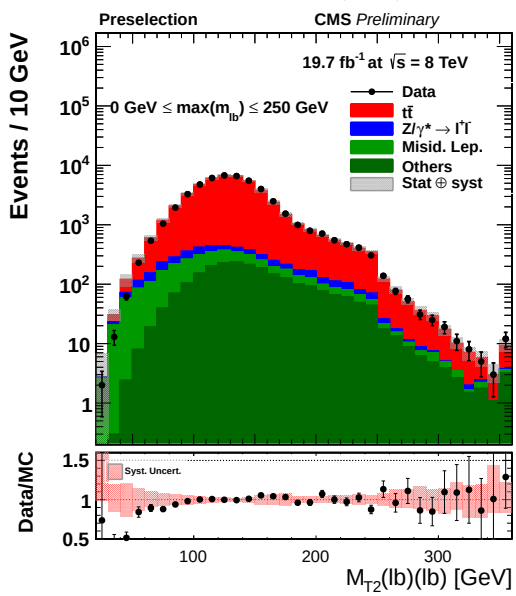

(c) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 250\,\mathrm{GeV}$.

Figure C.11: Distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for different cut requirements of $\max(m_{\ell\mathrm{b}})$, in a sample of events passing the basic preselection.

Some care has to be taken in comparing the sub-figures of Fig. C.11, as the far tail of the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution is nominally a signal region in this event sample. Taking into account the pre-existing $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ shape uncertainties that are already included, these comparisons indicate that, at worst, the cut on the $\max(m_{\ell\mathrm{b}})$ has an $\sim$ 2–3% effect on the ratio of data and simulation. This is negligible in comparison with the other shape uncertainties in this analysis.

As an additional check, we also considered a zero-b jet control region, where events are required to pass the basic preselection, **except** every event must contain no reconstructed b jets. Figure C.12 shows the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution for events falling into this subsample. As with Fig. C.11, the separate sub-figures show this distribution when different requirements are placed upon $\max(m_{\ell\mathrm{b}})$. Again, comparisons of the data and simulation indicate that the simulation is modeling this dependence with acceptable accuracy within the pre-existing systematic uncertainties.

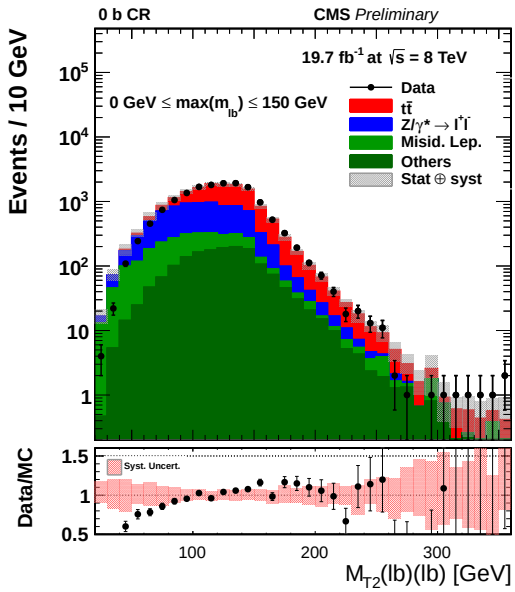

(a) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 150\,\mathrm{GeV}$.

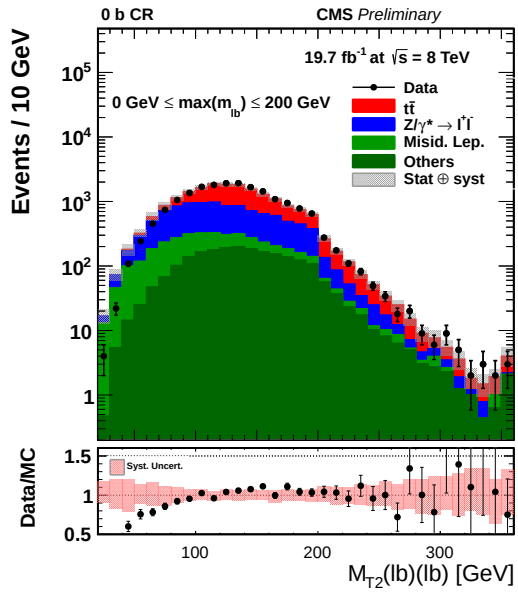

(b) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$.

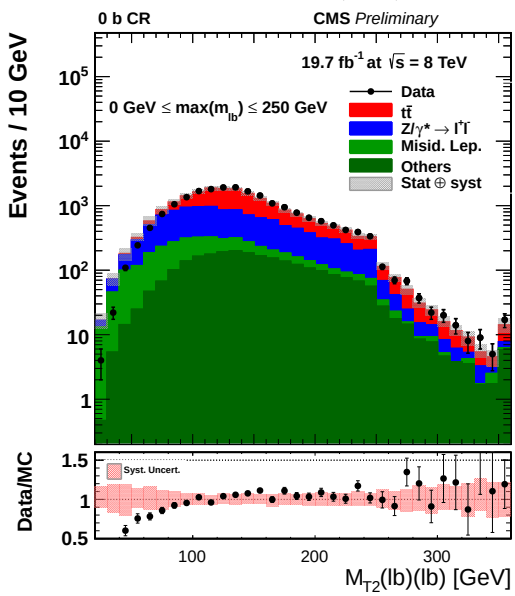

(c) The $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution when $\max(m_{\ell\mathrm{b}}) < 250\,\mathrm{GeV}$.

Figure C.12: Distributions of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ for different cut requirements of $\max(m_{\ell\mathrm{b}})$, in a sample of events passing the basic preselection, except that each event is required to have 0 reconstructed b-jets.

C.3.1.2 Checking the Post-fit Nuisance Parameters

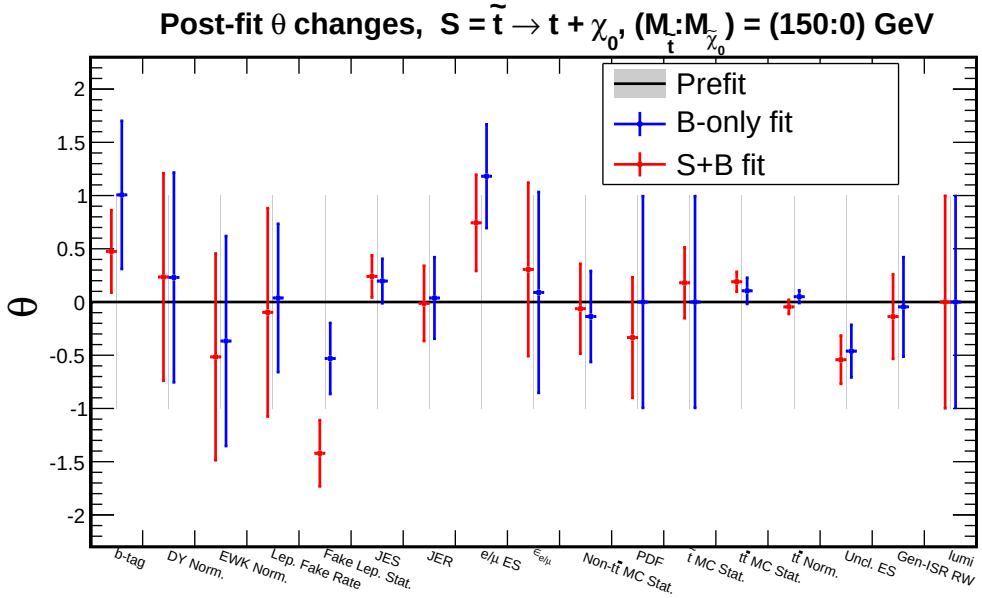

Figure C.13: The post-fit values for the nuisance parameters, θ, used as input into the dilepton-channel split 3D M_{T2} shape analysis. These post-fit values are shown for the background-only (blue) and unconstrained signal + background (red) hypotheses. The signal hypothesis used corresponds to the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ top-squark decay mode with $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0}) = (150{:}0)\,\mathrm{GeV}$. Pre-fit, the θ were parametrized such that they were dimensionless, centered at 0, and had width parameters of 1.

As part of limit-setting and significance calculations, Eqs. (G.5) and (G.9), we perform likelihood-maximizations of the nuisance parameters θ (e.g. systematic uncertainties). As a reminder, pre-fit, the θ are parametrized such that they are dimensionless, centered at 0, and have width parameters of 1 (e.g. the σ parameter for Gaussian uncertainties).

Both the mean value and width parameter are allowed to float in the fits. If the mean value of θ shifts, it corresponds to shifting (in the same direction) the overall expected contribution for the backgrounds (or signal) that are affected

by the corresponding systematic uncertainty. If the width's value is changed, it corresponds to changing (in a correlated fashion) the overall systematic uncertainty on the backgrounds (or signal) that are affected by the corresponding systematic uncertainty.

Figure C.13 shows the post-fit values of the θ parameters used as input into the split dilepton-channel 3D M_{T2} shape analysis. These post-fit values are shown for the background-only (blue) and unconstrained signal + background (red) hypotheses. The signal hypothesis used corresponds to the unpolarized T2tt top-squark decay mode with $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0}) = (150{:}0)\,\mathrm{GeV}$.

A number of the nuisance parameters are strongly constrained by the fits. For example, the widths of the 2 $\mathrm{t\bar{t}}$-normalization systematic uncertainties ($\mathrm{t\bar{t}}$ MC simulation statistics and the conservative, additional $\mathrm{t\bar{t}}$ normalization factor) are reduced significantly, down to $\approx 10 - 15\%$ of their initial values.

This is expected. In the control region, $M_{\mathrm{T2}}\,(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}\,(\mathrm{bb}) < 170\,\mathrm{GeV}$, which is included in the significance calculation, $\approx 4 \times 10^4$ $\mathrm{t\bar{t}}$ events are expected. This enables us to **strongly** constrain the $\mathrm{t\bar{t}}$ normalization and subsequently the magnitude of these systematic uncertainties. As well, because we see a similar amount of events in this region in data, the overall normalization of the $\mathrm{t\bar{t}}$ is not strongly shifted by the likelihood fit.

The widths of the systematics related to the misidentified lepton background are reduced to approximately 30% of their initial values. More notable, however, is the fact that the mean values are shifted downward by approximately 0.5 (1.5) in the background-only (signal + background) fits.

This is also expected. Numerous checks shown at various points in this dissertation have outlined the clear overestimate of the misidentified lepton background. However, this background tends to be fairly localized to a few regions of the 3D M_{T2} parameter space. Thus, as part of the statistical calculations, the large systematic uncertainties on this background enable a morphing of the misidentified lepton shape to better match the observed data, see Appendix H.3. As well, there is enough statistical precision in the observed data to also constrain the uncertainty on the misidentified lepton background.

Calculating the Significances with Fixed Nuisance Parameters

We performed a cross-check to validate that the significance calculations, e.g. Section 7.2.7.1, weren't being biased by overly constrained nuisance parameters. We repeated the calculation of the observed local significance for the high significance point, $m_{\tilde{\mathrm{t}}_1} = 200\,\mathrm{GeV}$, $m_{\tilde{\chi}_1^0} = 50\,\mathrm{GeV}$, from Fig. 7.24c, except we fixed one or more of the nuisance parameters to the pre-fit values (note that this included freezing any shape-based morphing related to the fixed nuisance parameters). Some of the nuisance parameters we tried fixing included the jet energy scale calibration (JES), the unclustered energy scale calibration (Uncl. ES), and the MC statistics for both the $\mathrm{t\bar{t}}$ and non-$\mathrm{t\bar{t}}$ simulation samples. We found that the significance varied by at most 5%, and in general, the variations were much closer to 1%. This strongly suggests that the large observed excesses shown in Section 7.2.7.1 are not due to systematic biases in the nuisance parameter constraint procedure.

C.3.2 Cross-checks on the Shape Analysis Results

One of the more interesting results of the 3D M_{T2} shape analysis was the presence of observed excesses in the 3D M_{T2} shape that fit with strong statistical significances to top-squark signal models. In this sub-section, we discuss a number of cross-checks on this result, designed to both develop our intuition about the result as well as to confirm the robustness of the statistical calculations.

C.3.2.1 The Indirect Effects of a Top-squark in the Higgs Sector

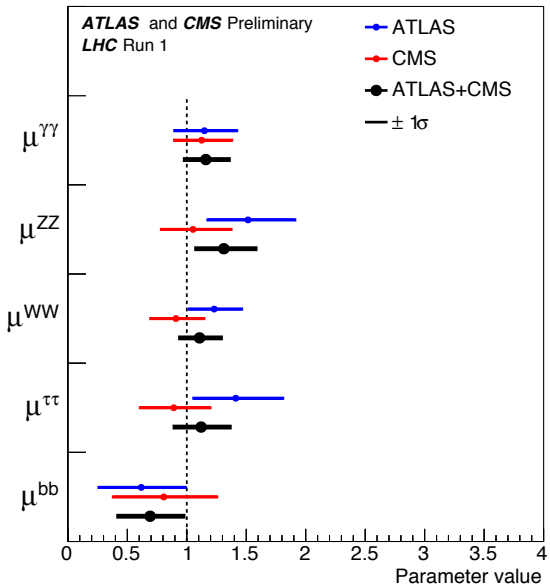

Figure C.14: Best-fit results for the observed signal strength μ for different decay modes of the Higgs-boson, as measured by the combination of ATLAS and CMS results. Also shown for completeness are the individual results for each experiment. The error bars indicate the $\pm 1\sigma$ intervals. Reprinted from Fig. 12 of [172].

One of the main channels used to discover the Higgs boson at the LHC was the decay, $\mathrm{H} \to \gamma\gamma$. Although the Higgs boson does not directly couple to photons, this decay can proceed through triangle-loop diagrams involving charged particles,

with the most relevant diagram (due to relative coupling strengths) involving top quarks running in the loop. Similarly, at the LHC, one of the main production modes for the Higgs boson is gluon-gluon fusion, again (primarily) involving top quarks running in a triangle-loop.

If the top-squark exists, there would be top-squark loop processes for both the gluon-gluon fusion production mode and the $\mathrm{H} \to \gamma\gamma$ decay mode. The overall effect of the top-squark would be to change the relative rates of the Higgs-boson production as well as its decay to photons. These changes would depend upon the masses of the top-squarks, $m_{\tilde{\mathrm{t}}_1}$ and $m_{\tilde{\mathrm{t}}_2}$. Measuring the rates (i.e. observed signal strengths), relative to the SM expectation, of Higgs boson production through gluon-gluon fusion and the $\mathrm{H} \to \gamma\gamma$ decay can therefore provide an indirect check on any possible top-squark signals, such as the possible signal observed in the 3D M_{T2} shape. The measurement of the observed signal strength in the $\mathrm{H} \to \gamma\gamma$ decay mode is particularly interesting as a nominal top-squark would enter into both the Higgs boson production and decay diagrams.

Figure C.14 shows the current ATLAS + CMS combined best measurements of the Higgs boson signal strength for a number of final states [172]. All of the observed signal strengths are consistent with the SM expectations within 1–2σ. This of course means that none of the observed signal strengths can definitively point towards the existence of new physics. However, they can still serve to rule out the points in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\mathrm{t}}_2})$ plane that would lead to statistically incompatible changes in μ.

Figure C.15 displays two contour maps, shown in the 2D plane of $m_{\tilde{\mathrm{t}}_1}$ and $m_{\tilde{\mathrm{t}}_2}$, of the expected change in the Higgs-boson's gluon-gluon fusion production

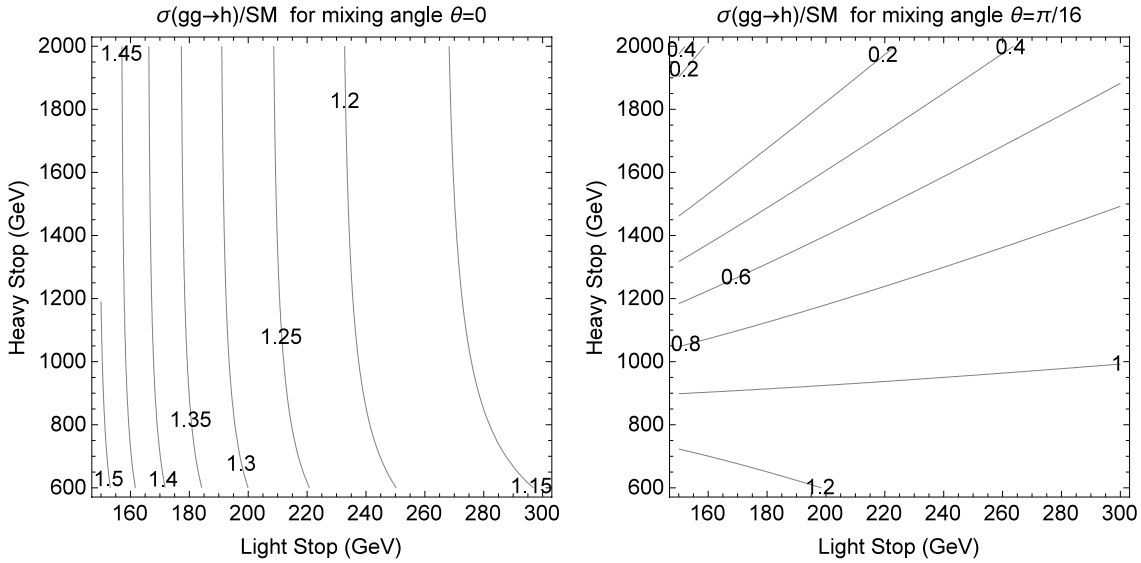

(a) The ratio, $\frac{\sigma^{\mathrm{MSSM}}(\mathrm{gg}\to\mathrm{H})}{\sigma^{\mathrm{SM}}(\mathrm{gg}\to\mathrm{H})}$, for the $\widetilde{\mathrm{t}}$ mixing angle, $\theta_t = 0$.

(b) The ratio, $\frac{\sigma^{\mathrm{MSSM}}(\mathrm{gg}\to\mathrm{H})}{\sigma^{\mathrm{SM}}(\mathrm{gg}\to\mathrm{H})}$, for the $\widetilde{\mathrm{t}}$ mixing angle, $\theta_t = \frac{\pi}{16}$.

Figure C.15: Contour plots showing the dependence of $\frac{\sigma^{\mathrm{MSSM}}(\mathrm{gg}\to\mathrm{H})}{\sigma^{\mathrm{SM}}(\mathrm{gg}\to\mathrm{H})}$ on the masses of the $\widetilde{\mathrm{t}}_1$ and $\widetilde{\mathrm{t}}_2$. Figures received from private communication [173].

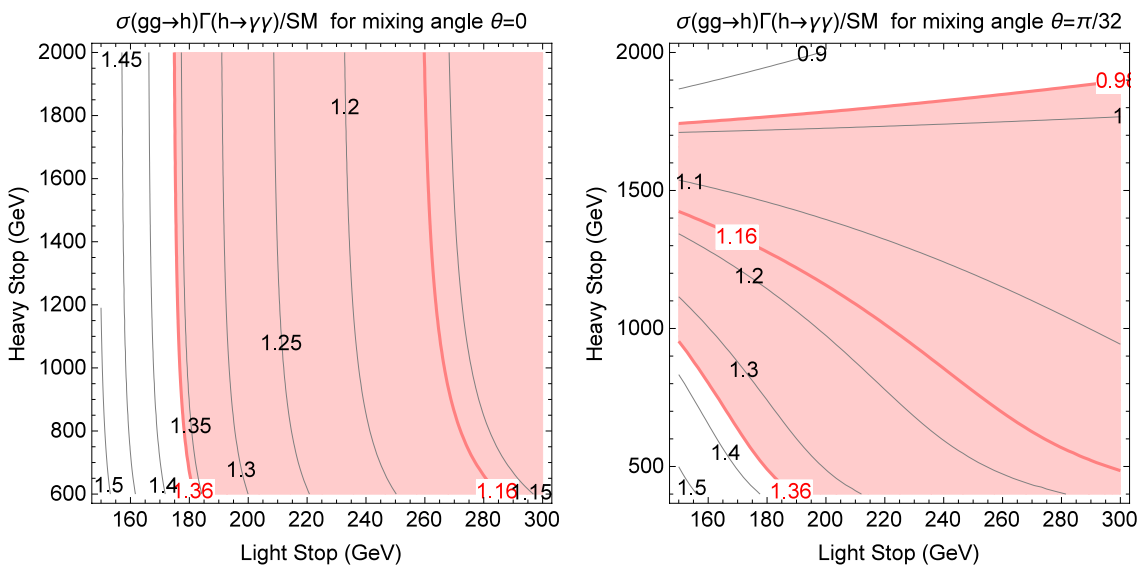

(a) The ratio, $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$, for the $\widetilde{\mathrm{t}}$ mixing angle, $\theta_t = 0$.

(b) The ratio, $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$, for the $\widetilde{\mathrm{t}}$ mixing angle, $\theta_t = \frac{\pi}{32}$.

Figure C.16: Contour plots showing the dependence of $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$ on the masses of the $\widetilde{\mathrm{t}}_1$ and $\widetilde{\mathrm{t}}_2$. The red contour shows the current best ($\pm 1\sigma$) combined ATLAS + CMS measurement of $\mu^{\gamma\gamma}$ (c.f. Fig. C.14). Figures received from private communication [173].

cross section, $\frac{\sigma^{\mathrm{MSSM}}(\mathrm{gg}\to\mathrm{H})}{\sigma^{\mathrm{SM}}(\mathrm{gg}\to\mathrm{H})}$. These contour maps were calculated in the MSSM for different values of the mixing angle θ_t between the $\widetilde{\mathrm{t}}_{\mathrm{L}}$ and $\widetilde{\mathrm{t}}_{\mathrm{R}}$. One can see that, for no mixing between the $\widetilde{\mathrm{t}}_{\mathrm{L}}$ and $\widetilde{\mathrm{t}}_{\mathrm{R}}$, the expected change in $\frac{\sigma^{\mathrm{MSSM}}(\mathrm{gg}\to\mathrm{H})}{\sigma^{\mathrm{SM}}(\mathrm{gg}\to\mathrm{H})}$ depends primarily upon $m_{\widetilde{\mathrm{t}}_1}$. However, once $\widetilde{\mathrm{t}}_{\mathrm{L}}$–$\widetilde{\mathrm{t}}_{\mathrm{R}}$ mixing is introduced, the expected change develops a (rather nontrivial) dependence on $m_{\widetilde{\mathrm{t}}_2}$ as well.

Figure C.16 shows analogous contour maps as Fig. C.15, except for the expected change in $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$. Also shown on these contour maps are red-shaded regions that display the measured ($\pm 1\sigma$) $\mu^{\gamma\gamma}$ from Fig. C.14. Again, in analogy with Fig. C.16, the expected change in $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$ can depend upon both $m_{\widetilde{\mathrm{t}}_1}$ and $m_{\widetilde{\mathrm{t}}_2}$, where the relative dependence varies with θ_t. We note, however, that none of the displayed regions show an expected change in $\mu^{\mathrm{MSSM}}_{\mathrm{H}\to\gamma\gamma}$ that is *inconsistent* with the current measurements of $\mu^{\gamma\gamma}$. Notably, for $\theta_t = 0$, i.e. $\widetilde{\mathrm{t}}_1 = \widetilde{\mathrm{t}}_{\mathrm{L}}$, the points with $m_{\widetilde{\mathrm{t}}_1} = 200\,\mathrm{GeV}$, which yielded the strongest significance when fit to the 3D M_{T2} excesses, c.f. Fig. 7.24c, are consistent the current measured $\mu^{\gamma\gamma}$ within 1σ.

C.3.2.2 The (Apparent) Lack of a Top-squark Signal in Other Analyses

Consider the hypothetical situation that the possible low-mass top-squark signal observed in the 3D M_{T2} shape is a real signature of MSSM-like top-squark pair-production and not just a statistical fluctuation of the background or a missing systematic uncertainty. In this case, we would expect to see it show up both in other CMS analyses looking at complementary final states, such as the semi-leptonic

top-squark search discussed in Section 7.1.1 and Appendix E, as well as the ATLAS analysis looking in the dileptonic final state [59].

We note that neither of these analyses observe an excess in this region of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane. Of course, an obvious explanation, albeit disappointing, is that the possible signal is just a statistical fluctuation in the background. For the rest of this section, however, we will adopt an optimistic viewpoint and discuss possible reasons for why the possible signal could be real, yet still not be observed by the other analyses.

One straightforward explanation for why the semi-leptonic top-squark search might not observe any excesses is that the new signal is not an MSSM top-squark at all, but is instead some SUSY-variant where the dileptonic final state of top-squark pair-production is heavily favored. This would alter the relative signal efficiencies to emphasize observing new physics in the dilepton channel over other channels. On that same note, this possible new signal might not be a top-squark at all, but instead some new, pair-produced particle that preferentially decays to a b quark, charged lepton, and (possibly) some invisible, new light particle. Of course, if either of these were the case, one might still expect that the ATLAS dilepton top-squark search would have observed something.

Here we bring up an important point in the discussion. Both the semi-leptonic top-squark search as well as the ATLAS dilepton top-squark search utilized BDT-based approaches to discriminate between the top-squark signal and the SM backgrounds. Multivariate approaches are certainly powerful but, as noted in Section 7.2.4, they can suffer from systematic biases and inefficiencies if the "true"

signal is notably different from the one used to train the BDTs. One can see an example of this in Fig. 7.2, where BDTs trained on the unpolarized T2tt decay mode were used to exclude top-squark pair-production when the chirality of the $\widetilde{\mathrm{t}}_1$–t coupling is entirely left- or right-handed. In this figure, we showed that this resulted in worse statistical sensitivities in the region $\Delta\mathrm{M} < m_{\mathrm{t}}$, where the kinematics of the top-squark signal change very quickly from one mass-point to another, and notably depend upon the chirality of the coupling between the $\widetilde{\mathrm{t}}_1$ and the t.

This highlights an advantage of the 3D M_{T2} shape approach, which is that it is fairly agnostic with respect to which signals it is sensitive to. Notably, because the SM backgrounds are so heavily concentrated in the control region, $M_{\mathrm{T2}}(\ell\ell) < 80\,\mathrm{GeV} \cap M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 170\,\mathrm{GeV} \cap M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) < 170\,\mathrm{GeV}$, the 3D M_{T2} shape analysis possesses some degree of sensitivity for a wide variety of possible new signals, simply through examination of the tail regions of the M_{T2} shape.

C.3.2.3 Dilepton Channel Compatibility

The large observed significances in Fig. 7.24a correspond to the region $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}_1^0} = 150\,\mathrm{GeV}$. For these significance calculations, the signal strength for all three dilepton channels were tethered into a single, global signal strength that was left as a floating parameter in the alternative hypothesis, see Eq. (G.9). However, if we are truly observing signatures of top-squark production in the data, then there should be consistency between all three dilepton channels.

We check for this consistency by allowing the signal strength to float indepen-

dently in each channel, so as to maximize the likelihoods independently for each channel. We then compare against the global signal strength hypothesis via the "χ^2"-like variable, $\tilde{\chi}^2 = -2 \ln \frac{\mathcal{L}(\text{data}|\text{nominal})}{\mathcal{L}(\text{data}|\text{alternate})}$. We calculate a null hypothesis p-value by comparing the observed value of $\tilde{\chi}^2$ against the expected distribution of $\tilde{\chi}^2$ that we generate from an ensemble of pseudo-data. Note that we allow the individual signal strengths to be slightly negative in order to improve the approximate equivalence between $\tilde{\chi}^2$ and a true χ^2.

Table C.1 shows the results of these channel compatibility checks for the unpolarized T2tt decay mode. For each mass point, we show the global $\hat{\mu}$ as well as the individual channel $\hat{\mu}$. As well, we should the observed value of $\tilde{\chi}^2$ and the corresponding null hypothesis p-value. One can see that, for the two lowest $m_{\tilde{\mathrm{t}}_1}$ points, the three dilepton channels are quite consistent with one another. Another notable point is that none of $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ points shows a deviation in the channel compatibility that's larger than 2σ.

Table C.2 shows three additional channel compatibility checks: A channel compatibility check for the unpolarized $x = 0.25$ T2bw decay mode, a channel compatibility check between the main 3D M_{T2} shape selection and a 0 b jet control sample, and a channel compatibility check between the main 3D M_{T2} shape selection and the semi-leptonic top-squark search.

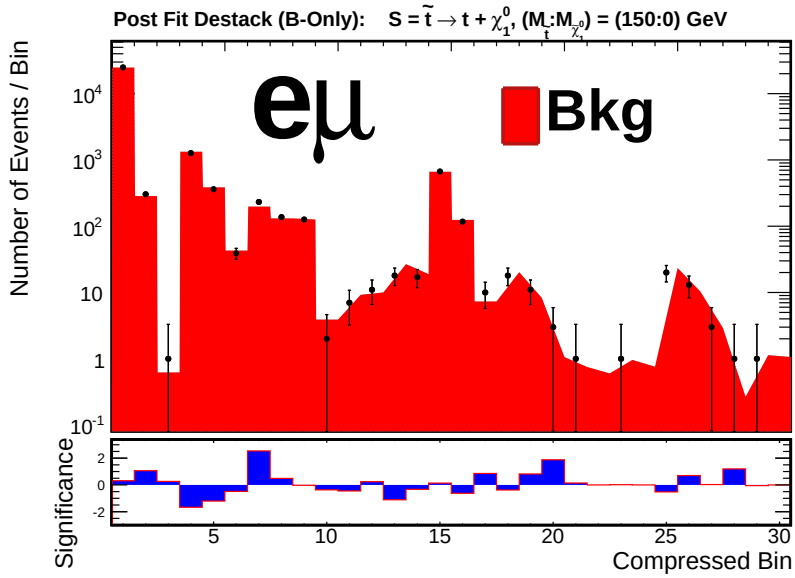

(a) The post-fit comparison for the background-only hypothesis.

