Title

Search for new physics using the MT2 variable in all-hadronic final states produced in 13 TeV proton-proton collisions at the CMS detector

Author

Marsh, Bennett

Publication Date

2020

University of California
Santa Barbara

Search for new physics using the M_{T2} variable in all-hadronic final states produced in 13 TeV proton-proton collisions at the CMS detector

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Bennett J. Marsh

Committee in charge:

Professor Claudio Campagnari, Chair
Professor David Stuart
Professor Nathaniel Craig

June 2020

The Dissertation of Bennett J. Marsh is approved.

Professor David Stuart

Professor Nathaniel Craig

Professor Claudio Campagnari, Committee Chair

May 2020

Search for new physics using the M_{T2} variable in all-hadronic final states produced in 13 TeV proton-proton collisions at the CMS detector

Acknowledgments

The work performed for this dissertation was only possible with the help and guidance of many mentors, colleagues, and friends, to whom I would like to express my gratitude. First, I would like to thank my advisor, Claudio Campagnari, for supporting me since my arrival five years ago and for teaching me much of what I know about experimental particle physics. I could not have asked for a better mentor.

Frank Golf and Indara Suarez provided much-needed guidance and direction during the early years of my graduate research. It was a joy working with them and everyone else in Claudio's group, including Nick Amin, Sicheng Wang, Franny Setti, Ulascan Sarica, Hualin Mei, Jonathan Guiang, Jerry Ling, and Japnidh Thakral. I would especially like to thank Nick and Sicheng for being great co-workers and friends from my first day at UCSB (and much of the work done here would have been significantly harder without Nick's various secret projects. Thank you for DIS and Metis.)

My work would not have been possible without the knowledge and technical support provided by the many professors, post-docs, and graduate students comprising the Surf-n-Turf collaboration. In particular, Dominick Olivito and Giovanni Zevi Della Porta showed me the ropes of the M_{T2} analysis, and Dylan Gilbert, Mario Masciovecchio, Slava Krutelyov, and Avi Yagil were invaluable partners and mentors in the push to complete its latest iteration.

I started work with the milliQan experiment in my first month as a graduate student, and finished with its first public result in my final week. I have greatly enjoyed working on this unique project throughout my time here with everyone involved at UCSB and elsewhere. In addition to those already mentioned above, I learned a tremendous amount from David Stuart, who has acted many times as almost a second advisor for this and other projects. Matthew Citron provided much guidance in our efforts to extract useful physics from the prototype detector.

Finally, I must thank my family for everything they have provided me. My parents Jim and Janice Marsh have unconditionally supported me always, and encouraged my passion for science from an early age. I would not have had the opportunities I have had without them. And to my lovely wife, Amber, who moved to California and put up with the grad school life for five years as I pursued my goals. I would not have been able to do this without your support, and I am so excited to start the next chapter of our lives together.

Curriculum Vitae

Bennett J. Marsh

Education

2020	Ph.D. in Physics, University of California, Santa Barbara.
2018	M.A. in Physics, University of California, Santa Barbara.
2015	B.S. in Physics and Mathematics, Purdue University, West Lafayette, IN.

Publications

- A. Ball *et al.*, "Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Submitted to Phys. Rev. D* (2020). [arXiv:2005.06518]
- CMS Collaboration, "Searches for physics beyond the standard model with the M_{T2} variable in hadronic final states with and without disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Eur. Phys. J. C* **80** (2020), 3. [arXiv:1909.03460].
- CMS Collaboration, "A MIP Timing Detector for the CMS Phase-2 Upgrade." Tech. Rep. CMS-TDR-020 (CERN, 2019).
- CMS Collaboration, "Constraints on models of scalar and vector leptoquarks decaying to a quark and a neutrino at $\sqrt{s} = 13$ TeV." *Phys. Rev. D* **98** (2018), 032005. [arXiv:1805.10228].
- CMS Collaboration, "Search for new phenomena with the M_{T2} variable in the all-hadronic final state produced in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Eur. Phys J. C* **77** (2017), 710. [arXiv:1705.04650].
- A. Ball *et al.*, "A Letter of Intent to Install a Milli-Charged Particle Detector at LHC P5." (2016). [arXiv:1607.04669].

Public Talks

- "The milliQan experiment: a search for milli-charged particles at the LHC". Light Dark Matter @ Accelerators Workshop, Nov. 20–22, 2019. Venice, Italy.
- "Monitoring tools for the muon system in the Compact Muon Solenoid (CMS) detector". Machine Learning in Science and Engineering Conference, June 6–8, 2018. Carnegie Mellon University, Pittsburgh, PA.

Abstract

Search for new physics using the M_{T2} variable in all-hadronic final states produced in 13 TeV proton-proton collisions at the CMS detector

by

Bennett J. Marsh

A search for phenomena beyond the Standard Model (BSM) is performed using events with hadronic jets and significant transverse momentum imbalance. The results are based on a sample of proton-proton collisions at a center-of-mass energy of 13 TeV, collected by the CMS experiment at the Large Hadron Collider in 2016–2018 and corresponding to an integrated luminosity of 137 fb^{-1}. The search is based on signal regions defined by the hadronic energy in the event, the jet multiplicity, the number of b-tagged jets, and the value of the kinematic variable M_{T2} for events with at least two jets. For events with exactly one jet, the transverse momentum of the jet is used instead. No significant excess event yield is observed above the predicted Standard Model background. This is used to constrain a range of BSM models that predict the following: the pair production of gluinos and squarks in the context of supersymmetry models conserving R-parity; the resonant production of a colored scalar state decaying to a massive Dirac fermion and a quark; and the pair production of scalar and vector leptoquarks each decaying to a neutrino and a top, bottom, or light-flavor quark. In most of the cases, the results obtained are the most stringent constraints to date. The analysis is published in the European Physical Journal C vol. 80, #3.

Additionally, the first search at a hadron collider for elementary particles with charges much smaller than the electron charge is presented. These results are based on a sample of proton-proton collisions at a center-of-mass energy of 13 TeV provided by the LHC in 2018, corresponding to an integrated luminosity of 37.5 fb^{-1}. A prototype scintillator-based detector is deployed near the CMS interaction point to conduct a search sensitive to particles with charges $\leq 0.3e$. The existence of new particles with masses between 20 and 4700 MeV is excluded at 95% confidence level for charges varying between $0.006e$ and $0.3e$, depending on their mass. New sensitivity is achieved for masses larger than 700 MeV.

Contents

Chapter 1

The Standard Model and Beyond

The Standard Model of particle physics describes all known fundamental particles and the interactions between them, with the exception of the gravitational force. Developed over decades in the latter half of the 20^{th} century, it has predicted experimental findings to an extraordinary degree of accuracy, and represents one of the crowning achievements of modern science. However, despite it successes, there are a number of known phenomena that cannot be explained with the current theory, giving physicists reason to look for an expanded model. In this chapter, we outline the Standard Model as it currently exists, discuss some of the challenges that the model faces, and give a brief summary of proposed extensions to the Standard Model that are relevant to this dissertation.

1.1 The Standard Model of particle physics

The Standard Model (SM) is at its core a quantum field theory defined by a local $SU(3) \times SU(2)_L \times U(1)$ gauge symmetry. We will explain what exactly this means shortly (essentially, each term gives rise to one of the three fundamental interactions), but for now we simply describe all of the known particles and their properties.

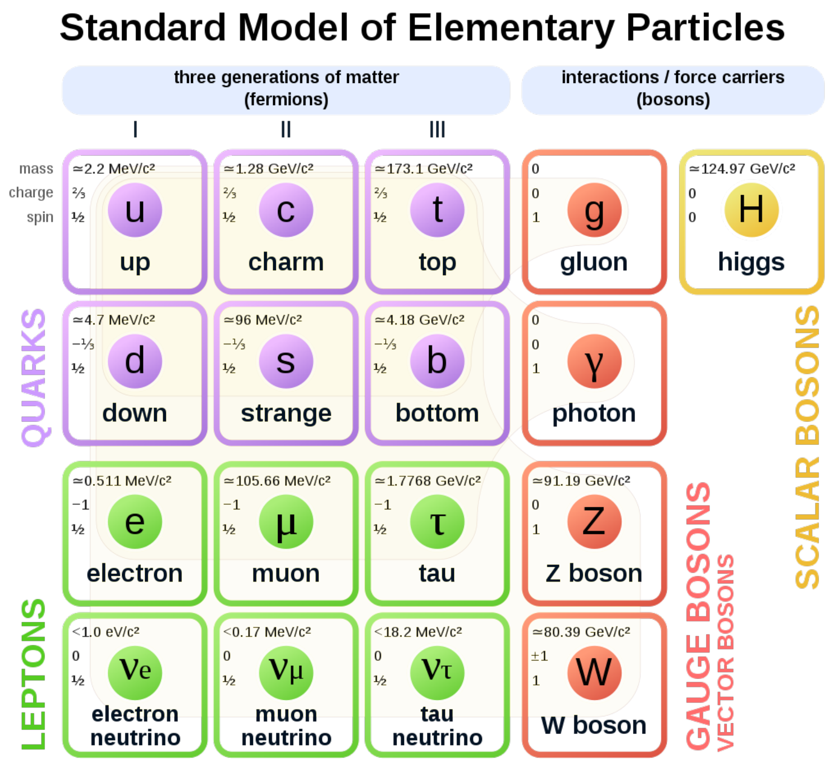

Figure 1.1: Diagram of all fundamental particles making up the Standard Model. There are 12 fermions (spin-$1/2$), consisting of six quarks (purple) and six leptons (green), and divided into three generations. There are additionally four force-carrying vector bosons (spin-1), that couple to the fermions and give rise to the strong, weak, and electromagnetic interactions. Finally, the scalar Higgs boson, proposed in 1964 and discovered in 2012, provides the mechanism by which the other particles acquire mass. (Image from [1])

1.1.1 Fundamental particles

Figure 1.1 shows all distinct particles in the SM, organized into groups. On the left, in purple and green, are the 12 fundamental fermions, defined by the value of their intrinsic angular momentum, or "spin", of $1/2$ (in units of $\hbar$). By the spin-statistics theorem, fermions obey the Pauli exclusion principle, meaning no two can occupy the same quantum state simultaneously.

The fermions are further defined by the various charges they carry (which determine how they interact with the bosons, as we will see). In green are the six leptons, three with electric charge of -1 (electron, muon and tau) and three that are electrically neutral (electron neutrino, muon neutrino, and tau neutrino). In purple are the 6 types of quarks. The up-type quarks (up, charm, and top) have electric charges of $+2/3$, and the down-type quarks (down, strange, and bottom) have electric charges of $-1/3$. The quarks, in contrast to the leptons, carry color charge, meaning they can interact via the strong interaction. The color charge can be one of three values, typically referred to as red, green, and blue. All quarks and leptons carry weak isospin, meaning they can interact via the weak interaction. Each of the fermions also has a corresponding antiparticle, which has the same mass but opposite charges.

The fermions can be classified into three generations as indicated in the diagram, with masses increasing in each generation. Due to various conservation laws arising from the allowed interactions, the first-generation particles are all stable, and hence form the building blocks of matter. The strong interaction allows up and down quarks to strongly bind to one another, forming protons (two ups and a down) and neutrons (two downs and an up). The strong interaction further binds these into nuclei, which themselves bind to electrons via the electromagnetic interaction to form atoms.

More generally, the strong interaction allows quarks to bind themselves into *hadrons*, either in quark-antiquark pairs called *mesons*, or in three-quark configurations called *baryons* (as well as more exotic structures such as *tetraquarks* [2], *pentaquarks* [3], and *glueballs* [4]). The net "color content" must be zero, which means mesons must contain e.g. a red quark and anti-red antiquark, and baryons must contain exactly one red, blue, and green quark. The only stable baryon is the proton (at least, its lifetime is longer than $\sim 10^{34}$ years [5]! The neutron decays to a proton, electron, and antineutrino with lifetime 15 minutes but becomes stable when confined in a nucleus). There are no stable mesons.

Moving to the right in the diagram, in red we have the four gauge bosons, which have spin-1 and which mediate the fundamental interactions. Their integer spin means they obey Bose statistics, and are not constrained by the Pauli exclusion principle as are fermions. The gluon is massless, electrically neutral and mediates the strong interaction. The photon is also massless and neutral and mediates the electromagnetic interaction. Finally, the W and Z bosons are massive and mediate the weak interaction. The Z is electrically neutral while the W carries charges of ± 1.

Finally, in yellow is the scalar (spin-0) Higgs boson. The Higgs mechanism was proposed in 1964 as an explanation for how gauge bosons can acquire mass [6–8]. A consequence of this mechanism is the prediction of a scalar boson of undetermined mass, that couples to all SM particles proportionally to their masses. The discovery of new boson of mass 125 GeV fitting these criteria (at least at the limits of current experimental precision) was announced by the CMS and ATLAS Collaborations in July 2012 [9, 10].

1.1.2 Fundamental interactions

Empirically, there are four known fundamental interactions, or "forces", between particles in our universe: the electromagnetic, strong, weak, and gravitational interactions. The electromagnetic force binds electrons to atomic nuclei and atoms into molecules, and hence is responsible for all of chemistry. It is also responsible for electromagnetic radiation such as visible light, radio waves, X-rays, etc. The strong force confines quarks into protons and neutrons and binds these together as atomic nuclei. The weak interaction is responsible for radioactivity and the nuclear fusion that powers stars, among other phenomena. Gravity, the mutual attraction of massive bodies, is understood through Einstein's theory of general relativity, and is not (yet) integrated with our quantum understanding of elementary particles. The other three interactions, on the other hand, are mathematically described by the SM.

Essentially, each of the interactions arises by requiring that the Lagrangian describing the particle dynamics is invariant under a certain *local gauge transformation*. This means that if the particle fields are transformed by some function that depends on spacetime position x, the Lagrangian is unchanged. The full details are given in any quantum field theory textbook (e.g. [11]), and we just give a brief summary of the main ideas here.

The simplest example is the electromagnetic interaction. Starting with the Lagrangian of a free fermion

$$\mathcal{L}_{\text{fermion}} = i\bar{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\bar{\psi}\psi, \tag{1.1}$$

where ψ is the spinor field of the fermion and m is its mass, we require that $\mathcal{L}$ is invariant under the local $U(1)$ transformation $\psi \rightarrow e^{-iq\theta(x)}\psi$. The Lagrangian in Eq. 1.1 is *not* invariant under this transformation, as we are left with an extra term proportional to $\partial_{\mu}\theta(x)$.

It turns out we can fix this by adding an extra term $-q\bar{\psi}\gamma^{\mu}\psi A_{\mu}$, where A_{μ} is a new vector field that transforms by $A_{\mu} \rightarrow A_{\mu} + \partial_{\mu}\theta$. The new terms introduced by the gauge transformation then cancel, and the Lagrangian is invariant. The field A_{μ} represents a new spin-1 particle, and to ensure that the Klein-Gordon equation is satisfied, we must also add a "free" term. Doing this, it turns out the boson must be massless to ensure that gauge invariance is maintained.

So by requiring that the Lagrangian for a free fermion is invariant under a local $U(1)$ transformation, we were forced to introduce a new massless, spin-1 boson that couples to the fermion via the term $q\bar{\psi}\gamma^{\mu}\psi A_{\mu}$. This new boson is the photon, and this fermion coupling is the fundamental interaction of quantum electrodynamics (QED)! The photon couples to any particle with electric charge, represented by q in the coupling term.

Feynman diagrams for the fundamental vertices of QED are shown in Fig. 1.2. On the

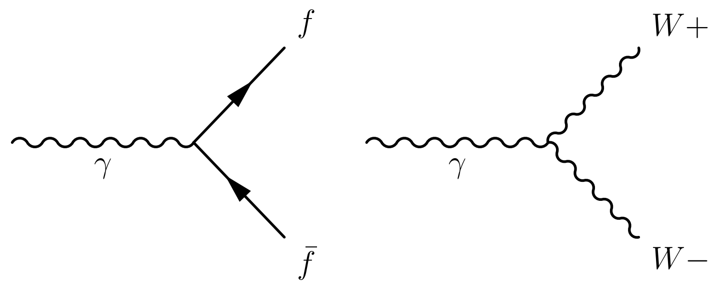

Figure 1.2: Fundamental vertices of the electromagnetic interaction. The photon couples to charged fermions (left) and the W boson (right), though the γWW vertex is best understood through electroweak unification discussed later in this section.

left is the coupling of the photon to any charged fermion (leptons, with $|q| = 1$, or quarks, with $|q| = 2/3$ or $1/3$). The W boson is also charged, and so couples to the photon, but this vertex is best understood through electroweak unification discussed below.

To derive the strong interaction, we start with the assumption that quarks have "color charge", which can be one of three values we refer to as red, green, and blue. Then the free quark Lagrangian can be written

$$\begin{aligned} \mathcal{L}_{\text{quarks}} &= \sum_{c=r,g,b} i\bar{q}_c\gamma^\mu\partial_\mu q_c - m\bar{q}_c q_c \\ &= i\bar{q}\gamma^\mu\partial_\mu q - m\bar{q}q \end{aligned} \tag{1.2}$$

where q is the quark spinor, c represents the color charge, and we have introduced the shorthand

$$q = \begin{pmatrix} q_r \\ q_g \\ q_b \end{pmatrix}, \quad \bar{q} = (\bar{q}_r\ \bar{q}_g\ \bar{q}_b)\,. \tag{1.3}$$

This can also be summed over the six flavors of quarks to describe all of them simultaneously.

The Lagrangian in Eq. 1.2 is already invariant under the global gauge transformation $q \to Uq$, where U is any 3×3 unitary matrix of determinant 1 (this group of matrices

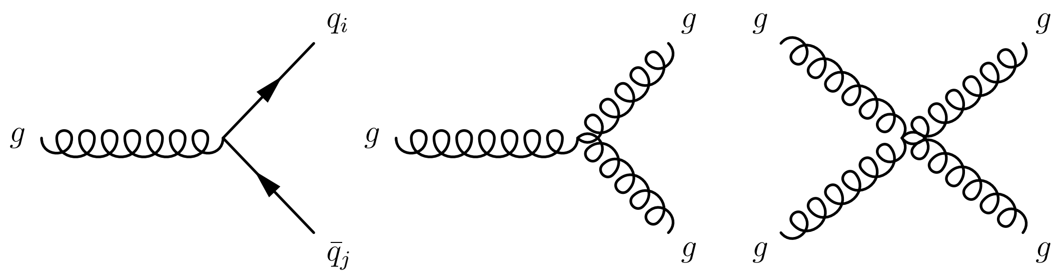

Figure 1.3: Fundamental vertices of the strong interaction. The gluon couples to a quark-antiquark pair, and also has 3- and 4-gluon self-interaction vertices. These are important for internal dynamics of hadrons as well as the hadronization of quarks and gluons at colliders.

is called $SU(3)$. It is not strictly necessary to require determinant 1, but we can factor out the determinant as a separate $U(1)$ transformation, which would just re-derive QED for quarks). But if we further require that the theory should be invariant under a *local* $SU(3)$ transformation $U(x)$, we are again required to add extra terms.

The details are quite a bit messier this time, as $SU(3)$ is a non-abelian group. But it turns out that any matrix in $SU(3)$ can be written as $e^{-ig_s \boldsymbol{\lambda}\cdot\boldsymbol{\alpha}(x)}$, where $\boldsymbol{\lambda}$ are the eight Gell-Man matrices and $\boldsymbol{\alpha}$ is an 8-dimensional vector, and we are forced to add a term to the Lagrangian $g_s\bar{q}\gamma^\mu(\boldsymbol{\lambda}\cdot\mathbf{A}_\mu)q$.

This time, $\mathbf{A}_\mu$ is a vector of eight new spin-1 bosons (gluons), and this term represents the coupling of gluons to quarks. In adding the kinetic terms for the gluons, we are again forced to make them massless, but this time, due to the non-abelian nature of $SU(3)$, we are also forced to add additional terms that represent three- and four-gluon self-interaction vertices. The fundamental interactions for quantum chromodynamics (QCD) are shown in Fig. 1.3. The two on the right are the boson self-interactions, which are not present in QED, and are a feature of non-abelian gauge theories.

The final interaction we must incorporate is the weak interaction. However, there are a number of difficulties that did not arise for the electromagnetic and strong interactions. First is the observation that charged weak currents (that is, W boson interactions), only involve left-handed chiral fermions states, or right-handed anti-fermion states (in the relativistic limit, or for massless particles, chirality is the same as helicity). The strong and electromagnetic interactions make no distinction between left- and right-handed states.

This is incorporated into the theory by introducing left-handed fermion doublets (e.g. $(\nu_e\ e)_L$, $(u\ d)_L$, etc.) and right-handed singlets, and then imposing a local $SU(2)$ symmetry on the doublet states. The charge corresponding to this symmetry is called *weak isospin*.

Running through the same procedure as above, we find that this $SU(2)_L$ symmetry generates two charged currents and one neutral current, which look deceivingly like the $W^{\pm}$ and Z bosons. However, there is a problem: the neutral Z boson is observed to couple to both left- and right-handed particles, while so far we have only dealt with the left-handed case.

The answer to this is tied to the solution for the other main problem of weak interactions: the fact that the weak bosons are massive. We have seen that local gauge invariance requires that the gauge bosons be massless, but in reality the weak gauge bosons are observed to be quite heavy. The solution to this in general was provided by the Higgs mechanism, mentioned above and discussed in the following. Then in the late 1960s, Glashow, Weinberg, and Salam [12–14] incorporated this into the SM and in the process unified the electromagnetic and weak interactions into a single *electroweak* interaction.

Before getting into the Higgs mechanism, we can understand this unification by introducing a weak hypercharge, defined as $Y = 2T_3 - 2Q$, where Q is the electric charge and T_3 is the third (neutral) component of weak isospin. This is an invariant quantity under weak isospin and the total symmetry group becomes $SU(2)_L \times U(1)$. Denoting $\mathbf{j}_\mu$ and j_μ^Y as the four currents associated with weak isospin and weak hypercharge, respectively, the electroweak interaction Lagrangian can be written

$$\mathcal{L}_{\text{EWK,int}} = -ig_w \mathbf{j}_\mu \cdot \mathbf{W}^\mu + \frac{ig'}{2} j^Y_\mu B^\mu, \tag{1.4}$$

where we have introduced coupling constants g_w and g'. This is written in terms of the *gauge fields* rather than the *physical fields*, to make it more manifestly invariant under our symmetry group, but it can be re-written in terms of the physical fields. The two W bosons are linear combinations of the first two (charged) components of $\mathbf{W}_\mu$, and the Z boson and photon are linear combinations of W^3_μ and B_μ:

$$W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu) \tag{1.5}$$

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_w & \sin\theta_w \\ -\sin\theta_w & \cos\theta_w \end{pmatrix} \begin{pmatrix} B_\mu \\ W^3_\mu \end{pmatrix}, \tag{1.6}$$

where θ_w is a free parameter that controls the degree of mixing, known as the Weinberg angle or weak mixing angle

From these and the known QED coupling, one can derive expressions for the weak coupling constants g_w and g' in terms of e and θ_w. Plugging Eqs. 1.5 and 1.6 into the Lagrangian Eq. 1.4, one can get an expression for the Lagrangian in terms of the physical fields and read off the allowed couplings. The fundamental vertices are shown in Fig. 1.4. The W couples to $\ell\nu$ or $q\bar{q}'$ pairs (though only to left-handed fermions and right-handed antifermions). The $\ell\nu$ pair must be within a single generation, but cross-generational $q\bar{q}'$ couplings are possible as the *weak eigenstates* of quark flavor are rotated slightly from the *mass eigenstates* that we have been referring to. The matrix that performs this rotation and determines the coupling strengths is known as the CKM matrix. The Z couples to

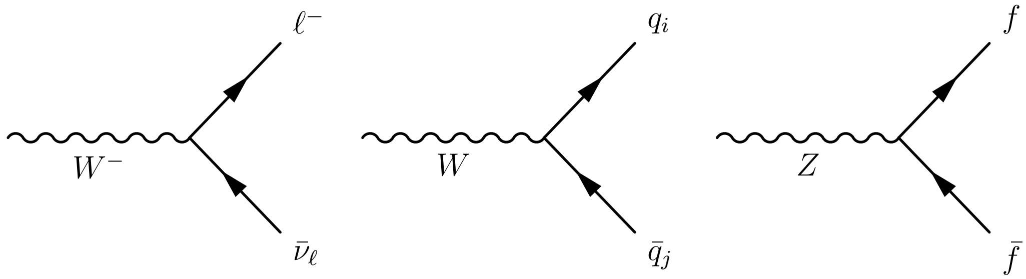

Figure 1.4: Fundamental vertices of the weak interaction. The W boson can couple to a lepton and a same-generation neutrino, or a quark and anti-quark (most strongly within the same generation, but cross-generational couplings are possible via the CKM matrix). The Z boson can couple to any fermion and its antiparticle. There are also boson self-interaction vertices (WWZ, $WWWW$, $WWZZ$, $\gamma\gamma WW$) necessary for the self-consistency of the theory, but they are less important practically.

any fermion-antifermion pair.

We have seen how the electromagnetic, strong, and weak interactions can arise by imposing a local $SU(3)\times SU(2)_L\times U(1)$ gauge symmetry on the SM Lagrangian. However, there remain the questions of how exactly the weak bosons acquire their masses without breaking local gauge invariance, and how the $SU(2)_L \times U(1)$ symmetry underlying the electroweak interaction is broken, such that the weak bosons are so heavy while the photon is massless.

The answer lies in the Higgs mechanism and spontaneous symmetry breaking, mentioned earlier. Again the details can be found in any quantum field theory textbook, but the basic idea is this: imagine that there exists a scalar complex field $\phi \equiv \phi_1 + i\phi_2$ governed by the Lagrangian

$$\mathcal{L} = \frac{1}{2}\partial_\mu\phi^*\partial_\mu\phi + \frac{1}{2}\mu^2\phi^*\phi - \frac{1}{4}\lambda^2(\phi^*\phi)^2. \tag{1.7}$$

This Lagrangian already has a global $U(1)$ symmetry, and requiring local symmetry again requires the addition of a massless gauge boson A_μ. Further, the potential (given here by the negative of the final two terms) normally has a minimum at 0, but in this case 0 is actually a local maximum and the minimum lies on a circle in the complex plane of radius μ/λ. Since the Feynman calculus is really a perturbation theory in the fields, we must expand around a local minimum of the potential. Choosing $(\mu/\lambda, 0)$ as the point to expand around ("spontaneously" breaking the symmetry), we define

$$\eta \equiv \phi_1 - \mu/\lambda, \quad \xi \equiv \phi_2. \tag{1.8}$$

Plugging these in to the Lagrangian, we find that η represents a scalar boson of mass $m_\eta = \sqrt{2}\mu$, and A_μ picks up a mass term $m_A = 2q\mu/\lambda$, where q is the charge generated by the $U(1)$ symmetry. So by introducing a complex scalar field and breaking the symmetry of its potential, we have given mass to the gauge boson at the cost of a new scalar boson η! There is a problematic term representing a direct ξA coupling, but this can be transformed away with the proper choice of gauge, thereby eliminating the non-physical "Goldstone boson" ξ.

It is a bit more complicated with the full $SU(2)_L \times U(1)$ symmetry of the electroweak Lagrangian (the scalar field must become a doublet of complex fields, and we have three gauge bosons), but the basic ideas are the same. We find that the W and Z bosons acquire masses proportional to μ/λ (and related by $M_W/M_Z = \cos\theta_W$) while the photon remains massless, and a new scalar "Higgs boson" is introduced. Its mass is undetermined by the theory and must be measured.

Masses of the fermions can also be added to the theory via the Higgs boson, by adding Yukawa coupling terms. We then get all of the fundamental interactions of the Higgs boson shown in Fig. 1.5. It directly couples to all massive particles, with strength proportional to the mass. While it does not directly couple to photons or gluons as they are massless, the Higgs can still decay to these particles through loops involving massive particles (most prominently virtual top quarks).

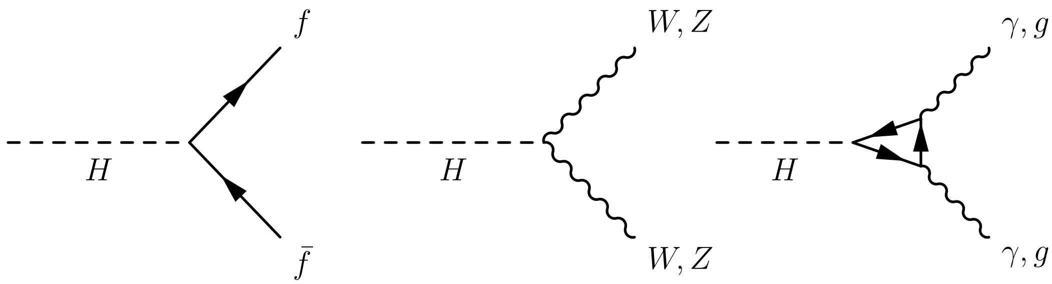

Figure 1.5: Interactions of the Higgs boson. It directly couples to any particle with mass, including fermions (left) and the W and Z bosons (center). It cannot couple directly to massless particles such as the photon and gluon, but it may decay to these through loops of massive particles (right). The most likely particle in the loop is the top quark, since it is the heaviest option (thought it needs to be virtual, since $2m_t > m_H$). There are also three- and four-Higgs self-interaction vertices.

We have now seen the basic ideas of how the Standard Model is constructed: starting from the Lagrangians of free fermions, we impose invariance under local $SU(3) \times SU(2)_L \times U(1)$ gauge transformations and are forced to add massless gauge bosons. By introducing a scalar Higgs field, we can spontaneously break the symmetry of the electroweak interaction and give masses to the weak bosons, as well as give mass terms to the fermions in the theory. The resulting theory has had remarkable success in predicting experimental findings to high degree of accuracy. However, there are a number of phenomena that the SM cannot explain, and it thus cannot be a complete "theory of everything". We examine some of the issues with the SM in the following section.

1.2 Shortcomings of the Standard Model

The Standard Model has a long history of making successful predictions about the sub-atomic world: the existence of the charm and top quarks, gluons, the weak bosons, and the Higgs boson; branching ratios and lifetimes of the fundamental particles and hadrons; the precise value of the electron's anomalous magnetic moment (down to one part in a billion, making it one of the most accurate verified predictions in the history of science); and more. Many experiments continue today to verify the SM's accuracy. However, it is not a perfect theory and cannot sufficiently explain a number of observed phenomena.

1.2.1 Unexplained phenomena and theoretical considerations

First and foremost among unexplained phenomena is the gravitational interaction. The SM provides a mathematical framework to describe three out of the four known fundamental interactions, but does not say anything about the mutual attraction of massive particles. Instead, gravitation is currently understood in terms of curvature of spacetime through Einstein's general theory of relativity. This is incompatible with our present understanding of particle physics, and the search for a quantum theory of gravity remains an active area of research among theoretical physicists.

Next, the SM does not naturally allow for neutrino masses (the Yukawa mechanism for adding masses to the charged leptons requires the existence of right-handed leptons, but only left-handed neutrinos have been observed). However, experiments have observed the phenomenon of "neutrino oscillation" [15], whereby a neutrino that starts as one flavor (e.g. an electron neutrino) is later measured after some propagation as a different flavor (e.g. a muon neutrino). This is only possible if the mass eigenstates of neutrinos are rotated from the flavor eigenstates that are produced in weak interactions, which necessarily requires that the neutrinos have small but non-zero masses. It is possible to modify the SM to incorporate neutrino masses, but such solutions are unsatisfactory as they offer no insight into the vastly different scales of neutrino and charged fermion masses. Theorists seek a more natural explanation for the origin of neutrino masses.

Another problem lacking a satisfactory explanation under the SM is that of matter-antimatter asymmetry, or the "baryon asymmetry" [16]. We observe that the universe is made almost entirely out of matter, as opposed to antimatter, and there is no good reason that this should be the case (the SM predicts that almost equal amounts of matter and antimatter should have been created at the beginning of the universe, after which most of it would have been annihilated). A necessary condition for explaining this asymmetry is a mechanism for violating charge-parity, or CP, symmetry (the symmetry that exchanges all particles for antiparticles and reverses chiralities). The weak interaction allows for small CP-violation via a complex phase in the CKM quark mixing matrix, but this is not big enough to explain the observed baryon asymmetry, so beyond the SM sources are necessary.

Finally, most relevant for this dissertation is the existence of *dark matter*. If one plots the velocity of matter as a function of radial distance from the center in a rotating spiral galaxy, Newton's laws applied to the distribution of visible matter say that the velocity should decrease with radius. However, observed galactic rotation curves are actually flat over large distances (see Fig. 1.6 for an example from the M33 galaxy). This suggests the presence of invisible, or "dark", matter permeating the galaxies that interacts gravitationally but not electromagnetically. Other observations support this same hypothesis, such as gravitational lensing and the cosmic microwave background power spectrum. The SM does not provide for any fundamental particles that can account for this dark matter, and explaining it is one of the central problems of modern physics.

There are also a few problems with the SM that are more theoretical in nature and motivated by "naturalness" concerns. That is, that the parameters of the SM are "fine-tuned" to specific values for no apparent reason (e.g. an underlying symmetry that would cause the parameter to take on that value). Extensions to the SM are sought that would provide more natural explanations.

One example is known as the *strong CP problem*. As mentioned above, the weak sector admits slight CP symmetry violation via a complex phase in the CKM quark

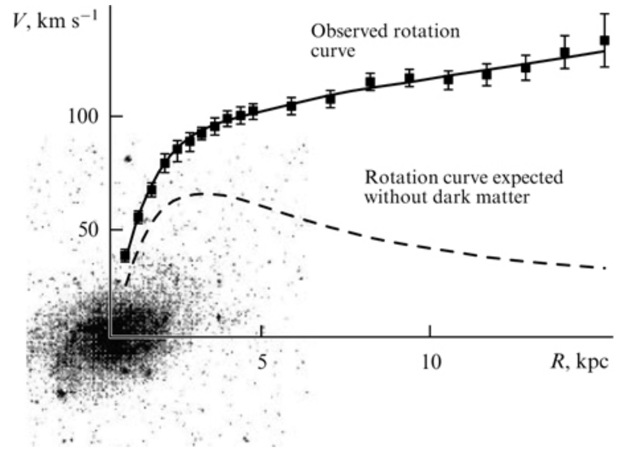

Figure 1.6: Observed and expected rotation curves of the M33 galaxy, suggesting a significant mass contribution from dark matter. (Image from [17])

mixing matrix. It is also possible to add CP-violation to the strong sector, but this is not observed experimentally. This means that the parameter controlling this CP-violating term in the QCD Lagrangian must be very close to 0. There is no known reason for this, and theorists have provided a number of models that can explain it more naturally. However, no evidence for any of these theories has yet been found experimentally.