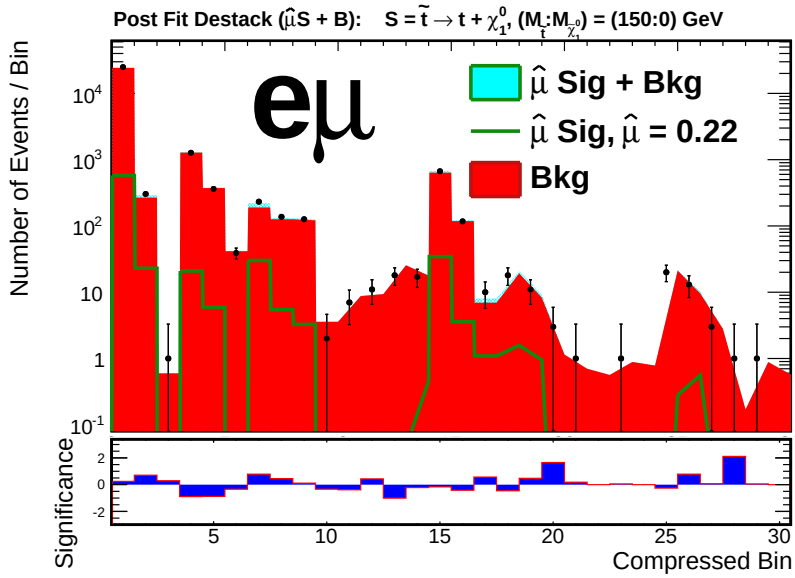

(b) The post-fit comparison for the unconstrained signal + background hypothesis.

Figure C.17: Post-fit comparisons between the data and the background-only and unconstrained signal ($\widetilde{t}_1 \to t\widetilde{\chi}_1^0$) + background hypotheses in the $e^{\pm}\mu^{\mp}$ channel of the 3D M_{T2} shape analysis. The upper panel shows the fits of the hypotheses to the data. Note that the linear interpolations between bin edges is purely visual. The bottom panels show the observed statistical significances of deviations between the data and the fit for the two hypotheses.

C.3.2.4 Post-fit Comparisons with Data

To better understand why certain mass points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane fit with such strong significance to the observed excesses in the 3D M_{T2} shape, it is useful to consider post-fit comparisons of the observed shape of the collision data against the shapes of the background-only and signal + background hypotheses. Figure C.17 shows the comparisons for the $\mathrm{e}^{\pm}\mu^{\mp}$ shape for the mass point $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (150{:}0)\,\mathrm{GeV}$. The upper panel shows the fits of the hypotheses to the data Note that we have explicitly removed the bins from this figure that were rebinned as part of the procedure described in Section 7.2.4.5, hence the reduced number of bins relative to Fig. 7.21b. Note that the linear interpolations between bins is purely visual. The bottom panel shows the observed statistical significances of the deviations of the fit from the data for each bin of the fit.

It is interesting to note that the observed significance of deviations are much lower for the signal + background hypothesis, particularly in bins with large statistics. This leads to a much stronger likelihood for this hypothesis, which subsequently yields a stronger calculated significance.

Table C.1: A compatibility check between the 3 dilepton channels for the 3D M_{T2} shape dileptonic top-squark search for various mass points in the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}^0_1})$ plane for the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}^0_1$ top-squark decay mode. The signal strength μ is floated ($\mu \to \hat{\mu}$) in order to best fit the observed data. Two separate fits are done. For the first, μ is kept fixed between all 3 channels. For the second, μ is allowed to individually float in each separate dilepton channel in order to maximize the fit likelihood for each channel separately. The compatibility of the channels is then quantified by a "χ^2-like" variable, $\tilde{\chi}^2 = -2\ln\frac{\mathcal{L}(\mathrm{data}|\mathrm{nominal})}{\mathcal{L}(\mathrm{data}|\mathrm{alternate})}$. The p-value of the nominal hypothesis (the three dilepton channels are consistent) is then quantified using the expected distribution of $\tilde{\chi}^2$ from an ensemble of pseudo-data.

$(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}^0_1})$ [GeV]	Global $\hat{\mu}$	Split-Channel $\hat{\mu}$		$\tilde{\chi}^2$	$\tilde{\chi}^2$ p-value
$\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}^0_1 \to \mathrm{bW}^+\widetilde{\chi}^0_1$					
(125:0)	$0.29^{+0.00}_{-0.14}$	$0.25^{+0.14}_{-0.15}$	$\mathrm{e}^\pm\mu^\mp$	1.58	0.38 ± 0.02
		$0.36^{+0.18}_{-0.17}$	$\mu^+\mu^-$		
		$-0.32^{+0.70}_{-0.74}$	$\mathrm{e}^+\mathrm{e}^-$		
(150:0)	0.22 ± 0.05	0.21 ± 0.00	$\mathrm{e}^\pm\mu^\mp$	1.38	0.65 ± 0.02
		0.28 ± 0.00	$\mu^+\mu^-$		
		0.38 ± 0.00	$\mathrm{e}^+\mathrm{e}^-$		
(200:0)	0.23 ± 0.10	$0.23^{+0.12}_{-0.12}$	$\mathrm{e}^\pm\mu^\mp$	5.58	0.07 ± 0.00
		$0.38^{+0.17}_{-0.18}$	$\mu^+\mu^-$		
		$0.96^{+0.50}_{-0.43}$	$\mathrm{e}^+\mathrm{e}^-$		
(225:0)	0.22 ± 0.15	0.27 ± 0.19	$\mathrm{e}^\pm\mu^\mp$	5.35	0.08 ± 0.00
		0.68 ± 0.26	$\mu^+\mu^-$		
		0.02 ± 0.58	$\mathrm{e}^+\mathrm{e}^-$		
(175:25)	0.27 ± 0.11	0.27 ± 0.02	$\mathrm{e}^\pm\mu^\mp$	4.99	0.11 ± 0.01
		0.51 ± 0.04	$\mu^+\mu^-$		
		1.15 ± 0.05	$\mathrm{e}^+\mathrm{e}^-$		
(225:75)	1.52 ± 0.47	$1.54^{+0.47}_{-0.43}$	$\mathrm{e}^\pm\mu^\mp$	5.97	0.07 ± 0.01
		$2.23^{+0.68}_{-0.59}$	$\mu^+\mu^-$		
		$4.64^{+1.57}_{-1.54}$	$\mathrm{e}^+\mathrm{e}^-$		
(250:100)	1.32 ± 0.28	$1.10^{+0.59}_{-0.56}$	$\mathrm{e}^\pm\mu^\mp$	1.86	0.34 ± 0.01
		$1.64^{+0.74}_{-0.70}$	$\mu^+\mu^-$		
		$2.83^{+1.75}_{-1.56}$	$\mathrm{e}^+\mathrm{e}^-$		

Table C.2: Additional results on channel compatibility presented in the same format as Table C.1. Notably, these additional results include channel compatibility checks between the main 3D M_{T2} shape selection (Main) and a 0 b jet (0 b) control sample, as well as a channel compatibility check between the main 3D M_{T2} shape selection and the semi-leptonic top-squark search (e/μ).

$(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ [GeV]	Global $\hat{\mu}$	Split-Channel $\hat{\mu}$		$\tilde{\chi}^2$	$\tilde{\chi}^2$ p-value
$\tilde{\mathrm{t}}_1 \to \mathrm{t}\tilde{\chi}_1^0 \to \mathrm{bW}^+\tilde{\chi}_1^0$					
(150:0)	$0.13^{+0.05}_{-0.04}$	$0.11^{+0.06}_{-0.06}$	$\mathrm{e}^{\pm}\mu^{\mp}$ Main	3.00	0.54 ± 0.01
		$0.05^{+0.09}_{-0.05}$	$\mu^+\mu^-$ Main		
		$0.30^{+0.23}_{-0.15}$	$\mathrm{e}^+\mathrm{e}^-$ Main		
		0.16 ± 0.08	$\mathrm{e}^{\pm}\mu^{\mp}$ 0 b		
		$0.16^{+0.24}_{-0.09}$	$\mu^+\mu^-$ 0 b		
		$0.36^{+0.18}_{-0.12}$	$\mathrm{e}^+\mathrm{e}^-$ 0 b		
$\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+ \to \mathrm{bW}^+\tilde{\chi}_1^0$, $x = 0.25$					
(225:75)	$1.80^{+0.44}_{-0.43}$	$1.60^{+0.48}_{-0.49}$	$\mathrm{e}^{\pm}\mu^{\mp}$	0.81	0.54 ± 0.01
		$2.10^{+0.61}_{-0.60}$	$\mu^+\mu^-$		
		$1.20^{+1.20}_{-0.72}$	$\mathrm{e}^+\mathrm{e}^-$		
(175:25)	0.40 ± 0.01	0.29 ± 0.03	$\mathrm{e}^{\pm}\mu^{\mp}$	2.30	0.28 ± 0.01
		0.49 ± 0.04	$\mu^+\mu^-$		
		$0.00^{+0.28}_{-0.00}$	$\mathrm{e}^+\mathrm{e}^-$		
		$0.00^{+0.17}_{-0.00}$	e/μ		

Appendix D: Reweighting to Different Top-squark Coupling Chiralities

As noted in Section 3.2.6.3, the daughter particles produced in top-squark decays can have non-trivial polarizations. The signal simulation samples generated by CMS were made with arbitrary fixed choices of the relevant SUSY parameters, and thus represent the expected signal only for fixed polarization scenarios.

However, it is known that both the dileptonic top-squark search, Chapters 6 and 7, as well as the semi-leptonic top-squark search, Section 7.1 and Appendix E, can be strongly dependent upon the top-squark coupling-polarization scenarios.

Notably, both of these analyses depend upon the reconstructed object selection efficiencies as well as the relative angular correlations between objects, both of which can be strongly affected by the considered polarization scenarios.

The authors of Ref. [174] developed techniques to reweight the signal simulation to match any expected top-squark coupling-polarization scenario, for both the $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ (T2tt) and $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ (T2bw) decay modes. We will provide a quick overview of the technique here; more details can be found in Ref. [174].

D.1 General Discussion on Reweighting Signal MC

The basic idea of reweighting, for a general top-squark decay mode is to take our signal MC sample (our "reference" sample) and apply an event weight to each event given by,

$$\mathcal{W}^i(\text{ref} \to \text{target}) = \frac{\sigma^{\text{total}}_{\text{ref}}}{\sigma^{\text{total}}_{\text{target}}} \frac{\left|\mathcal{M}^i_{\text{target}}\right|^2}{\left|\mathcal{M}^i_{\text{ref}}\right|^2}, \tag{D.1}$$

where $\mathcal{M}^i_{\text{ref}}$ and $\mathcal{M}^i_{\text{target}}$ are the matrix element terms for the two considered scenarios, summed over helicities of initial and final state particles. The normalization factor, $\sigma^{\text{total}}_{\text{ref}}/\sigma^{\text{total}}_{\text{target}}$, accounts for possible differences in the total cross section between the two scenarios, and ensures that the total number of events (i.e. without cuts) does not change on average after reweighting (the total cross section times branching fraction is an external input to the analysis).

There are some additional simplifications that can be applied to Eq. (D.1) in the context of the $\text{pp} \to \widetilde{\text{t}}_1\widetilde{\text{t}}_1^*$ process. The top-squark particle is a scalar, and furthermore, in the signal generation, is assumed to be on shell. Thus, the squared matrix element term for the whole process, $\text{pp} \to \widetilde{\text{t}}_1\widetilde{\text{t}}_1^* \to \text{bW}^+\widetilde{\chi}^0_1\ \overline{\text{b}}\text{W}^-\widetilde{\chi}^0_1$. can be split into three independent factors. The first of these is related just to the production of the $\text{pp} \to \widetilde{\text{t}}_1\widetilde{\text{t}}_1^*$ system directly; this part includes relevant EWK/QCD corrections, including things such as ISR.

Because, however, the "reference" and "target" scenarios use the same top-squark masses, the ratio of this component for the "reference" and "target" scenarios is 1.

The other two matrix element factors describe the decays of the top-squark and its antiparticle, $|\mathcal{M}_{\widetilde{\mathrm{t}}_1\to\ldots}|^2$ and $|\mathcal{M}_{\widetilde{\mathrm{t}}_1^*\to\ldots}|^2$. In terms of these two pieces, the weight can be calculated as,

$$\mathcal{W}^i = \left[\frac{\Gamma^{\widetilde{\mathrm{t}}_1}_{\mathrm{ref}}}{\Gamma^{\widetilde{\mathrm{t}}_1}_{\mathrm{target}}}\right]^2 \left[\frac{\left|\mathcal{M}^i_{\widetilde{\mathrm{t}}_1,\mathrm{target}}\right|^2}{\left|\mathcal{M}^i_{\widetilde{\mathrm{t}}_1,\mathrm{ref}}\right|^2}\right] \left[\frac{\left|\mathcal{M}^i_{\widetilde{\mathrm{t}}_1^*,\mathrm{target}}\right|^2}{\left|\mathcal{M}^i_{\widetilde{\mathrm{t}}_1^*,\mathrm{ref}}\right|^2}\right]. \tag{D.2}$$

In Eq. (D.2), the ratio of the corresponding partial widths ensures the maintenance of the normalization. Furthermore, it is worth noting that the interactions between the stop and its direct decay products are expected to CP conserving (even if, through neutralino-mixing, a $\widetilde{\mathrm{t}}_1$ might decay to either a $\widetilde{\mathrm{t}}_\mathrm{L}$ or $\widetilde{\mathrm{t}}_\mathrm{R}$). This means that in Eq. (D.2) the second and third term should be equivalent after performing the exchange of particles with their corresponding anti-particles.

We will now discuss the reweighting in more detail for the T2tt decay mode. For the T2bw decay mode, the resulting equations for the event weights are similar [174], albeit slightly more complicated due to the additional SUSY decay chain.

D.2 Details of the Reweighting for the T2tt Decay Mode

The general form for the $\widetilde{\mathrm{t}}_1$ decay's $|\mathcal{M}|^2$, including when the top is produced off-shell in the stop decay, i.e. when $\Delta\mathrm{M} = m_{\widetilde{\mathrm{t}}_1} - m_{\widetilde{\chi}^0_1} \leq m_\mathrm{t}$, can be written as,

$$\begin{aligned}|\mathcal{M}_{\widetilde{\mathrm{t}}_1\to\widetilde{\chi}^0_1\mathrm{b}\bar{\mathrm{d}}\mathrm{u}}|^2 \propto\ & \frac{1}{(p_\mathrm{t}^2 - m_\mathrm{t}^2)^2 + m_\mathrm{t}^2\Gamma_\mathrm{t}^2}\cdot\frac{1}{((p_\mathrm{u}+p_{\bar{\mathrm{d}}})^2 - m_\mathrm{W}^2)^2 + m_\mathrm{W}^2\Gamma^2_{\mathrm{W}^+}}\cdot \\ & [\ \left(2\ s_\theta^2(p_{\widetilde{\chi}^0_1}\cdot p_\mathrm{t}) + 2 s_\theta c_\theta m_\mathrm{t} m_{\widetilde{\chi}^0_1}\right)(p_{\bar{\mathrm{d}}}\cdot p_\mathrm{t})\ (p_\mathrm{u}\cdot p_\mathrm{b}) \\ & +\ \left(c_\theta^2 {m_\mathrm{t}}^2 - s_\theta^2 p_\mathrm{t}^2\right)(p_{\bar{\mathrm{d}}}\cdot p_{\widetilde{\chi}^0_1})\ (p_\mathrm{u}\cdot p_\mathrm{b})\ ],\end{aligned}$$

where $s_\theta \equiv \sin\theta^{\rm t}_{\rm eff}$ and $c_\theta \equiv \cos\theta^{\rm t}_{\rm eff}$ (and $\theta_{\rm eff}$is defined in Section 3.2.6.3). In the expression, $p_{\rm t}$ denotes the four-momentum of the $\rm b\bar{d}u$ system (i.e. the product decays expected from a top quark). The off-shell nature of the top manifests in the dependence on $p_{\rm t}^2$ (on-shell, $p_{\rm t}^2 = m_{\rm t}^2$). Note that the $m_{\rm t}$ in Eq. (D.3) refers to the **on-shell** mass of the top quark ($\approx 173\,\rm GeV$).

For the remaining piece, the partial width, $\Gamma_{\tilde{\rm t}_1\to\tilde{\chi}^0_1{\rm t}}$, of the stop decay for a given SUSY scenario has the following simple form,

$$\Gamma_{\tilde{\rm t}_1\to\tilde{\chi}^0_1{\rm t}} \propto (p_{\tilde{\chi}^0_1}\cdot p_{\rm t}) + 2s_\theta c_\theta m_{\rm t} m_{\tilde{\chi}^0_1}. \tag{D.3}$$

Thus, the weight for the $\tilde{\rm t}_1$ is given by,

$$\mathcal{W}^i_{\tilde{\rm t}_1} = \frac{\left[2\ s_\theta^2(p_{\tilde{\chi}^0_1}\cdot p_{\rm t}) + 2s_\theta c_\theta m_{\rm t} m_{\tilde{\chi}^0_1}\right](p_{\bar{\rm d}}\cdot p_{\rm t}) + \left[c_\theta^2 {m_{\rm t}}^2 - s_\theta^2 p_{\rm t}^2\right](p_{\bar{\rm d}}\cdot p_{\tilde{\chi}^0_1})}{\left[(s_\theta^2 + c_\theta^2\frac{m_{\rm t}^2}{p_{\rm t}^2})(p_{\tilde{\chi}^0_1}\cdot p_{\rm t}) + 2s_\theta c_\theta m_{\rm t} m_{\tilde{\chi}^0_1}\right](p_{\bar{\rm d}}\cdot p_{\rm t})}. \tag{D.4}$$

So, in summary, if one wishes to assign a weight for a given event to reweight it to match a given "polarization" scenario, one simply calculates the weight specified in Eq. (D.4) for the $\tilde{\rm t}_1$ (along with an additional weight for the $\tilde{\rm t}^*_1$), plugging in the correct $\theta_{\rm eff}$to correspond to the desired "polarization" scenario. Listing the correct $\theta_{\rm eff}$ for some common "polarization" scenarios below,

- "Fully Right-handed" tops: $\theta_{\rm eff} = 0$
- "Unpolarized" tops: $\theta_{\rm eff} = \frac{\pi}{4}$
- "Fully Left-handed" tops: $\theta_{\rm eff} = \frac{\pi}{2}$

Where it is worth noting that the notion of the "polarization" of the top-quark only makes complete sense in the "on-shell" case, and that further more, the top-quark can only be fully left- or right-handed when $m_{\tilde{\chi}^0_1} = 0$, as per Eq. (3.28).

Appendix E: Overview of the Semi-leptonic top-squark Search

The majority of the information in this section is paraphrased from Ref. [37].

Table E.1 contains the relevant details on the preselection utilized by the semi-leptonic top-squark search.

At the preselection level, the two primary backgrounds are $\mathrm{t\bar{t}}$ (90%) and W + jets (7%). Because of the $M_{\mathrm{T}} > 100$ GeV cut, the majority of the $\mathrm{t\bar{t}}$ background is dileptonic $\mathrm{t\bar{t}}$ where one of the leptons is not reconstructed.

For the signal selection, the semi-leptonic analysis applies a multivariate approach based on BDTs (Boosted Decision Trees, [164]).

Separate BDTs are trained for each top-squark decay mode and for various regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane, primarily demarcated by $\Delta\mathrm{M} = m_{\tilde{\mathrm{t}}_1} - m_{\tilde{\chi}_1^0}$.

The full list of variables fed into each separate BDT training are provided in Ref. [37].

The output of these BDTs, when applied to a collision event (simulated or real) is a BDT discriminant, a variable that represents how "signal-like" a given event is. The distribution of this variable tends to be *approximately* exponentially distributed for both signal and background. Relative to the background, the signal tends to have a stronger tail.

Table E.1: Preselection requirements applied for the semi-leptonic top-squark search, Section 7.1.1.

Object	Selection
N_ℓ	exactly 1
p_{T}^{ℓ}	$> 30\,(20)$ GeV for e (μ)
$\lvert\eta^{\ell}\rvert$	$< 1.4442\,(2.1)$ for e (μ)
$\not{E}_{\mathrm{T}}$	> 80 GeV
M_{T}	> 100 GeV
N_{jets}	≥ 4
N_{b}	≥ 1
Isolated Track Veto	
Tau Veto	

Once these BDT discriminant distributions are produced for both signal and background, a cut-and-count experiment is performed in the tail of the BDT discriminant distribution. The exact BDT discriminant cut values used depend upon the BDT, and are optimized based on minimizing the median expected upper limit produced using an asymptotic method, see Appendix G.2 or Ref. [168].

Figure E.1 shows examples of two of these BDT discriminant distributions. Both BDTs were trained on the T2bw, $x = 0.75$ top-squark decay mode, but in different regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ 2D plane. For both sub-figures the final selection cut defining the 1-bin counting experiment is shown as a vertical, dotted black line with an arrow pointing in the direction of integration.

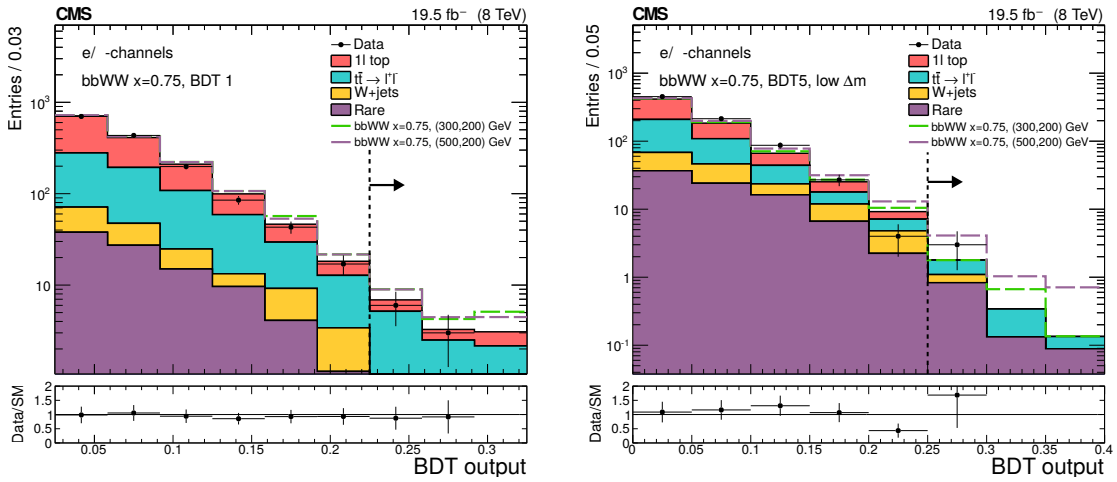

(a) (b)

Figure E.1: The distribution of the BDT discriminant for two different BDTs used by the semi-leptonic top-squark search. Both BDTs were trained on the $x = 0.75$ $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ top-squark decay mode, but in different regions of the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ 2D plane. In addition to collision data (points with error bars) and the expected background (filled histograms), two representative signal mass points are shown as dashed lines: $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ [GeV] = (300, 200), and (500, 200). For both sub-figures the final selection cut defining the 1-bin counting experiment is shown as a vertical, dotted black line with an arrow pointing in the direction of integration. Reprinted from Fig. 5 of [37].

E.1 Background Estimation in the Semi-leptonic top-squark Search

The majority of the backgrounds in the semi-leptonic top-squark search are estimated in a data-driven fashion.

E.1.1 Normalizing the M_{T} Tail

For example, the modeling of M_{T} tails for the 1ℓ and W + jets backgrounds was checked in a 0 b-jet control region. Figure E.2a shows the distribution of the M_{T} variable in this control region. There is a clear discrepancy between data and simulation in the tail of this distribution. In order to correct this, separate templates

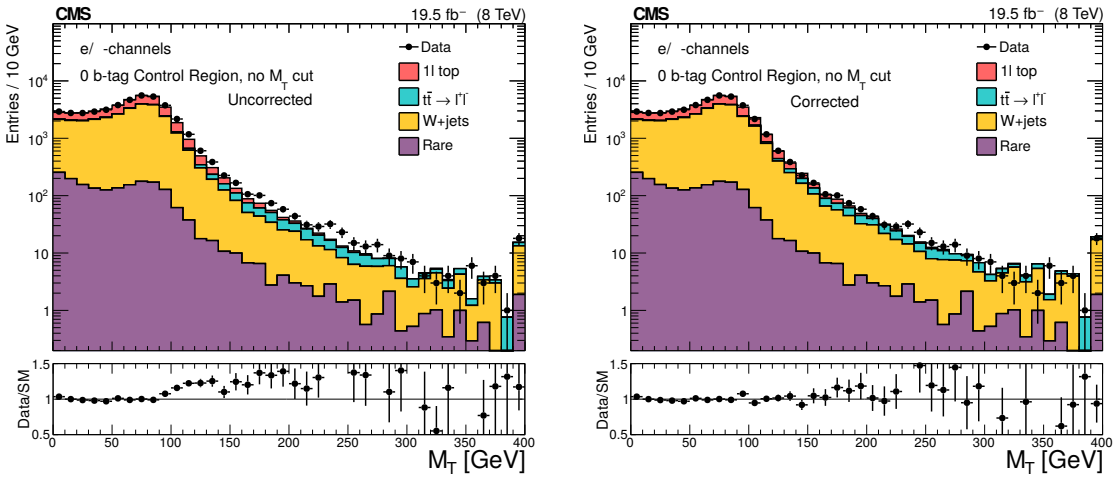

(a) The M_{T} distribution in the semi-leptonic 0-bjet control region, *without* the $\mathrm{SF}_{1\ell}$ and SF_{W} scale factors applied.

(b) The M_{T} distribution in the semi-leptonic 0-bjet control region, *with* the $\mathrm{SF}_{1\ell}$ and SF_{W} scale factors applied.

Figure E.2: The full M_{T} distribution for the semi-leptonic analysis, Section 7.1.1, in a 0 b-jet control region. The left sub-figure shows the distribution without any scale factors applied to correct the 1ℓ $\mathrm{t\bar{t}}$ or W + jets backgrounds. The right sub-figure is the same plot after the application of correction scale factors. Reprinted from Fig. 6 of [37].

for the W + jets and 1ℓ $\mathrm{t\bar{t}}$ distributions were fit to data in the peak M_{T} region ($50 < M_{\mathrm{T}} < 80$ GeV) and in the tail ($M_{\mathrm{T}} > 100$ GeV). Individual scale factors for each background were derived from comparing the template fits for these two regions. Using separate templates for the two backgrounds account for the relative increased amount of high-mass off-shell Ws in the W + jet events.

Figure E.2b shows the M_{T} distribution in this control region after the application of the derived scale factors. As expected, the agreement between data and simulation has been substantially improved.

Table E.2: Selection requirements applied in the semi-leptonic top-squark search to generate a sample enriched in dileptonic $t\bar{t}$ events, Section 7.1.1.

Object	Selection
e, μ	exactly two, oppositely charged
$\lvert\eta^{e}\rvert$	< 1.4442
$\lvert\eta^{\mu}\rvert$	< 2.1
p_{T}^{ℓ}	$> 20\,\text{GeV}$
$M_{\ell^+\ell^-}$	$M_{\ell^+\ell^-} < 76$ GeV $\cup$ $M_{\ell^+\ell^-} > 106$ GeV
$\not{E}_{T}$	> 50 GeV
N_{jets}	≥ 1
N_{b}	≥ 1

E.1.2 Dileptonic $t\bar{t}$ with a "Lost" Lepton

In order to check the modeling of the dileptonic $t\bar{t}$ background, the semi-leptonic top-squark search utilized a dileptonic $t\bar{t}$ control region. Table E.2 shows the event selection for this control region. Direct comparisons of the distribution of M_{T} (constructed between the $\not{E}_{T}$ and the leading lepton) found that the simulation accurately modeled the expected shape and normalization of the M_{T} variable in this control region.

E.1.3 Systematics in the Semi-leptonic Top-squark Search

Table E.3 contains a list of the relevant systematics affecting the total background estimate in the semi-leptonic top-squark search. The relative magnitude at the preselection level, as well as the range of variation over the BDT selections, are both shown. The majority of the systematics listed in Table E.3 are derived

Table E.3: Summary of the relative systematic uncertainties on the total background estimate for the semi-leptonic top-squark search, Section 7.1.1. The preselection level values are shown on the left, while the range of their variations over the BDT selections is shown on the right. Reprinted from Table 4 of [37].

	Uncertainty (%) at preselection	**Uncertainty (%)** range over BDTs
$\mathrm{SF}_{1\ell}$ uncertainty	16.4	0 - 24
SF_{W} uncertainty	1.4	0 - 5
Modeling of M_{T} tail in $\mathrm{t\bar{t}} \to \ell^+\ell^-$	1.6	7 - 39
Modeling of N_{jets} in $\mathrm{t\bar{t}}$	1.1	1 - 4
Modeling of the 2nd lepton veto	1.2	1 - 4
M_{T} peak normalization	0.7	3 - 37
MC statistics in SR	0.4	3 - 38
Cross section uncertainties	2.0	4 - 34
Total	16.8	23 - 58

as part of data-driven background estimations, e.g. the uncertainties on $\mathrm{SF}_{1\ell}$ and SF_{W} are related to the statistical uncertainties from the template fits described in Section E.1.

E.2 Results of the Semi-leptonic top-squark Search

Figure E.3 displays maps of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane that show the demarcation where different BDTs are used to separate the signal and background in the semi-leptonic top-squark search. Notably, for the T2tt decay mode, in the off-shell top quark region, $\Delta\mathrm{M} < m_{\mathrm{t}}$, the kinematics of the top-squark events change very quickly from one point to another. Thus, three different BDTs are utilized in this region in order to ensure a maximal sensitivity to top-squark production in this region of the

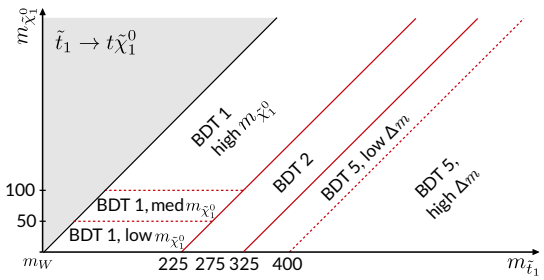

(a) The separate BDT regions for the unpolarized T2tt decay mode.

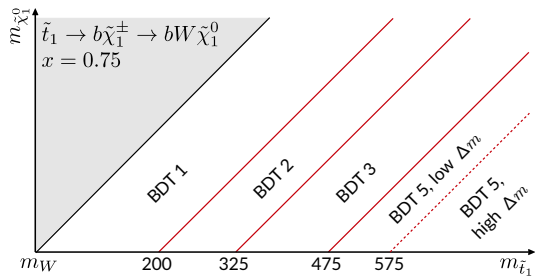

(b) The separate BDT regions for the unpolarized T2bw decay mode with $x = 0.75$.

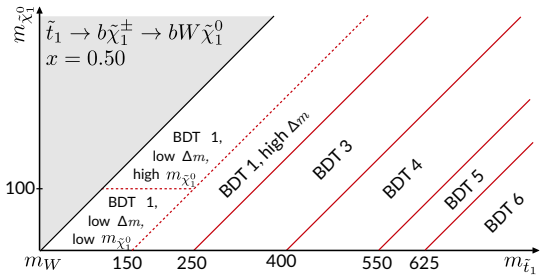

(c) The separate BDT regions for the unpolarized T2bw decay mode with $x = 0.50$.

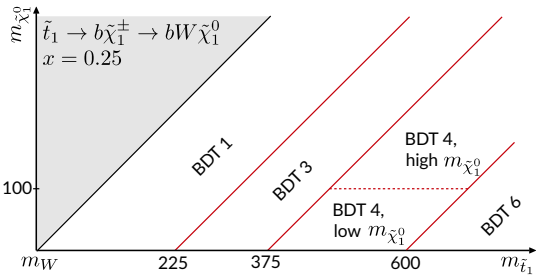

(d) The separate BDT regions for the unpolarized T2bw decay mode with $x = 0.25$.

Figure E.3: Maps of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane showing the demarcation where different BDTs are used to separate the signal and background in the semi-leptonic top-squark search. Reprinted from Fig. 4 of [37].

$(m_{\widetilde{t}_1}, m_{\widetilde{\chi}_1^0})$ plane.

Table E.4: Background prediction, without signal contamination, and observed data for the BDT selections in the semi-leptonic top-squark search, see Section 7.1.1 or Appendix E. Reprinted from Table 5 of [37].

T2tt	**BDT 1**	**BDT 1**	**BDT 1**	**BDT 2**	**BDT 5**	**BDT 5**
	Low $m_{\tilde{\chi}_1^0}$	**Medium** $m_{\tilde{\chi}_1^0}$	**High** $m_{\tilde{\chi}_1^0}$		**Low** Δ**M**	**High** Δ**M**
Background	363.3 ± 35.0	46.3 ± 15.5	19.0 ± 6.7	37.0 ± 12.9	5.7 ± 2.4	3.6 ± 1.7
Data	286	33	17	33		1

T2bw (x**=0.25**)	**BDT 1**	**BDT 3**	**BDT 4**	**BDT 4**	**BDT 6**
			Low $m_{\tilde{\chi}_1^0}$	**High** $m_{\tilde{\chi}_1^0}$	
Background	41.6 ± 11.2	28.5 ± 7.2	19.9 ± 5.4	5.0 ± 1.9	5.8 ± 2.5
Data	27	23	19	5	6

T2bw (x**=0.50**)	**BDT 1**	**BDT 1**	**BDT 1**	**BDT 3**	**BDT 4**	**BDT 5**	**BDT 6**
	Low Δ**M**	**Low** Δ**M**	**High** Δ**M**				
	Low $m_{\tilde{\chi}_1^0}$	**High** $m_{\tilde{\chi}_1^0}$					
Background	14.0 ± 4.8	2.7 ± 1.5	91.0 ± 25.4	7.1 ± 2.4	0.8 ± 0.3	0.7 ± 0.4	2.7 ± 1.3
Data	16	1	85	4	1	2	5

T2bw (x**=0.75**)	**BDT 1**	**BDT 2**	**BDT 3**	**BDT 5**	**BDT 5**
				Low Δ**M**	**High** Δ**M**
Background	13.2 ± 3.9	23.1 ± 7.1	10.5 ± 3.1	2.3 ± 1.1	0.4 ± 0.2
Data	9	15	6	3	0

Table E.4 shows the background prediction, without signal contamination, and observed data for the BDT selections in the semi-leptonic top-squark search.

Figure E.4 shows the expected and observed mass exclusions of the semi-leptonic top-squark search for the four top-squark decay modes considered.

For the T2tt decay mode, depending upon the top-squark mass and the mass splitting between the top-squark and the LSP, the semi-leptonic top-squark search can probe top-squark masses up to 700 GeV and LSP masses up to 260 GeV.

As with the dileptonic analysis, the final results of semi-leptonic top-squark search for the T2bw decay mode depend strongly upon the value of the chargino mass-splitting parameter. In contrast to the dileptonic cut-and-count analysis, Figs. 6.38b through 6.38d, the relative loss in sensitivity as x gets lower is not nearly as large. This is primarily due to two effects.

The first of these is the typically lower lepton p_{T} in the low-x T2bw decay modes, c.f. Fig. 6.27a. Because the semi-leptonic top-squark search only requires one lepton, it has a higher preselection efficiency for the low-x T2bw decay modes. even with a slightly higher lepton p_{T} requirement ($p_{\mathrm{T}} > 30$ GeV).

The second, which is perhaps more straightforward, is that the semi-leptonic top-squark analysis's use of BDTs strongly improves the signal efficiency in regions where, nominally, a simple cut-and-count approach would suffer. For example, in the regions where $m_{\widetilde{\chi}_1^{\pm}} - m_{\widetilde{\chi}_1^0} < m_{\mathrm{W}}$, i.e. where the W is produced off-shellrecovers the signal efficiency for relative to the dilepton cut-and-count approaches (i.e. the dileptonic top-squark search) will tend to, in general, have notably lower statistical sensi of the dilepton analysis lisemi-leptonic analysis utilized a

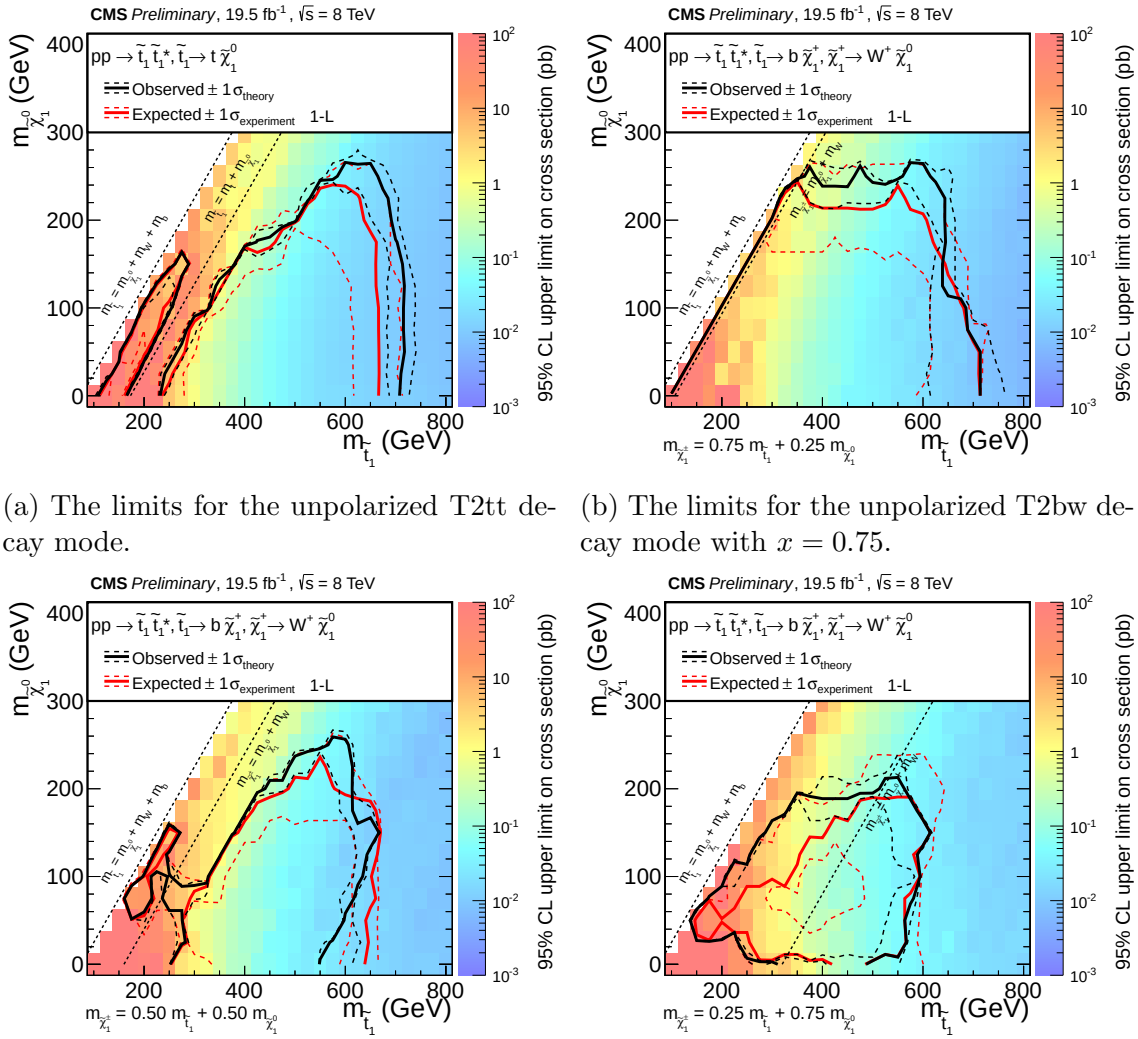

(a) The limits for the unpolarized T2tt decay mode.

(b) The limits for the unpolarized T2bw decay mode with $x = 0.75$.

(c) The limits for the unpolarized T2bw decay mode with $x = 0.50$.

(d) The limits for the unpolarized T2bw decay mode with $x = 0.25$.

Figure E.4: 2D $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ maps of the median expected 95% CL upper-limit on the cross section for top-squark pair-production for the single lepton analysis. Overlaid on top of these maps are contours denoting the regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane where top-squark pair-production has been excluded at the 95% CL. The two solid contour lines were made using the median expected cross-section limit (red) and the observed cross-section limit (black). The dashed red lines represent the exclusion region when using the 0.16 and 0.84 percentile expected limits, while the dashed black lines represent the exclusion region when varying the reference top-squark pair-production cross section within its theoretical uncertainties. A small amount of kernel-based smoothing, Appendix H.2.3, has been applied in order to aid the visual interpretation of the results.

Given that the dilepton analysis requires two leptons to pass the selection requirements, this hurts the preselection efficiency of the dilepton analysis relative to the semi-leptonic top-squark search. while First, for the low x T2bw top-squark decay modes, the lepton p_{T}, and correspondingly tends to be much lower, c.f. Fig. 6.27a. Because the dileptonic search requires two leptons to pass selection requirements, this hurts the efficiency In the $x = 0.75$ scenario, the semi-leptonic analysis can probe top-squark masses up to 475 GeV and $\widetilde{\chi}^0$ masses up to 150 GeV. For the $x = 0.50$ scenario, the analysis only probes as small region around $m_{\widetilde{\mathrm{t}}_1} = 350$ GeV, $m_{\widetilde{\chi}_1^0} = 75$ GeV. Our analysis has no observed sensitivity to the $x = 0.25$ scenario.

Appendix F: Further explanation of the "Tight-to-Loose" Method

In this section we provide a more detailed description of the "Tight-to-Loose" method that was introduced in Section 6.3.3 as a method for estimating the contribution of misidentified leptons to our relevant signal regions.

To explain the method, first consider two nominal lepton selection criteria, a "tight" selection with relatively strict selection criteria that is designed to primarily select for "prompt" leptons – leptons coming directly from the decay of a W boson or Z boson — and a "loose" selection, with relaxed selection criteria relative to the tight selection. It should be noted that in a given event, the set of leptons that pass the "loose" selection is a strict superset of those that pass the "tight" selection — i.e. all "tight" leptons are by definition, also "loose" leptons.

F.1 The "Tight-to-Loose" method in 1-lepton events

Now, imagine we are performing an analysis looking for events with only a single "prompt" lepton e.g. a W+jets cross-section measurement. One of the notable backgrounds for this kind of analysis would be the contribution from non-prompt leptons from QCD multijet events that were misidentified as prompt leptons — where the misidentified lepton could be a real lepton embedded in a heavy-flavor

hadron decay or a more traditional, "fake lepton" such as a jet reconstructed incorrectly as an electron. As analysts, we would likely be interested in estimating the total contribution of QCD multijet events to a final, "tight" lepton selection sample.

We could first perform the usual suite of object selections/cuts, except using the **"loose"** lepton selection in place of the tight lepton selection. Doing this, we will end up with N_ℓ events. Of these N_ℓ events, N_p will be *prompt loose lepton* events – events where the reconstructed loose lepton is a genuine, prompt lepton. while N_f will be *misidentified non-prompt, loose lepton* events — events where the reconstructed loose lepton is coming from a misidentified non-prompt lepton (defined before).

We cannot directly measure either N_p or N_f. However, N_p and N_f are directly related to the number of events where there is one (no) loose lepton candidate that passes the stricter tight selection, N_{t1} (N_{t0}). Using p to represent the prompt rate, the efficiency with which a prompt loose lepton also passes the tight lepton selection, and f to represent misidentification rate (or "fake" rate), the efficiency with which a non-prompt loose lepton is misidentified because it also passes the tight lepton selection, we can display these relations analytically,

$$\begin{aligned} N_\ell &= N_p + N_f = N_{t0} + N_{t1} \\ N_{t0} &= (1-p)N_p + (1-f)N_f, \\ N_{t1} &= pN_p + fN_f, \end{aligned} \tag{F.1}$$

The relations between N_p/N_f and N_{t0}/N_{t1} in Eq. (F.1) can be inverted,

$$\begin{aligned} N_p &= \frac{1}{p-f}\left[(1-f)N_{t1} - fN_{t0}\right] \\ N_f &= \frac{1}{p-f}\left[pN_{t0} - (1-p)N_{t1}\right] \end{aligned} \tag{F.2}$$

From Eq. (F.2) we can subsequently (and simply) determine the total number of misidentified leptons passing our tight selection criteria (e.g. the estimate of the misidentified lepton contribution), $N_f^{pass} = fN_f$.

Using the power of hindsight, we could instead rewrite p and f into a different (ultimately more convenient) form, $\eta = \frac{1-p}{p}$, $\epsilon = \frac{f}{1-f}$.

Side note: this form for p and f comes from direct relations between N_p^{pass}/N_p^{fail} and N_f^{pass}/N_f^{fail},

$$\begin{aligned} N_f^{pass} &= fN_f = \frac{f}{1-f}N_f^{fail} = \epsilon N_f^{fail} \\ N_p^{fail} &= (1-p)N_p = \frac{1-p}{p}N_p^{pass} = \eta N_p^{pass} \end{aligned} \tag{F.3}$$

Rewriting Eq. (F.2) in terms of η and ϵ,

$$\begin{aligned} N_p^{pass} &= p\frac{1-f}{p-f}\left[N_{t1} - \frac{f}{1-f}N_{t0}\right] = \frac{1}{1-\epsilon\eta}\left[N_{t1} - \epsilon N_{t0}\right] \\ N_f^{pass} &= f\frac{p}{p-f}\left[N_{t0} - \frac{1-p}{p}N_{t1}\right] = \frac{\epsilon}{1-\epsilon\eta}\left[N_{t0} - \eta N_{t1}\right] \end{aligned} \tag{F.4}$$

F.1.1 Possible Systematic Biases

The results derived above rely intrinsically on the assumption that p and f are accurate representations of the lepton prompt and fake rates. However, in practice, we do not know p and f *a priori*, and thus have to measure them (likely in two different control samples).

Because, in practice, we have to measure p and f in separate samples from the

one where we apply them, we possibly have a systematic bias from this extrapolation. To minimize the bias coming from extrapolating the measured p and f into our signal sample, a reasonable prescription is to measure the dependence of p and f on relevant object quantities, presumably in a binned sense; the two most straightforward (and probably most important) variables are the lepton's p_{T} and η.

F.1.2 The Estimated Yield of Misidentified Leptons in 1-lepton Events

To estimate the final expected contribution from misidentified leptons, one would then calculate N_f^{pass}, á la Eq. (F.4), in each bin of relevant quantities (e.g. lepton p_{T} and η) and then sum up all of these individual estimated N_f^{pass} to arrive at a final estimated number.

Alternatively, a mathematically equivalent way to arrive at this answer is to iterate over the sample of loose lepton events and for each event, calculate the p and f for the selected loose lepton (or, alternatively, calculate the η and ϵ for the selected loose lepton). Then, for each event, multiply said event's pre-existing event weight (which is presumably 1 if we're dealing with collision data) by an additional weight based on whether the selected loose lepton passes or fails the tight lepton selection.

$$\begin{aligned} w^{fail} &= \frac{\epsilon(p_{\mathrm{T}},\, \eta^{\ell})}{1-\epsilon(p_{\mathrm{T}},\, \eta^{\ell})\eta(p_{\mathrm{T}},\, \eta^{\ell})} \\ w^{pass} &= \frac{-\epsilon(p_{\mathrm{T}},\, \eta^{\ell})\eta(p_{\mathrm{T}},\, \eta^{\ell})}{1-\epsilon(p_{\mathrm{T}},\, \eta^{\ell})\eta(p_{\mathrm{T}},\, \eta^{\ell})}. \end{aligned} \tag{F.5}$$

The advantage of this second method, is that one is also free to generate the expected contribution of misidentified leptons to the distribution of various event kinematic

variables simply by creating the distribution from the loose lepton sample, weighted by the calculated weights from Eq. (F.5).

F.2 Misidentified Leptons in Dilepton Events

We turn now, to considering the case where we have a dilepton analysis, e.g. the top-squark search described in this dissertation.

First, a quick description of notation: $N_{\ell\ell}$ represents the total number of selected dilepton events; for all other N quantities, the first subscript index represents the leading (p_{T}) lepton while the second subscript index represents the sub-leading lepton — e.g. N_{pf} represents the number of events where the leading (p_{T}) lepton is a prompt lepton and the sub-leading lepton is a misidentified lepton.

With this notation in mind, we can immediately write the relations, ła Eq. (F.1) for the 1-lepton case, between $N_{\ell\ell}$ and its relevant demarcations,

$$
\begin{aligned}
N_{\ell\ell} &= N_{pp} + N_{fp} + N_{pf} + N_{ff} = N_{t11} + N_{t01} + N_{t10} + N_{t00} \\
N_{t11} &= p_1 p_2 N_{pp} + p_1 f_2 N_{pf} + f_1 p_2 N_{fp} + f_1 f_2 N_{ff} \\
N_{t10} &= p_1(1-p_2) N_{pp} + p_1(1-f_2) N_{pf} + f_1(1-p_2) N_{fp} + f_1(1-f_2) N_{ff} \\
N_{t01} &= (1-p_1) p_2 N_{pp} + (1-p_1) f_2 N_{pf} + (1-f_1) p_2 N_{fp} + (1-f_1) f_2 N_{ff} \\
N_{t00} &= (1-p_1)(1-p_2) N_{pp} + (1-p_1)(1-f_2) N_{pf} + (1-f_1)(1-p_2) N_{fp} + (1-f_1)(1-f_2) N_{ff}
\end{aligned}
\tag{F.6}
$$

We can then write these linear relations between the various N in terms of a matrix relation. Rather than do this, we will just quote the inverted matrix relation instead

— i.e. relating N_{pp} to N_{t00}, N_{t01},... etc.,

$$\begin{pmatrix} N_{pp} \\ N_{pf} \\ N_{fp} \\ N_{ff} \end{pmatrix} = \frac{A}{(p_1 - f_1)(p_2 - f_2)} \begin{pmatrix} N_{t00} \\ N_{t10} \\ N_{t01} \\ N_{t11} \end{pmatrix} \tag{F.7}$$

In the 1-lepton case, at this point we were able to write down analytic expressions — Eq. (F.4) — for the two possible kinds of 1-lepton events passing our tight selection. We can write down an analogous expression for the 2-lepton case, except we now have four possible kinds of 2-lepton events passing our tight selection (because the leading and sub-leading tight leptons could each respectively be prompt or misidentified leptons),

$$\begin{aligned} N_{pp}^{pass} &= \frac{1}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \left[N_{t11} - \epsilon_2 N_{t10} - \epsilon_1 N_{t01} + \epsilon_1\epsilon_2 N_{t00}\right] \\ N_{ff}^{pass} &= \frac{1}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \left[N_{t00} - \eta_1 N_{t10} - \eta_2 N_{t01} + \eta_1\eta_2 N_{t00}\right] \\ N_{pf}^{pass} &= \frac{-\epsilon_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \left[\eta_2 N_{t11} - N_{t10} - \epsilon_1\eta_2 N_{t01} + \epsilon_1 N_{t00}\right] \\ N_{fp}^{pass} &= \frac{-\epsilon_1}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \left[\eta_1 N_{t11} - N_{t10} - \epsilon_2\eta_1 N_{t01} + \epsilon_2 N_{t00}\right] \end{aligned} \tag{F.8}$$

The events where at least one of the two tight leptons is a misidentified lepton comprise the total contribution of misidentified leptons to our final signal sample. So, we would derive our final estimate of our misidentified lepton background by calculating and adding together N_{ff}^{pass}, N_{pf}^{pass}, and N_{fp}^{pass} from Eq. (F.8).

F.2.1 Possible Systematic Biases in Dilepton Events

Similar to the discussion in Section F.1.1 above, the main systematic bias of this estimation method comes in from the measurement of p and f. The solution to minimize this systematic bias is the same as before.

F.2.2 The Estimated Yield of Misidentified Leptons in Dilepton Events

In analogy with the discussion in Section F.1.2, one can derive weights to assign each event in a loose-loose dilepton sample based on the calculated p and f for each selected loose lepton and whether the leading and sub-leading loose leptons in each event also pass the tight lepton selection. This was already shown in Eq. (6.14), but for posterity, we include it here as well,

$$\begin{aligned}
\text{Pass} - \text{Pass}: \quad w^{\text{pass pass}} &= -\frac{\epsilon_1\eta_1 + \epsilon_2\eta_2 - \epsilon_1\epsilon_2\eta_1\eta_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Fail} - \text{Fail}: \quad w^{\text{fail fail}} &= -\frac{\epsilon_1\epsilon_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Pass} - \text{Fail}: \quad w^{\text{pass fail}} &= \frac{\epsilon_2}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)} \\
\text{Fail} - \text{Pass}: \quad w^{\text{fail pass}} &= \frac{\epsilon_1}{(1-\epsilon_1\eta_1)(1-\epsilon_2\eta_2)}
\end{aligned} \tag{F.9}$$

Appendix G: $\mathrm{CL_s}$-based Limit-Setting and Null p-values

The end result for the various analyses described in this dissertation are statistical limits and null p-values. The calculations to derive these results utilize the machinery and methodology originally designed by the ATLAS and CMS Collaborations for the purpose of combining searches for the Higgs boson across multiple channels [175, 176].

The basis of the limit-setting methodology is to construct the $\mathrm{CL_s}$ test criterion [177, 178] using likelihood ratios of signal and background hypotheses. When calculating null p-values, likelihood ratios are also utilized to quantify any observed excesses in the data. Commonly, these likelihood ratios are constructed in a fully frequentist fashion, but asymptotic formulae [168] are sometimes used for approximate answers. We provide below a general prescription of the methodology below.

Following the notation of Ref. [176], we denote the expected signal and background yields as s and b respectively. If there are multiple bins for the given analysis or analyses in question — e.g. the yields of the separate 1-lepton and 2-lepton analyses, Section 7.1, or the individual bins of the 2-lepton shape analysis, Section 7.2.4 — the subscript index i designates the individual "bin".

G.1 Systematic Uncertainties

Systematic uncertainties represent the intrinsic, quantified sources of uncertainty in the estimates of the expected signal and background. This includes things such as the uncertainty on the integrated luminosity and theoretical cross-sections; detector-modeling based uncertainties, such as the calibration of reconstructed object energy scales; and the intrinsic statistical uncertainty that arises in estimating the expected yield of signal and background events, whether via data-driven techniques or through the use of finite-precision simulation samples.

Following the notation of Ref. [176], we term systematic uncertainties as "nuisance parameters"[1] and designate them mathematically with the general symbol θ. Consequently, the expected signal and background, including their yield and overall shape, can be thought of as being functions of these θ parameters: $s(\theta)$ and $b(\theta)$, where it should be understood that each separate, independent, systematic uncertainty is designated by its own θ.

In order to quantify the effect of systematic uncertainties on the signal and background expectations, we parameterize the θ such that their probability density function (*p.d.f.*) is a standard normal distribution (Gaussian with mean zero and unit variance). Because in these statistical calculations we primarily deal with positive-definite quantities, we then take an exponential parameterization e^{θ}, so that

[1] This is a rather (humorously) apt term, given how much effort it often requires to properly characterize and quantify the relevant systematic uncertainties in a given analysis.

the resulting *p.d.f.* for θ is a *log-normal* distribution,

$$p(\theta|\tilde{\theta}) = \frac{1}{\theta\sqrt{2\pi}\ln\kappa_\theta}\exp\left(-\frac{\left(\ln(\theta-\tilde{\theta})\right)^2}{2(\ln\kappa_\theta)^2}\right), \tag{G.1}$$

where $\tilde{\theta}$ is the mean value of the Gaussian and the effective uncertainty is dictated by $\ln\kappa_\theta$ (e.g. $\kappa_\theta = 1.10$ signifies that 32% of the time, θ will be within 10% of $\tilde{\theta}$).

These *p.d.f.s* for the systematic uncertainties are then reinterpreted as Bayesian posteriors for θ arising from some hypothetical measurements $\tilde{\theta}$, a la Bayes theorem,

$$p(\theta|\tilde{\theta}) \sim p(\tilde{\theta}|\theta)\cdot\pi_\theta(\theta), \tag{G.2}$$

where π_θ are flat priors for the parameters θ. This reinterpretation enables us to handle the systematic uncertainties in a fully frequentist fashion, as the *p.d.f.s* $p(\tilde{\theta}|\theta)$ for these hypothetical auxiliary measurements can be used to constrain the likelihood of the primary measurement. As well, we can use $p(\tilde{\theta}|\theta)$ to generate sampling distributions of our test statistics (discussed later) following a pure frequentist procedure.

G.2 Setting Limits on the Signal Strength

When we do not observe any clear excess in the collision data relative to the expected background in our respective signal regions, we follow the customary procedure and set 95% CL limits on μ, the *signal strength parameter*,

$$\mu = \frac{\sigma^{\text{obs.}}_{\tilde{\text{t}}_1\tilde{\text{t}}_1*}}{\sigma^{\text{calc.}}_{\tilde{\text{t}}_1\tilde{\text{t}}_1*}},$$

as a function of the decay mode considered and nominal SUSY particle masses/mass splittings, where σ_{obs} represents the observed cross-section limit we can set using our

data and $\sigma_{\tilde{t}_1\tilde{t}_1}$ is the nominal 8 TeV pair production cross-section for top-squarks. In this representation, scenarios – e.g. $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ points, top-squark decay modes, etc. – where the calculated upper limit on μ is less than 1 are excluded at a 95% CL.

These limits are calculated in a modified-frequentist fashion. We will explain what this means below.

G.2.1 Construction of Likelihood Functions

The first step of limit-setting is to construct likelihood functions, $\mathcal{L}(\text{data}|\mu,\, \theta)$, which is typically written as a product over the Poisson probabilities for each bin multiplied by the probability densities for the nuisance parameters,

$$\mathcal{L}(\text{data}|\mu,\, \theta) = \left[\prod_i \frac{[\mu s_i(\theta) + b_i(\theta)]^{n_i}}{n_i} e^{-[\mu s_i(\theta)+b_i(\theta)]}\right] \cdot \prod_j p(\tilde{\theta}_j|\theta_j), \tag{G.3}$$

where "data" can refer to the actual collision data or pseudo-data generated from one of the hypotheses being tested.