Finally, with what partially motivates the kinds of new physics sought in the present analysis, we have what is known as the *hierarchy problem*. The observed Higgs mass of 125 GeV is the sum of the bare Higgs mass, a free parameter in the theory, plus terms arising from loop-level corrections to the free propagator. The most important of these corrections is shown in Fig. 1.7. As the Higgs couples to all massive fermions, the virtual particle loop can be anything. The most important contribution however is from the top quark, since the top quark is the heaviest particle and couples most strongly to the Higgs. Computing amplitudes, one finds (to leading order in Λ)

$$m_{H,\mathrm{obs}}^2 \approx m_{H,\mathrm{bare}}^2 - \frac{|\lambda|^2}{8\pi^2}\Lambda^2, \tag{1.9}$$

where λ is the top quark Yukawa coupling and Λ is the scale up to which the SM is valid,

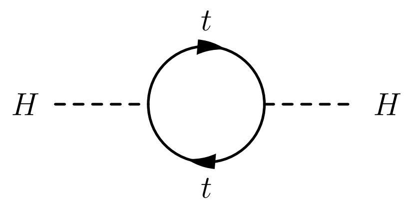

Figure 1.7: Diagram representing the loop-level correction to the Higgs mass. The loop can be any charged fermion, but the top quark is most important since it has the highest mass.

typically taken to be the Planck scale $\Lambda \approx 10^{19}$ GeV.

The observed Higgs mass is of order $m^2_{H,\text{obs}} \sim 10^4$ GeV^2, and the correction term is of order 10^{36} GeV^2. That means that the unobservable bare Higgs mass must be fine-tuned to one part in $\mathcal{O}(10^{32})$ in order to cancel the correction almost perfectly, and there is no reason why this must be so. In other words, we ask *why is the Higgs mass so much smaller than the Planck scale?* The fact that this must be a complete accident in the SM is highly unnatural, and leads to the question of whether there is some kind of physics beyond the SM that could provide a more satisfactory explanation.

1.2.2 Signs of experimental deviation

While the SM has been extensively tested and probed to astounding degrees of accuracy, there are have been a few signs of deviation between theory and experimental results. Many of these have come and gone over the years, shown to be merely statistical fluctuation, but a small number remain open questions today.

First is the anomalous magnetic moment of the muon. Particles with intrinsic spin also have an intrinsic magnetic moment caused by that angular momentum. The dimensionless proportionality constant relating this moment to the particle's spin and the Bohr magneton is known as the "g-factor". In ordinary quantum mechanics, this is exactly equal to 2 for fermions such as the electron and muon. However, QED predicts small corrections to this from virtual particle loops. The g-factor of the electron has been measured as $g_e = 2.002\,319\,304\,361\,46(56)$ [18], agreeing with theory to around one part in 10^{10} and making it one of the most precisely verified predictions in all of physics.

The muon g-factor has been measured to around 1 part in 10^6, but is higher than the predicted theory value by 3–4 standard deviations [19]. There are various models of beyond the SM physics that could in principle explain this, but it is still an open question whether the discrepancy is due to new physics, statistical fluctuation, or poorly modeled experimental or theoretical uncertainties. The Muon g–2 experiment at Fermilab is currently taking data and should reduce the experimental uncertainty by more than a factor of three.

Another potential sign of deviation is a collection of flavor physics anomalies in the decays of B mesons [20]. The SM predicts the principle of *lepton flavor universality* (LFU), which states that the three generations of leptons are identical in all but their masses (i.e., their couplings to the electroweak bosons are the same). New particles could in principle couple differently to different generations, changing branching fractions for mesons that decay via the weak interaction into leptons. The BaBar, Belle, and LHCb Collaborations have measured the ratios

$$R_{D^{(*)}} = \frac{\mathcal{B}(B \to D^{(*)}\tau\bar{\nu}_\tau)}{\mathcal{B}(B \to D^{(*)}\ell\bar{\nu}_\ell)}, \qquad \text{with } \ell = \mu, e \tag{1.10}$$

and have found values consistently higher than the SM predictions, with the averages being 2–3 standard deviations larger than the SM values. Fig. 1.8 shows the measured averages compared to SM expectations.

The experiments have also probed LFU via the rarer loop-level neutral current $b \to s\ell\ell$ processes, in the form of $R_{K^{(*)}} = \mathcal{B}(B \to K^{(*)}\mu^+\mu^-)/\mathcal{B}(B \to K^{(*)}e^+e^-)$. This shows similar 2σ-level deviation from the SM prediction. Analysis of data from LHCb and the

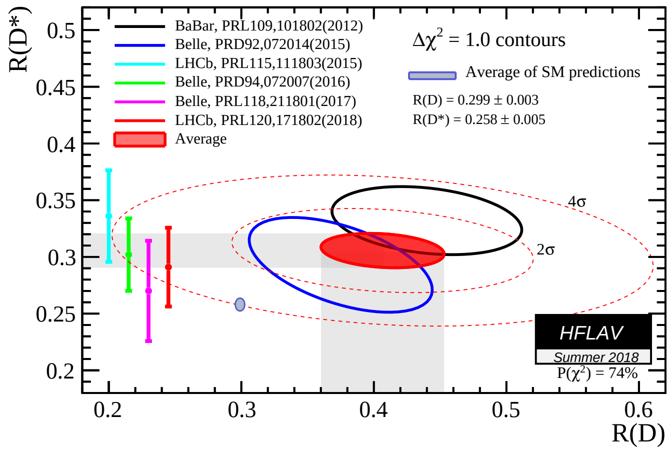

Figure 1.8: Averages of R_D and R_{D^*} as measured by the BaBar, Belle, and LHCb Collaborations. Both ratios are 2–3 standard deviations higher than the SM expectation. (Image from [21])

recently begun Belle 2 run should improve the experimental precision and give a clearer picture of whether or not the anomalies are real.

1.3 Theories of physics beyond the Standard Model

There is a huge variety of models of physics beyond the SM (BSM), each with their own predictions of new particles and phenomena. The list is far too vast and diverse to summarize completely here, so we only discuss briefly the ones most relevant to the analysis presented in this dissertation.

1.3.1 Supersymmetry

Supersymmetry (SUSY) is a proposed extension to the SM dating to the 1970s that posits that every SM fermion has a bosonic "superpartner", and every SM boson has a fermionic superpartner. While a theory that doubles the number of fundamental particles and adds over 100 parameters to the SM is not the most elegant of solutions, it has become popular as it can solve the hierarchy problem and provide a natural candidate for dark matter. We give a brief overview of the main ideas here; a more complete summary can be found at [22].

First, a word on notation and terminology. Supersymmetric particles in general are referred to as sparticles. The superpartners to the SM fermions are referred to by prepending an ‘s’ to the name of the fermion (e.g. slepton, squark, selectron, stau, stop, sbottom). The superpartners of the SM bosons are referred to by appending an “ino” to the name of the boson (e.g. wino, higgsino). Symbolically, we add a tilde above the symbol of the SM particle to refer to its superpartner (e.g. $\widetilde{t}$, $\widetilde{e}$, $\widetilde{W}$).

In the minimal possible supersymmetric extension to the SM (MSSM) [23], there actually must be four higgsinos due to the way spontaneous symmetry breaking works with the added SUSY degrees of freedom. Two of these are charged and two are neutral. Then due to mixing between bosons with identical quantum numbers, the charged bosinos $\widetilde{H}^{\pm}$ and $\widetilde{W}^{\pm}$ can in general mix into charginos $\widetilde{\chi}^{\pm}_{1,2}$, and the neutral zino, photino, and higgsinos can mix into neutralinos $\widetilde{\chi}^{0}_{1,2,3,4}$.

In raw, unbroken supersymmetry, sparticles have the same mass and charge as their SM counterparts, and spin differing by $1/2$. However, if this were the case we would have already discovered supersymmetry, as the sparticles would be produced copiously at existing colliders. Since we have not seen them, we know that if it exists supersymmetry must be a spontaneously broken symmetry, such that the sparticles have masses too large to have been seen with existing experiments.

While not strictly necessary for the theory, it is a common starting point to assume that there is a conserved multiplicative quantum number called R-parity that is equal to $+1$ for SM particles and -1 for sparticles. This means that sparticles must be produced in pairs, and that the lightest supersymmetric particle (LSP) must be stable (as there is no lighter particle with $R = -1$ that it can decay into). A neutral LSP is a natural candidate

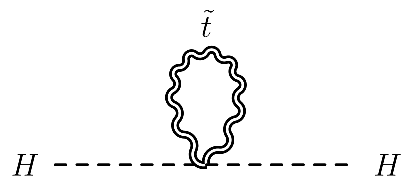

Figure 1.9: Diagram representing the SUSY correction to the Higgs mass from the $\widetilde{t}\widetilde{t}HH$ coupling. The term has opposite sign to the typical fermion correction terms (due to the bosonic nature of the stop squark), and depending on the stop mass can be of similar magnitude and largely cancel them, thus solving the hierarchy problem.

for dark matter, as it is massive, stable, and does not interact electromagnetically. The LSP is typically taken to be $\widetilde{\chi}_1^0$. R-parity conservation furthermore guarantees that baryon and lepton numbers are still conserved in supersymmetric interactions, a desirable feature as these conservation laws have been precisely tested experimentally.

In addition to providing a dark matter candidate, SUSY can also solve the hierarchy problem by adding additional terms that largely cancel the problematic term proportional to Λ^2 in Eq. 1.9. Superpartners to the SM fermions (themselves bosons) couple to the Higgs with vertices like that shown in Fig. 1.9. Computing the amplitudes, one finds the correction from one of these diagrams to be

$$\Delta m^2_{H,\mathrm{SUSY}} \approx +\frac{\lambda_S}{8\pi^2}\Lambda^2, \tag{1.11}$$

where λ_S is the coupling strength of the vertex. Note that this correction term is *positive*, in contrast to the negative contribution from SM fermion loops. This is due to the spin-statistics theorem and the fact that the superpartners are bosons.

So if $\lambda_S \approx \lambda^2$, then the terms largely cancel and the problem is solved. Since the most problematic term comes from a top quark loop, the SUSY correction from stop squarks is the most important for the cancellation. This means that the stop cannot be too heavy, or the cancellation will not work. While this is fairly subjective, the commonly cited criterion is that the stop must have mass $\mathcal{O}$(TeV) or smaller in order to provide an acceptable solution.

The MSSM has 120 new parameters, making an interpretation of experimental data nearly impossible. It is possible to reduce this to 19 parameters by adding some phenomenological constraints (e.g. no flavor changing neutral currents) [23], but interpretations would still be highly complicated in a 19-dimensional parameter space. To facilitate interpretation, physicists have developed a set of *simplified models* [24–26]. These are models consistent with the MSSM but with a constrained phase space, so that they represent one specific SUSY process and contain only 1–3 free parameters. Typically sparticles not involved in the process are assumed to be heavy enough that they are decoupled from the process and do not affect it in a significant way. An example are the so-called "type 1" models, in which gluinos are pair produced and decay with 100% branching fraction to a quark-antiquark pair and LSP (e.g. $pp \to \widetilde{g}\widetilde{g}$, $\widetilde{g} \to t\bar{t}\widetilde{\chi}^0_1$). In this case the mass of the gluino and LSP are free parameters. A full description of all of the simplified models considered in this dissertation can be found in Sec. 9.3.

1.3.2 Leptoquarks

A variety of BSM theories, such as grand unified theories [27–29], technicolor models [30–32], compositeness scenarios [33,34], and R-parity violating SUSY [35] predict the existence of particles called *leptoquarks* (LQ) that carry quantum numbers of both quarks and leptons, allowing them to interact directly. Details vary among theories, allowing either quark-charged lepton or quark-neutrino vertices or both, and some allowing cross-generational couplings. The LQs can either be spin-0 (scalar, $\mathrm{LQ_S}$) or spin-1 (vector, $\mathrm{LQ_V}$). Some example vertices involving leptoquarks are shown in Fig. 1.10

Leptoquarks have become more heavily discussed in recent years as a possible expla-

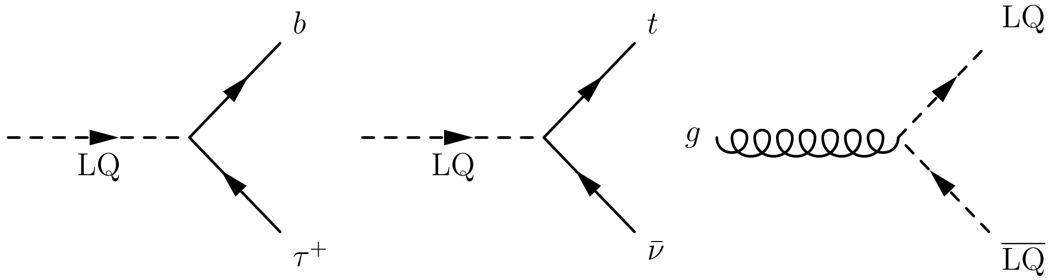

Figure 1.10: Some example vertices involving leptoquarks. (left) A third-generation LQ decaying into a bottom quark and tau lepton, (center) a LQ decaying into a top quark and neutrino, and (right) a gluon-LQ-LQ vertex illustrating the color charge that LQs carry.

nation of the various B-physics anomalies discussed in Sec. 1.2.2. In particular, Ref. [36] predicts a best-fit model consisting of a vector leptoquark with mass of $\mathcal{O}$(TeV), decaying with 50% branching fraction to each of $\mathrm{LQ_V} \to b\tau$ and $\mathrm{LQ_V} \to t\nu$. The analysis in this dissertation is interpreted in terms of this and other LQ models in Sec. 9.3.

1.4 Hadron collider physics

We have seen in the previous sections how the SM is constructed, what interactions are allowed, and a few examples of possible theoretical extensions. What we are mainly interested in for this dissertation, however, is how exactly this all manifests itself at hadron colliders.

Since interesting phenomena involve relatively heavy particles (e.g. the weak bosons, the Higgs, theorized SUSY particles, etc.), physicists produce them by colliding together particles at very high energies. This is generally done with either e^+e^- or proton-proton colliders, which can either be pp or $p\bar{p}$. The Large Hadron Collider (LHC), discussed in the following chapter, is a pp collider.

In contrast to electrons, protons are composite particles, which makes the collisions inherently messy. This is both a benefit and a complication: a complication because the

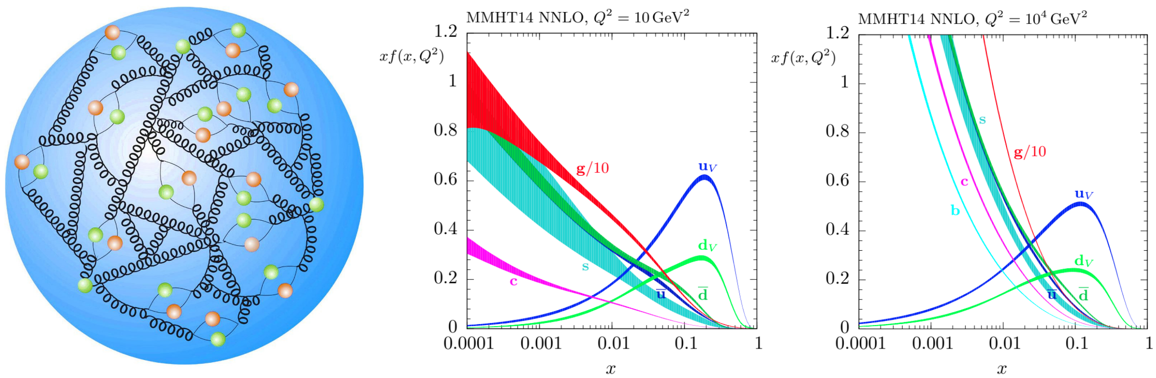

Figure 1.11: On the left, an illustration of the internal structure of the proton (image from [37]). Gluons and quark-antiquark pairs constantly pop in and out of existence. The right two plots show MMHT2014 NNLO proton parton distribution functions, for momentum transfer scales $Q^2 = 10$ GeV^2 (center) and $Q^2 = 10^4$ GeV^2 (right). The quantity x is the fraction of proton momentum carried by the parton. (Image from [38])

underlying collision energies are ill-defined, with the constituents of the protons carrying an unknown fraction of their total energy, and a benefit because this allows physicists to probe a wide range of energy scales and processes simultaneously.

We stated in the previous section that the proton consists of three quarks: two ups and one down. However, the real story is more complicated than this. The *uud* quarks are only what are referred to as the *valence quarks*, which determine the quantum numbers of the proton. The strong force, mediated by gluons, binds these quarks together, and the gluons can produce virtual quark-antiquark pairs that pop in and out of existence. So in reality, the proton consists of a "sea" of quarks and gluons, excitations of the quantum field between the valence quarks (actually, the bare mass of the valence quarks only accounts for around 1% of the proton's mass. The other 99% comes from the energy of this quark-gluon "sea"). Fig. 1.11 (left) shows an illustration of the inner dynamics of a proton.