G.2.2 Constructing a Test Statistic

We then construct a test-statistic that facilitates comparisons of the compatibility of the "data" with the two respective hypotheses being tested,

$$\tilde{q}_\mu \;=\; -2\ln\frac{\mathcal{L}(\text{data}|\mu,\, \hat{\theta}_\mu^{\text{obs}})}{\mathcal{L}(\text{data}|\hat{\mu},\, \hat{\theta})}, \quad \text{with the constraint } 0 \le \hat{\mu} \le \mu \tag{G.4}$$

where $\hat{\theta}_\mu^{\text{obs}}$ is the maximum likelihood estimator of θ given a fixed μ and "data", and $\hat{\mu}$ and $\hat{\theta}$ are the values of μ and θ when the likelihood has its global maximum. In the construction of $\tilde{q}_\mu$, fixing the nuisance parameters to their maximum-likelihood

estimates has been found to have good coverage properties [179].

It is worth noting that physics considerations motivate the additional constraint on $\hat{\mu}$ – namely, from $0 \leq \hat{\mu}$, downward fluctuations in the data cannot imply a negative signal and, from $\hat{\mu} \leq \mu$, upward fluctuations in the data do not weaken the strength of the μ-strength signal hypothesis.

We then calculate values of the nuisance parameters that best describe the data – by maximizing the likelihood in Eq. (G.3) – under the B-only and μS+B hypotheses; we use $\hat{\theta}_0^{\text{obs}}$ and $\hat{\theta}_\mu^{\text{obs}}$ to specify these likelihood maximizing values for the two respective hypotheses.

We then construct the effective distribution of the test statistic under the two hypotheses by generating ensembles of pseudo-data and calculating the test statistic for each member of this ensemble. For each of these members, before generating the pseudo-data, new values of the nuisance parameters are generated according to their respective *p.d.f.s*, $p(\tilde{\theta}|\theta)$, where the input θ are the likelihood maximizing values, e.g. $\hat{\theta}_0^{\text{obs}}$ and $\hat{\theta}_\mu^{\text{obs}}$. Figure G.1a shows example distributions $f\left(q_\mu \,|\, b(\hat{\theta}_0^{\text{obs}})\right)$ and $f\left(q_\mu \,|\, \mu s(\hat{\theta}_\mu^{\text{obs}}) + b(\hat{\theta}_\mu^{\text{obs}})\right)$, along with an observed $\tilde{q}_\mu^{\text{obs}}$.

With these example distributions in hand, we can calculate a p-value associated with each distribution,

$$\begin{aligned} \text{CL}_{\text{s+b}}(\mu) = p_\mu = & \ \mathcal{P}\left(\tilde{q}_\mu \geq \tilde{q}_\mu^{\text{obs}} \,|\, \mu s(\hat{\theta}_\mu^{\text{obs}}) + b(\hat{\theta}_\mu^{\text{obs}})\right) \\ = & \int_{\tilde{q}_\mu^{\text{obs}}}^{\infty} f\left(q_\mu \,|\, \mu s(\hat{\theta}_\mu^{\text{obs}}) + b(\hat{\theta}_\mu^{\text{obs}})\right) d\tilde{q}_\mu \,, \end{aligned} \tag{G.5}$$

$$
\begin{aligned}
\mathrm{CL_b} = 1 - p_b = & \ \mathcal{P}\left(\tilde{q}_\mu \geq q_\mu^{obs} \,|\, b(\hat{\theta}_0^{obs})\right) \\
= & \int_{\tilde{q}_\mu^{\mathrm{obs}}}^{\infty} f\left(q_\mu \,|\, b(\hat{\theta}_0^{\mathrm{obs}})\right) d\tilde{q}_\mu \,,
\end{aligned}
\tag{G.6}
$$

and then calculate $\mathrm{CL_s}$ as the ratio of these two,

$$
\mathrm{CL_s}(\mu) = \frac{\mathrm{CL_{s+b}}}{\mathrm{CL_b}} \,. \tag{G.7}
$$

In order to quote the $1 - \alpha$ (e.g. 95%) confidence level (CL) upper-limit on μ, we scan through values of μ until we find the first value that yields $\mathrm{CL_s} \leq \alpha$ (e.g. 5%).

We note that this prescription is strictly more statistically conservative than the commonly-used alternative, $\mathrm{CL_{s+b}} \leq \alpha$. In particular, the $\mathrm{CL_s}$-based method becomes much more statistically conservative for strong downwards fluctuations of the background, as those correlate to small values of $\mathrm{CL_b}$.

G.2.3 Calculation of Expected Limits

A common type of limit calculated for CMS analyses is the *expected* median upper-limit on μ for the *B-only* hypothesis. This type of limit gives an analyst a quantitative idea the expected statistical power of the analysis under the assumption that there is no signal.

Comparisons between the expected median upper-limit and the observed upper-limit can directly point to "interesting" regions of the parameter space being explored; for example, in the case of the Higgs boson searches, the observed upper-limits on the Higgs production for $m_{\mathrm{H}} = 125$ GeV were notably weaker relative to the median expected upper-limits, implying the possible presence of signal in the

data. Alternatively, comparisons between the expected median upper-limit for variants of an analysis (e.g. the same search performed at ATLAS and CMS) can give indications on which analysis is "stronger" in a statistical power sense.

To calculate an expected median upper-limit, we generate a large number of samples of pseudo-data from the background-only hypothesis. For each pseudo-data sample, values of the nuisance parameters are sampled from their respective *p.d.f.s*, $p(\tilde{\theta}|\theta)$. Similar to the observed limits, final calculations involve utilizing the post-fit nuisance parameters, i.e. $\hat{\theta}_0^{\text{obs}}$, as the input θ for the *p.d.f.s*.

For each generated set of pseudo-data, we can then calculate the CL_s parameter, record the result, and combine the results to create an overall probability distribution for the CL_s parameter for the analysis in question. We can integrate this probability distribution to create the cumulative distribution function (CDF) for the analysis CL_s. The median expected upper-limit is the μ upper-limit value where the CDF first crosses 0.5. Analogously, one can generate the $\pm 1\sigma$ expected upper-limit band using the μ region, $0.84 > \text{CDF} > 0.16$ (or, for the $\pm 2\sigma$ band, the μ region corresponding to $0.025 > \text{CDF} > 0.975$).

Figure G.2 shows an example CL_s *p.d.f.* and CDF and the subsequent median, $\pm 1\sigma$, and $+2\sigma$ upper limits on μ (there were not enough events generated in the relevant region to exactly pin down the μ value for the -2σ band edge).

G.2.3.1 Data-blind Expected Limits

When performing analysis optimizations that, to avoid bias, should be "blind" with respect to the data, we utilize the pre-fit nuisance parameters, which are the θ values prior to the likelihood maximization. This is known as the *data-blind* median expected limit. An example of where this was used in this dissertation is in Section 6.8.2, where we optimized the choice of $M_{\mathrm{T2}}(\ell\ell)$ threshold based upon the calculated data-blind median expected limits for each of the 5 considered thresholds.

G.3 Quantifying Observed Excesses

In the shape analysis described in Chapter 7, we observed a number of excesses with respect to the SM expectations. We quantify these excesses in terms of a local, null p-value, i.e. the probability for the background to fluctuate and give an excess of events as "signal-like" or more so than the observed one. In order to calculate these null p-values, we construct a test statistic, q_0,

$$q_0 \;=\; -2\ln\frac{\mathcal{L}(\mathrm{data}|0,\,\hat{\theta}_0^{\mathrm{obs}})}{\mathcal{L}(\mathrm{data}|\hat{\mu},\,\hat{\theta})}, \quad \text{with the constraint } \hat{\mu} \geq 0. \tag{G.8}$$

In the denominator, in the alternative, signal + background hypothesis, the signal strength is floated so as to maximize the likelihood. The motivation for this choice is that its distribution asymptotically (in the large sample-size limit) converges to half of a χ^2 distribution with one degree of freedom plus 0.5 * $\delta(q_0)$ [168].

With this definition of q_0, there are three general qualitative cases:

- There is a deficit in the data relative to the background-only hypothesis: $q_0 =$

0.

- There is an excess in the data (of any size), but it does not look very much like the expected signal: $q_0 > 0$ but is "small".

- There is a "large", **signal-like** excess in the data: $q_0 > 0$ and is "large".

Observed excesses in the data are transformed into values for p_0, the "null p-value with respect to the signal hypothesis," by calculating the probability to obtain a value of q_0 as large as the one observed in the data, $q_0^{\rm obs}$,

$$\begin{aligned} p_0 &= \mathcal{P}\left(q_0 \geq q_0^{\rm obs} \,|\, b(\hat{\theta_0^{\rm obs}}) \right) \\ &= \int_{q_0^{\rm obs}}^{\infty} f\left(q_0 \,|\, b(\hat{\theta}_0^{\rm obs}) \right) dq_0 \,, \end{aligned} \tag{G.9}$$

The distribution $f\left(q_0 \,|\, b(\hat{\theta}_0^{\rm obs}) \right)$ is calculated either by an analogous method as above for $\tilde{q}_\mu$, or, in the asymptotic regime, by using the analytic approximation described above. An example distribution of q_0 is shown in Fig. G.1b. The solid blue line represents a χ^2 with one degree of freedom. As can be seen, this approximation is quite accurate, especially for large values of q_0.

The statistical "significance", i.e. the commonly quoted $n\sigma$, is calculated via a one-sided Gaussian tail integral,

$$p_0 = \int_{\rm Z}^{\infty} \frac{1}{\sqrt{2\pi}} \exp(-x^2/2) dx. \tag{G.10}$$

or, in the asymptotic regime, by using,

$$\rm Z = \sqrt{q_0^{\rm obs}}. \tag{G.11}$$

G.3.1 Quantifying the *Global* Significance of Observed Excesses: the "Look-Elsewhere" Effect

In the top-squark searches described in this dissertation, the masses of the SUSY particles, e.g. $m_{\tilde{t}_1}$ and $m_{\tilde{\chi}_1^0}$, are parameters that are undefined for the background-only hypothesis, unlike the signal strength μ which is set to 0. Consequently, the conditions of Wilks' theorem [180] are not applicable, and one cannot construct a unique test statistic that can account for all possible signals while still maintaining asymptotic χ^2 behavior.

Thus, although the observed excesses in the 3D M_{T2} shape yielded large significances for multiple mass points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, the quoted *local* significances for each of these mass points overestimate the true, *global* significances of the observed excesses. In particular, they fail to account for the probability for the background to fluctuate in a coherent fashion so as to mimic the signal at any of the mass points we scanned through in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane; this is colloquially known as the "look-elsewhere" effect.

This problem has been studied in detail [181,182,183], and solutions to correct for this effect exist for an arbitrary number of additional signal-only parameters. We provide here a quick overview of one of these solutions, summarized from Ref. [183].

For notation purposes, we denote the set of additional signal-only parameters as θ'. As an example, for the $\tilde{t}_1 \to t\tilde{\chi}_1^0$ decay mode of the top-squark, these additional parameters would be the unknown masses, $m_{\tilde{t}_1}$ and $m_{\tilde{\chi}_1^0}$. The test statistic q_0 can then be thought of as a function of these additional parameters, $q_0(\theta')$. These

additional parameters can be mapped onto a $\mathcal{D}$-dimensional manifold, such as the $(m_{\widetilde{\mathrm{t}}_1}, m_{\widetilde{\chi}_1^0})$ plane for the aforementioned $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode. For any fixed point θ' on this manifold, the conditions of Wilks' theorem apply, and the $q_0(\theta')$ at that fixed point will follow a χ^2 distribution with one degree of freedom, as per the discussion above regarding Eq. (G.8).

Viewed in the context of the entire manifold, $q_0(\theta')$ is a χ^2 random field, a set of random variables that can be continuously mapped onto a single manifold. In this context, then, the *global* significance of an observation can be characterized in terms of the *excursion probability* — the probability for the maximum of $q_0(\theta')$ to be above an arbitrary threshold.

Closed-form expressions exist for the excursion probabilities for large values of the threshold. These expressions are related to the topological structures of the manifold and the intrinsic covariance structure of the random field. They capitalize on the fact that, asymptotically, the global p-value is given by,

$$E\left[\phi(A_{\sigma_{\mathrm{local}}})\right] \approx p_{\mathrm{global}} \tag{G.12}$$

where $E\left[\phi(A_{\sigma_{\mathrm{local}}})\right]$ is the expected value of $\phi(A_{\sigma_{\mathrm{local}}})$, the Euler characteristic of $A_{\sigma_{\mathrm{local}}}$, the observed manifold when requiring the local points of $q_0(\theta')$ to be above a certain threshold significance, $\sigma_{\mathrm{local}} > \sigma_{\mathrm{test}}$.

For a search with one signal-only parameter, e.g. the mass of the Higgs boson in the Higgs-boson search, the Euler characteristic of $A_{\sigma_{\mathrm{local}}}$ is just the number of disjoint regions that are above the threshold significance, $N(\sigma_{\mathrm{local}} > \sigma_{\mathrm{test}})$. The formula for the global p-value of the observed excess at a particular local point in

$q_0(\theta')$ is given by [182],

$$p_{\text{global}} \sim p_{\text{local}} + N(\sigma_{\text{local}} > \sigma_{\text{test}}) e^{-\left(\sigma^2_{\text{local}} - \sigma^2_{\text{test}}\right)/2}. \tag{G.13}$$

For a search with two signal-only parameters, e.g. the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}^0_1})$ plane for the top-squark searches described in this dissertation, the Euler characteristic of $A_{\sigma_{\text{local}}}$ is just the number of disjoint regions of the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}^0_1})$ plane that are above the threshold significance, minus the number of "holes" in these regions. The formula for the global p-value is slightly more complicated [183],

$$E\left[\phi(A_{\sigma_{\text{local}}})\right] = p_{\text{local}} + e^{-\sigma^2_{\text{local}}/2}\left(c_0 + c_1 \sigma_{\text{local}}\right). \tag{G.14}$$

The above equation is the one we utilized in Section 7.2.7.1 to estimate the global significance of the observed excesses in the 3D M_{T2} shape.

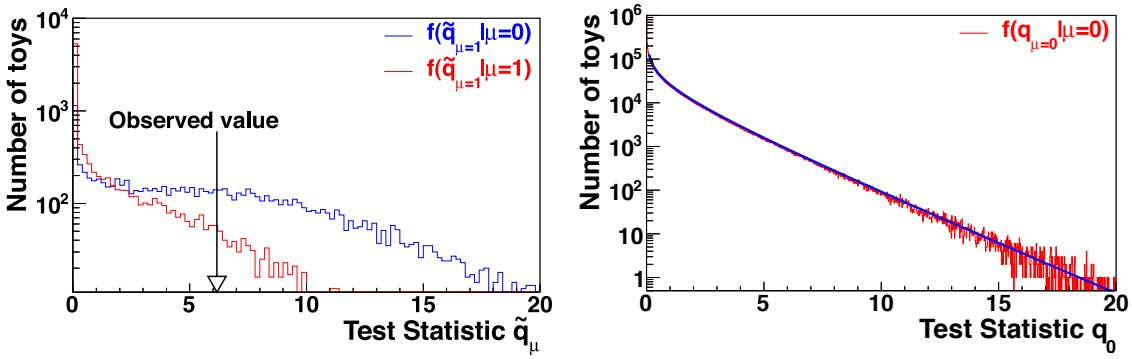

(a) The distribution of the $\tilde{q}_\mu$ test statistic, Eq. (G.4) for pseudo-data generated under the *signal + background* and *background-only* hypotheses. The observed value of $\tilde{q}_\mu$ is indicated by the arrow. Reprinted from Fig. 1 of [175].

(b) The distribution of the q_0 test statistic, Eq. (G.8) for pseudo-data generated using the *background-only* hypothesis. The solid blue line represents a χ^2 distribution with one degree of freedom. Reprinted from Fig. 3 of [175].

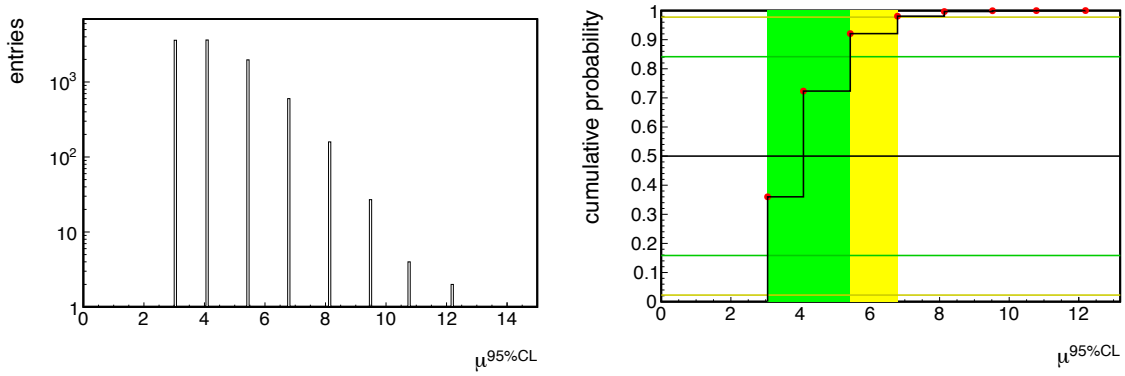

(a) (b)

Figure G.2: (Left) An example *p.d.f.* for the 95% CL upper-limit on μ using pseudo-data generated from a toy analysis ($s = 1$, $b = 1$, no systematic errors – note for pseudo-data generation, $s = 0$). (Right) The cumulative distribution function constructed by integrating the *p.d.f.* from the left. The colored horizontal lines represent the 2.5th, 16th, 50th, 84th, and 97.5th percentiles. The μ values where the CDF crosses these thresholds define the median, $\pm 1\sigma$ (68%), and $\pm 2\sigma$ (95%) expected upper limits on μ for the *B-only* hypothesis. Reprinted from Fig. 2 of [175].

Appendix H: Additional Details on Statistical Techniques

In this appendix, we provide additional details on some of the statistical techniques that were utilized in the analyses described in Chapters 6 and 7.

H.1 Defining the "Pull" of a Fit

"Pull" distributions are a common analysis item in particle physics. We provide here a very quick overview of them. Readers that are interested in additional information are encouraged to refer to the excellent review in Ref. [184].

Definition of a "Pull"

For a random variable X that is Gaussian distributed with parameters (μ, σ), the "pull",

$$g = \frac{X - \mu}{\sigma} \tag{H.1}$$

will be distributed as a standard Gaussian with mean zero and unit width.

Pulls are commonly used in particle physics data analysis to check the quality of fits to experimental data or the validity of the value for a parameter extracted from a fit to data. When the pull distribution is properly constructed, any deviations of

the mean of the pull distribution from zero indicate a possible bias in the fit, while deviations of the standard deviation of the pull distribution from unity indicate either under- or over-coverage of the relevant uncertainties.

H.2 Dealing with Stochastic Data

Experimental particle physics is a field that is filled with stochastic data — data that is being sampled from some "true" distribution that is subject to stochastic fluctuations. It can often be the case that the analyst is interested in information contained in the underlying true distribution. For example,

1. In Section 6.8.3 of Chapter 6, we showed 95% CL upper-limits on the top-squark signal strength parameter μ. We, in turn, utilized these upper-limits to exclude regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane by using the individual scan points as an estimate for the regions where $\mu < 1$.

2. As a simpler example, a physicist studying the lifetime of a particular excited state of an atom might be interested in the time constant τ_0 that governs the decay of this excited state.

There are a number of techniques that can aid the analyst in extracting these parameters of the "true" underlying distribution (e.g. in Example 1, the regions of $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ where the upper limit on μ is less than 1 or τ_0 in Example 2).

One of the simplest techniques is known as parametric functional fitting, or parametric estimation. In Example 2 above, in order to extract τ_0, the hypothetical physicist can make a histogram of the observed decay times of the excited state and

then fit this histogram with an exponential distribution. The extracted value of the exponential constant, one of the *parameters* of the fitted function, is then a direct estimate of τ_0.

Parametric estimation is a commonplace technique because it is both simple and often quite robust. However, in order for it to be as accurate as possible, a necessary requirement of parametric estimation is an optimal parametric fit-function.

This requirement limits the application of parametric fits; as an illustration, in Example 1, *a-priori* there is no "reasonable" parametric form that one can assume the observed signal strengths, i.e. the "sampled data," follows. One could certainly find parametric forms that approximately fit the data — after all, arbitrarily high-order polynomials can usually fit a given dataset with reasonable "accuracy" — but this immediately raises concerns about possible systematic biases coming from the choice of parametric form. Moreover, the physical interpretation of the parametric fit might not be straightforward.

One recourse is to utilize a technique known as non-parametric functional estimation. The gist of this technique is that, rather than using the data to provide a statistical estimator for some particular parameter of the underlying distribution (e.g. τ_0 in Example 2), one instead uses the data to provide a statistical estimator of the underlying distribution itself. We now illustrate this technique through an example.

H.2.1 Non-parametric Functional Estimation using the Priestley-Chao Estimator

Consider a sample of n observations taken in an fixed interval in some one-dimensional parameter space, x,

$$y(x_i) = f(x_i) + \epsilon_i, \tag{H.2}$$

where $f(x_i)$ is some unknown, non-random true function and ϵ_i represents stochastic fluctuations, each uncorrelated with one another, with zero mean and homoscedastic (constant) variance σ_ϵ^2.

Assuming that the n measurement points x_i are equally spaced along this fixed interval, Priestley and Chao propose the following non-parametric functional estimator (NPE) for $f(x_i)$ [185],

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right) y(x_i), \tag{H.3}$$

, where $K(\cdot)$ is the *kernel* function, which satisfies the following qualities

$$\begin{aligned} \int_{-\infty}^{\infty} K(u)du &= 1 \\ \int_{-\infty}^{\infty} K(u)^2 du &= R_K < \infty \\ \int_{-\infty}^{\infty} u^2 K(u)du &= \sigma_K^2 < \infty, \end{aligned} \tag{H.4}$$

and h is the *bandwidth* of the kernel function.

Using the Priestley-Chao (PC) estimator, the best estimate for the unknown function at a point x is given by a weighted average over all the measurements y_i, based upon the proximity of their respective x_i to the chosen x.

It was shown by Priestley and Chao in their original paper [185] that asymptotically, as the number of measurements, $n \to \infty$, and the bandwidth, $h \to 0$, the bias of the PC estimator is equal to,

$$B\left[\hat{f}(x)\right] = \left(E\left[\hat{f}(x)\right] - f(x)\right) \approx \frac{h^2}{2}\frac{d^2 f}{dx^2}\sigma_K^2, \tag{H.5}$$

and correspondingly, the PC estimator's variance is, asympotically,

$$V\left[\hat{f}(x)\right] \approx \frac{\sigma_\epsilon^2}{hn} R_K. \tag{H.6}$$

Thus, the PC estimator is an asymptotically unbiased estimator of $f(x)$ and, assuming that $n \to \infty$ faster than $h \to 0$, it is also (again, asympotically) a consistent estimator of $f(x)$.

The mean-squred error (MSE) of the PC estimator is given by[1],

$$\begin{aligned} MSE\left[\hat{f}(x)\right] &= B\left[\hat{f}(x)\right]^2 + V\left[\hat{f}(x)\right] \\ &\approx \frac{h^4}{4}\left(\frac{d^2 f}{dx^2}\right)^2 \sigma_K^4 + \frac{\sigma_\epsilon^2}{hn} R_K. \end{aligned} \tag{H.7}$$

Optimal Kernel Choice

One can attempt to minimize the MSE of the PC estimator by varying the choice of kernel function $K(\cdot)$; this issue was first solved by Benedetti via calculus of variations in Ref. [186]. The optimal choice of kernel is the "Epanechnikov kernel,"

$$K(t) = \frac{3}{4}\left(1 - t^2\right) \qquad |t| < 1. \tag{H.8}$$

It turns out, however, that the loss in efficiency for the PC estimator when choosing a non-optimal kernel is rather small [187].

[1] As a general rule of thumb, when working with statistical estimators, the smaller the MSE, the better.

Optimal Bandwidth

Given a fixed kernel choice, one can instead attempt to minimize the MSE of the PC estimator by choosing an optimal bandwidth, h_{opt}. Performing this optimization leads to,

$$h_{\text{opt}}(x) = n^{-\frac{1}{5}} \left[\frac{\sigma_\epsilon^2 R_K}{\left(\frac{d^2 f}{dx^2}\right)^2 \sigma_K^4} \right]^{\frac{1}{5}} . \tag{H.9}$$

From Eq. (H.9), we find the perhaps intuitive result that the optimal choice of bandwidth for minimizing the MSE depends upon both how noisy, σ_ϵ, the data is and the relative smoothness, $\frac{d^2 f}{dx^2}$, of the function $f(x)$. If one is dealing with noisy data, then the bandwidth should be larger in order to average out the fluctuations as much as possible; however, for regions of x where the underlying true function has sharp curvature, the optimal choice of bandwidth is a small one in order to minimize the bias that comes from sampling $f(x)$ too far from the desired x point of interest. Of course, the expression in Eq. (H.9) is, in some respects, not applicable in practice as one needs to know the value of $\frac{d^2 f}{dx^2}$ (and if one knew that already, for most reasonable $f(x)$, the functional form could be derived by simple integration).

H.2.1.1 Testing the PC Estimator: Proof in Practice

In order to demonstrate the relative strength of NPEs, we performed a comparison between a parametric estimator (PE) — specifically a least-squares based functional fit — and the PC estimator by fitting each respectively to two chosen true functions, $f(x)$.

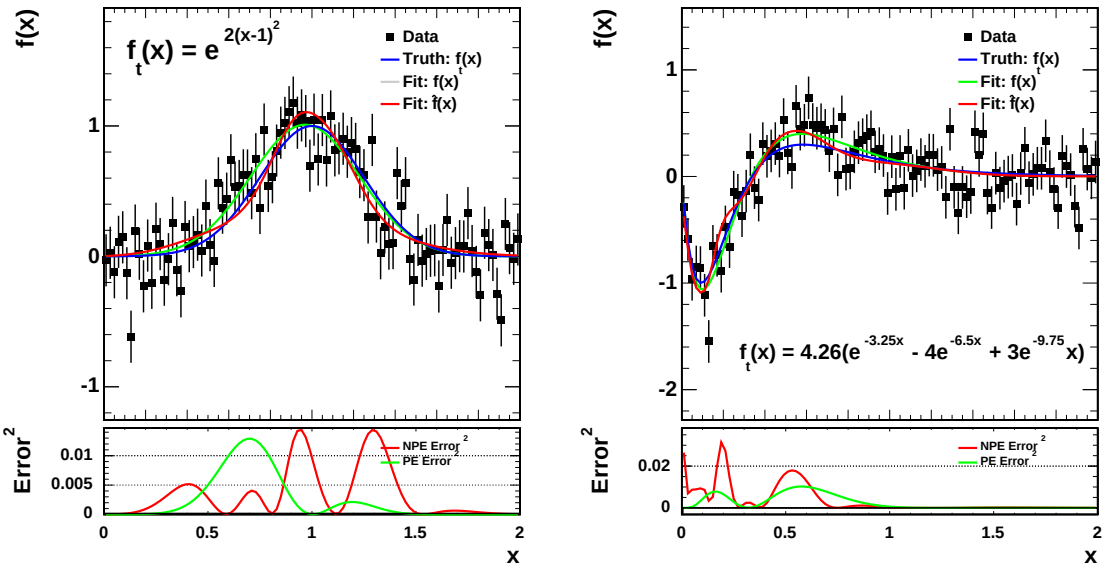

(a) Examples of parametric and non-parametric fits to a standard Gaussian function.

(b) Examples of parametric and non-parametric fits to a "triple exponential," c.f. Eq. (H.10).

Figure H.1: Examples of parametric and non-parametric fits to two separate functions. For every "data" point (black markers), the value is given by the true underlying function, $f(x)$ plus a Gaussian-distributed random variable with parameters ($\mu = 0$, $\sigma = 0.2$), where these Gaussian random variables are independently sampled for each point, c.f. Eq. (H.2). The blue solid line represents the actual functional form of $f(x)$, while the green and red solid lines represent respectively the parametric and non-parametric functional fits to the "data". For the parametric fit, the true parametric form is used. For the non-parametric fit, the Priestley-Chao estimator, Eq. (H.3), is applied to the data, using a Gaussian kernel function, Eq. (H.11), where the bandwidth parameter has been chosen to be the optimal one, as per Eq. (H.9). The bottom panels of each sub-figure show the "squared error", the squared difference between the true function and the fits of the two estimators.

The two functions used were a standard Gaussian function, $f(x) \sim e^{-x^2/c}$, and a "triple exponential,"

$$f(x) = N \sum_{i=1}^{3} A_i e^{-\lambda_i x}. \tag{H.10}$$

The function shown in Eq. (H.10) has been used before [188, 189, 190] for testing NPEs because, with the right choice of parameters, it can have both regions of strong curvature and regions of relative smoothness, thus facilitating testing the robustness of a given NPE. The choices of the parameters, N, A_i, and λ_i, were chosen in concordance with Ref. [189].

For each function, we generated pseudo-"data" by sampling at 100 fixed points in x. For each point, x_i, the "data" is equal to $y(x_i)$ as defined by Eq. (H.2), where ϵ_i is a number drawn from a Gaussian-distributed random variable with parameters ($\mu = 0$, $\sigma = 0.2$). These samplings to generate the ϵ_i for each x_i are statistically independent of one another.

We then fit this pseudo-data with both a PE and NPE. For the parametric fit, the true, nominal parametric form was used. For the non-parametric fit, the PC estimator is applied to the data, using a Gaussian kernel function,

$$K\left(\frac{x - x_i}{h}\right) = \frac{1}{\sqrt{2\pi h_{\text{opt}}^2(x)}} e^{-\left(\frac{x - x_i}{h_{\text{opt}}(x)}\right)^2} \tag{H.11}$$

where the bandwidth parameter has been optimized dynamically for each x, as per Eq. (H.9).

Figure H.1 shows the results of these fits. In the top panel of each sub-figure, the blue solid line represents the actual functional form of $f(x)$, while the green (red) solid line represents the (non)parametric functional fits to the "data". The bottom

panels of each sub-figure show the squared difference between the true function and the two respective attempted fits (the "squared error").

As you can see from comparisons of the squared error, the PC estimator can perform better than a parametric fit in terms of estimating the value of $f(x)$[2]. This mostly happens in regions of relatively low curvature in $f(x)$ where the data experienced relatively strong statistical fluctuations, e.g. $0.5 \lesssim x \lesssim 0.8$ in Fig. H.1a.

Moreover, the PC estimator can also "regress to the mean" faster than a PE, as it heavily weights local information relative to non-local information (in contrast to a PE, which utilizes all of the available information). So, when a sequence of strong, coherent statistical fluctations occur in the data, e.g. $0.4 \lesssim x \lesssim 0.7$ in Fig. H.1b or the aforementioned $0.5 \lesssim x \lesssim 0.8$ in Fig. H.1a, while both estimators are often biased by these fluctuations, assuming the following sub-samples of data are approximately centered around the true function, the PC estimator more quickly "returns" to tracking the true function.

That said, the performance of the PC estimator suffers in regions of large curvature, e.g. $x \sim 1$ in Fig. H.1a or $x \lesssim 0.25$ in Fig. H.1b, because, from Eq. (H.9), the optimal bandwidth chosen to be relatively lower in these regions in order to not bias the estimation. Consequently, in these regions, the PC estimator is more subject to statistical fluctations in the data (to word this statement differently, when using the optimal bandwidth in a region of large curvature, the local bias of the estimator — Eq. (H.5) — is reduced but the local variance — Eq. (H.6) — is increased).

[2]Of course, by construction, the PC estimator, or any NPE for that matter, cannot yield estimates of any distinctive parameters of the true $f(x)$.

H.2.1.2 Testing the PC Estimator: Checking the "Pull" of the Fit

As another test of the PC estimator's robustness, we calculated the overall "pull" (Section H.1) of the PC fit, and compared it against the "pull" for the sampled data as well as the PE fit.

Specifically, for every x_i, we calculated the "pull" — i.e. the difference of the PC fit and the sampled data, $y(x_i) - \hat{f}_(x_i)$, divided by the Gaussian σ with which the sampled data was generated (0.2 for the example above). We then created the overall pull distribution and checked the consistency of its mean and RMS with those expected from a standard pull distribution (0 and 1, respectively)[3]. We repeated this process for the PE fit as well as for the sampled data itself, where, for the sampled data we compared the $y(x_i)$ values against the true $f(x_i)$.

The results of these comparisons for both true $f(x)$ that we attempted to fit — i.e. the Gaussian and the triple exponential from Eq. (H.10) — are shown in Table H.1. In addition to showing the values of the pull distribution mean and standard deviation, the table also shows the respective "consistencies" of the calculated mean and standard deviations with the "standard pull" hypothesis — that is, the difference of the expected mean (0) or standard deviation (1) from the calculated value, divided by the uncertainty on the calculated value. Finally, shown in the last column is the percentage difference of the calculated standard deviation for the given

[3]Remember that the sample mean and sample variance are the maximum-likelihood estimators for the true mean and variance of the underlying distribution governing the sampled data. Thus, for sampled data that nominally should be coming from a fixed Gaussian distribution, comparisons of the sample mean and sample variance (standard deviation) to the nominal μ and σ^2 (σ) of the distribution serve as a straightforward but relatively accurate means of checking the consistency of the data with the hypothesis that it actually is governed by the Gaussian in question.

fit and the comparison between the true distribution and the sampled data. As you can see, for both functions and for the sampled data as well as the two attempted fits, the mean value of the pull distribution is consistent within uncertainties with the expected mean of a "standard pull". The standard deviations however, are not consistent within uncertainties. Further investigation showed that this effect was spurious, specifically that it was due to pull distributions only being calculated within a limited range (-2 to 2); This explanation was corroborated by attempting to fit each individual pull distribution with a Gaussian distribution; each of these respective fits yielded an estimated mean and variance for the Gaussians that was easily consistent within uncertainties with the expected values for a standard pull.

Checking the Pull for Varying σ_ϵ

As a further check of the robustness of the PC estimator, we also repeated the above cross-check using pull distributions, but we allowed the individual σ_ϵ for each x point to vary. Specifically, when generating the sampled data for a given x_i, the σ_ϵ was chosen randomly from a uniform distribution on the interval, [0, σ_ϵ]

The results of this are shown in Table H.2, which uses the same relative formatting as Table H.1. Direct comparisons of the two tables shows that the calculated pull distributions shown in Table H.2 are even more consistent with standard pull distributions (again, with the caveat that the relative inconsistency of the calculated pull's standard deviations is due to a limited bin range).

Table H.1: Results from calculating the pull distribution for the distributions shown in Fig. H.1a. The pull distribution mean and σ refer to the sample mean and sample standard deviation. The "Value" column refers to the calculated value of the quantity in question. The "Consistency" column refers to how consistent (using the calculated uncertainties on the mean and σ) the calculated value is with the value expected from a standard pull distribution (mean of zero, standard deviation of one). Note that the pull distributions were calculated within a limited sub-range, hence biasing the calculated value of the sample standard deviation to be slightly lower than normal, c.f. the discussion in Section H.2.1.1. The Δ(data, fit) shows the relative percentage difference in the calculated sample standard deviation between the pull distribution calculated with the pseudo-data and the pull distribution calculated with the fit in question.

Fit Type	Pull distribution statistic				Δ(data, fit) (%)
	Pull distribution mean		Pull distribution σ		
	Value	Consistency	Value	Consistency	
$f_1(x)$ "Data"	0.07	0.72	0.90	-1.58	-
$f_1(x)$ Parametric Fit	-0.03	-0.30	0.87	-2.13	3.55
$f_1(x)$ Non-Parametric Fit	-0.02	-0.25	0.83	-2.75	7.99
$f_2(x)$ "Data"	0.06	0.72	0.90	-1.58	-
$f_2(x)$ Parametric Fit	0.01	0.14	0.90	-1.60	0.11
$f_2(x)$ Non-Parametric Fit	0.07	0.77	0.91	-1.47	-0.70

Checking the Pull for Non-optimal Bandwidths

As a cross-check, we also tried fitting a variant of the PC estimator where, instead of using the optimal kernel bandwidth for each point, h_{opt}, we instead fixed the bandwidth to a set value for all x. We tested several values for h, including discrete steps through a range of values for h, utilizing the average value of the optimal bandwidth across the entire x range, and utilizing the maximum or minimum optimal bandwidths in the entire x range. Notably as well, we also tested the "best-estimate" optimal *fixed*-bandwidth, where the second derivative term, $\frac{d^2 f}{dx^2}$,

Table H.2: Analogous information as Table H.1, but calculated using a slight variant of the toy experiment discussed in Section H.2.1.1. See the text for more details on this variant.

Fit Type	Pull distribution statistic				Δ(data, fit) (%)
	Pull distribution mean		Pull distribution σ		
	Value	Consistency	Value	Consistency	
$f_1(x)$ "Data"	-0.04	-0.40	0.94	-0.92	-
$f_1(x)$ Parametric Fit	-0.00	-0.04	0.94	-0.82	-0.66
$f_1(x)$ Non-Parametric Fit	0.01	0.07	0.88	-1.88	7.25
$f_2(x)$ "Data"	-0.04	-0.40	0.94	-0.92	-
$f_2(x)$ Parametric Fit	-0.05	-0.56	0.93	-1.05	0.87
$f_2(x)$ Non-Parametric Fit	-0.04	-0.44	0.90	-1.46	4.52

from Eq. (H.9) is estimated through data-driven methods[4].

For all fixed-values of h, we repeated the functional estimation of $f_1(x)$ and $f_2(x)$ and the subsequent calculation of the resulting pull distribution. We found that, unanimously, the resulting pull distributions were quite statistically inconsistent with standard pull distributions. This often manifested in the form of large statistical inconsistencies for both the calculated pull distribution mean and standard deviation. Checking the actual fits to data, these large statistical inconsistencies in the calculated pull distributions were due to large systematic biases in the non-parametric fit in certain regions of x.

This test points to an underlying guiding principle for NP estimation: The bandwidth parameter must be allowed to vary with x, as fixed bandwidths often result in systematically biased fits of the underlying data.

[4]This method for bandwidth selection is commonly utilized with NPEs in real situations, due to the aforementioned typical lack of knowledge regarding $\frac{d^2 f}{dx^2}$.

Choosing Locally Optimal Bandwidths using Pull Distributions

The difficulty, then, is determining how best to vary the bandwidth with the data in order to optimize the NP estimation. We propose the following approach: For each local region of x that one wishes to estimate with an NPE, one iterates through multiple values of h, calculating the resulting *local* NPE, $\hat{f}(x)$, for each. In order to create a global $\hat{f}(x)$, one stitches together the various locally calculated $\hat{f}(x)$, iterating through permutations of the respective fixed bandwidths h for each.

For each stitched-together global $\hat{f}(x)$, one calculates the overall pull distribution from comparing the resulting $\hat{f}(x)$ with the observed data. Referencing prior discussions, we observed that the calculated pull distribution of an NPE is seemingly statistically consistent with a standard pull distribution *only* when the NPE's bandwidth is optimized for each local region in x. Assuming this relation is biconditional, a $\hat{f}(x)$ that has a standard pull distribution should have a minimal (i.e. optimal) MSE. Thus, the global $\hat{f}(x)$ that yields a pull distribution that is most statistically consistent with a standard pull distribution is then chosen as the final global $\hat{f}(x)$.

One issue with this method is that it is computationally intensive, as it involves the combinatoric choices for the local bandwidths of each individual local region in x. We thus propose a computationally simpler alternative. The underlying principle is the same; for each local region of x, one iterates through fixed values of h and calculates the local $\hat{f}(x)$. However, instead of permuting through each local $\hat{f}(x)$ from a given x region, one first ranks these calculated local $\hat{f}(x)$ by calculating a

local pull distribution for each. Similar to the discussion above, the local $\hat{f}(x)$ that yields a pull distribution that is most statistically consistent with a standard pull distribution is then chosen as the local $\hat{f}(x)$ for the x region in question. One then repeats this procedure for each local x region and builds a global $\hat{f}(x)$ by stitching together the locally-optimized $\hat{f}(x)$.

At first glance, one might expect that the two methods have notable systematic differences between them. However, we note that as long as the pull distribution for each local $\hat{f}(x)$ is a standard pull distribution, then the pull distribution of the global $\hat{f}(x)$ will be as well. We dedicate some additional discussion to this method in the context of a 2D kernel-based NPE in Section H.2.3.2.

H.2.2 Summary

In this section, we have presented a detailed discussion on non-parametric estimators (NPEs), which can be used to model "noisy" data when no suitable parametric form is available. We presented several toy experiments where we fit parametric and non-parametric estimators to pseudo-data drawn from some underlying true distribution, where stochastic fluctuations were added to simulate statistical uncertainties. Direct comparisons of any biases in the resulting fits, as well as the calculated pull distributions for the respective fits, showed that NPEs can achieve similar levels of performance as PEs, so long as the bandwidth parameter of the NPE is optimized for each local region in x.

Because this bandwidth optimization can often be quite difficult without addi-

tional knowledge of the underlying distribution, we included a proposal for a method for bandwidth optimization. This method for determining optimal local bandwidths involves finding which fixed, local bandwidths result in calculated pull distributions for the NPE that are most statistically consistent with a standard pull distribution.

In the next section, we provide details on a 2D NPE that we utilized to smooth experimental results from the analyses described in this dissertation.

H.2.3 Two-dimensional Gaussian Kernel Smoothing

At several points in this dissertation, we discussed Gaussian kernel-based smoothing. In this section we provide more a specific description of this kernel-based smoothing procedure.

H.2.3.1 Introduction to the Idea

The basic idea of the two-dimensional kernel-based smoothing is extremely similar to the nonparametric functional estimation described in Section H.2.1. However, for this purpose, the functional form of the NPE is slightly altered from Eq. (H.3),

$$\hat{f}(x,\, y) = \frac{\int_{-a_{x'}}^{a_{x'}} \int_{-a_{y'}}^{a_{y'}} dy' dx' f(x',\, y') K(x'\, y')}{\int_{-a_{x'}}^{a_{x'}} \int_{-a_{y'}}^{a_{y'}} dy' dx' K(x'\, y')}, \tag{H.12}$$

but in practice, both equations utilize the same underlying approach: the individual data points that we have calculated/measured provide collective, local information about the "true" expectations of the data. In the case of Eq. (H.12), the "smoothed" results are calculated as a weighted average over the local information

using the kernel function $K(x'\,y')$. The main issue that arises in applying Eq. (H.12) is determining the optimal kernel bandwidth(s).

H.2.3.2 Optimizing the Choice of Kernel Widths

When discussing the PC estimator, we noted in Section H.2.1 how having access to additional information, namely, the functional form of the second-derivative of the true function, $f(x)$, enabled the calculation of an "optimal" kernel bandwidth, where the "optimal" bandwidth minimized the mean-squared error from the function estimation. Of course, as already noted when we discussed this bandwidth optimization, this type of information is typically not available in most situations where smoothing is needed.

More general, data-driven methods for optimizing kernel bandwidths have been studied at length, and there is a copious amount of literature on the subject. Rather than utilize one of those methods, however, we propose instead a novel method. The motivation for this method is the following. Nominally, we would like to choose a set of bandwidths that minimize the overall mean-squared error, — MSE, top half of Eq. (H.7) — of the smoothed data, as this subsequently minimizes the combined expected loss of accuracy from any systematic and statistical biases in the smoothing.

Although we cannot directly calculate the mean-squared error for a given choice of kernel bandwidths, as long as there is another statistical quantity that we can calculate and that is correlated with the mean-squared error, we can utilize

this other statistical quantity to guide the choice of kernel bandwidths.

In Section H.2.1.2 we showed how, for the optimal choice of kernel bandwidth, fitting (e.g. smoothing) a function with the Priestley-Chao estimator leads to an overall pull distribution between the calculated fit function and the initial data that is consistent with a standard pull distribution.

Moreover, in that same section, we demonstrated how other choices of kernel bandwidth led to pull distributions that were relatively less consistent with a standard pull distribution. We thus proposed a method for optimizing the local bandwidth parameters by capitalizing on these differences. We describe below the two-dimensional version of our proposed method.

Specifically, for a given x:y point that we wish to smooth, we scan through a range of individual pairs of σ_x:σ_y values. For a given pair, σ_{x_i}:σ_{y_i}, we perform the two-dimensional smoothing using these σ_x:σ_y values in a "local" area around the chosen x:y point. We then calculate the pull distribution that results from this smoothing (again, only calculating it locally — using the points where we actually applied the smoothing).

Scanning over the range of σ_x:σ_y values thus yields a set of "local" pull distributions. We then rank these pull distributions based on how consistent (we will define this ranking shortly) they are with a standard pull distribution — i.e. $(\mu = 0, \sigma = 1)$ — and choose the σ_x:σ_y pair that yielded the most consistent local pull distribution.

Ranking the pull distributions' consistency

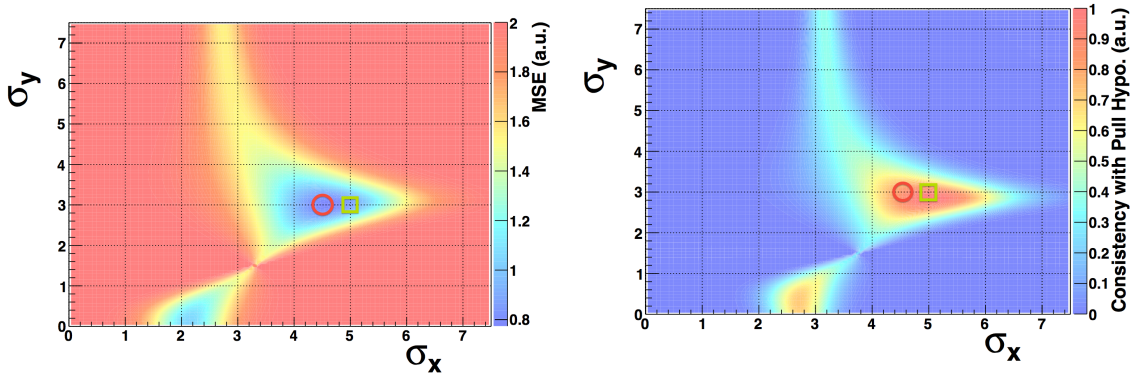

(a) The MSE value versus σ_x:σ_y.

(b) The "standard pull consistency", versus σ_x:σ_y, of the calculated local pull distribution.

Figure H.2: Cartoon of a hypothetical two-dimensional optimization scan over individual axis kernel widths. This hypothetical scan would be performed for each point in the x:y plane where one is attempting to apply kernel-based smoothing using Gaussian kernels. The grid line intersection points represent the scanned points in the σ_x:σ_y plane. The σ_x:σ_y pair that yields the most "consistent", "local" pull distribution (the definitions of "consistent" and "local" are in the text) is chosen as the optimal pair, and is represented on both sub-figures by the point in the center of the green square. The expectation is that this point is approximately close to the true optimal choice (red circle) of kernel widths, based on minimizing the mean-squared error of the smoothing.

1. Within the set of calculated pull distributions, we find the distribution that has a sample RMS closest to 1. Call this distribution A.

2. Distribution A is added into a subset. For each other calculated pull distribution, B_i, if the difference $\mathrm{RMS}(A) - \mathrm{RMS}(B_i)$ is less than the estimated total error on the difference, $\sigma_{\mathrm{RMS}(A)} \oplus \sigma_{\mathrm{RMS}(B_i)}$, then B_i is also added to aforementioned subset.

3. From the subset constructed in Step 2, we choose the pull distribution that has the smallest bias — the distribution mean is closest to 0.

The expectation, from the studies shown in Section H.2.1.2, is that the σ_x:σ_y pair that this method chooses will be approximately close to the optimal pair that minimizes the MSE of the smoothing.

Figure H.2 shows a cartoon detailing this procedure. The underlying assumption is that the MSE for a given choice, σ_x:σ_y, of kernel bandwidths is strongly correlated with the overall consistency of the calculated pull distribution with a standard pull distribution. If this is true, then finding the point in the σ_x:σ_y plane that maximizes the consistency directly correlates to minimizing the MSE of the resulting 2D NPE.

H.2.4 Smoothing Example Results

In this section, we show some examples of how the 2D NPE-based smoothing discussed in Section H.2.3 visually affects the data. Figure H.3b (H.3a) shows the (un)smoothed signal efficiency with which mass-points from the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay

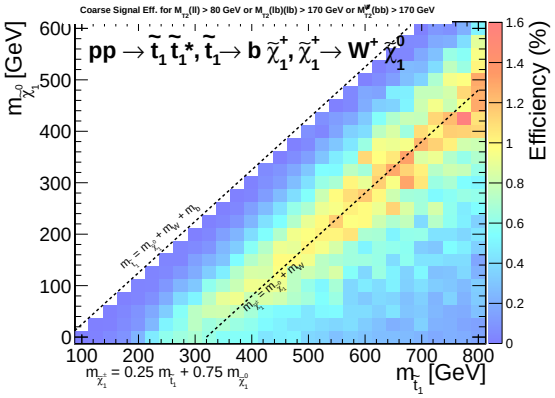

(a) The unsmoothed signal efficiency for $x = 0.25$.

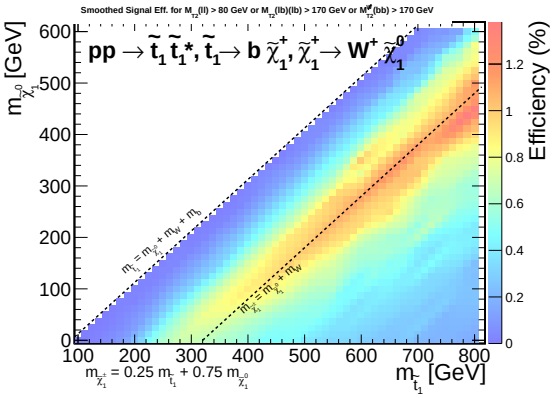

(b) The smoothed signal efficiency for $x = 0.25$.

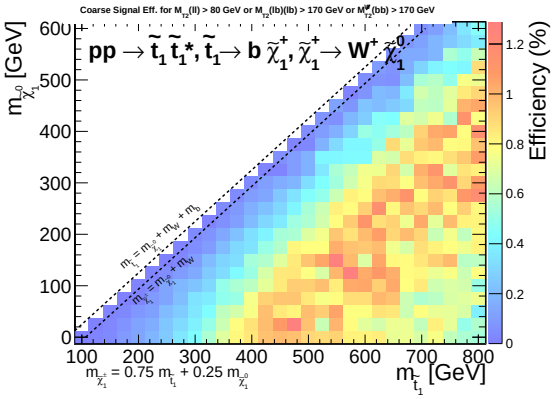

(c) The unsmoothed signal efficiency for $x = 0.75$.

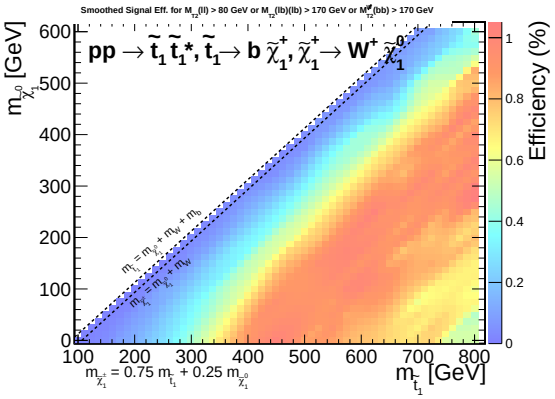

(d) The smoothed signal efficiency for $x = 0.75$.

Figure H.3: The signal efficiency, shown pre- and post-smoothing for the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV} \cup M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) > 170\,\mathrm{GeV} \cup M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) > 170\,\mathrm{GeV}$, for two variants of the $\tilde{\mathrm{t}}_1 \to \mathrm{b}\tilde{\chi}_1^+$ decay mode. See Section H.2.3 for details on the smoothing procedure.

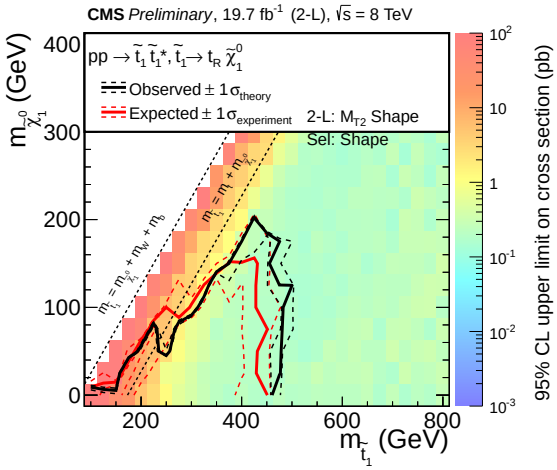

(a) The unsmoothed exclusion limits for the $\widetilde{t}_1 \to t_R \widetilde{\chi}_1^0$ decay mode.

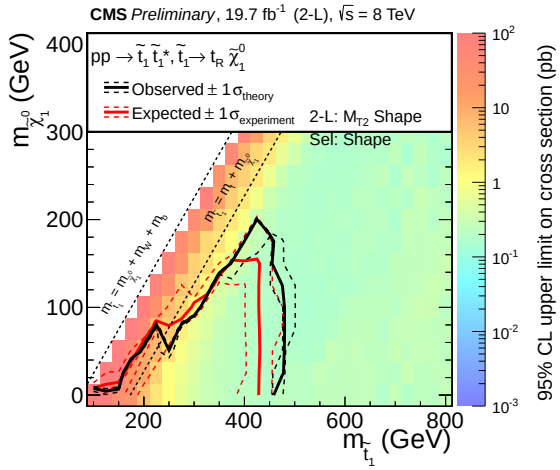

(b) The smoothed exclusion limits for the $\widetilde{t}_1 \to t_R \widetilde{\chi}_1^0$ decay mode, i.e. a copy of Fig. 7.16b

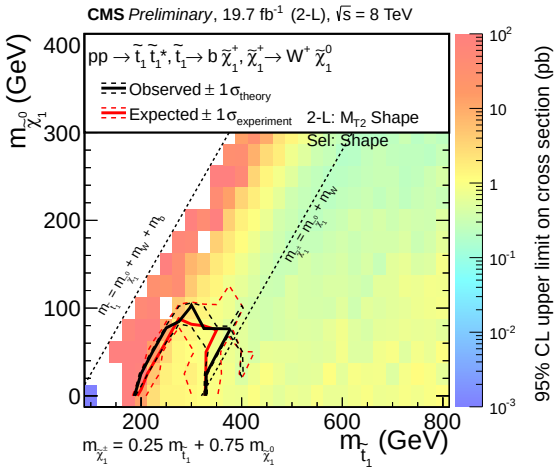

(c) The unsmoothed exclusion limits for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with $x = 0.25$.

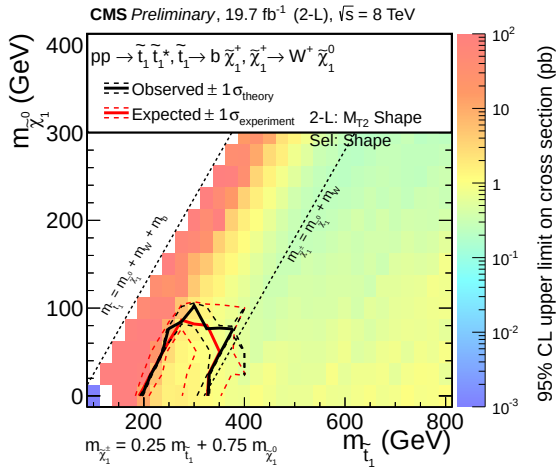

(d) The smoothed exclusion limits for the $\widetilde{t}_1 \to b\widetilde{\chi}_1^+$ decay mode with $x = 0.25$. i.e. a copy of Fig. 7.15d

Figure H.4: The impact of the kernel-based smoothing on exclusion limits for top-squark pair production. On the left are the unsmoothed upper limits, with associated 95% CL exclusion contours, on top-squark pair-production. The maps on the right show the results of applying the kernel-based smoothing, Section H.2.3, to the maps on the left.

mode, with $x = 0.25$, pass the cut, $M_{\mathrm{T2}}(\ell\ell) > 80\,\mathrm{GeV} \cup M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) > 170\,\mathrm{GeV} \cup M_{\mathrm{T2}}^{\mathrm{W}}(\mathrm{bb}) > 170\,\mathrm{GeV}$. Figures H.3c and H.3d show the analogous results for the case when $x = 0.75$. Figure H.4 shows examples of how the smoothing affects the calculated exclusion limits for top-squark pair-production in the $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0})$ plane.

For both figures, one can see that in general, the smoothing performs quite well in terms of removing "rough" edges in the associated 2D distributions. The smoothing also is able to fill in "holes" in the 2D distributions — e.g. the point $(m_{\tilde{\mathrm{t}}_1}, m_{\tilde{\chi}_1^0}) = (250{:}100)$ in Fig. H.4c — by utilizing local surrounding information to best estimate the expected value for those points. Finally, in this figures, there is no clear indication that the smoothing introduces any notable systematic biases. We have confirmed this last point in other cases where we applied the 2D NPE-based smoothing.

H.3 Maximum-Likelihood Morphing of Shapes

There are a number of salient differences between the 3D M_{T2} dilepton shape analysis, Section 7.2.4, and the cut-and-count dileptonic top-squark search, Chapter 6. One of the more notable ones is the fact that, in the 3D M_{T2} shape analysis, a number of key systematic uncertainties are treated as *shape* uncertainties. This is relevant when calculating the test statistics used to either set exclusion limits, Eq. (G.4), or quantify observed excesses, Eq. (G.8).

For both the cut-and-count as well as the 3D M_{T2} shape analysis, the systematic uncertainties on both the signal and background are constrained further

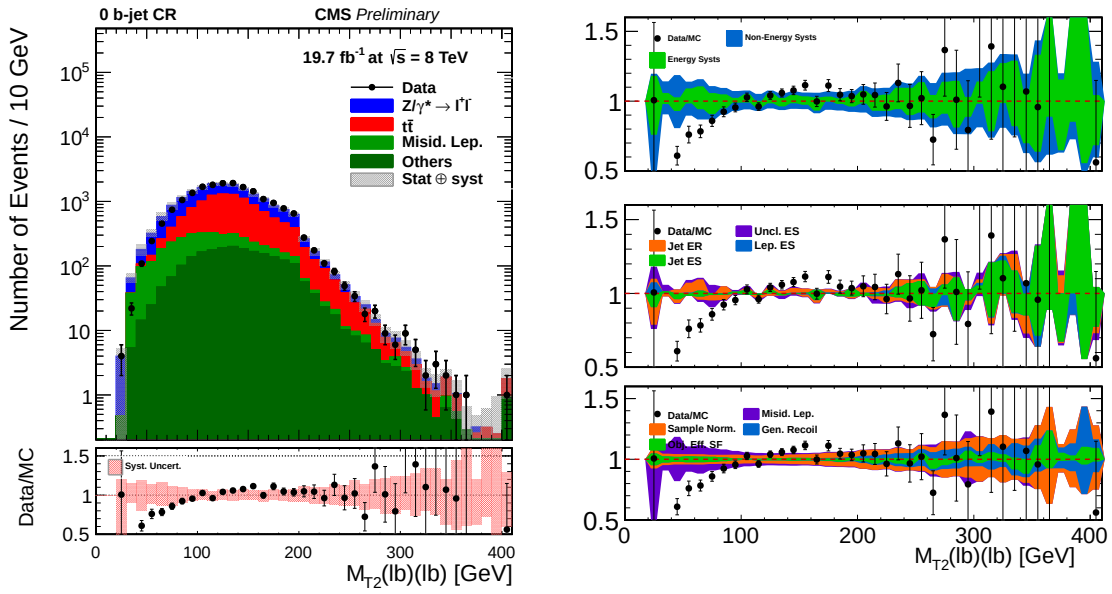

(a) (b)

Figure H.5: The distribution of $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ in a sample of events passing the 3D M_{T2} shape selection, except each event is required to have no reconstructed b-jets. The left sub-figure shows the individual expected background contributions, along with the expected total impact of systematic uncertainties. The right sub-figure shows an itemized breakdown of the contribution of individual systematic uncertainties to the ratio of data and simulation.

based upon maximum-likelihood fits to the observed data. For the 3D M_{T2} shape analysis, the fact that some of the systematic uncertainties are shape uncertainties also enables a re-evaluation, morphing, and subsequent constraining of the 3D M_{T2} shapes for the signal and background.