When two protons collide, collisions can happen between any combination of valence quarks and sea quarks/gluons. The higher the collision energy, the smaller the distance scale probed and the more "visible" the sea quarks/gluons become. This is quantified in what are called "parton distribution functions" (PDFs), which are probability densities of finding a particular particle with longitudinal momentum fraction x, given an energy scale Q^2. Examples are shown in Fig. 1.11 for $Q^2 = 10$ GeV^2 and 10^4 GeV^2. One sees that the PDFs for valence u and d quarks peak near 1/3, as these are disproportionately likely to carry a significant fraction of the proton's momentum. The sea quarks and gluons peak lower, and become more prominent at higher Q^2. These PDFs must be taken into account when computing cross sections for the various pp production processes.

Now knowing that collisions can happen between any valence quarks or sea quarks and gluons, we can put together the vertices presented in Sec. 1.1.2 to make diagrams of various interesting final states possible to achieve at pp colliders. A few examples are shown in Fig. 1.12.

First, in the top left, is the production of a Z boson. From the PDFs, we know that the most likely quarks are a u valence quark and a $\bar{u}$ sea quark, followed by a d valence quark and a $\bar{d}$ sea quark. Also note that since one is a valence quark and one is a sea quark, there is likely a momentum imbalance, so the entire event will be boosted in the longitudinal direction. The transverse momentum, however, will be very close to 0.

In this case we have shown the Z decaying into neutrinos, which will be invisible to the detector. The Z can also decay into charged leptons, giving the Drell-Yan process, or quarks. We have also shown a gluon radiating off one of the initial state quarks, called "initial state radiation", or ISR.

Now the strong force is a short range force, and energy actually increases without bound as a colored particle is separated from its bound state (in contrast to electromagnetism, where the energy goes to 0 at infinite separation). This means that isolated quarks and gluons cannot exist (called *color confinement*), as eventually it becomes energetically

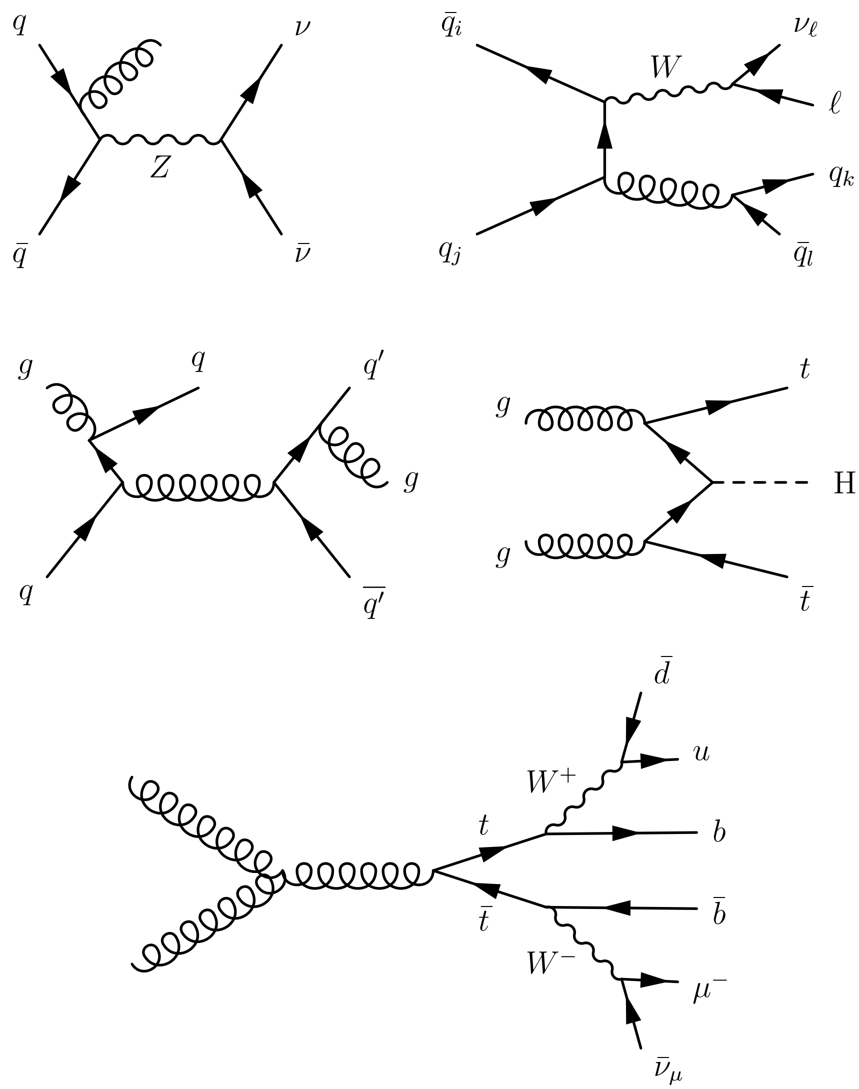

Figure 1.12: Diagrams of the production of a few important processes at pp colliders. From left-to-right, top-to-bottom: a SM monojet signature, caused by a Z decaying invisibly to $\nu\bar{\nu}$ and an initial quark radiating a gluon; W+jets production, in this case $W \to \ell\nu$ plus two jets; QCD multijet production, in this case with three quark jets and one gluon jet; $t\bar{t}H$ production; and a full $t\bar{t}$ decay chain, where one of the W's from a top decays leptonically and one decays hadronically, producing a final state with a prompt lepton, missing energy, and four jets, two of which are b jets.

favorable to produce new quark-antiquark pairs. In hadron colliders, a high-energy quark or gluon thus undergoes a process called *hadronization*, in which it produces a narrow cone of hadrons, called a jet, as it is expelled from the collision point. The signature represented by the first diagram in Fig. 1.12 is thus a "monojet": a single high-energy jet, recoiling against invisible energy from the neutrinos from Z decay.

Continuing to the right in Fig. 1.12, we have the production of a W boson which decays to a $\ell\nu$ pair. We've also shown an ISR gluon splitting into two quarks, which show up as two jets. This diagram thus represents a leptonic W+jets event.

Next in the middle left is a generic diagram representing QCD multijet production, the most common type of high-p_{T} event at hadron colliders. In this case we have an interaction between a valence quark and sea gluon, which results in a final state of three quarks and a gluon, giving a four jet event. By tacking on more quark/gluon vertices to this or similar diagrams, we can create events with any number of jets. Dijets are the most common type of event, and the rate falls exponentially with the number of jets (each jet reduces the cross section roughly by a factor of 10).

In the middle right we have a diagram representing $t\bar{t}H$ production from two gluons. This will not be very relevant for the analysis in this dissertation, but it is an important process under study at the LHC to better understand the nature of the Higgs boson and its coupling to top quarks.

And finally, the bottom diagram shows the production of a $t\bar{t}$ pair through gluon-gluon fusion. Top quarks are the only quarks that do not undergo hadronization, since they decay too quickly due to their very high mass. Since the on-diagonal V_{tb} element of the CKM matrix is very nearly one, top quarks decay nearly 100% of the time to a W boson and b quark. The b quarks hadronize into b jets (identifiable due to the long lifetime of B hadrons), and the W bosons can decay either hadronically (70% of the time) or leptonically (10% each to e, μ, τ).

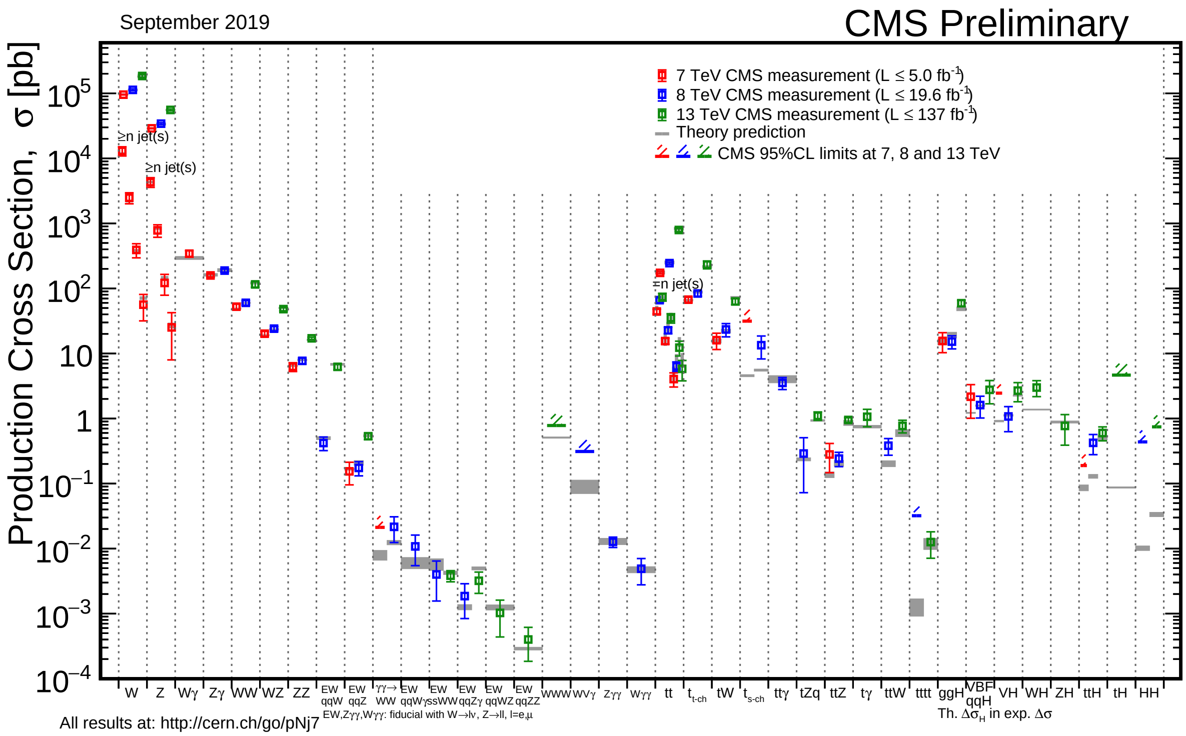

Figure 1.13: Summary of SM production cross sections measured at CMS, compared to values predicted from theory. (Image from [39])

Using the PDFs, theorists have computed cross sections for a huge variety of SM processes at the LHC. These have then been measured by the various experiments. A summary plot from CMS is shown in Fig. 1.13; one sees that the experimentally measured cross sections agree well with the theoretical predictions in all cases.

Theorists can also calculate the cross sections of hypothetical BSM production modes, such as SUSY. Fig. 1.14 shows the cross sections for the pair production of gluinos and squarks at $\sqrt{s} = 13$ TeV. One sees that the cross section exponentially falls with sparticle mass. Comparing with Fig. 1.13, one can get a sense for how much rarer these SUSY processes are than the SM processes that will constitute backgrounds for the search. For example, 2 TeV gluino pair production has a cross section of $\mathcal{O}(10^{-3})$ pb, five to six orders of magnitude smaller than that for W, Z, or $t\bar{t}$ production. Similar theoretical

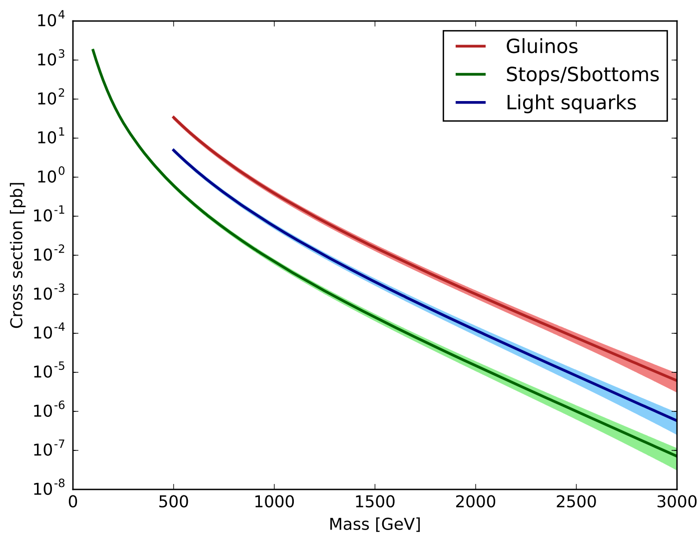

Figure 1.14: Cross sections for pair production of SUSY particles at the LHC at $\sqrt{s} = 13$ TeV, computed at NNLO(approx)+NNLL order [40] and assuming squarks and gluinos are decoupled. The x axis is the mass of the relevant particle. The light squark cross section assumes 8-fold degeneracy between left- and right-handed $\widetilde{u}$, $\widetilde{d}$, $\widetilde{s}$, and $\widetilde{c}$ squarks.

cross sections for the pair production of leptoquarks are shown in Fig. 9.14 in Sec. 9.3.3.

Chapter 2

The CMS Experiment

In the previous chapter we summarized the theoretical underpinnings of modern particle physics, and discussed possible extensions that point towards future research. We now turn our attention to the machines that make such research possible. The Large Hadron Collider (LHC) is the largest particle collider ever built, smashing protons together at record energies and luminosities. On it are located four major experiments, designed to record electrical snapshots of the collisions and allow physicists to reconstruct the details of each event. Sec. 2.1 describes the design and operation of the LHC, and Sec. 2.2 introduces the Compact Muon Solenoid (CMS) experiment, one of the four at the LHC and the one used for the analysis in this dissertation.

2.1 The Large Hadron Collider

Underneath the French-Swiss border near the city of Geneva lies the LHC, the world's largest proton collider. With a circumference of 26.7 kilometers and depth of up to 175 meters below ground, it is designed to accelerate protons to an energy of 7 TeV, over 7,000 times their rest mass (and corresponding to a speed 99.9999991% the speed of light). Utilizing the tunnel of the older Large Electron–Positron collider (LEP) at CERN, it was constructed to supersede the Tevatron at Fermilab, which accelerated protons to 1 TeV.

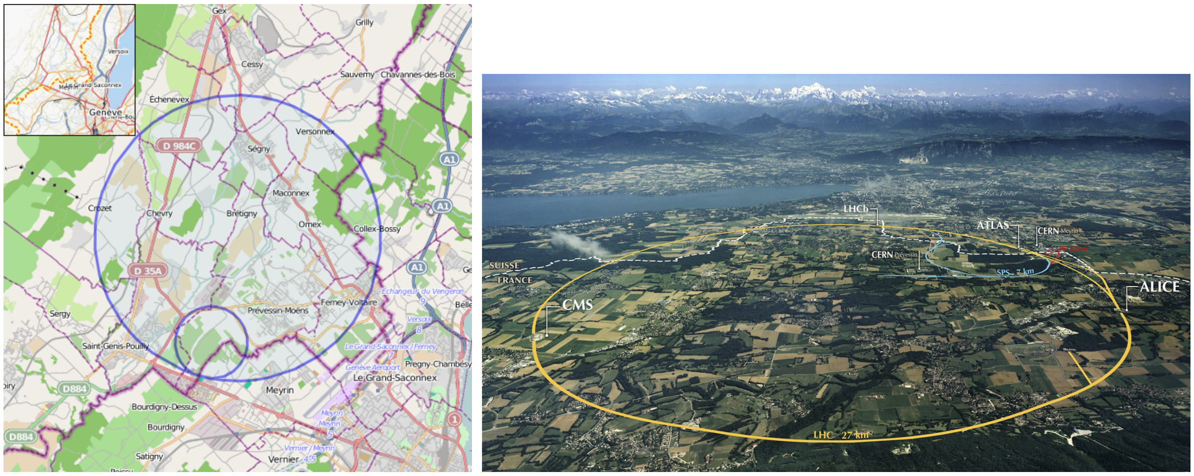

Figure 2.1: (left) The SPS and LHC overlaid on a map of the Franco-Swiss border near Geneva. The diameter of the LHC is 8.5 km (5.3 miles). (right) The LHC, the four major experiments, and CERN's campus overlaid on a photo of the Geneva area, looking southeast. Lake Geneva and the French Alps can be seen in the background. (Images from [41, 42])

What was the motivation for such a machine? The LEP collider, which accelerated electrons and positrons to up to 209 GeV, had allowed for precise measurements of many SM quantities including the W and Z masses, but was not able to find definitive proof of a Higgs boson or any evidence for BSM physics. Similarly, the Tevatron had discovered the top quark and made many measurements but failed to find evidence of the Higgs or new physics. Physicists thus sought a next-generation machine that could find these things.

Since the Higgs and potential new physics were at energy scales above the reach of present experiments, a higher energy collider was needed. Circular lepton colliders like LEP are limited in energy, as the energy radiated by an charged particle traveling in a circular path falls as $1/m^4r^2$, and so the size of the necessary electron collider would be prohibitive. The much heavier proton is far easier to accelerate to high energies, and an accelerator in the pre-existing LEP tunnel was sufficient to achieve the necessary energies (and additionally, cost-effective). The downside is that collisions of composite particles

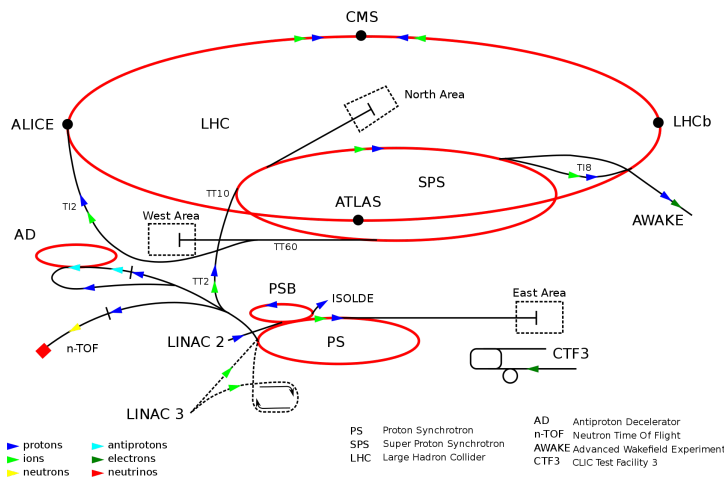

Figure 2.2: The CERN accelerator complex. Protons originate in the LINAC 2 where they are injected at 50 MeV into the PSB. The PSB, PS, and SPS subsequently accelerate them to 1.4 GeV, 26 GeV, and 450 GeV, respectively, before they enter the LHC where they reach up to 6.5 TeV. (Image from [43])

like protons are messy, and the actual parton collision energy is indeterminate. However, the LHC was meant do be a “discovery machine”, so probing a wide range of collision energies was desirable. For precision studies of any interesting physics uncovered by the LHC, a future linear e^+e^- collider would be a good option.

From these considerations, the particle physics community decided on a *pp* collider built in the LEP tunnel, and construction on the LHC began in 1998. A map of its location can be seen in Fig. 2.1. It sits in a tunnel between 50 and 175 meters underneath the suburbs of Geneva, Switzerland, straddling the France-Switzerland border. To the west are the Jura mountains, and to the east Lake Geneva. The CERN campus is on the south end.

Before entering the main ring, the protons progress through a series of increasingly large accelerators. The steps in the chain are as follows; the CERN accelerator complex is illustrated in Fig. 2.2.

- The protons begin as nuclei in hydrogen gas, and electric fields are used to strip the electrons away.

- The bare protons are fed into the linear LINAC 2 accelerator, where they are boosted to 50 MeV kinetic energy and injected into the 160 meter Proton Synchrotron Booster (PSB).
- The PSB accelerates the protons to 1.4 GeV and feeds them into the 628 meter Proton Synchrotron (PS).
- The PS accelerates them to 26 GeV and injects them into the 6.9 kilometer Super Proton Synchrotron (SPS).
- The SPS accelerates them to 450 GeV before finally injecting them into the main 26.7 kilometer LHC ring (it inserts two beams, traveling in opposite directions)
- Over a period of around 20 minutes, the LHC accelerates the protons to 6.5 TeV (this is the present maximum; the LHC is designed to eventually go to 7 TeV).

In order to steer the beams in a circle, the LHC makes use of powerful superconducting magnets. Dipole magnets are used as the main steering mechanism. From the formula for a relativistic charged particle in a uniform magnetic field, the necessary average magnetic field is $B = \gamma m\beta/qR = 5.1$ T. However, the magnets are not all the way around the ring so the peak dipole magnetic field is 7.74 T. In addition to the dipole magnets, quadrupole magnets are used for beam focusing, and higher multipole magnets are used for finer corrections. In total there are nearly 10,000 individual magnets. In order to maintain such a strong field, the superconducting magnets operate at a temperature of only 1.9 K, achieved using 96 tons of superfluid helium [44].

Magnets can only steer the beam, not increase its energy. In order to perform the acceleration from 450 GeV to 6.5 TeV, the LHC makes use of 8 radio-frequency (RF) cavities per direction that produce an oscillating electric field at a frequency of 400 MHz. This naturally produces a beam of "bunches", spaced 25 ns apart: the electric fields boost the protons, and the gradient of the field is constructed in such a way that any protons that arrive early or late receive a slightly different kick, pushing them back towards the bunch center.

Together, the magnetic and electric fields produce two counter-rotating beams of 6.5 TeV protons, organized into bunches 25 ns apart (corresponding to around 7.5 m of spatial separation). There are about 1.2×10^{11} protons per bunch, and up to 2808 bunches per beam, giving 3.4×10^{14} protons in each beam. At 6.5 TeV each, the total kinetic energy in both beams together is over 700 million joules, equivalent to a Boeing 737 traveling at 200 mph!

The "collision rate" at colliders is measured with a quantity known as instantaneous luminosity, defined by the rate N of a given type of interaction with cross section (roughly, likelihood of interaction) σ as $\mathcal{L} = N/\sigma$. This luminosity is a function of the bunch crossing rate (at the LHC, once every 25 ns), the number of protons per bunch, and the effective area of the beam (i.e., how tightly packed the protons in the beam are). The LHC can reach an instantaneous luminosity of around 2×10^{34} cm^{-2}s^{-1}.

The total proton-proton inelastic cross section is measured to be around 78 mb [45]. Multiplying by the LHC instantaneous luminosity, that means we expect 1.5 billion inelastic collisions per second, or around 40 per bunch crossing. Most of these will be relatively uninteresting, and a major challenge in analyzing experimental data is disentangling the interesting event from the $\sim$40 other interactions (referred to as *pileup* interactions).

The LHC began beam operations in September 2008. Just nine days after the first beams circulated, an electrical fault led to the loss of six tons of liquid helium, causing a magnet quench and an explosion that damaged 53 magnets [46]. This delayed the start of physics by a year until November 2009, and limited the collision energy in the LHC's first run to 8 TeV. This Run 1 lasted through the beginning of 2013, before a two year shutdown and the start of 13 TeV collisions in Run 2 in 2015.

The LHC steers the counter-rotating beams to collide at four pre-defined experimental interaction points, corresponding to the ATLAS, ALICE, CMS, and LHCb experiments (illustrated in Figs. 2.1 (right) and 2.2). ALICE is optimized to study collisions of Pb nuclei (another capability of the LHC, and the reason for the use of "hadron collider" instead of "proton collider"), which produce quark–gluon plasma. LHCb is an asymmetric forward detector meant to study processes involving B hadrons. CMS and ATLAS are both general-purpose hermetic detectors, designed to study a wide range of physics. We now turn our attention to the CMS experiment, which was used to collect the data utilized in the present analysis.

2.2 The CMS detector

The Compact Muon Solenoid (CMS) detector is one of two general-purpose detectors at the LHC, designed to capture the details of proton-proton collisions as completely as possible in order to study the SM and search for any signs of physics beyond the SM.

CMS is 21 meters long and 15 meters in diameter (relatively "compact", compared to ATLAS). It is built around a central solenoidal magnet, which provides a powerful magnetic field that bends the trajectories of charged particles. The detector consists of various layers of sub-detectors and components, illustrated in the cutaway diagram in Fig. 2.3. The layers, working from the inside out, are: the beampipe; silicon tracker; electromagnetic calorimeter; hadronic calorimeter; magnet; muon detectors; and steel return yokes (these last two are interspersed in alternating layers). Moreover, these components are organized into a central cylindrical barrel, providing coverage out to roughly 25° from the beamline, and two circular endcaps that can be pulled away for access to the detector. Photos of an opened CMS (i.e. with the endcaps pulled away)

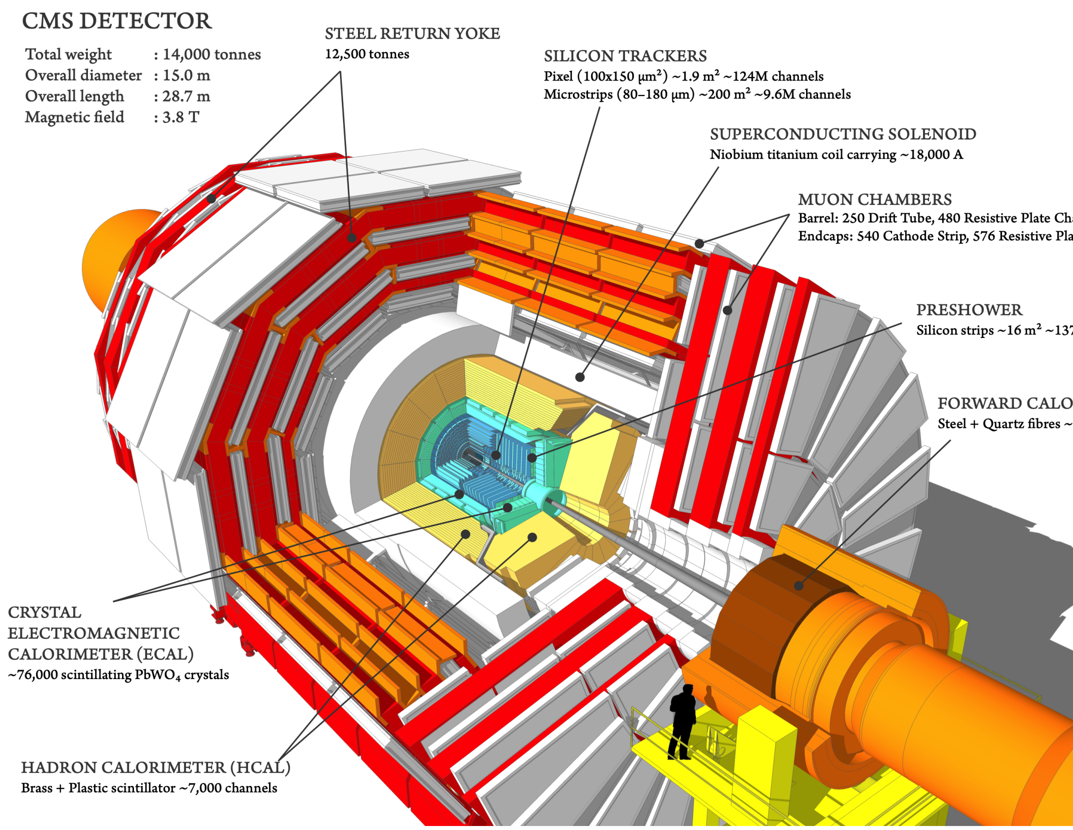

Figure 2.3: Cutaway illustration of the CMS detector. Note the outline of a human in the lower right for scale. (Image from [47])

are shown in Fig. 2.4.

We now go through each of the components and briefly describe their design and operation. More detailed accounts and the full technical details can be found in the CMS Technical Design Reports (volume I contains detector performance and software [48], and volume II describes physics performance [49]. There are also more recent reports describing various detector upgrades).

First, a word on coordinates and notation (used in this chapter and throughout this dissertation): the standard CMS coordinate system is centered on the nominal collision point, with the x axis pointing radially inward toward the center of the LHC, the y axis pointing vertically upward, and the z axis pointing along the beamline, counter-

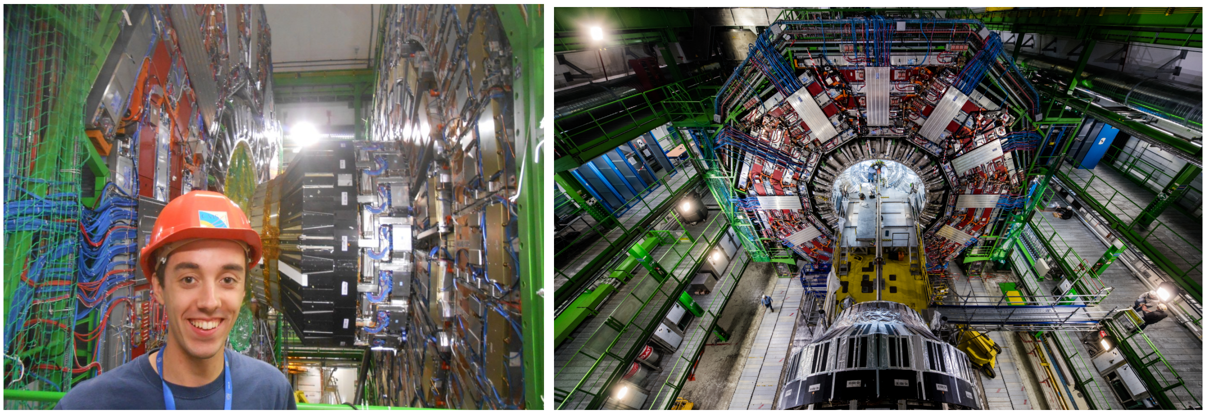

Figure 2.4: (left) The author next to an opened CMS detector in summer 2014. On the right side is the muon endcap detector, with the endcap hadronic calorimeter protruding. On the left is the barrel, with the muon detectors and iron yokes on the outside, and the magnet forming the inner ring. (right) A broader view of the opened CMS detector, during the installation of the new phase-1 upgrade pixel detector in 2017. Note the technicians inside the barrel for scale. (Image via CERN)

clockwise if looking at the LHC from above. The coordinate r refers to the cylindrical radius $r = \sqrt{x^2 + y^2}$, and ϕ the azimuthal angle $\tan(\phi) = y/x$. Instead of the polar angle $\theta = \tan^{-1}(r/z)$, particle physicists generally use pseudorapidity $\eta = -\log(\tan(\theta/2))$. The reason is that for relativistic particles, η is a very good approximation of rapidity, and differences in rapidity are invariant under Lorentz boosts along the z axis. A pseudorapidity of 0 is central (i.e. orthogonal to the beamline), and a pseudorapidity high in absolute value indicates something nearly parallel to the beamline. A subscript T on a quantity (e.g. p_{T}) indicates the transverse (i.e. xy) component.

2.2.1 Magnet and return yoke

The core of CMS is a large solenoid designed to produce a strong magnetic field. Charged particle trajectories are curved by this field, allowing reconstruction of their momenta (the radius of curvature is given by $R = p_{\mathrm{T}}/qB$). The solenoid is 13 meters long and has an inner radius of around 3 meters, making it the largest solenoidal magnet

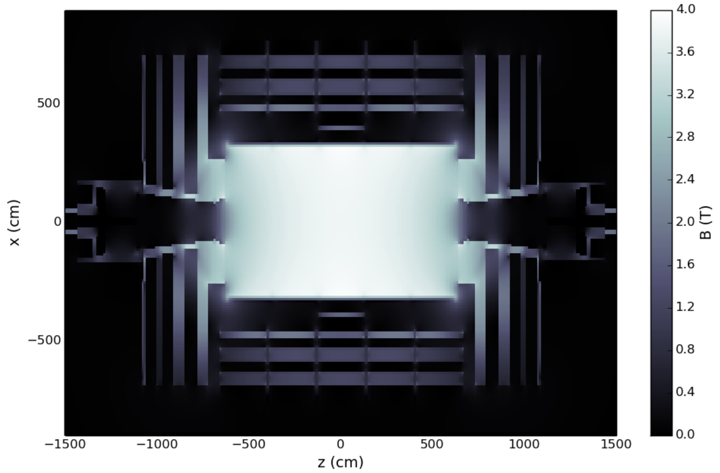

Figure 2.5: Map of the CMS magnetic field magnitude in an rz cross-sectional plane. The magnitude is a roughly constant 3.8 T in the core inside the magnet coils, and weaker outside. The steel return yokes that carry back most of the magnetic flux are visible.

ever built. The coils are made of superconducting niobium-titanium, and carry a current of 18,160 A to produce a magnetic field of 3.8 T.

Outside of the magnet coils are three layers of steel (called the "return yoke"), which provide support and guide the exterior return field (to limit the field strength outside of the physical detector). A map of the magnetic field magnitude in a cross-sectional plane of CMS is shown in Fig. 2.5. The central 13 m $\times$ 6 m core is visible, along with the layers of steel return yoke that carry most of the exterior field. Outside of these structures, the field is relatively small.

2.2.2 Tracker

The innermost detector, just outside of the beampipe and inside of the calorimeters and magnet, is the silicon tracking detector. The aim of this detector is to accurately reconstruct the curved trajectories of charged particles in order to reconstruct their momenta.

A central challenge in the design of this detector is the radiation level so close to the beam: at 10 cm from the beampipe, the incident particle flux is around 10 million per cm^2 per second. Hence, the detector must be radiation-hard. Additionally, material budget should remain low so as to minimally affect the trajectories of throughgoing particles, and the detector should give high spatial resolution and have fast response time, to allow accurate and timely reconstruction of particle tracks.

From these considerations, a silicon detector was chosen, as silicon is relatively resilient to radiation and has a fast response time. The detector works by reverse-biasing narrow strips of silicon; throughgoing charged particles then ionize the atoms to create electron-hole pairs, which can be collected via a voltage gradient and detected as a small electrical pulse lasting a few nanoseconds.

The tracker consists of two main sub-modules. The first, called the pixel detector, is close to the beamline and consists of arrays of 100 μm $\times$ 150 μm silicon pixels, with 124 million pixels in total. There are four barrel layers at radii 2.9, 6.8, 10.9, and 16.0 cm, and three endcap layers. The detector provides coverage out to $|\eta| < 2.4$. The small size of the pixels allows for highly accurate identification of particle location.