Figures H.5 and H.6 show example results from this procedure. The pre-morphing $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ distribution, for a 0 b jet control sample, is shown in Fig. H.5a. In the region where the discrepancies between data and simulation are the largest, i.e. $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b}) < 100\,\mathrm{GeV}$, the total systematic uncertainty is dominated by the uncertainty on the misidentified lepton contribution (purple in bottom panel of Fig. H.5b), the electroweak vector boson (e.g. WW, ZZ, etc.) sample nor-

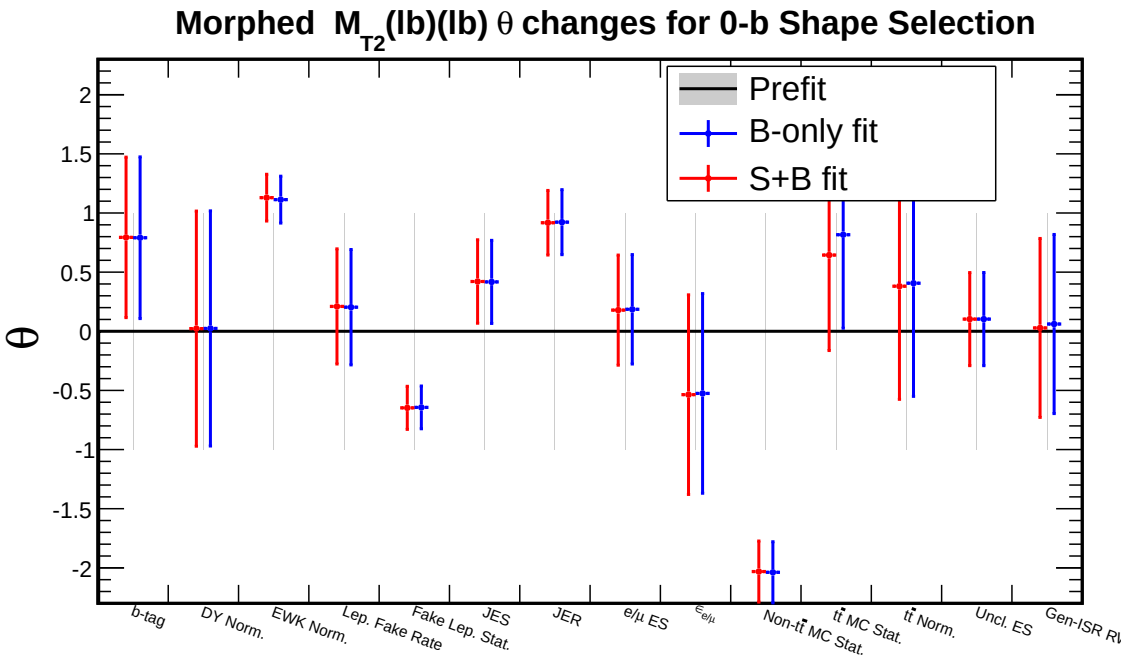

(a)

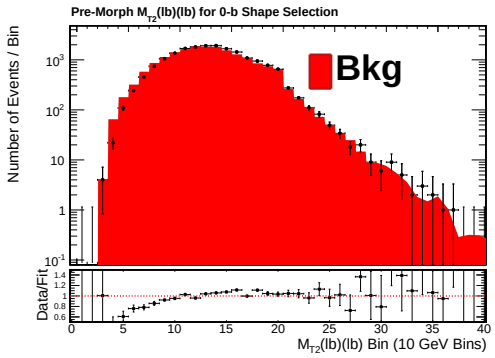

(b)

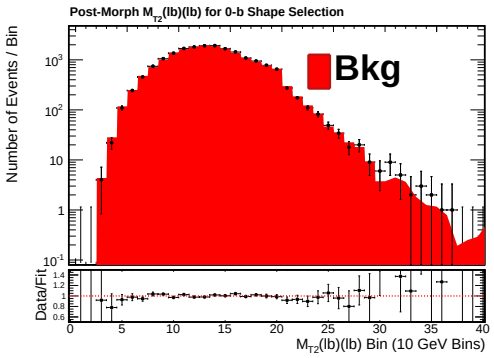

(c)

Figure H.6: (Bottom) Two example distributions of $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ in a sample of events passing the 3D M_{T2} shape selection, except that each event is required to have no reconstructed b-jets. The left sub-figure shows the pre-fit distribution, i.e. Fig. H.5a, that is input into the likelihood-based morphing, see Section H.3. The right sub-figure shows the post-fit distribution. (Top) the post-fit constraints of the associated nuisance parameters, see Fig. H.5b, of the fit. Pre-fit, the θ were parametrized such that they were dimensionless, centered at 0, and had width parameters of 1.

malization (orange in bottom panel of Fig. H.5b), and (partially) the uncertainty on the $\not{E}_{\mathrm{T}}$ reconstruction due to JER (orange in middle panel of Fig. H.5b).

This distribution is input, Fig. H.6b, into the likelihood-maximization and, as mentioned above, the background shape is morphed within the shape-based systematic uncertainties on the background. The resulting distribution, Fig. H.6c, clearly shows much better agreement between data and simulation. The post-fit nuisance parameters (systematic uncertainties) are shown in Fig. H.6a. One can see that the nuisance parameters that were most altered by the likelihood-maximization are the misidentified lepton uncertainty, the electroweak normalization, the JER, and the non-$\mathrm{t\bar{t}}$ MC statistics (i.e. all simulated samples that are not the direct $\mathrm{t\bar{t}}$ production sample). This makes sense given that these were the dominant uncertainties in the region with the largest discrepancies between data and simulation. We note that the uncertainty related to the non-$\mathrm{t\bar{t}}$ MC statistics was shifted downward by $\sim 2\sigma$ relative to its initial position. This, along with the relative shift downward of the statistical uncertainty on the lepton misidentification rate, reflects the fact that the overall background had to be shifted downward in order for data and simulation to match. The fact that the non-$\mathrm{t\bar{t}}$ MC statistics had to be shifted by 2σ is likely indicative of an underestimate of this particular uncertainty in this region.

Appendix I: Optimizing Experimental Selections with the Punzi Parameter

At several points in the dissertation we discussed optimizing a particular analysis choice or set of analysis choices using the "Punzi parameter," a metric for optimizing analysis choices originally defined in Ref. [191].

In this chapter, we provide important details on the Punzi parameter, and discuss more specifically how we utilized it for optimizing selection choices. including a paraphrasing of the salient details from Ref. [191].

I.1 Optimization in Data Analysis

Optimization is one of the more prevalent tasks underlying data analysis in particle physics experiments. How should reconstructed objects (e.g. leptons, jets) be selected? What experimental variables should be used for discrimination between signal and background? One necessary, but not sufficient, guiding principle, at least with respect to searches for new physics, is that after tuning all the various "knobs" of the analysis, there must be acceptable modeling of the relevant parameters of interest by the simulation.

In the context of searches for new physics, a guiding principle is the maxi-

mization of the "statistical power" of the experiment. Stated generally, if the new physics does exist, we want to be able to maximize our ability to see it. If the new physics does not exist, then we want the maximal exclusion ranges for relevant parameters of the new theory (e.g. production cross sections, new particle masses, couplings, etc.).

In the search for signatures of top-squark pair-production described in Chapter 6, we used $M_{\mathrm{T2}}(\ell\ell)$ as a discriminating variable between our top-squark signal and our SM backgrounds. The relative shape of the $M_{\mathrm{T2}}(\ell\ell)$ distribution depends upon both $m_{\tilde{\mathrm{t}}_1}$ and the mass difference between the top-squark and neutralino, $\Delta\mathrm{M} = m_{\tilde{\mathrm{t}}_1} - m_{\tilde{\chi}_1^0}$. In order to optimize the choice of $M_{\mathrm{T2}}(\ell\ell)$ threshold used for defining our signal region, we utilized the "data-blind" median expected upper-limit on the cross section determined for each signal point individually.

This procedure is robust in the sense that the optimal $M_{\mathrm{T2}}(\ell\ell)$ thresholds that we derived are guaranteed to represent, within the statistical uncertainties of the simulation, the "true" optimal $M_{\mathrm{T2}}(\ell\ell)$ thresholds in the case that the top-squark signal does not exist. However, this procedure is an extremely involved one because it requires, among other things, running the very computationally demanding frequentist-based limit setting procedure (Section 6.8.2 and Appendix G) It is often useful, especially in the earlier stages of the development of an analysis, to have a simpler, faster method for determining analysis choices like the aforementioned $M_{\mathrm{T2}}(\ell\ell)$ cut.

One method commonly used in statistical data analysis is to maximize a *figure of merit* (F.O.M.), a quantitative measure of how "powerful" the experiment is with

the hypothetical choice relative to other choices.

Two of the more common F.O.M.s are defined below,

$$F_1(\vec{x}) = \frac{S(\vec{x})}{\sqrt{B(\vec{x})}}$$
$$F_2(\vec{x}) = \frac{S(\vec{x})}{\sqrt{B(\vec{x}) + S(\vec{x})}} \tag{I.1}$$

where $S(\vec{x})$ and $B(\vec{x})$ respectively represent the expected amount of signal and background events passing the set of cuts defined by the set $\vec{x}$, There are a few issues with the F.O.M.s defined above [191]. For example, in order to maximize F_2, the experimentalist must explicitly know what the expected signal yield is, which introduces a source of possible systematic bias in the optimization. To contrast this, maximizing F_1 only requires knowledge of the dependence of S on $\vec{x}$. However, the robustness of F_1 suffers greatly at low B. As an example, naively requiring F_1 to be maximized would favor an experiment with $S = 10^{-1}$, $B = 10^{-5}$ over $S = 10$, $B = 1$.

Giovanni Punzi proposes a more robust, albeit more complicated metric in Ref. [191]. We have paraphrased below the important points from his description and discussion of the parameter.

I.2 The Punzi Parameter

Consider the situation of a single-bin counting experiment, where we are comparing two hypotheses, the background-only hypotheses, H_0, and the signal + background hypothesis, H_S.

In this experiment, we perform a sequence of cuts, count the number of ex-

pected events passing these cuts for our background, B, and also the analogous number of events for signal, S. From these numbers, we can construct the probability to observe n events for each hypothesis,

$$\begin{aligned} p\left(n|H_0\right) &= e^{-B}B^n/n! \\ p\left(n|H_S\right) &= e^{-(B+S)}\left(B+S\right)^n/n! \end{aligned} \tag{I.2}$$

It is straightforward to see that, for some desired level of significance, the observed number of events that pass our cuts must be greater than some threshold, $n > n_{\min}$, where $n_{\min}$ depends upon the desired level of significance. It is also clear to see that for this single bin counting experiment, the power of the test with respect to the signal hypothesis H_S will grow monotonically with the number of signal events, and so, just as there is some number of minimum number of observed events, $n_{\min}$, there is an analogous minimum number of signal events, $S > S_{\min}$, which depends, again, on the desired levels of significance or confidence level (CL) at which one desires to set the experimental limits.

If one performs a Gaussian approximation of the probabilities in Eq. (I.2), then there is an analytic expression for $S_{\min}$,

$$S_{\min} = a\sqrt{B} + b\sqrt{B + S_{\min}}, \tag{I.3}$$

where a and b correspond to the number of Gaussian σ for a one-sided Gaussian tail integral at significances α and β respectively. For a typical search experiment, $a = 5$ (5σ discovery) and $b = 2$ (95% confidence level).

Solving Eq. (I.3) for $S_{\min}$,

$$S_{\min} = \frac{b^2}{2} + a\sqrt{B} + \frac{b}{2}\sqrt{b^2 + 4a\sqrt{B} + 4B}. \tag{I.4}$$

To compare the performance of this analytic estimation method against other

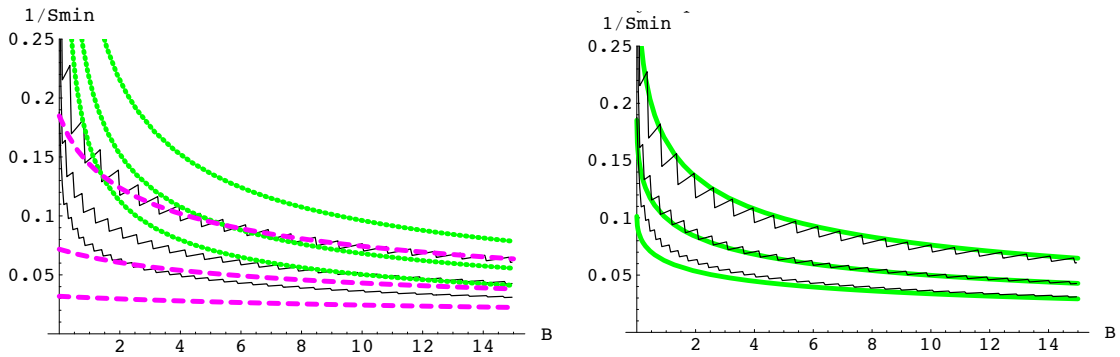

(a) Comparison of the $1/S_{\min}$ calculated with two common F.O.M.s, F_1 (dotted) and F_2 (dashed), Moving from the top downwards, the three curves represent significances [CL] of 1.96σ [95%], 3σ [95%], 5σ [90%]. Reprinted from Fig. 3 of [191].

(b) Comparison of the $1/S_{\min}$ calculated with the Punzi parameter, Moving from the top downwards, the three curves represent significances [CL] of 1.96σ [95%], 3σ [95%], 5σ [90%]. Reprinted from Fig. 5 of [191].

Figure I.1: The dependence of the accuracy of different experimental sensitivity metrics on the number of expected background events as well as the desired significance and CL.

F.O.M.s, Fig. I.1a shows a comparison between the true dependence of $1/S_{\min}$ on B and the estimated dependence using F_1 and F_2 from Eq. (I.1). Figure I.1b shows the analogous comparison using Punzi's method, in the approximation that the two parameters, a and b, are approximately equal. As can be seen from Fig. I.1a, the accuracy of both F_1 and F_2 clearly suffers, with F_1 overestimating the sensitivity at low backgrounds, and conversely, F_2 underestimating the sensitivity. Punzi's method, however, is clearly much more accurate, although its accuracy also suffers somewhat at low backgrounds, particularly when high significance or confidence levels are desired. These inaccuracies in Punzi's method stem from the Gaussian approximation to the Poisson probabilities – Eq.(I.3) – and can be improved by

accounting for the remaining differences between Gaussian and Poisson tail integrals,

$$S_{\min} = \frac{a^2}{8} + \frac{9b^2}{13} + a\sqrt{B} + \frac{b}{2}\sqrt{b^2 + 4a\sqrt{B} + 4B}. \tag{I.5}$$

Figure I.2 shows the analogous comparison as Fig. I.1b, but using the $S_{\min}$ expres-

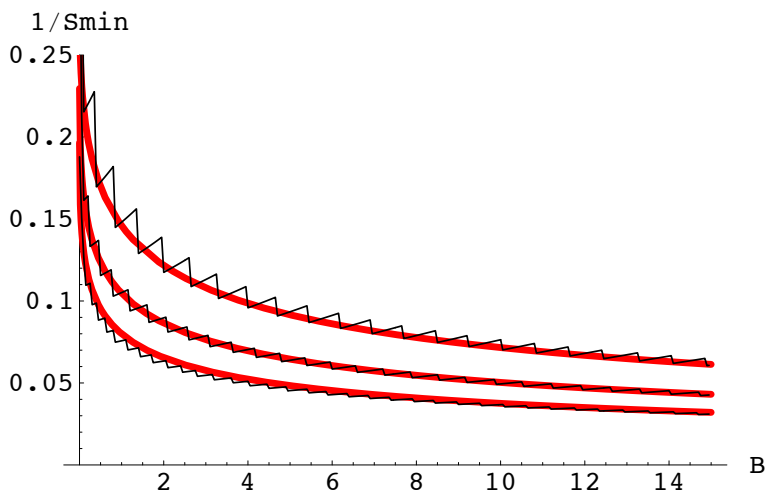

Figure I.2: Comparison of the $1/S_{\min}$ calculated with the Punzi parameter, Moving from the top downwards, the three curves represent significances [CL] of 1.96σ [95%], 3σ [95%], 5σ [90%]. Reprinted from Fig. 6 of [191].

sion from Eq. (I.5). As can be seen, the accuracy of Punzi's method further improves and is quite robust throughout a wide range of expected background yields and desired significances or confidence levels.

Optimizing Experimental Selections with the Punzi Parameter

To summarize what has been discussed thus far, as long as the counting experiment's expected yield of signal events is larger than this $S_{\min}$, i.e. $S > S_{\min}$, we can expect that the counting experiment will be sensitive at the desired level of significance or confidence level.

However, in analogy with Eq. (I.1), both the expected signal and background

yields in our counting experiment will depend upon the set of experimental cuts chosen. Consequently, we can rewrite the above inequality as a function of x, $S(x) > S_{\min}(x)$. We can simplify this further by slightly rewriting the expression for $S(x)$ in terms of the experimental luminosity, L, the signal efficiency, $\epsilon(x)$, and the signal cross section σ,

$$S(x) = \epsilon(x) \cdot L \cdot \sigma \tag{I.6}$$

We can then define a "minimum-detectable" cross section

$$\sigma_{\min} = \frac{\frac{b^2}{2} + a\sqrt{B} + \frac{b}{2}\sqrt{b^2 + 4a\sqrt{B} + 4B}}{\epsilon(x) \cdot L} \tag{I.7}$$

It should be straightforward to see that the statistical power of our counting experiment is maximized when this $\sigma_{\min}$ is minimized. Alternatively, defining the Punzi parameter $\mathcal{P}$ as the inverse of $\sigma_{\min}$ multiplied by the experimental luminosity,

$$\begin{aligned} \mathcal{P}(x) &= (\sigma_{\min}(x) \cdot L)^{-1} \\ &= \frac{\epsilon(x)}{\frac{b^2}{2} + a\sqrt{B(x)} + \frac{b}{2}\sqrt{b^2 + 4a\sqrt{B(x)} + 4B(x)}} \end{aligned} \tag{I.8}$$

our experimental sensitivity is maximized when we maximize the Punzi parameter.

I.2.1 Testing Punzi Parameter-based Optimization

We performed a toy experiment in order to test the above claim. In this toy experiment, we are performing a hypothetical 1-bin counting experiment. The signal efficiency, ϵ_S, has a specific analytic dependence on B, the background yield remaining after the hypothetical cut has been applied,

$$\epsilon_S = \epsilon_{S0} + \Delta\epsilon_S \frac{B^C}{\Delta B} \tag{I.9}$$

where $\epsilon_{S0} = 0.5\%$, $\Delta\epsilon_S = 1\%$, $\Delta B = 2$, and C is varied.

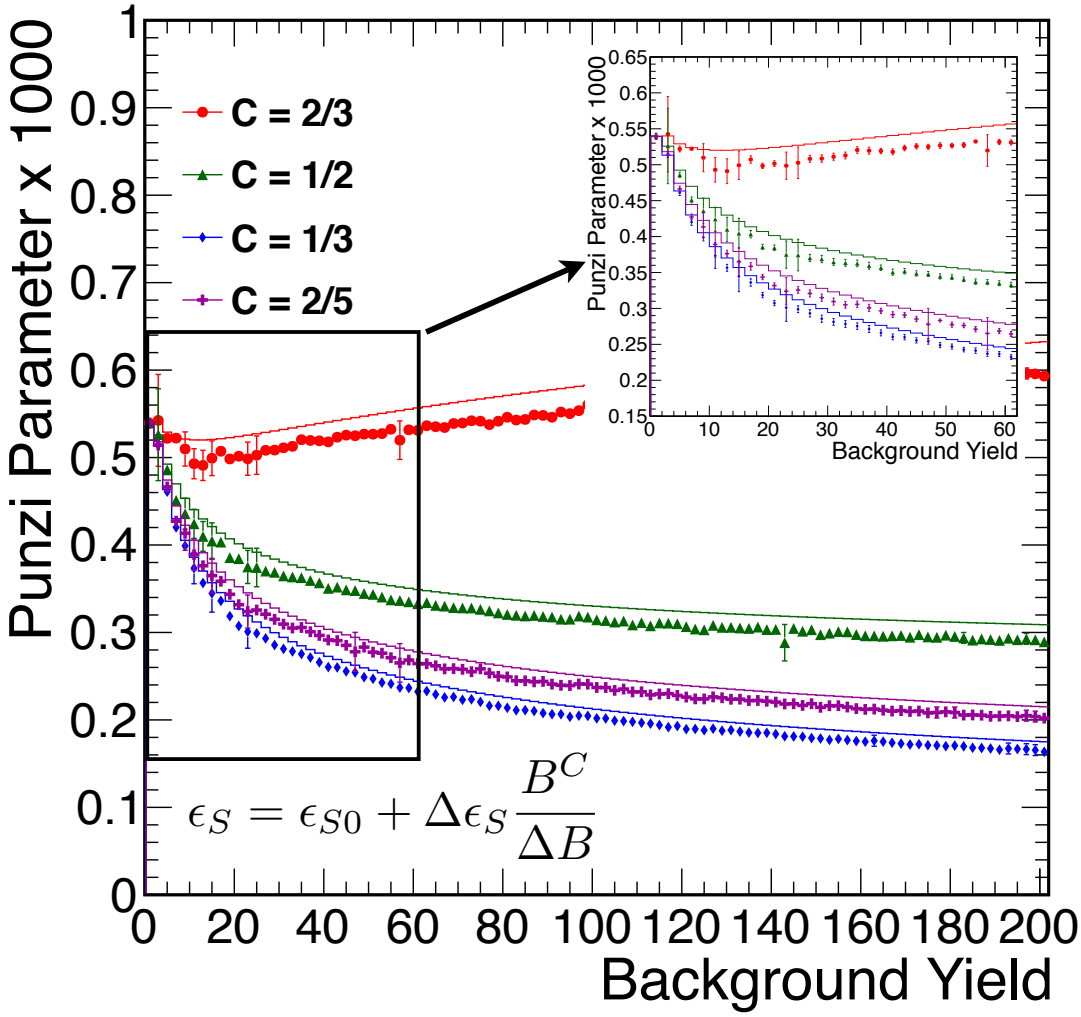

Figure I.3: The dependence of the (solid lines) Punzi parameter, Eq. (I.8), and the (points with error bars) inverse of the fully-frequentist median expected upper limit on the total background yield. Four different signal efficiency functions were tested, all of the form of Eq. (I.9) with differing values for the power parameter C.

We scan through background yields from 0 to 200 and also test four different C values: $C = \frac{2}{3}$, $\frac{1}{2}$, $\frac{1}{3}$, and $\frac{2}{5}$. For each of these, we calculate the Punzi parameter and also calculate the fully frequentist median expected upper-limit on the signal cross section. For all of the calculations of the expected limits, a nominal signal cross section of $\sigma = 100$ pb is used, but the specific value does not significantly alter

the upper-limit.

Figure I.3 shows the result of this toy experiment. The solid colored lines represent the values of the Punzi parameter, while the points with errors bars show the inverse of the upper-limits on the signal cross section. The different colors reference the different C values that we tested. The Punzi parameter tracks along fairly well with the inverse of the upper-limits, showing that the Punzi parameter captures the expected qualitative behavior of the sensitivity of 1-bin counting experiments.

I.2.1.1 Testing the Punzi Parameter for the $M_{\mathrm{T2}}(\ell\ell)$ Optimization

The specific form of the signal efficiency for the toy test, Eq. (I.9), is, admittedly a bit contrived. Specifically, the direct tethering between the signal efficiency and the background yield does not reflect the reality of most searches. Thus, we also tested the Punzi Parameter as an optimization tool for the choice of $M_{\mathrm{T2}}(\ell\ell)$ threshold in the cut-and-count dileptonic top-squark search, Section 6.8.2.

We compared the five choices of $M_{\mathrm{T2}}(\ell\ell)$ threshold for three $m_{\tilde{\mathrm{t}}_1}$ values for all four of the considered top-squark decay modes. The three $m_{\tilde{\mathrm{t}}_1}$ values were chosen independently for each top-squark decay mode. For each $m_{\tilde{\mathrm{t}}_1}$ value for the four top-squark decay modes, we calculated the Punzi parameter for each $M_{\mathrm{T2}}(\ell\ell)$ threshold. The optimal $M_{\mathrm{T2}}(\ell\ell)$ threshold was chosen based on which cut value yielded the maximum value for Punzi parameter. We then compared these optimal $M_{\mathrm{T2}}(\ell\ell)$ thresholds against the corresponding thresholds calculated from comparisons of the data-blind median expected limits.

Table I.1: Comparisons between two methods for calculating the optimal $M_{\mathrm{T2}}(\ell\ell)$ threshold for the dileptonic cut-and-count top-squark search, Chapter 6. The two tested methods were: finding which $M_{\mathrm{T2}}(\ell\ell)$ threshold yielded the maximimum value for the Punzi parameter, Eq. (I.8), and finding which $M_{\mathrm{T2}}(\ell\ell)$ threshold yielded the strongest, data-blind median expected limit, Section 6.8.2. The results for several $m_{\widetilde{\mathrm{t}}_1}$, all with $m_{\widetilde{\chi}_1^0} = 0\,\mathrm{GeV}$, for four different top-squark decay modes are shown: The $\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$ decay mode and the $\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$ decay mode with chargino mass-splitting parameter values, $x = 0.75$, 0.50 and 0.25. For all of the decay modes, the decay products are assumed to be unpolarized.

Optimal $M_{\mathrm{T2}}(\ell\ell)$ [GeV]					
$m_{\widetilde{\mathrm{t}}_1}$ [GeV]	$\widetilde{\mathrm{t}}_1 \to \mathrm{t}\widetilde{\chi}_1^0$		$m_{\widetilde{\mathrm{t}}_1}$ [GeV]	$\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$, $x = 0.75$	
	Punzi	Limit		Punzi	Limit
150	80	90	**200**	100	100
300	110	110	**300**	110	110
450	120	120	**400**	120	120
$m_{\widetilde{\mathrm{t}}_1}$ [GeV]	$\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$, $x = 0.50$		$m_{\widetilde{\mathrm{t}}_1}$ [GeV]	$\widetilde{\mathrm{t}}_1 \to \mathrm{b}\widetilde{\chi}_1^+$, $x = 0.25$	
	Punzi	Limit		Punzi	Limit
200	80	110	**425**	100	100
300	100	110	**525**	90	110
400	120	110	**625**	100	110

Table I.1 shows the results of this comparison. The relative consistency between the two optimal choices of $M_{\mathrm{T2}}(\ell\ell)$ threshold depend upon both the decay mode and the $m_{\widetilde{\mathrm{t}}_1}$ value considered. The observed discrepancies between the two optimal choices for each $m_{\widetilde{\mathrm{t}}_1}$ point are driven by two effects that were not accounted for in the Punzi parameter calculation.

The first of these is the intrinsic systematic uncertainties on the signal efficiency and background yields, see Section 6.4. The second, and far more important effect, is the correction to the signal yield to account for signal contamination in the $\mathrm{t}\bar{\mathrm{t}}$ control region, see Section 6.8.1. The magnitude of the correction to the signal

yield tends to go down with increasing $M_{\mathrm{T2}}(\ell\ell)$ threshold. This naturally leads to higher optimal $M_{\mathrm{T2}}(\ell\ell)$ cut values when accounting for this effect.

I.2.2 Extending the Punzi Approach for Multiple Bins

The above derivations, discussions, and validations were all for the example of a 1-bin counting experiment. In this section, we provide a simple extension to Punzi's method that can be utilized for the optimization of multi-bin counting experiments, such as the 3D M_{T2} binned-shape analysis described in Section 7.2.4.

We will motivate this extension in the context of the 3D M_{T2} shape analysis. In that analysis, a 3D binned shape is constructed by choosing particular bins along each M_{T2} axis. After nominal bin choices are made, there are regions of the 3D M_{T2} parameter space where, in the full, data-driven background estimate, there is no expected contribution from backgrounds.

Using these background-empty bins would hurt the robustness of the statistical calculations used to set limits and derive p-values, see Appendix G. A nominal solution is to ignore them when performing the statistical calculations. However, there could be a relatively large expected signal component in these regions. Thus, it is almost assuredly better if, instead of throwing the information away, these low-background bins are added to other bins. The issue, then, is deciding which bin a low-background bin is added to.