Outside of the pixel detector are 10 layers of silicon strips in the barrel, and 12 in the endcaps. The four inner barrel layers consist of 10 cm $\times$ 180 μm strips, and the next six consist of 25 cm $\times$ 180 μm strips. There are 10 millions strips in total, and the strip tracker reaches out to a radius of 130 cm.

2.2.3 Electromagnetic calorimeter

Between the tracker and the magnet lie two calorimeters. The first, called the electromagnetic calorimeter (ECAL), is designed to measure the energy of electrons and photons. It works by triggering an electromagnetic shower: in the presence of mat-

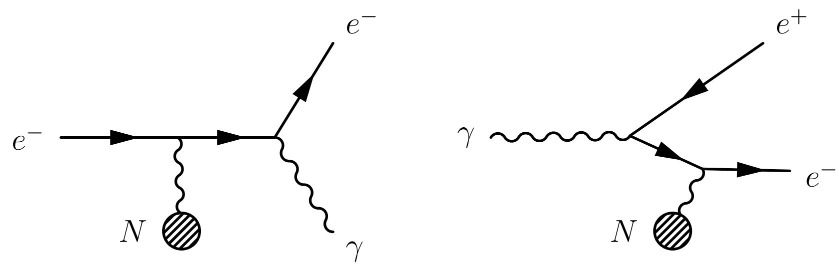

Figure 2.6: Diagrams for bremsstrahlung ($e^{\pm} \to e^{\pm}\gamma$, left) and photon conversion ($\gamma \to e^{+}e^{-}$, right), the processes that cause electromagnetic showers. The presence of a nucleus N to absorb/provide momentum is necessary for the conservation of momentum.

ter, a high-energy electron will lose energy by emitting a photon in a process called bremsstrahlung, and a high-energy photon will "convert" into an $e^{+}e^{-}$ pair (see Fig. 2.6 for diagrams of these processes). These processes then produce a cascade of electrons and photons, which cause the material of the calorimeter to emit scintillation light that can be detected and measured.

The ECAL is a homogeneous calorimeter (i.e., the material that produces the shower is also used for measuring the energy). It is constructed from crystals of lead tungstate ($PbWO_4$). This material is very dense, good for producing showers and stopping particles in a short distance, and is also optically clear and scintillates when electrons and photons pass through it. The scintillation timescale is also short, making for fast, well-defined pulses.

The crystals are individually 22 mm $\times$ 22 mm $\times$ 230 mm, giving the detector a depth of 26 radiation lengths ($\chi_0 = 0.89$ cm for $PbWO_4$). The barrel consists of 61,200 crystals organized into 36 supermodules, and the flat endcaps consist of an additional 15,000 crystals.

The inner sides of the ECAL endcaps are covered by a "preshower detector", consisting of two layers of lead interleaved with two layers of silicon strip detectors, that are used to distinguish between high-energy prompt photons and less interesting closely-spaced photon pairs from $\pi^0 \to \gamma\gamma$ decay. The lead layers cause photons to shower, and the silicon strips provide a high-granularity snapshot of the shower. Two closely spaced photons from π^0 decay, while often indistinguishable in the ECAL, can be distinguished in the higher resolution silicon detector.

2.2.4 Hadronic calorimeter

The final sub-detector inside of the magnet volume is the hadronic calorimeter, or HCAL. The purpose of this calorimeter is to measure the energy of charged and neutral hadrons (most often charged pions, charged and neutral kaons, protons, and neutrons). These hadrons, in contrast to electrons and photons, are not stopped by the ECAL, and must be measured by a separate detector.

In contrast to electromagnetic showers, hadronic showers develop primarily via the strong interaction, and are less contained (i.e., the their characteristic longitudinal and lateral sizes are bigger). In order then to provide the maximal amount of absorption power within the confines of the magnet radius, a "sampling calorimeter" design was chosen. In this case, brass absorption layers, very efficient at stopping particles and causing showers, are interleaved with layers of plastic scintillator, which produce detectable light as particles pass through. The light from the scintillators is captured by wavelength-shifting optical fibers and converted into electrical signals by magnetic field-resistant hybrid photomultipliers. This design has a worse energy resolution than a homogeneous calorimeter design like the ECAL, as some fraction of the energy is lost in the absorption material and must be estimated.

The HCAL barrel and endcaps, each made up of 36 wedges, make up the bulk of the HCAL detector. However, there is an additional "outer barrel" detector that sits just outside of the magnet coil, used to record energy from any particles that manage to get through. Finally, there are separate forward hadronic calorimeters at either end of CMS that record energies of particles very close to the beamline ($3.0 < |\eta| < 5.0$). Due to the high radiation doses and higher particle density these detectors face, they have a slightly different design with steel absorbers and Cerenkov light-producing quartz readout fibers.

2.2.5 Muon detectors

The only SM particles produced at the LHC that can make it through CMS's calorimeters are neutrinos and muons. Neutrinos are only weakly interacting, so there is no feasible way to detect them and their presence can only be inferred through transverse momentum imbalance. Muons, on the other hand, are minimum-ionizing due to their heavier mass and hence easily pass through the calorimeters, but are still electrically charged so can be detected with dedicated components.

CMS employs three different types of muon detectors, all located outside of the calorimeters and magnet where muons are likely to be the only detectable particles. Drift tubes (DTs) in the barrel and cathode strip chambers (CSCs) in the endcaps are used for precise position measurements from which muon trajectories are inferred. Additionally, resistive plate chambers (RPCs) provide fast measurements in both the barrel and endcaps and are used for triggering.

The DTs utilize 4 cm wide gas-filled tubes containing a stretched wire with a positive voltage. Throughgoing muons ionize atoms in the Ar/CO_2 gas mixture, and the electrons travel to the central wire and register a detectable current. By utilizing information from perpendicular wires, the DTs provide a 2D position measurement; resolution is on the order of 100 μm. There are four layers of DTs, at radii 4.0, 4.9, 5.9, and 7.0 m from the beam axis.

DTs do not function as well with high occupancy rates or uneven magnetic fields, and are more susceptible to neutron backgrounds, so in the endcaps CSC detectors are used. CSCs contain a gaseous volume, similar to the DTs, but with both positively-charged anode wires and perpendicular negatively-charged cathode strips. Muons ionize atoms in the gas, electrons travel to the anode, and the resultant positive ions travel to the cathode, inducing detectable currents. The perpendicular nature of the anode wires and cathode strips provide a 2D position measurement. Each chamber is composed of six layers that provide independent measurements, which are averaged together to give a more precise global fit. The resultant spatial resolution is on the order of 200 μm, and the angular resolution in ϕ is of order 10 mrad.

The DTs and CSCs are both fast enough to be used in the trigger system. However, a third even faster type of muon detector is additionally used in both the barrel and endcaps for redundancy and to allow unambiguous identification of the correct bunch crossing. These RPCs consist of two parallel plates made of a high-resistivity plastic, one positively and one negatively charged. They are separated by a gaseous volume. When a muon passes through, it ionizes electrons from the gas, which pass through the plastic plates and are detected by external metallic strips. The pattern on the strips is able to give a quick muon measurement, with time resolution of only 1 ns. This is used by the triggering system to quickly decide whether or not the event may have something "interesting" enough to keep.

2.2.6 Trigger system

With a bunch spacing time of 25 ns, the LHC delivers 40 million bunch crossings per second. This translates to over 1 billion individual *pp* interactions per second, once accounting for the total *pp* cross section and LHC luminosity. Recording full detector information from a single bunch crossing takes on average 460 kB of disk space. At 40

MHz, this means that recording every event would require writing around 18 TB/s to disk! This is completely infeasible, so a drastic reduction in the number of recorded events is necessary.

To do this, CMS (and other experiments) make use of a trigger system, which is able to quickly decide which events are potentially interesting and worth saving. In CMS's case, it is a two-tiered system, with the "Level-1" (L1) trigger making very fast decisions based on rough low-level event information, and the "high-level" trigger (HLT) examining the output of the L1 trigger at a finer level of detail to achieve further reduction.

The L1 trigger consists of fast custom electronics that look for the existence of various "trigger primitive" objects such as photons, electrons, jets, and muons from low-level information from the calorimeter and muon systems (tracking is too slow to be used in the L1 trigger). For example, the L1 trigger might look for a sum of calorimeter deposits (H_{T}) larger than 300 GeV, or a muon track with $p_{\mathrm{T}} > 18$ GeV. The L1 trigger lowers the event rate by a factor of around 1000, to about 100 kHz, and the resulting events are then passed to the HLT.

The HLT is a software-based trigger that runs a fuller reconstruction of the events passing the L1 trigger, allowing for more detailed analysis of the events and tighter restrictions on what passes. For example, one trigger path might look for a muon with $p_{\mathrm{T}} > 12$ GeV that is isolated (i.e., no other high-p_{T} tracks near it), along with an electron with $p_{\mathrm{T}} > 23$ GeV that is isolated and passes a loose identification requirement. The HLT reduces the event rate to around 1000 Hz, and the passing events are then saved to disk for offline analysis.

Often physicists are interested in studying a particular kind of event that occurs much too often to be recorded 100% of the time. To allow for these types of events, trigger paths can be "prescaled", which means that only one in every N events is saved. As an example, in Chapter 8 we make use of an inclusive selection of events with $H_{\mathrm{T}} > 250$ GeV. At peak luminosity these events happen at around 3 kHz, which would completely overwhelm the HLT. So the trigger path for these events is prescaled by a factor of $\sim$4000, reducing their contribution to the overall rate to below 1 Hz.

2.2.7 Computing and reconstruction

Events that pass the triggers are then fed into the worldwide CMS computing grid for processing and analysis. First, the Tier-0 computing center on-site at CERN performs initial reconstruction, and distributes data between $\sim$10 Tier-1 computing centers across the world. These sites store the large, lower-level reconstructions of the data and can re-process the data when necessary. Finally, $\sim$150 Tier-2 centers around the world store smaller high-level reconstructions, and facilitate analysis of data by end-users.

CMS utilizes a hierarchy of data formats, as follows:

- RAW: full set of low-level "raw" detector information, such as hits on individual detector elements. A typical RAW event from the JetHT data set is 770 kB.
- RECO: the output of the initial processing at the Tier-0, containing reconstructed but still very detailed physics objects. A typical JetHT event size is 2–3 MB, due to the additional information.
- AOD: a much slimmed-down version of the RECO event information, eliminating much of the space-consuming low-level information not necessary for a large majority of analysis. A typical JetHT event size is 330 kB.
- MINIAOD: a further reduction of the AOD tier, eliminating much rarely-used information and adding some relevant high-level information used by many analyses (e.g. b-tagging information, jet/lepton IDs, etc.) A typical JetHT event size is 40 kB. This is the data format used for the analysis in this dissertation.
- NANOAOD: a newly introduced format, consisting of a lightweight columnar data

set built from MINIAOD. Much low-level particle information is dropped, keeping only higher level variables used for a majority of analyses. A typical JetHT event size is 0.7 kB.

In addition to processing real data from the detector, the CMS computing network also manages the production of simulated data, or Monte Carlo (MC). For a given process (e.g. $Z \to \ell^+\ell^-$), the initial events are generated with an event generator such as MadGraph [50]; parton showering, fragmentation, and decays of semi-stable particles are done with a tool such as pythia [51]; and then all of the resulting particles are fed into a full Geant4 simulation [52] of the CMS detector. This produces realistic detector hits that can then be processed with the same code as is used for real data, producing the same data formats as above. For certain samples where computation time is prohibitive (usually for BSM models), CMS has developed a fast simulation ("FastSim") framework [53] that forgoes the full Geant4 simulation in favor of a much faster parametric approach. Separate simulation is generated for each year of data collection (2016, 2017, and 2018 for the analysis presented here), as pileup and detector conditions change from year to year.

2.3 Phase-2 CMS MIP timing detector

The LHC is currently in the midst of "Long Shutdown 2", which will be followed by Run 3 which lasts through 2024. This is then followed by Long Shutdown 3, during which the LHC and experiments will be upgraded and prepared for the high-luminosity LHC (HL-LHC) era starting in 2027.

The HL-LHC will provide a leveled luminosity of 5.0×10^{34} cm^{-2}s^{-1}, two and a half times larger than the present maximum, and corresponding to an average pileup (PU) of 140 *pp* interactions per bunch crossing. Ultimately, the plan is to go all the way to

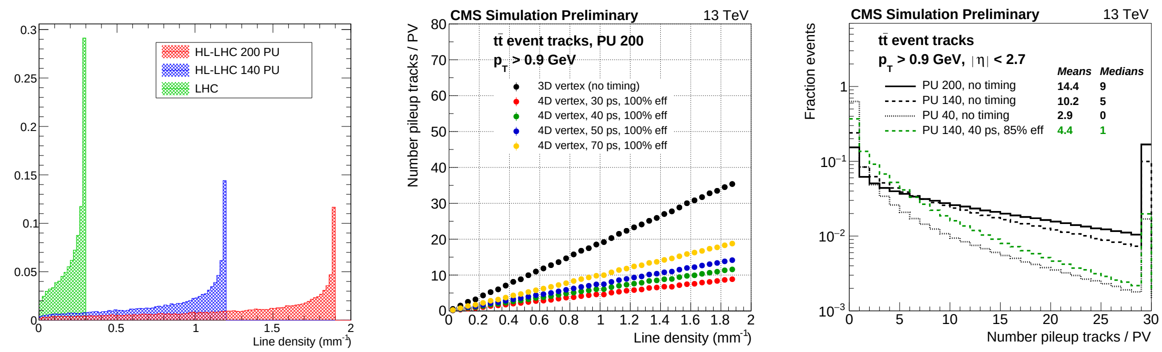

Figure 2.7: (left) Distribution of line density of PU vertices for different $\langle N_{\rm PU}\rangle$ scenarios. The right edges of the distributions correspond to a vertex in the center of the Gaussian spatial spread of vertices. (center) The mean number of PU tracks associated to the primary vertex as a function of line density, with no timing information and with 30, 40, 50, and 70 ps timing resolutions, assuming $\langle N_{\rm PU}\rangle = 200$. A resolution of $\sim$30 ps reduces the PU contamination to near present-LHC levels. (right) Distributions of the number of PU tracks associated to the primary vertex, assuming no timing and $\langle N_{\rm PU}\rangle$ of 40, 140, and 200, and under a realistic HL-LHC scenario of $\langle N_{\rm PU}\rangle = 140$ with a timing layer with $\sigma_t = 40$ ps and an efficiency of 85%.

7.5×10^{34} cm^{-2}s^{-1}, corresponding to an average of 200 pp interactions per bunch crossing. The current CMS experiment is ill-equipped to handle this many simultaneous collisions, as the increased spatial overlap of tracks, interaction vertices, and calorimeter deposits will degrade particle identification and reconstruction capabilities below acceptable levels (not to mention the greatly increased radiation levels that current components cannot handle).

In order to maintain performance during the HL-LHC era, CMS plans to under go a "Phase-2 upgrade" [54], in which components will be upgraded to withstand higher levels of radiation and allow more discrimination between overlapping PU tracks. A key component of this is the addition of a minimum-ionizing particle (MIP) timing detector between the tracker and calorimeters [55].

The idea is to add a "fourth dimension" in the tracking of particles, which will allow discrimination between spatially overlapping vertices. Pileup interactions have not only

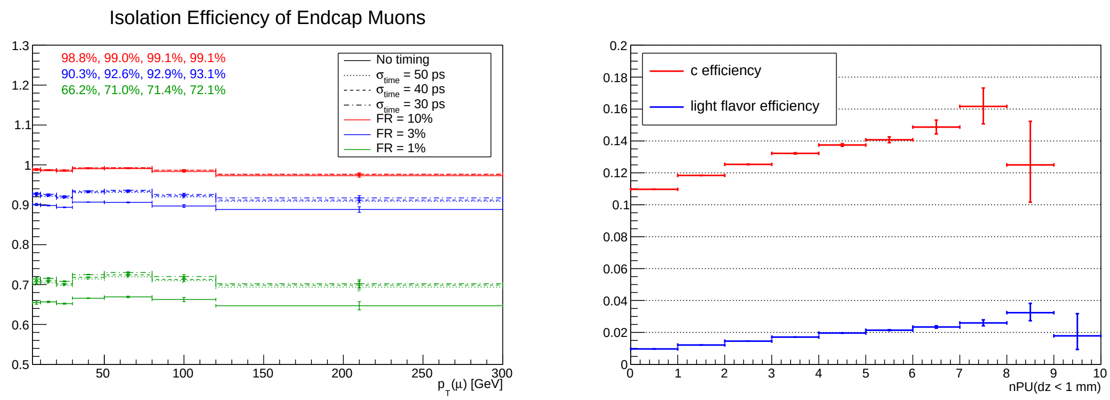

Figure 2.8: (left) Track isolation efficiency for prompt muons under different timing scenarios, for various fixed fake rates. (right) b-tagging fake rate for charm and light flavor quarks, as a function of the number of PU vertices within 1 mm of the primary vertex. Adding timing information effectively reduces this number.

a spatial spread with an RMS of around 4.3 cm, but also a spread in time with an RMS of 180–200 ps. An ability to measure the time of tracks with a resolution of around 40 ps would allow an effective reduction of HL-LHC pileup to current levels. This is illustrated in Fig. 2.7, which shows the distribution of PU vertex line density under different $\langle N_{\mathrm{PU}} \rangle$ scenarios, and the number of PU tracks associated to the primary vertex as under different timing capability assumptions.

In addition to reducing the raw number of PU vertices near the primary vertex, adding timing information has visible effects on more tangible detector performance metrics such as prompt lepton identification and b tagging, as shown in Fig. 2.8. These figures were made using generator-level tracks from $t\bar{t}$ simulation, overlaying 200 random minimum-bias interaction vertices, and emulating a timing layer by smearing each track's time by the desired resolution.

The design of the timing layer is separated into barrel and endcap components, due to their different physics requirements. The barrel timing layer, situated between the tracker and the ECAL, will consist of arrays of scintillating Lutetium Yttrium Orthosilicate

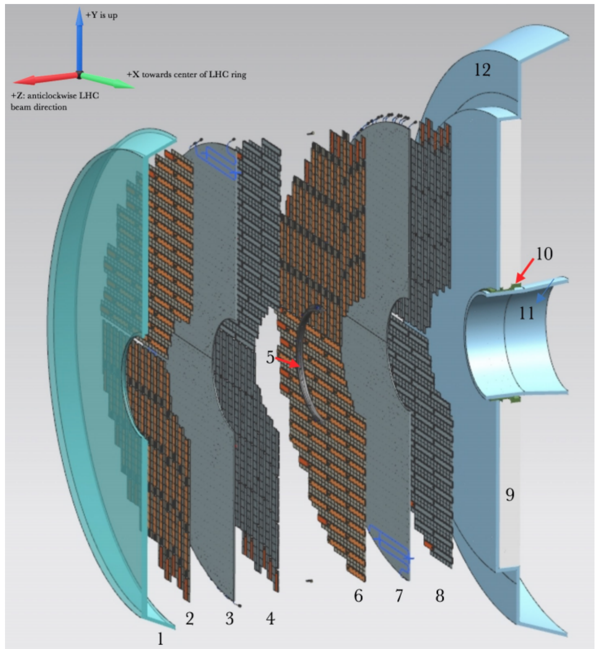

Figure 2.9: Layout of the proposed endcap timing layer. Rows of LGAD sensors are mounted on either side of two support disks, with the rows on either side offset such that full coverage is attained. (Image from [55])

crystals, read out by silicon photomultipliers. Charged particles passing through generate optical photons, which are measured to give a timestamp of the passing particle.

The endcap detectors are closer to the beamline and hence receive much larger radiation doses. This would degrade the performance of scintillating crystals too quickly, so silicon low-gain avalanche detectors (LGADs) [56] are used instead, which add a thin p-type multiplication layer to traditional silicon sensors, and can achieve time resolutions of up to 20 ps. Rows of LGAD sensors will be mounted on either side of two support disks, as shown in Fig. 2.9. Gaps are left between the rows to allow for power and cooling service channels, but the rows on either side of the disks are offset to provide full coverage. The disks are mounted on the inside of the calorimeter endcap structure.

Chapter 3

Jets, Missing Energy, and Jet Response Templates

Understanding the production and measurement of hadronic jets is of central importance to any hadron collider experiment. Not only is pure multijet production the most common type of high-p_{T} event, but *any* process can be accompanied by one or more jets arising from initial- or final-state radiation of quarks and gluons. Moreover, any mis-measurement of jet energies harms the experiment's ability to measure the imbalance in transverse momentum, a key component of many analyses. Understanding jet mis-measurement is thus highly important for any analysis sensitive to this fake missing transverse momentum, such as the one presented here. Secs. 3.1 and 3.2 describe how jets and missing transverse momentum are measured at CMS, Sec. 3.3 summarizes the various sources of jet mis-measurement, and Sec. 3.4 details how the jet response of the detector is actually measured so that it can be accounted for.

3.1 Jets at CMS

The production of jets, narrow cones of hadrons produced in pp collisions, was described in Sec. 1.4. Due to color confinement, isolated colored states cannot exist, so high-energy quark and gluons that are expelled from the collision point will *hadronize* into a collection of hadrons in a cone around the initial momentum direction.

Jets consist predominantly of a mixture of charged an neutral hadrons, though they may also contain photons or leptons from the decay of heavy-flavor hadrons. The charged hadrons leave tracks in the tracker, and both the charged and neutral hadrons deposit the majority of their energy in the HCAL.

3.1.1 Jet clustering

A central task in event reconstruction is identifying individual jets among the mess of tracks and energy deposits left in the detector. To do this there exist various clustering algorithms that attempt to reasonably group objects reconstructed by the detector into separate jets.

As almost everything in a jet leaves most of its energy in the calorimeters, one option is to just try to cluster the calorimeter deposits into jets. However, calorimeters have relatively low energy resolution, so CMS has developed a "particle-flow" (PF) algorithm [57] that combines calorimeter information with that from the tracker and muon systems to try to reconstruct each individual particle. By using information from the higher-resolution tracker and muon systems, a superior jet energy measurement is achieved. A rough outline of the algorithm is as follows:

- Match tracks in the tracker to tracks in the muon system. These are PF muons.
- Match remaining tracks to calorimeter deposits in the HCAL and ECAL. These are PF charged hadrons (HCAL) or electrons (ECAL).
- Any remaining calorimeter energy is assigned to PF neutral hadrons (HCAL) or photons (ECAL).

Once particles are identified, they are clustered into "PF jets". There are a variety of ways to do this, but the most commonly used at CMS and the one used for this analysis is the anti-k_{T} algorithm [58]. Essentially, this defines a distance metric parametrized by a characteristic size R and iteratively combines objects until a certain criterion is met. The parameter R roughly corresponds to the radius of the jet in (η, ϕ) space. For this analysis we use a parameter of $R = 0.4$. For analyses looking at boosted heavy objects such as top quarks or Higgs bosons, resulting jets are often merged into one large jet and a parameter of $R = 0.8$ is commonly used.

3.1.2 Pileup removal

Blindly applying the above clustering procedures over all particles in an event will lead to many particles from pileup interactions being clustered into jets originating from the primary vertex, artificially increasing their energies. There are a number of ways to attempt to account for this. The first, used for jets in this analysis, is known as "charged hadron subtraction" (CHS). One simply removes all of the charged hadrons whose tracks are unambiguously associated with a pileup vertex before clustering the remaining particles into jets. This method is simple, but does not account for any pileup contribution from particles that are not charged hadrons.

Another commonly used algorithm is called "pileup per particle identification" (PUPPI). In this case, each particle is assigned a weight based on the likelihood that it originates from pileup, by considering the shape of the energy distribution in the immediate vicinity of the particle. The particle energies are then scaled by this weight when computing jet energies. The PUPPI algorithm gives slightly better resolution in jet p_{T} and mass, and is used by analyses that are particularly sensitive to these quantities.

3.1.3 Jet energy corrections

The measurement of jet energies is imperfect, as there are uncertainties from pileup contribution, calorimeter inefficiencies and fluctuations, undetectable neutrinos, etc. CMS attempts to correct for this by calibrating the jet energy scale as a function of various event and jet variables, and correcting measured jet energies to give a more uniform response. These corrections are together referred to as *jet energy corrections*, or JECs.

The procedure is factorized into a number of discrete stages. First, the L1 correction (unrelated to the L1 trigger) attempts to remove any residual energy coming from pileup interactions. The corrections are derived from a MC sample of dijet events, processed once with pileup interactions overlaid on the main event, and once with no pileup added. The quantity $p_{\mathrm{T}}^{\mathrm{PU}}/p_{\mathrm{T}}^{\mathrm{no\ PU}}$ is fit in bins of jet η as a function of mean event energy density ρ (correlated with the number of pileup interactions N_{PU}), jet area A, and jet p_{T}.

Next, the L2 corrections account for imperfect and non-uniform detector response. It is derived again on a MC sample of dijet events, with L1 corrections already applied to the reconstructed jet energies. The quantity $p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{gen}}$, where $p_{\mathrm{T}}^{\mathrm{reco}}$ is the reconstructed jet energy and $p_{\mathrm{T}}^{\mathrm{gen}}$ is the generator-level jet energy (with neutrinos removed), is fit in bins of jet η as a function of jet p_{T}.

Finally, it is known that detector response is not modeled exactly correctly in MC, so data-only L2 residual corrections are derived on a control sample of real data. *Relative* corrections as a function of jet η are derived from a dijet sample, by comparing the jet p_{T} to a reference jet. *Absolute* corrections, which correct the η-independent absolute jet energy scale as a function of p_{T}, are derived on a sample of Z+jets or γ+jets events.

3.1.4 b-jet tagging

Many important processes contain final states with b quarks, including the decays of top quarks, Higgs bosons, and SUSY particles. Identifying the jets that originate from b quarks is a key method for reducing background.

Hadrons involving b quarks tend to be heavier and have longer lifetimes than hadrons made from lighter quarks (for example, B mesons haves masses of $\sim$5 GeV and $c\tau$ values of $\mathcal{O}$(1 mm) when Lorentz-boosted). These properties can be used to discriminate b jets from jets produced by light-flavor quarks or gluons. Some variables/properties that provide discriminating power are:

- Track impact parameters d_{z} and d_{xy}. Charged hadrons from b hadron decay tend to be slightly displaced from the primary vertex, as the b hadron propagates a measurable amount before decaying. Related to this is the impact parameter significance $\mathrm{SIP}_{z,xy}$, computed by dividing the impact parameter by the uncertainty in its measurement.
- Reconstructed secondary vertices. If there are enough high-quality tracks from the decay of a displaced b hadron, a *secondary vertex* may be reconstructed which can indicate a b hadron was present in the jet.
- Track p_{T} relative to the jet axis. Since b hadrons are heavy, their decay products tend to have more transverse momentum relative to the original direction of propagation.
- Presence of a charged lepton. Around 20% of decays of b hadrons result in an electron or muon, much higher than for non-b hadrons. Hence, the presence of a charged lepton in a jet can indicate b decay.

These and other variables can be fed into a variety of multivariate algorithms in order to identify b jets. A number of these algorithms have been designed and tested by

CMS [59]. For this analysis, we use the DeepCSV algorithm [60], a deep neural network trained on track-level variables to discriminate between b, c, bb, cc, and light-flavor jets.

Three nominal working points are set for the trained algorithm, corresponding to light-flavor fake rates of roughly 0.1%, 1%, and 10%. We use the medium working point, which achieves an efficiency of 68% and fake rates of 12% and 1.1% for c and light-flavor jets, respectively, for jets with $p_{\mathrm{T}} > 20$ GeV.

The efficiency and fake rates are not identical between data and simulation. To correct for the known differences, CMS derives b-tagging scale factors that are applied to MC. The procedure is as follows: for each MC event, compute the probabilities of the observed b-tagging configuration as

$$P(\mathrm{MC}) = \prod_{i \text{ tagged}} \varepsilon_i \prod_{j \text{ not tagged}} (1 - \varepsilon_j) \tag{3.1}$$

$$P(\mathrm{data}) = \prod_{i \text{ tagged}} SF_i \varepsilon_i \prod_{j \text{ not tagged}} (1 - SF_j \varepsilon_j), \tag{3.2}$$

where ε_i is the measured efficiency to tag a particular jet as a b jet (a function of p_{T}, η, and generator-level jet flavor, which is one of b, c, or light-flavor), and SF_i is the derived data/MC scale factor, also a function of jet p_{T}, η, and generator-level flavor. The entire event is then weighted by a factor

$$w = \frac{P(\mathrm{data})}{P(\mathrm{MC})}. \tag{3.3}$$

3.2 Missing transverse momentum at CMS

As protons are composite particles, the underlying collisions at the LHC are between partons carrying an undetermined fraction of the total momentum. Hence, the lab frame is in general not the center-of-mass frame, and the total momentum of the colliding system is boosted along the beam direction. The *transverse* momentum, orthogonal to the beamline, on the contrary should be zero (technically, the partons can have a transverse momentum on the scale of the proton mass of a few hundred MeV, but this is negligible at the LHC energy scales).

If all produced particles were perfectly measured, then, the vector sum of their $\vec{p}_{\mathrm{T}}$ values should be zero. This is not the case in practice, however. Any measured imbalance in the transverse momentum is called *missing transverse momentum/energy* (though "energy" is a bit of a misnomer, as energy is a scalar with no directionality), referred to variously as $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$, $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, $\vec{\not{E}}_{\mathrm{T}}$, or MET. Real $p_{\mathrm{T}}^{\mathrm{miss}}$ comes from undetectable particles such as neutrinos or hypothetical LSPs. Fake $p_{\mathrm{T}}^{\mathrm{miss}}$ can arise from effects like pileup contamination or imperfect detector response; a central part of this analysis will be measuring the CMS jet response distributions to estimate contribution from events with fake $p_{\mathrm{T}}^{\mathrm{miss}}$ from mis-measured jets.

At CMS, the primary way of computing $p_{\mathrm{T}}^{\mathrm{miss}}$, called PFMET, is to simply sum the momenta of all reconstructed particle-flow candidates. The raw $p_{\mathrm{T}}^{\mathrm{miss}}$ can thus be defined

$$\vec{\not{E}}_{\mathrm{T}}^{\mathrm{raw}} = -\sum_{i\in\mathrm{PFcands}} \vec{p}_{\mathrm{T},i}. \tag{3.4}$$

However, we know that detector response is imperfect, and we have already calibrated for this when deriving jet energy corrections. So the $p_{\mathrm{T}}^{\mathrm{miss}}$ measurement is easily improved by simply applying the same jet energy corrections to the jets before computing $p_{\mathrm{T}}^{\mathrm{miss}}$. We can re-write the raw $p_{\mathrm{T}}^{\mathrm{miss}}$ as

$$\vec{\not{E}}_{\mathrm{T}}^{\mathrm{raw}} = -\sum_{i\in\mathrm{jets}} \vec{p}_{\mathrm{T},i}^{\mathrm{raw}} - \sum_{i\in\mathrm{uncl}} \vec{p}_{\mathrm{T},i}, \tag{3.5}$$

where the first sum is over all jets and the second is over any particle-flow candidates not clustered into jets. Then the JEC-corrected, or type 1-corrected, $p_{\mathrm{T}}^{\mathrm{miss}}$ can be computed by simply replacing the jet momenta with their corrected values:

$$\begin{aligned}\vec{\not{E}}_{\mathrm{T}}^{\text{type-1}} &= -\sum_{i\in\text{jets}} \vec{p}_{\mathrm{T},i}^{\,\text{corr}} - \sum_{i\in\text{uncl}} \vec{p}_{\mathrm{T},i} \\ &= \vec{\not{E}}_{\mathrm{T}}^{\text{raw}} - \sum_{i\in\text{jets}} \left(\vec{p}_{\mathrm{T},i}^{\,\text{corr}} - \vec{p}_{\mathrm{T},i}^{\,\text{raw}}\right).\end{aligned} \tag{3.6}$$

This type 1-corrected $p_{\mathrm{T}}^{\text{miss}}$ is the version used in this analysis.

3.3 Sources of jet mis-measurement

For the purposes of estimating background from QCD multijet events (Chapter 8), we will be interested in modeling the "response" of the CMS detector to jets. That is, for a jet with a true p_{T} of $p_{\mathrm{T}}^{\text{true}}$, what will be the measured $p_{\mathrm{T}}^{\text{reco}}$? Jet measurement is an inherently random process, so the response will be given in the form of a probability density function in the variable $p_{\mathrm{T}}^{\text{reco}}/p_{\mathrm{T}}^{\text{true}}$.

Fig. 3.1 shows a few examples of these functions measured in simulation (called jet response templates, or JRTs). It is found that the templates are described well by a central Gaussian core (red), with larger, non-Gaussian tails (green). Details on the exact derivation of these functions, and the procedure for fitting the core and tails, are given in Sec. 3.4. For now, it is sufficient to know that they are measured in simulation by matching jets clustered from reconstructed PF candidates ("reco jets") to jets clustered from generator-level particles ("gen jets") and comparing their p_{T} values.

The size and shape of the core are due to standard stochastic smearing in the calorimeters. The width of this Gaussian core relative to the jet p_{T}, referred to as the "jet resolution", generally falls as $1/\sqrt{E}$. This is illustrated in Fig. 3.2 (left), which shows the measured resolution improving as jet p_{T} increases.