Consider a scenario where one has a 3-bin counting experiment. Using i to denote the bin index, we have expected background yields B_i and expected signal

efficiencies ϵ_i. Imagine now that we have to add the expected contents of the third bin to one of the first two. The Punzi parameter can be utilized to facilitate this decision.

We use $\mathcal{P}_{i+j}$ to denote the Punzi parameter calculated from adding the contents of bins i and j together. The proposed method for optimizing the rebinning is to compare the relative values of $\mathcal{P}_{1+3}$ and $\mathcal{P}_{2+3}$ and utilize this to decide whether to add bin 3 to bin 1 or bin 2,

$$\begin{aligned} \mathcal{F}\left(\mathcal{P}_{i+3}, \mathcal{P}_j\right) &> \mathcal{F}\left(\mathcal{P}_{j+3}, \mathcal{P}_i\right) \\ \rightarrow \mathcal{S}_{i+3} &> \mathcal{S}_{j+3}, \end{aligned} \tag{I.10}$$

where, in the bottom line, $\mathcal{S}_{i+3}$ ($\mathcal{S}_{j+3}$) denotes the scenario where bin 3 has been added to bin i (j), leaving a 2-bin counting experiment consisting of bin $i+3$ ($j+3$) and bin j (i), and the top line indicates that we compare some function $g(x)$ that depends upon the calculated Punzi parameters for each scenario (i.e. the individual Punzi parameters for the individual bins in the resulting 2-bin counting experiment). If $\mathcal{F}\left(\mathcal{P}_{i+3}, \mathcal{P}_j\right) > \mathcal{F}\left(\mathcal{P}_{j+3}, \mathcal{P}_i\right)$, then we choose $\mathcal{S}_{i+3}$ — i.e. we add bin 3 to bin i.

We tested four comparison algorithms, $g(x)$ to try to relatively rank $\mathcal{S}_{1+3}$ and $\mathcal{S}_{2+3}$. The first two involve directly comparing the sums of the Punzi parameters,

$$\mathcal{F}\left(\mathcal{P}_{i+3}, \mathcal{P}_j\right) = \left(\mathcal{P}_{i+3}\right)^m + \left(\mathcal{P}_j\right)^m, \tag{I.11}$$

where m is a tuning parameter. We considered two values for m, $m = 1$ (the "Linear" algorithm), and $m = 2$ (the "Quadratic" algorithm).

The "Ellipse" Intersection Method

The next two comparison algorithms were slightly more involved, as they involved a rederivation of the Punzi parameter for two bins. The basic operating assumption is that in the two bin case, there is total minimum number S_{min} of signal events that can be written as a linear sum of some minimum number for each bin. This, in turn, can be written in similar terms as Eq. (I.3),

$$\begin{aligned} S_{\text{min}} &= S_{\text{min, 1}} + S_{\text{min, 2}} \\ &= a_1\sqrt{B_1} + b_1\sqrt{B_1 + S_{\text{min, 1}}} + a_2\sqrt{B_2} + b_2\sqrt{B_2 + S_{\text{min, 2}}}, \end{aligned} \tag{I.12}$$

where B_i represents the expected yield in bin i, $\mathcal{S}_{\text{min, i}}$ is the analogous quantity, but for the signal, and a_i and b_i are parameters used, for now, to rank the relative contribution of each bin. If one wishes to push the analogy with Eq. (I.3) further, a_i and b_i, in some sense, can represent the desired statistical significances for the individual bins at significances α_i and β_i respectively.

Defining $x = \sqrt{S_{\text{min}} + B}$, Eq. (I.3) defines the intersection of a line, $y \propto x$ with a parabola, $y \propto x^2$, The solution involved finding the smaller of the two intersection points. Similarly then, defining $x_i = \sqrt{S_{\text{min, i}} + B_i}$, Eq. (I.12) is the equation for the intersection of a plane, $z \propto x_1 + x_2$, and a paraboloid, $z \propto x_1^2 + x_2^2$. Solving for the minimum point of this intersection,

$$S_{\text{min, 1}} + S_{\text{min, 2}} = a_1\sqrt{B_1} + a_2\sqrt{B_2} + b_{\text{tot}}\sqrt{C} \tag{I.13}$$

where $b_{\text{tot}} = b_1 \oplus b_2$ and $C = a_1\sqrt{B_1} + B_1 + a_2\sqrt{B_2} + B_2 + \frac{1}{4}b_{\text{tot}}^2$.

To progress further, we utilize a metric by which the relative contributions of

the two bins, via their parameters a_i and b_i, can be ranked in terms of the individual bins' Punzi parameters.

$$\begin{aligned} a_i &= a\frac{(\mathcal{P}_1)^n}{(\mathcal{P}_1)^n + (\mathcal{P}_2)^n}, \\ b_i &= b\frac{(\mathcal{P}_1)^n}{(\mathcal{P}_1)^n + (\mathcal{P}_2)^n}, \end{aligned} \tag{I.14}$$

where the a and b are the parameters from Eq. (I.3) and n is an additional parameter, similar to m in Eq. (I.11). The idea of Eq. (I.14) is that the calculated Punzi parameters for each bin determine their relative importance in contributing to the 2-bin experiment. Specifically, the solution of Eq. (I.13) then places a relative importance on the individual bin that, on its own, would yield a more statistically powerful counting experiment relative to the other bin. After calculating Eq. (I.13) for each scenario, we construct the ranking function $g(x)$, defining $r(\mathcal{P}_1, \mathcal{P}_2) \equiv S_{\text{min, 1}} + S_{\text{min, 2}}$,

$$\mathcal{F}\left(\mathcal{P}_{i+3}, \mathcal{P}_j\right) = r^{-1}(\mathcal{P}_{i+3}, \mathcal{P}_j). \tag{I.15}$$

We call this type of ranking function the "Ellipse Intersection" rank, and we tested two values for the n parameter from Eq. (I.14), $n = 1$ ("Linear Punzi Ellipse Intersection" or "Linear Ellipse" algorithm), and $n = 2$ ("Quadratic Punzi Ellipse Intersection" or "Quadratic Ellipse" algorithm).

I.2.2.1 Ranking the Different Multi-bin Punzi-based Optimizations

We developed a toy experiment in order to determine which $\mathcal{F}(x)$ statistically provides the best relative ranking of the two rebinning scenarios, $\mathcal{S}_{1+3}$ and $\mathcal{S}_{2+3}$. For this toy experiment, we iterated through various combinations of the B_i and ϵ_i

focusing on scenarios where B_3 was low, as these are the scenarios we are interested in[1].

For each of these permutations of the B_i and ϵ_i, we tried both possible rebinning combinations, and calculated a relative rank as per Eq. (I.10), trying the two previously defined $\mathcal{F}(x)$ (Eqs. (I.11) and (I.15)) with different values for n and m. For each rebinning combination, we also calculated the fully-frequentist median expected upper-limits for the resulting 2-bin counting experiment.

For each scenario, the rebinning combination that yielded the strongest expected upper-limits was considered the "optimal" rebinning combination. We subsequently quantified the performance of the $\mathcal{F}(x)$ (i.e. the "Linear", "Quadratic", "Linear Ellipse", and "Quadratic Ellipse" algorithms) by considering statistical metrics such as how often their ranking reflected the true optimal ranking, how often their ranking was unambiguously wrong (i.e. more than 2σ away from the correct ranking, where σ is the uncertainty from the fully-frequentist limit calculation).

Ranking Results

In general, all four algorithms performed quite well. For all of the scenarios tested, each $\mathcal{F}(x)$ correctly ranked the two rebinning choices $\mathcal{O}\,(80\%)$ of the time and was unambiguously wrong only $\mathcal{O}\,(10\%)$ of the time. However, relative to one another, we found that the "Linear" algorithm performed the best, outperforming the other algorithms in terms of its correct-selection efficiency and relative rate of unambiguous mistakes in $\sim 90\%$ of the permutations of the B_i and ϵ_i that we

[1]Remember, the motivation is to try and optimize which bin to add bin 3 to, bin 1 or bin 2.

considered. To quote some quantitative metrics, we found that for all of the considered permutations of the B_i and ϵ_i, the "Linear" algorithm's selection efficiency was $\sim (86 \pm 2)\%$, and its unambiguous mistakes occurred $\sim (8 \pm 2)\%$ of the time, where the $\pm N\%$ term indicates the typical spread over the considered permutations of the B_i and ϵ_i. The "Quadratic" algorithm had similar metrics, albeit slightly worse by $\sim 2\%$ for both the selection efficiency and unamiguous mistakes. Both of the "Ellipse Intersection" algorithms had selection efficiencies $\sim (78 \pm 5)\%$, with unambiguous mistakes occurring $\sim (11 \pm 2)\%$ of the time.

Conclusions and Applications of the Ranking Comparisons

We have found via toy experiments that the "Linear" algorithm — a simple linear comparison of the calculated sums of Punzi parameters for hypothetical rebinning scenarios — has the optimal performance in terms of correctly ranking these hypothetical rebinning scenarios relative to one another.

Consequently, when performing the dynamic rebinning of the low-background bins in the 3D M_{T2} shape, Section 7.2.4.5, we utilized the "Linear" algorithm.

I.3 A Novel Usage Example of the Punzi Parameter

Another, novel usage of the Punzi parameter is to decide which variable, among a set of correlated discriminating variables, is the optimal variable for the purposes of signal-background discrimination.

The maximum value of the Punzi parameter directly correlates with the sta-

tistical strength of a hypothetical "analysis"; thus, for the set of discriminating variables under consideration, whichever variable achieves the highest maximum Punzi parameter value is expected to be the statistically strongest discriminating variable.

In the 3D M_{T2} shape analysis, this scenario arose when trying to decide which $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ flavor to utilize in the 3D M_{T2} shape, Section 7.2.2.2.

As a reminder, in constructing these M_{T2} variants – $M_{\mathrm{T2}}\,(\ell\ell)$, $M_{\mathrm{T2}}^{\mathrm{W}}\,(\mathrm{bb})$, and the chosen flavor of $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$– the larger analysis goals were to utilize a combination of information from all these M_{T2} variants to maximize the sensitivity to top-squark signatures in the data. Given that the low-x T2bw decay modes, relative to the T2tt and high x T2bw decay modes, performed the worst with the $M_{\mathrm{T2}}\,(\ell\ell)$ variable alone, the relative ranking will favor the $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ flavor that, for the low-x T2bw decay mode, performs the best relative to the other $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ flavors.

Figures I.4 through I.6 show the calculated values of the Punzi parameter in these hypothetical experiments for some of the aforementioned flavors of $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$. Events are required to pass the basic preselection and the top-squark signal efficiency is compared against the full, data-driven SM background estimate, Section 6.3.

Directly comparing these maximum Punzi values in Figs. I.4a and I.5a shows that $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ performs better than $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ for the T2tt decay mode, as it reaches a maximum Punzi value of ~ 2 for the former, relative to ~ 0.9 for the latter. This relative ranking of these two $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ flavors also holds for the T2bw, $x = 0.25$ decay mode for $\Delta\mathrm{M} > 200\,\mathrm{GeV}$, as can be seen from comparing

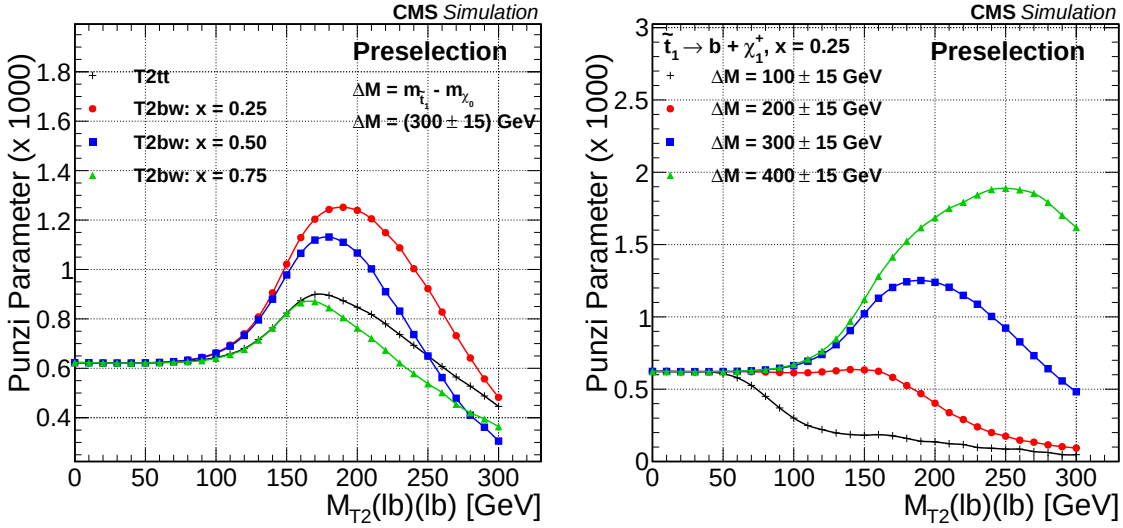

(a) The Punzi parameter dependence on the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different top-squark decay modes for a ΔM value of 300 ± 25 GeV.

(b) The Punzi parameter dependence on the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure I.4: The value of the Punzi parameter, Appendix I, calculated using the tuning parameter values $a = 2$, $b = 5$, when performing a hypothetical cut-and-count using the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ variable to separate the full data-driven SM background estimate against the top-squark signal, for the sample of events passing the preselection.

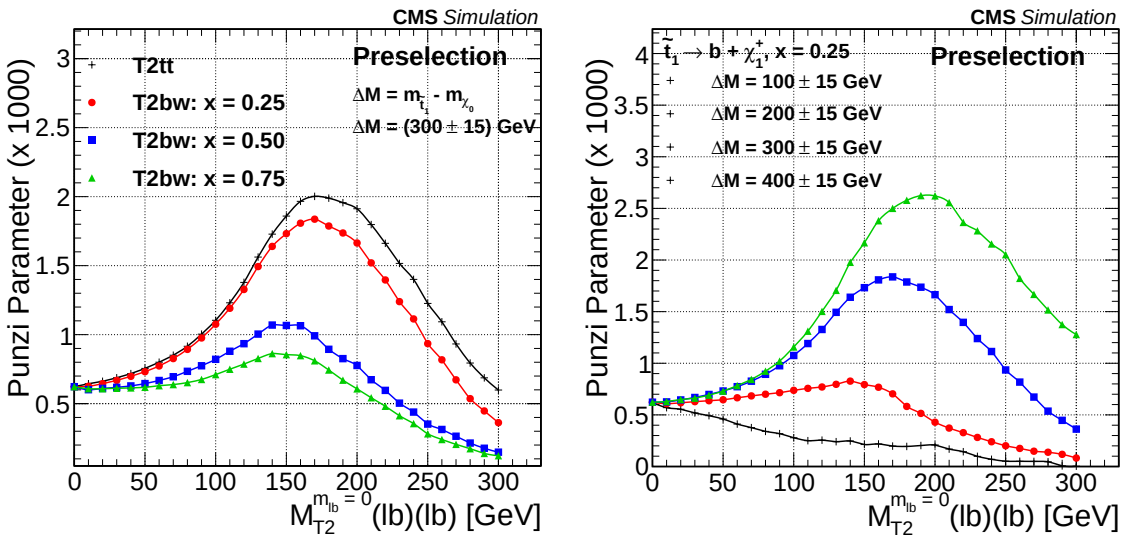

(a) The Punzi parameter dependence on the $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different top-squark decay modes for a ΔM value of 300 ± 25 GeV.

(b) The Punzi parameter dependence on the $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure I.5: The value of the Punzi parameter [Appendix I], calculated using the tuning parameter values $a = 2$, $b = 5$, when performing a hypothetical cut-and-count using the $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}(\ell\mathrm{b})(\ell\mathrm{b})$ variable to separate the full data-driven SM background estimate against the top-squark signal, for the sample of events passing the preselection.

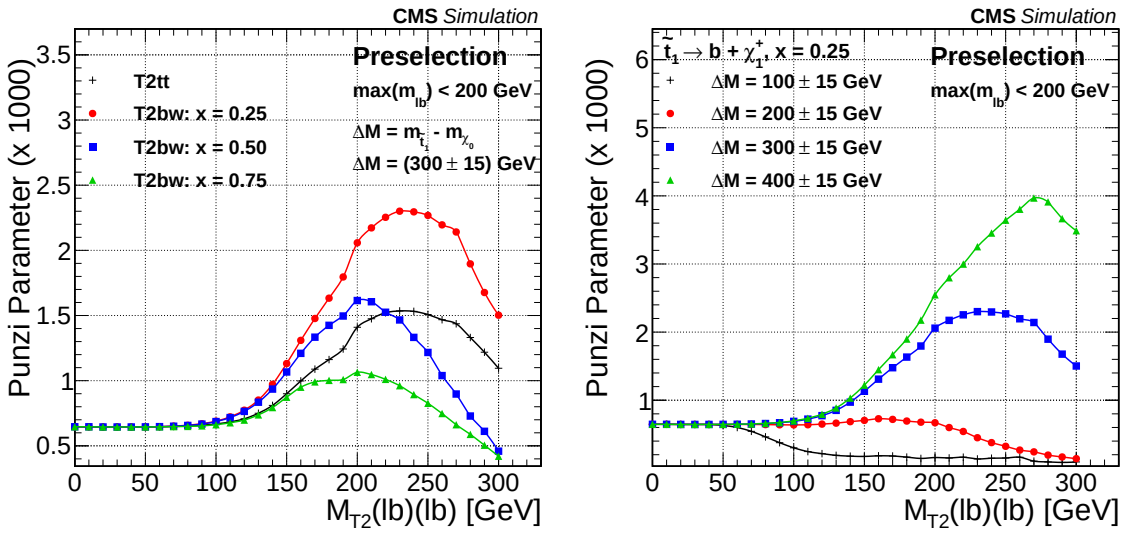

(a) The Punzi parameter dependence on the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different top-squark decay modes for a ΔM value of 300 ± 25 GeV.

(b) The Punzi parameter dependence on the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ cut for several different ΔM regions for the T2bw decay mode with chargino mass-splitting parameter $x = 0.25$.

Figure I.6: The value of the Punzi parameter [Appendix I], calculated using the tuning parameter values $a = 2$, $b = 5$, when performing a hypothetical cut-and-count using the $M_{\mathrm{T2}}(\ell\mathrm{b})(\ell\mathrm{b})$ variable to separate the full data-driven SM background estimate against the top-squark signal, for the sample of events passing a modified version of the preselection where an additional requirement, $\max(m_{\mathrm{b}\ell}) < 200$ GeV, has been applied.

Figs. I.4b and I.5b.

Analogous comparisons between $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ constructed with the additional requirement, $\max(m_{\ell\mathrm{b}}) < 200\,\mathrm{GeV}$ – i.e. Fig. I.6, and $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ show that the cut on $\max(m_{\ell\mathrm{b}})$ improves the performance of $M_{\mathrm{T2}}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ relative to $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$ for all of the T2bw decay mode scenarios for $\Delta\mathrm{M} \geq 200\,\mathrm{GeV}$. The performance in the T2tt decay mode is better for $M_{\mathrm{T2}}^{m_{\ell\mathrm{b}}=0}\,(\ell\mathrm{b})\,(\ell\mathrm{b})$, but as noted above, the focus for this relative ranking is on the low-x T2bw decay modes.

Bibliography

[1] Particle Data Group Collaboration, "Review of Particle Physics", *Chin. Phys.* **C38** (2014) 090001, doi:10.1088/1674-1137/38/9/090001. 3, 20

[2] C. G. Lester and D. J. Summers, "Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders", *Phys. Lett.* **B463** (1999) 99–103, doi:10.1016/S0370-2693(99)00945-4, arXiv:hep-ph/9906349. 6, 165

[3] A. Barr, C. Lester, , and P. Stephens, "m(T2): The truth behind the glamour", *Jour. of Phys.* **G2910** (2003) doi:10.1088/0954-3899/29/10/304, arXiv:hep-ph/0304226. 6, 165

[4] L. Boyle, "Standard Model of Particle Physics: Most Complete Diagram". Creative Commons Attribution 3.0 Unported License, Feb, 2014. https://commons.wikimedia.org/wiki/File:Standard_Model_Of_Particle_Physics,_Most_Complete_Diagram.jpg. xviii, 13, 18

[5] G. P. Salam, "Towards jetography", *Eur. Phys. J. C* **67** (2010) 637, doi:10.1140/epjc/s10052-010-1314-6, arXiv:0906.1833. xx, 14, 126, 127

[6] A. Bilal, "Lectures on Anomalies", arXiv:0802.0634. 16

[7] D. A. Dicus and V. S. Mathur, "Upper Bounds on the Values of Masses in Unified Gauge Theories", *Phys. Rev. D* **7** (May, 1973) 3111–3114, doi:10.1103/PhysRevD.7.3111. 16

[8] B. W. Lee, C. Quigg, and H. B. Thacker, "Weak interactions at very high energies: The role of the Higgs-boson mass", *Phys. Rev. D* **16** (Sep, 1977) 1519–1531, doi:10.1103/PhysRevD.16.1519. 16

[9] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons", *Phys. Rev. Lett.* **13** (Oct, 1964) 508–509, doi:10.1103/PhysRevLett.13.508. 16

[10] G. Bhattacharyya, "A Pedagogical Review of Electroweak Symmetry Breaking Scenarios", *Rept. Prog. Phys.* **74** (2011) 026201, doi:10.1088/0034-4885/74/2/026201, arXiv:0910.5095. 16

[11] F. Zwicky, "On the Masses of Nebulae and of Clusters of Nebulae", *Astrophys. J.* **86** (October, 1937) 217, doi:10.1086/143864. 21

[12] D. Clowe et al., "A direct empirical proof of the existence of dark matter", *Astrophys. J.* **648** (2006) L109–L113, doi:10.1086/508162, arXiv:astro-ph/0608407. 21

[13] L. A. Page et al., "Cosmological parameters from pre-planck cosmic microwave background measurements", *Phys.Rev.* **D87** (2013), no. 10, 103012, doi:10.1103/PhysRevD.87.103012, arXiv:1302.1841. 21, 23

[14] Atacama Cosmology Telescope Collaboration, "The Atacama Cosmology Telescope: Cosmological parameters from three seasons of data", *JCAP* **1310** (2013) 060, doi:10.1088/1475-7516/2013/10/060, arXiv:1301.0824. 21, 23

[15] Planck Collaboration, "Planck 2013 results. I. Overview of products and scientific results", *Astron.Astrophys.* **571** (2014) A1, doi:10.1051/0004-6361/201321529, arXiv:1303.5062. 21

[16] L. A. Page et al., "The Atacama Cosmology Telescope: cross correlation with Planck maps", *Journal of Cosmology and Astroparticle Physics* **2014** (2014), no. 07, 016. 21

[17] N. Suzuki et al., "The Hubble Space Telescope Cluster Supernova Survey. V. Improving the Dark-energy Constraints above z > 1 and Building an Early-type-hosted Supernova Sample", *Ap.J* **746** (February, 2012) 85, doi:10.1088/0004-637X/746/1/85, arXiv:1105.3470. 22

[18] L. Canetti, M. Drewes, and M. Shaposhnikov, "Matter and Antimatter in the Universe", *New J. Phys.* **14** (2012) 095012, doi:10.1088/1367-2630/14/9/095012, arXiv:1204.4186. 22

[19] H. Georgi, H. R. Quinn, and S. Weinberg, "Hierarchy of Interactions in Unified Gauge Theories", *Phys. Rev. Lett.* **33** (Aug, 1974) 451, doi:10.1103/PhysRevLett.33.451. 23

[20] S. P. Martin, "A Supersymmetry Primer", *Adv. Ser. Direct. High Energy Phys.* **21** (2010) 1, doi:10.1142/9789814307505_0001, arXiv:hep-ph/9709356. xiv, xviii, 24, 25, 26, 30, 31, 39

[21] ATLAS, CMS Collaboration, "Combined Measurement of the Higgs Boson Mass in *pp* Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments", *Phys. Rev. Lett.* **114** (2015) 191803, doi:10.1103/PhysRevLett.114.191803, arXiv:1503.07589. 25

[22] J. Wess and B. Zumino, "Supergauge Transformations in Four-Dimensions", *Nucl. Phys.* **B70** (1974) 39–50, doi:10.1016/0550-3213(74)90355-1. 27

[23] P. Fayet, "Supersymmetry and Weak, Electromagnetic and Strong Interactions", *Phys. Lett.* **B64** (1976) 159, doi:10.1016/0370-2693(76)90319-1. 29

[24] P. Fayet, "Spontaneously Broken Supersymmetric Theories of Weak, Electromagnetic and Strong Interactions", *Phys. Lett.* **B69** (1977) 489, doi:10.1016/0370-2693(77)90852-8. 29

[25] S. Dimopoulos, S. Raby, and F. Wilczek, "Supersymmetry and the scale of unification", *Phys. Rev. D* **24** (Sep, 1981) 1681–1683, doi:10.1103/PhysRevD.24.1681. 34

[26] L. Ibez and G. Ross, "Low-energy predictions in supersymmetric grand unified theories", *Physics Letters B* **105** (1981), no. 6, 439 – 442, doi:http://dx.doi.org/10.1016/0370-2693(81)91200-4. 34

[27] W. J. Marciano and G. Senjanović, "Predictions of supersymmetric grand unified theories", *Phys. Rev. D* **25** (Jun, 1982) 3092–3095, doi:10.1103/PhysRevD.25.3092. 34

[28] A. Djouadi, "The Anatomy of electro-weak symmetry breaking. II. The Higgs bosons in the minimal supersymmetric model", *Phys. Rept.* **459** (2008) 1–241, doi:10.1016/j.physrep.2007.10.005, arXiv:hep-ph/0503173. 37

[29] J. F. Gunion and H. E. Haber, "Higgs bosons in supersymmetric models (I)", *Nuclear Physics B* **272** (1986), no. 1, 1 – 76, doi:http://dx.doi.org/10.1016/0550-3213(86)90340-8. 38

[30] CMS Collaboration, "CMS Supersymmetry Physics Results", (June, 2015). xviii, 40, 44, 49, 358

[31] ATLAS Collaboration, "ATLAS Supersymmetry (SUSY) searches", (February, 2015). 40, 44, 358

[32] S. Dimopoulos and G. Giudice, "Naturalness constraints in supersymmetric theories with non-universal soft terms", *Phys. Lett. B* **357** (1995), no. 4, 573, doi:10.1016/0370-2693(95)00961-J. 41

[33] R. Barbieri, G. Dvali, and L. J. Hall, "Predictions from a U(2) flavour symmetry in supersymmetric theories", *Phys. Lett. B* **377** (1996), no. 13, 76, doi:10.1016/0370-2693(96)00318-8. 41

[34] M. Papucci, J. T. Ruderman, and A. Weiler, "Natural SUSY Endures", *JHEP* **09** (2012) 035, doi:10.1007/JHEP09(2012)035, arXiv:1110.6926. 41, 42, 43

[35] A. Delgado et al., "The light stop window", *Eur. Phys. J.* **C73** (2013), no. 3, 2370, doi:10.1140/epjc/s10052-013-2370-5, arXiv:1212.6847. 41

[36] N. Craig, "The State of Supersymmetry after Run I of the LHC", in *Beyond the Standard Model after the first run of the LHC Arcetri, Florence, Italy, May 20-July 12, 2013*. 2013. arXiv:1309.0528. 41

[37] CMS Collaboration, "Search for direct stop pair production in the single and double lepton channel at $\sqrt{(s)} = 8$ TeV", CMS Physics Analysis Summary CMS-PAS-SUS-14-015, CERN, 2015. xiv, xvii, xxxiv, 45, 283, 415, 417, 418, 420, 421, 423

[38] I. Low, "Polarized Charginos (and Tops) in Stop Decays", arXiv:1304.0491. 46, 47, 48, 261

[39] M. Perelstein and A. Weiler, "Polarized Tops from Stop Decays at the LHC", *JHEP* **0903** (2009) 141, doi:10.1088/1126-6708/2009/03/141, arXiv:0811.1024. 46, 47, 261

[40] G. Mahlon and S. J. Parke, "Angular correlations in top quark pair production and decay at hadron colliders", *Phys. Rev.* **D53** (1996) 4886–4896, doi:10.1103/PhysRevD.53.4886, arXiv:hep-ph/9512264. 47

[41] ATLAS Collaboration, "ATLAS Run 1 searches for direct pair production of third-generation squarks at the Large Hadron Collider", arXiv:1506.08616. xix, 50

[42] DELPHI Collaboration, "Search for supersymmetric partners of top and bottom quarks at $\sqrt{s} = 189$ GeV", *Phys. Lett B 496, 59* (2000). 49

[43] L3 Collaboration, "Searches for Scalar Quarks in e^+e Interactions at $\sqrt{s} =$ 189 GeV", *Phys. Lett B 471, 308* (1999). 49

[44] OPAL Collaboration, "Search for Scalar Top and Scalar Bottom Quarks at LEP", *Phys. Lett B 545, 272* (2002). 49

[45] ALEPH Collaboration, "Search for scalar quarks in e^+e^- collisions at $\sqrt{s}$ up to 209 GeV", *Phys. Lett B 537, 5* (2002). 49

[46] D0 Collaboration, "Search for scalar top quarks in the acoplanar charm jets and missing transverse energy final state in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV", *Phys. Lett. B 665, 1* (2008). 49

[47] CDF Collaboration, "Search for Scalar Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV", *JHEP* **10** (2012) 158, doi:10.1007/JHEP10(2012)158, arXiv:1203.4171. 49

[48] ATLAS Collaboration, "Search for pair-produced top squarks decaying into charm quarks and the lightest neutralinos using 20.3 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector at the LHC", ATLAS CONF Note 2013-068, 2013. 49

[49] A. Djouadi, Y. Mambrini, "Decays of the lightest top squark", *Phys. Rev. D.Vol. 61, 095006* (2000). 50

[50] A. Djouadi, Y. Mambrini, "Three-body decays of top and bottom squarks", *Phys. Rep. D (2001) 115005* (2001). 50

[51] ATLAS Collaboration, "Search for a supersymmetric partner to the top quark in final states with jets and missing transverse momentum at $\sqrt{s} = 7$ TeV with the ATLAS detector", `arXiv:1208.1447`. Submitted to *Phys. Rev. Lett.* 50

[52] ATLAS Collaboration, "Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 7$ TeV pp collisions using 4.7 fb^{-1} of ATLAS data", `arXiv:1208.2590`. Submitted to *Phys. Rev. Lett.* 50

[53] ATLAS Collaboration, "Search for light scalar top quark pair production in final states with two leptons with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions", `arXiv:1208.4305`. Submitted to *Eur. Phys. J C.* 50

[54] ATLAS Collaboration, "Search for light top squark pair production in final states with leptons and b-jets with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions", `arXiv:1209.2102`. Submitted to *Phys. Lett. B.* 50

[55] ATLAS Collaboration, "Search for a heavy top-quark partner in final states with two leptons with the ATLAS detector at the LHC", `arXiv:1209.4186`. Submitted to *JHEP.* 50

[56] ATLAS Collaboration, "Searches for direct scalar top pair production in final states with two leptons using the stransverse mass variable and a multivariate analysis technique in $\sqrt{s} = 8$ TeV pp collisions using 20.3 fb^{-1} of ATLAS data", ATLAS CONF Note 2013-065, 2013. 50

[57] ATLAS Collaboration, "Search for direct production of the top squark in the all-hadronic $t\bar{t} + E_T^{miss}$ final state in 21 fb^{-1} of p-p collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector", ATLAS CONF Note 2013-024, 2013. 50

[58] ATLAS Collaboration, "Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions using 21 fb^{-1} of ATLAS data", ATLAS CONF Note 2013-037, 2013. 50

[59] ATLAS Collaboration, "Search for direct top-squark pair production in final states with two leptons in pp collisions at $\sqrt{s} = 8$TeV with the ATLAS detector", *JHEP* **06** (2014) 124, doi:10.1007/JHEP06(2014)124, arXiv:1403.4853. xxix, 50, 329, 404

[60] CMS Collaboration, "Search for top squarks in multijet events with large missing momentum in proton-proton collisions in 8 TeV", CMS Physics Analysis Summary CMS-PAS-SUS-13-015, CERN, 2013. arXiv:1207.7235, http://cdsweb.cern.ch/record/1635353. 51, 281

[61] CMS Collaboration, "Searches for 3rd generation squark production in fully hadronic final states in pp collisions at 8 TeV", CMS Physics Analysis Summary CMS-PAS-SUS-14-011, CERN, 2015. arXiv:1503.08037. 51

[62] CMS Collaboration, "Search for top-squark pair production in the single lepton final state in pp collisions at 8 TeV", CMS Physics Analysis Summary CMS-PAS-SUS-13-011, CERN, 2013. arXiv:1308.1586. 51, 197, 198, 281

[63] D0 Collaboration Collaboration, "Observation of the Top Quark", *Phys. Rev. Lett.* **74** (Apr, 1995) 2632–2637, doi:10.1103/PhysRevLett.74.2632. 53

[64] CDF Collaboration Collaboration, "Observation of Top Quark Production in $\overline{p}p$ Collisions with the Collider Detector at Fermilab", *Phys. Rev. Lett.* **74** (Apr, 1995) 2626–2631, doi:10.1103/PhysRevLett.74.2626. 53

[65] L. Evans and P. Bryant, "LHC Machine", *JINST* **3** (2008) S08001, doi:10.1088/1748-0221/3/08/S08001. xix, 54, 56, 59, 60

[66] J.-L. Caron, "The LHC injection complex", (May, 1993). AC Collection. Legacy of AC. Pictures from 1992 to 2002. xix, 57

[67] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004. xix, xx, 62, 65, 66, 67, 72, 73, 79, 90, 91, 95, 96, 101

[68] A. Herve et al., "Status of the CMS magnet", *IEEE Trans.Appl.Supercond.* **12** (2002) 385–398, doi:10.1109/TASC.2002.1018425. 66

[69] CMS Collaboration, "Precise mapping of the magnetic field in the CMS barrel yoke using cosmic rays", *JINST* **5** (2010) T03021, doi:10.1088/1748-0221/5/03/T03021. xix, 68

[70] V. Klyukhin et al., "The CMS Magnetic Field Map Performance", *IEEE Trans.Appl.Supercond.* **20** (2010) 152–155, doi:10.1109/TASC.2010.2041200, arXiv:1110.0607. 68

[71] CMS Collaboration, "Performance of the missing transverse energy reconstruction by the CMS experiment in $\sqrt{s} = 8$ TeV pp data"", *JINST* **10** (2014) P02006, doi:doi:10.1088/1748-0221/10/02/P02006, arXiv:1411.0511. xiv, xxi, xxii, 69, 145, 150, 151, 154, 155, 156, 157, 158, 159, 160, 164, 191, 229, 297, 362

[72] R. Gluckstern, "Uncertainties in track momentum and direction, due to multiple scattering and measurement errors", *Nucl.Instrum.Meth.* **24** (1963) 381–389, doi:10.1016/0029-554X(63)90347-1. 71

[73] V. Karimki, "Explicit Covariance Matrix for Particle Measurement Precision",. 71

[74] CMS Collaboration CollaborationV. Karimäki, et al., "The CMS tracker system project: Technical Design Report". Technical Design Report CMS. CERN, Geneva, 1997. 73

[75] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014), no. 10, P10009, doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569. xix, 75, 114, 116

[76] K. A. Olive et al., "The Review of Particle Physics: Passage of Particles Through Matter", *Chin. Phys. C* **C38** (2014) 090001. 77

[77] A. Annenkov, M. Korzhik, and P. Lecoq, "Lead tungstate scintillation material", *Nucl.Instrum.Meth.* **A490** (2002) 30–50, doi:10.1016/S0168-9002(02)00916-6. 77

[78] X. Qu, L. Zhang, and R. Zhu, "Radiation induced color centers and light monitoring for lead tungstate crystals", *IEEE Trans.Nucl.Sci.* **47** (2000) 1741–1747, doi:10.1109/TNS.2000.914439. 77

[79] CMS Collaboration, "2012 ECAL detector performance plots", CMS Detector Performance Summary CMS-DP-2013-007, CERN, 2013. xix, 81, 83

[80] G. Akopdzhanov, A. Inyakin, and P. Shuvalov, "Photomultiplier short-term instability", *Nuclear Instruments and Methods* **161** (1979), no. 2, 247 – 257, doi:http://dx.doi.org/10.1016/0029-554X(79)90391-4. 82

[81] CMS Collaboration Collaboration, "Electromagnetic calorimeter calibration with 7 TeV data", Technical Report CMS-PAS-EGM-10-003, CERN, Geneva, 2010. xix, 83

[82] CMS Collaboration, "Energy Calibration and Resolution of the CMS Electromagnetic Calorimeter in *pp* Collisions at $\sqrt{s} = 7$ TeV", *JINST* **8** (2013) P09009, doi:10.1088/1748-0221/8/09/P09009, arXiv:1306.2016. xix, 83

[83] P. Adzic, "Energy resolution of the barrel of the CMS Electromagnetic Calorimeter", *Journal of Instrumentation* **2** (2007), no. 04, P04004. 84

[84] CMS Collaboration, "Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at sqrt(s) = 8 TeV", *JINST* **10** (2015), no. 08, P08010, doi:10.1088/1748-0221/10/08/P08010, arXiv:1502.02702. xix, 85, 124, 125, 228

[85] CMS Collaboration, "Performance of the CMS hadron calorimeter with cosmic ray muons and LHC beam data", *JINST* **5** (2010) T03012, doi:10.1088/1748-0221/5/03/T03012, arXiv:0911.4991. 89, 90

[86] J. Freeman, "Silicon photomultipliers for the CMS hadron calorimeter", *Nucl. Instrum. Meth. A* **617** (2010) 393, doi:10.1016/j.nima.2009.10.132. Proceedings of the 11th Pisa Meeting on Advanced Detectors. 90

[87] CMS Collaboration, "Technical proposal for the upgrade of the CMS detector through 2020", Technical Report CERN-LHCC-2011-006. LHCC-P-004, CERN, Geneva, Jun, 2011. xx, 93

[88] E. Gatti, A. Longoni, R. Boie, and V. Radeka, "Analysis of the position resolution in centroid measurements in MWPC", *Nuclear Instruments and Methods in Physics Research* **188** (1981), no. 2, 327 – 346, doi:http://dx.doi.org/10.1016/0029-554X(81)90513-9. 95

[89] CMS Collaboration, "CMS Luminosity Based on Pixel Cluster Counting - Summer 2012 Update", CMS Physics Analysis Summary CMS-PAS-LUM-12-001, CERN, Geneva, 2012. 97, 98, 99

[90] CMS Collaboration, "CMS Luminosity Based on Pixel Cluster Counting - Summer 2013 Update", CMS Physics Analysis Summary CMS-PAS-LUM-13-001, CERN, Geneva, 2013. 97, 98, 99

[91] V. Balagura, "Notes on van der Meer Scan for Absolute Luminosity Measurement", *Nucl. Instrum. Meth. A* **654** (2011) 634, doi:10.1016/j.nima.2011.06.007, arXiv:1103.1129. 100

[92] L. Agostino et al., "Commissioning of the CMS High Level Trigger", *JINST* **4** (2009) P10005, doi:10.1088/1748-0221/4/10/P10005, arXiv:0908.1065. 103

[93] D. Trocino, "The CMS High Level Trigger", *J. Phys. Conf.* **513** (2014) 012036, doi:10.1088/1742-6596/513/1/012036. 103, 105

[94] CMS Collaboration, "Performance of the CMS Level-1 Trigger", *PoS* **ICHEP2012** (2013) 508, arXiv:1302.2469. 105

[95] G. Bauer et al., "Automating the CMS DAQ", *J.Phys.Conf.Ser.* **513** (2014) 012031, doi:10.1088/1742-6596/513/1/012031. 105

[96] J. C. Collins, D. E. Soper, and G. F. Sterman, "Factorization of Hard Processes in QCD", *Adv. Ser. Direct. High Energy Phys.* **5** (1989) 1–91, doi:10.1142/9789814503266_0001, arXiv:hep-ph/0409313. 106

[97] T. Sjöstrand, S. Mrenna, and P. Skands, "PYTHIA 6.4 physics and manual", *JHEP* **05** (2006) 026, doi:10.1088/1126-6708/2006/05/026, arXiv:hep-ph/0603175. 107, 171

[98] J. Alwall et al., "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations", *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301. 107, 171

[99] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with parton shower simulations: the POWHEG method", *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092. 107, 171

[100] Z. Was, "TAUOLA the library for τ lepton decay, and KKMC/KORALB/KORALZ... status report", *Nucl. Phys. B, Proc. Suppl.* **98** (2001) 96, doi:10.1016/S0920-5632(01)01200-2. 108, 174

[101] J. Allison et al., "Geant4 a simulation toolkit", *Nucl. Instrum. Meth. A* **506** (2003), no. 3, 250 – 303, doi:10.1016/S0168-9002(03)01368-8. 108

[102] J. Allison et al., "Geant4 developments and applications", *IEEE Trans. Nucl. Sci.* **53** (February, 2006) 270 –278, doi:10.1109/TNS.2006.869826. 108

[103] S. Abdullin et al., "The Fast Simulation of the CMS Detector at LHC", *Journal of Physics: Conference Series* **331** (2011), no. 3, 032049. 108, 109

[104] A. Giammanco, "The Fast Simulation of the CMS Experiment", *Journal of Physics: Conference Series* **513** (2014), no. 2, 022012. 108

[105] CMS Collaboration, "Comparison of the Fast Simulation of CMS with the first LHC data",. 109

[106] CMS Collaboration, "Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus, and MET", CMS Physics Analysis Summary CMS-PAS-PFT-09-001, CERN, 2009. 110, 129

[107] CMS Collaboration Collaboration, "Commissioning of the Particle-flow Event Reconstruction with the first LHC collisions recorded in the CMS detector", Technical Report CMS-PAS-PFT-10-001, 2010. 110

[108] CMS Collaboration, "Commissioning of the Particle-Flow Reconstruction in Minimum-Bias and Jet Events from pp Collisions at 7 TeV", CMS Physics Analysis Summary CMS-PAS-PFT-10-002, CERN, 2010. 110, 129, 177

[109] CMS Collaboration, "Commissioning of the particle-flow event reconstruction with leptons from J/Ψ and W decays at 7 TeV", CMS Physics Analysis Summary CMS-PAS-PFT-10-003, CERN, 2010. 111

[110] F. Beaudette, "The CMS Particle Flow Algorithm", in *Proceedings, International Conference on Calorimetry for the High Energy Frontier (CHEF 2013)*, p. 295. École Polytechnique, Paris, 2013. `arXiv:1401.8155`. 111

[111] CMS Collaboration, "CMS Physics: Technical Design Report Volume 1: Detector Performance and Software", Technical Report CMS-TDR-8-1, CERN, Geneva, 2006. 112

[112] P. Billoir, "Progressive track recognition with a Kalman like fitting procedure", *Comput. Phys. Commun.* **57** (1989) 390–394, `doi:10.1016/0010-4655(89)90249-X`. 115

[113] D. Giordano and G. Sguazzoni, "CMS reconstruction improvements for the tracking in large pile-up events", *J. Phys. Conf.* **396** (2012), no. 2, 022044, `doi:10.1088/1742-6596/396/2/022044`. xiv, 117

[114] CMS Collaboration, "Tracking and Primary Vertex Results in First 7 TeV Collisions", CMS Physics Analysis Summary CMS-PAS-TRK-10-005, CERN, Geneva, 2010. 118

[115] CMS Collaboration, "Performance of CMS muon reconstruction in cosmic-ray events", *Journal of Instrumentation* **5** (2010), no. 03, T03022. 118

[116] CMS Collaboration Collaboration, "Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV", *JINST* **7** (2012) P10002, `doi:10.1088/1748-0221/7/10/P10002`, `arXiv:1206.4071`. 119, 179, 228

[117] CMS Collaboration, "The performance of the CMS muon detector in proton-proton collisions at sqrt(s) = 7 TeV at the LHC", *JINST* **8** (2013) P11002, `doi:10.1088/1748-0221/8/11/P11002`, `arXiv:1306.6905`. 119, 179, 228

[118] CMS Collaboration, "Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at s = 8 TeV", *JINST* **10** (2015), no. 06, P06005, `doi:10.1088/1748-0221/10/06/P06005`, `arXiv:1502.02701`. xx, 120, 122, 124, 180, 228

[119] G. P. Salam et al., "The Anti-k_{T} jet clustering algorithm", *JHEP* **0804** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189. 126

[120] CMS Collaboration, "Jet Energy Scale and Resolution in the CMS Experiment", *JINST* (2015) To be submitted. xxi, 129, 130, 132, 133, 134, 228

[121] CMS Collaboration Collaboration, "Identification of b-quark jets with the CMS experiment", *JINST* **8** (2013) P04013, doi:10.1088/1748-0221/8/04/P04013, arXiv:1211.4462. 136, 138

[122] CMS Collaboration, "Performance of b tagging at sqrt(s)=8 TeV in multijet, ttbar and boosted topology events",. xxi, 136, 141, 142, 143, 144

[123] CMS Collaboration, "Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **06** (2013) 81, doi:10.1007/JHEP06(2013)081, arXiv:1303.4571. 146

[124] M. P. Marionneau and C. Veelken, "Study of missing transverse energy and missing transverse energy corrections in events containing one Z boson decaying into two muons with the 2012 data", *CMS Note* **2012/333** (2012). xxi, 149

[125] CMS Collaboration, "Determination of jet energy calibration and transverse momentum resolution in CMS", *JINST* **6** (2011) P11002, doi:10.1088/1748-0221/6/11/P11002, arXiv:1107.4277. 150

[126] CMS Collaboration, "Studies of Tracker Material", CMS Physics Analysis Summary CMS-PAS-TRK-10-003, 2010. 150

[127] CMS Collaboration, "Jet Energy Resolution in CMS at $\sqrt{s} = 7$ TeV", *CMS PAS* **JME-10-014** (2010). 151

[128] MET Working Group., "Performance of the Missing Transverse Energy Reconstruction in the 8 TeV Data", *CMS PAS* **JME/2012/002** (2012). 151, 154

[129] B. Calvert et al., "Study of missing energy performance with photons in 8 TeV data", *CMS Analysis Note* (2012), no. CMS-AN-2012-341,. xxi, 154

[130] B. Calvert et al., "Study of missing energy performance with photons in 8 TeV data", *CMS Analysis Note* (2013), no. CMS-AN-2013-373,. xxi, 154

[131] CMS Collaboration, "Missing transverse energy performance of the CMS detector", *JINST* **6** (2011) P09001, doi:doi:10.1088/1748-0221/6/09/P09001, arXiv:1106.5048. 155

[132] A. Apresyan, "Identification and mitigation of anomalous signals in CMS HCAL", CMS Conference Report CMS-CR-2012-238, 2012. 155

[133] CDF Collaboration, "Precise measurement of the W-boson mass with the CDF II detector", *Phys.Rev.Lett.* **108** (2012) 151803, doi:10.1103/PhysRevLett.108.151803, arXiv:1203.0275. 165

[134] D0 Collaboration, "Measurement of the W boson mass with the D0 detector", *Phys.Rev.* **D89** (2014), no. 1, 012005, doi:10.1103/PhysRevD.89.012005, arXiv:1310.8628. 165

[135] CDF, D0 Collaboration, "2012 Update of the Combination of CDF and D0 Results for the Mass of the W Boson", arXiv:1204.0042. 165

[136] S. Alioli, P. Nason, C. Oleari, and E. Re, "NLO single-top production matched with shower in POWHEG: s- and t-channel contributions", *JHEP* **09** (2009) 111, doi:10.1088/1126-6708/2009/09/111, arXiv:0907.4076. [Erratum: doi:10.1007/JHEP02(2010)011]. 171

[137] E. Re, "Single-top Wt-channel production matched with parton showers using the POWHEG method", *Eur. Phys. J. C* **71** (2011) 1547, doi:10.1140/epjc/s10052-011-1547-z, arXiv:1009.2450. 171

[138] S. Frixione et al., "The MC@NLO 4.0 Event Generator", arXiv:1010.0819. 171, 175

[139] M. Czakon, P. Fiedler, and A. Mitov, "Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\frac{4}{S})$", *Phys.Rev.Lett.* **110** (2013) 252004, doi:10.1103/PhysRevLett.110.252004, arXiv:1303.6254. 174

[140] M. Aliev et al., "HATHOR: HAdronic Top and Heavy quarks crOss section calculatoR", *Comput.Phys.Commun.* **182** (2011) 1034–1046, doi:10.1016/j.cpc.2010.12.040, arXiv:1007.1327. 174

[141] P. Kant et al., "HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions", *Comput.Phys.Commun.* **191** (2015) 74–89, doi:10.1016/j.cpc.2015.02.001, arXiv:1406.4403. 174

[142] N. Kidonakis, "Top Quark Production", doi:10.3204/DESY-PROC-2013-03/Kidonakis, arXiv:1311.0283. 174

[143] K. Melnikov and F. Petriello, "Electroweak gauge boson production at hadron colliders through $\mathcal{O}(\alpha_s^2)$", *Phys. Rev. D* **74** (2006) 114017, doi:10.1103/PhysRevD.74.114017, arXiv:hep-ph/0609070. 174

[144] J. M. Campbell, R. K. Ellis, and C. Williams, "Vector boson pair production at the LHC", *JHEP* **07** (2011) 018, doi:10.1007/JHEP07(2011)018, arXiv:1105.0020. 174

[145] W. Beenakker et al., "Supersymmetric top and bottom squark production at hadron colliders", *JHEP* **1008** (2010) 098, doi:10.1007/JHEP08(2010)098, arXiv:1006.4771. 175

[146] M. Kramer et al., "Supersymmetry production cross sections in *pp* collisions at $\sqrt{s} = 7$ TeV", arXiv:1206.2892. 175

[147] CMS Collaboration, "Performance of muon identification in pp collisions at s**0.5 = 7 TeV", Technical Report CMS-PAS-MUO-10-002, CERN, 2010. Geneva, 2010. 179

[148] CMS Collaboration Collaboration, "Electron reconstruction and identification at $\sqrt{(s)} = 7$ TeV", Technical Report CMS-PAS-EGM-10-004, CERN, Geneva, 2010. 180

[149] N. Saoulidou, "Particle Flow Jet Identification Criteria", CMS Analysis Note CMS-AN-2010-003, CERN, Geneva, Jun, 2010. 181

[150] CMS Collaboration, "Jet Identification", (2012). 181

[151] C. S. Rogan, "Searches for New Symmetries in pp Collisions with the Razor Kinematic Variables at $\sqrt{\mathbf{s}} = 7$ TeV". PhD thesis, Caltech, 2013. 188, 219

[152] CMS Collaboration, "Measurement of differential top-quark pair production cross sections in *pp* colisions at $\sqrt{s} = 7$ TeV", *Eur.Phys.J.* **C73** (2013), no. 3, 2339, doi:10.1140/epjc/s10052-013-2339-4, arXiv:1211.2220. 196

[153] CMS Collaboration, "Measurement of the differential cross section for top quark pair production in pp collisions at sqrt(s) = 8 TeV", arXiv:1505.04480. 196

[154] D. Barge et al., "Hadronic recoil studies of heavy boosted systems", *CMS-AN2013-059* (2013). 197, 198

[155] H.-C. Cheng and Z. Han, "Minimal Kinematic Constraints and m(T2)", *JHEP* **0812** (2008) 063, doi:10.1088/1126-6708/2008/12/063, arXiv:0810.5178. 199

[156] CMS Collaboration Collaboration, "Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{(s)} = 8$ TeV", *CMS-PAS-TOP-12-007* (2012). 208, 217

[157] CMS Collaboration Collaboration, "Measurement of the $t\bar{t}$ production cross section and the top quark mass in the dilepton channel in pp collision at $\sqrt{(s)} = 7$ TeV", *JHEP* **1107** (2011) 049. 208, 217

[158] CMS Collaboration, "Search for new physics in events with same-sign dileptons and jets in pp collisions at $\sqrt{s} = 8$ TeV", *JHEP* **1401** (2014) 163, doi:10.1007/JHEP01(2015)014,10.1007/JHEP01(2014)163, arXiv:1311.6736. 218

[159] H. Bakhshian et al., "Computing the contamination from fakes in leptonic final states", CMS Analysis Note CMS-AN-2010-261, CERN, Geneva, 2010. 218

[160] S. Alekhin et al., "The PDF4LHC Working Group Interim Report", arXiv:1101.0536. 232

[161] M. Botje et al., "The PDF4LHC Working Group Interim Recommendations", arXiv:1101.0538. 232

[162] K. A. Olive et al., "The Review of Particle Physics: The Mass and Width of the W Boson", *Chin. Phys. C* **C38** (2014) 090001. 234

[163] ATLAS Collaboration, "Search for direct top-squark pair production in final states with two leptons in pp collisions at $\sqrt{s} = 8$TeV with the ATLAS detector", *JHEP* **06** (2014) 124, doi:10.1007/JHEP06(2014)124, arXiv:1403.4853. 277

[164] Rokach, Lior and Maimon, O., "Data mining with decision trees: theory and applications", *World Scientific Pub Co Inc.* (2008). 277, 283, 415

[165] ATLAS Collaboration, "Measurement of Spin Correlation in Top-Antitop Quark Events and Search for Top Squark Pair Production in pp Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector", *Phys. Rev. Lett.* **114** (2015), no. 14, 142001, doi:10.1103/PhysRevLett.114.142001, arXiv:1412.4742. 278

[166] ATLAS Collaboration, "Measurement of the $t\bar{t}$ production cross-section using $e\mu$ events with b -tagged jets in pp collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS detector", *Eur. Phys. J.* **C74** (2014), no. 10, 3109, doi:10.1140/epjc/s10052-014-3109-7, arXiv:1406.5375. 278

[167] A. Aubin, "Search for direct stop pair production in the semi-leptonic channel at 8 TeV", CMS Analysis Note CMS-AN-2014-067, CERN, Geneva, Feb, 2015. 281

[168] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics", *The European Physical Journal C* **71** (2011), no. 2, doi:10.1140/epjc/s10052-011-1554-0. 283, 416, 434, 441

[169] S. Baker and R. D. Cousins, "Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms", *Nucl. Instrum. Meth.* **221** (1984) 437–442, doi:10.1016/0167-5087(84)90016-4. xxx, 343, 344, 346

[170] O. Vitells and E. Gross, "Estimating the significance of a signal in a multi-dimensional search", *Astroparticle Physics* **35** (December, 2011) 230–234, doi:10.1016/j.astropartphys.2011.08.005, arXiv:1105.4355. 351

[171] C. Borschensky et al., "Squark and gluino production cross sections in pp collisions at $\sqrt{s}$ = 13, 14, 33 and 100 TeV", *Eur. Phys. J.* **C74** (2014), no. 12, 3174, doi:10.1140/epjc/s10052-014-3174-y, arXiv:1407.5066. xxxi, 361

[172] "Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s}$ = 7 and 8 TeV", Technical Report ATLAS-CONF-2015-044, CERN, Geneva, Sep, 2015. xxxiii, 400, 401

[173] C. Verhaaren. Private communication, 2015. xxxiii, 402

[174] J. Alcaraz Maestre and B. de la Cruz Martinez, "Simulating arbitrary SUSY-mixing scenarios in direct stop production via event reweighting", CMS Analysis Note CMS-AN-2013-128, CERN, Geneva. 411, 413

[175] ATLAS Collaboration, "Procedure for the LHC Higgs boson search combination in summer 2011",. xxxv, 434, 446

[176] CMS Collaboration, "Combined results of searches for the standard model Higgs boson in *pp* collisions at $\sqrt{s}$ = 7 TeV", *Phys. Lett.* **B710** (2012) 26–48, doi:10.1016/j.physletb.2012.02.064, arXiv:1202.1488. 434, 435

[177] A. L. Read, "Presentation of search results: the CLs technique", *J. Phys.* **G28** (2002) 2693, doi:10.1088/0954-3899/28/10/313. 434

[178] T. Junk, "Confidence level computation for combining searches with small statistics", *Nucl. Instrum. Meth.* **A434** (1999) 435, arXiv:9902006. 434

[179] K. Cranmer, "Statistical Challenges for Searches for New Physics at the LHC", in *Statistical Problems in Particle Physics, Astrophysics and Cosmology*, L. Lyons and M. Karagöz Ünel, eds., p. 112. 2006. arXiv:physics/0511028. doi:10.1142/9781860948985_0026. 438

[180] S. S. Wilks, "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses", *Ann. Math. Statist.* **9** (03, 1938) 60–62, doi:10.1214/aoms/1177732360. 443

[181] R. B. Davies, "Hypothesis Testing when a Nuisance Parameter is Present Only Under the Alternatives", *Biometrika* **74** (1987), no. 1, 33–43. 443

[182] E. Gross and O. Vitells, "Trial factors or the look elsewhere effect in high energy physics", *Eur. Phys. J.* **C70** (2010) 525–530, doi:10.1140/epjc/s10052-010-1470-8, arXiv:1005.1891. 443, 445

[183] O. Vitells and E. Gross, "Estimating the significance of a signal in a multi-dimensional search", *Astroparticle Physics* **35** (December, 2011) 230–234, doi:10.1016/j.astropartphys.2011.08.005, arXiv:1105.4355. 443, 445

[184] L. Demortier and L. Lyons, "Everything you always wanted to know about pulls", (2002). CDF/DOC/STATISTICS/PUBLIC/5776. 447

[185] M. B. Priestley and M. T. Chao, "Non-Parametric Function Fitting", *Journal of the Royal Statistical Society. Series B (Methodological)* **34** (1972), no. 3, pp. 385–392. 450, 451

[186] J. K. Benedetti, "On the Nonparametric Estimation of Regression Functions", *Journal of the Royal Statistical Society. Series B (Methodological)* **39** (1977), no. 2, pp. 248–253. 451

[187] B. W. Silverman, "Density Estimation for Statistics and Data Analysis". Chapman & Hall, London, 1986. 451

[188] G. Wahba and S. Wold, "A completely automatic french curve: fitting spline functions by cross validation", *Communications in Statistics* **4** (1975), no. 1, 1–17, doi:10.1080/03610927508827223, arXiv:http://dx.doi.org/10.1080/03610927508827223. 454

[189] M. G. Schimek, "Estimation and inference in partially linear models with smoothing splines", *Journal of Statistical Planning and Inference* **91** (2000), no. 2, 525 – 540, doi:http://dx.doi.org/10.1016/S0378-3758(00)00197-X. Perspectives in modern statistical inference: parametric semiparametrics, nonparametrics. 454

[190] X.-W. Chang and L. Qu, "Wavelet estimation of partially linear models", *Computational Statistics & Data Analysis* **47** (2004), no. 1, 31 – 48, doi:http://dx.doi.org/10.1016/j.csda.2003.10.018. 454

[191] G. Punzi, "Sensitivity of searches for new signals and its optimization", *eConf* **C030908** (2003) MODT002, arXiv:physics/0308063. xxxvi, 473, 475, 477, 478