The tails, on the other hand, come from rarer events in which the measured jet p_{T} is

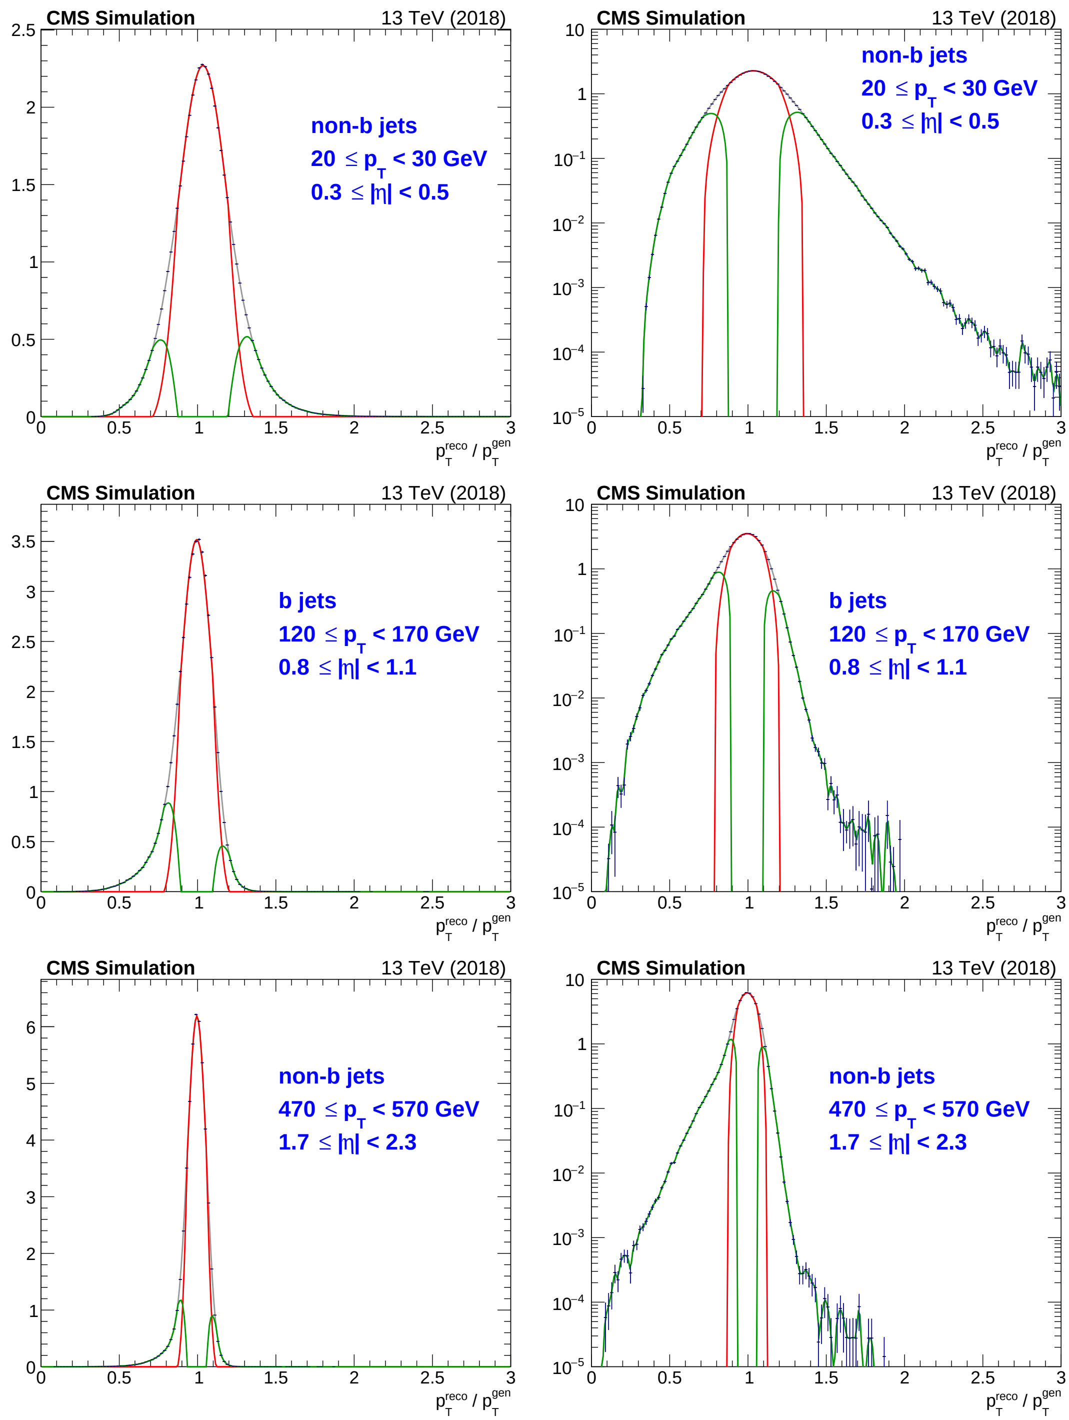

Figure 3.1: A selection of three example jet response templates for various p_T/η/b-jet bins, shown in linear (left) and log (right) scales. The dark blue points are the raw templates. The red curves are the fitted Gaussian “cores” of the templates, and the green curves are the “tails”, as described in Section 3.4.2.

much further from the true value. These types of events are harder to model, but as we will see are critically important for the QCD estimate method to work. In order for a multijet event to populate the high-$p_{\mathrm{T}}^{\mathrm{miss}}$ signal regions, generally one or more jets must be badly mis-measured (i.e., reside in the tails of the response functions).

The sources of the measured tails fall into two categories. The first, "real sources", are due to real effects that are present in the data we are interested in modeling. The second, "fake sources", are due to features of the simulation or template-derivation methodology that are not reflected in the actual data. We seek to model the first category as accurately as possible, and remove events from the second category so that they do not artificially enhance the tails of the templates.

- Real sources
  - Neutrinos from heavy-flavor decay (mostly in jets from b quarks). This enhances the *left* tails, and is the reason for deriving separate templates for b jets (see Fig. 3.2 (right)).
  - Reconstruction errors, e.g. a badly reconstructed high-p_{T} track that greatly increases the reconstructed jet p_{T}. This generally enhances the *right* tails.
  - Holes or cracks in the detector that cause part of the reconstructed jet to go "missing". This enhances the *left* tails.
  - Overlap with a pileup jet. This enhances the *right* tails.
- Fake sources
  - Gen or reco jet "splitting". i.e. for a single gen (reco) jet, the corresponding reco (gen) jet is clustered as two different jets, and only one gets matched. Depending on the direction, this enhances *both* tails.

– Mis-matching a gen jet to a reconstructed pileup jet. This generally enhances the *left* tails, but can also enhance the right tail for low-p_T jets.

– Holes in the simulated calorimeter that are not present in the real data. This enhances the *left* tails.

3.4 Derivation of jet response templates

The jet response templates are measured in simulation, by matching gen and reco jets using the distance measure $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$. Section 3.4.1 describes the gen/reco jet matching procedure, and measures taken to prevent fake matches that artificially enhance the tails of the templates. Section 3.4.2 describes the methodology for fitting the core/tails of the derived templates, and the procedure to correct the template resolutions for known differences in jet resolution between data and MC.

3.4.1 Gen/reco jet matching

The process for matching gen and reco jets in order to derive jet response templates begins with QCD Monte Carlo, in which both the generator-level and reconstructed particles have been clustered into jets. It is important to use gen jets that include neutrinos, as we are interested in modeling jet mis-measurement due to energy carried away by such neutrinos. Reco jet energies are fully corrected with the same jet energy corrections as used in the main analysis. Gen jet flavor is determined by identifying the flavor of hadrons within the jet cone. Once this is all done, all reco and gen jets with $p_\mathrm{T} > 10$ GeV are selected and saved.

After selecting the jets, gen and reco jets are matched and JRTs are constructed by filling histograms with $p_\mathrm{T}^\mathrm{reco}/p_\mathrm{T}^\mathrm{gen}$. The templates are constructed in bins of gen-level p_T and $|\eta|$, and separately for b and non-b jets. The p_T and η binning are as follows:

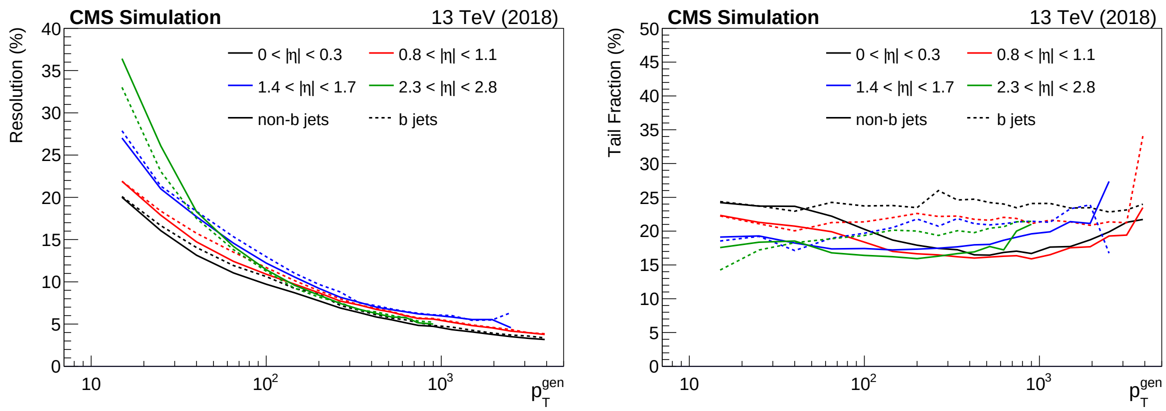

Figure 3.2: (left) Jet energy resolution (defined as the width of the fitted Gaussian core) as a function of generator-level jet p_{T}, as measured in 2018 MC. Resolution improves as jet p_{T} increases, and generally degrades as η increases. (right) The fraction of the jet response function that resides in the fitted tails (as opposed to the Gaussian core) as a function of generator-level jet p_{T}, as measured in 2018 MC. The main feature to notice is that at intermediate p_{T}, b jets have a higher probability to be in the tail. This is due to the presence of neutrinos from heavy-flavor decay in the jet, which cause the jet energy to be under-measured and hence enhance the left tail of the templates. At low p_{T}, the tails are driven by other effects so differences between b and non-b jets are not visible.

- p_{T} bin edges: [0, 20, 30, 50, 80, 120, 170, 230, 300, 380, 470, 570, 680, 800, 1000, 1300, 1700, 2200, 2800, 3500, 4300, 5200, 6500] GeV
- $|\eta|$ bin edges: [0, 0.3, 0.5, 0.8, 1.1, 1.4, 1.7, 2.3, 2.8, 3.2, 4.1, 5.0]

The algorithm for matching is outlined here; the reasoning for and effect of various steps are described following.

1. Skip events that fail any of the $p_{\mathrm{T}}^{\mathrm{miss}}$ filters used in the main analysis
2. Find all "tight pairs" of reco/gen jets that have $\Delta R(\mathrm{reco}, \mathrm{gen}) < 0.3$
3. Find all "loose pairs" of reco/gen jets that have $\Delta R(\mathrm{reco}, \mathrm{gen}) < 0.5$
4. Identify the reco jets and gen jets that are in more than one loose pair
5. Throw away any tight pairs that contain a jet identified in the previous step

6. Throw away pairs in which the reco jet is within one of the manually-identified calorimeter holes
7. Throw away pairs in which the reco jet fails tight jet ID
8. Throw away pairs in which the reco jet fails loose pileup jet ID
9. For all remaining tight pairs, identify the correct p_{T}/η-binned histograms and fill:
   - b jet histogram with weight

     $$w_{\mathrm{b}} = \varepsilon_{\text{b-tag}}(p_{\mathrm{T}},\, \eta,\, \text{gen flavor}) \times SF_{\text{b-tag}}(p_{\mathrm{T}},\, \eta,\, \text{gen flavor}),$$
   - non-b jet histogram with weight

     $$w_{\text{non-b}} = 1 - w_{\mathrm{b}},$$

   where $\varepsilon_{\text{b-tag}}$ is the efficiency for tagging a particular flavor of gen jet, as measured in MC, and $SF_{\text{b-tag}}$ is the CMS-derived scale factor to correct this efficiency for known differences between data and MC.

Step 1 rejects entire events that fail one of the MET filters (see Sec. 5.1.4) applied in the main analysis. These guard against certain kinds of mis-reconstruction that can enhance the tails of the jet response functions. Most notably, the EcalDeadCellTriggerPrimitive and HBHENoise filters reject events where there is energy in calorimeter regions known to be down/noisy (leading to fake energy, and enhancement of the right tails), and the BadChargedCandidate and BadPFMuon filters reject events containing badly reconstructed high-p_{T} tracks (which also enhance the right tails). Since we reject these events in the main analysis, we also reject jets from such events here to avoid false tails in the templates (see Fig. 3.3 (right) for an example of the effect on a template)

Steps 2 through 5 identify pairs of gen and reco jets with $\Delta R < 0.3$, but with the crucial caveat that *there are no additional matches within* $\Delta R < 0.5$. This protects

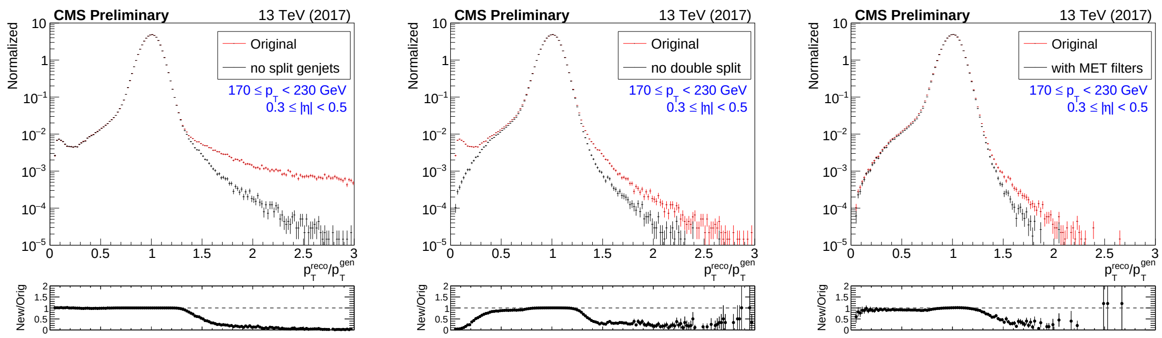

Figure 3.3: The effect of various cuts applied during the matching on an example template (non-b jet, $170 \leq p_{\mathrm{T}} < 230$ GeV, $0.3 \leq |\eta| < 0.5$). (left) Requiring *exactly* 1 gen jet within dR < 0.3 of a reco jet, instead of at least one. This removes "split" gen jets. (center) Requiring that there are no secondary nearby gen (reco) jets within dR < 0.5 of the reco (gen) jet. This protects against splitting in both directions, so reduces both tails. (right) Applying MET filters and rejecting jets from events that fail. This removes bad reconstructions (e.g. fake high-p_{T} track) and primarily reduces the right tail.

against one of the largest fake sources of tails in the measured templates: the "splitting" of gen and/or reco jets. This can happen in three distinct ways:

- A single gen jet is reconstructed as multiple reco jets, and gets matched to one. This enhances the left tail of the templates, as the matched reco jet only has a fraction of the energy.
- Multiple gen jets are reconstructed as a single reco jet. This enhances the right tail of the templates, as the matched gen jet only has a fraction of the energy.
- There are multiple nearby gen and reco jets, but energy/particles are distributed differently among the gen jets compared to reco. This enhances both tails, depending on how the energy distribution and matching happens.

Requiring that there are no secondary nearby gen (reco) jets within $\Delta R < 0.5$ of the reco (gen) jet protects against these scenarios. The effect on a sample template is shown in Fig. 3.3. The reduction in the tails after applying this cut is seen to be quite large.

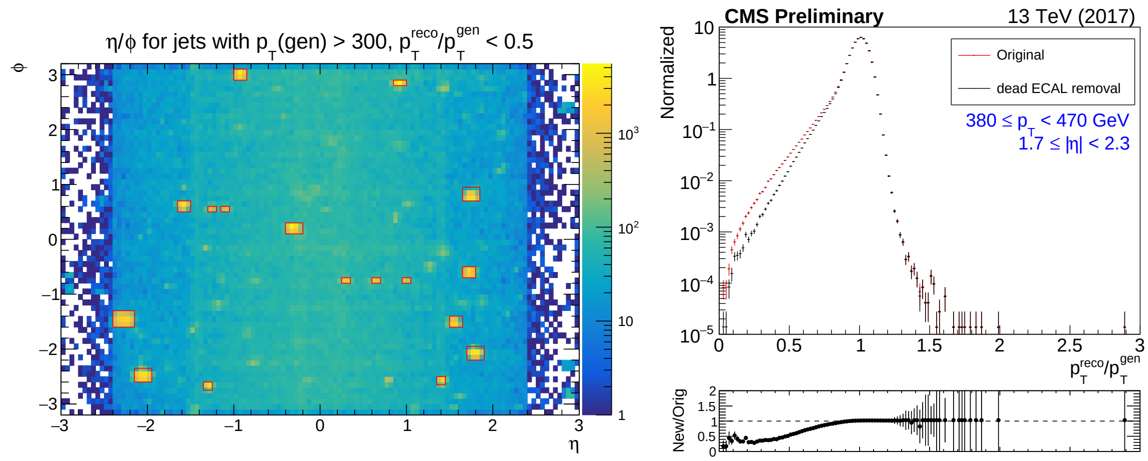

Figure 3.4: (left) an $\eta - \phi$ map for jet pairs with $p_{\mathrm{T}}^{\mathrm{gen}} > 300$ GeV and $p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{gen}} < 0.5$ from 2017 simulation. Identified hot-spots are outlined in red. Pairs with reco jets in these regions are not used. (right) The effect of this removal on an example template. Relatively large reductions in the left tail are seen for templates containing affected regions.

Step 6 throws away pairs involving a reco jet in one of a number of manually-identified "dead" spots in the detector. When plotting an $\eta - \phi$ map of the jets with small (<0.5) $p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{gen}}$ values, we observe a number of "hot-spots", corresponding to calorimeter regions that are dead or off in the simulation. These do not generally correspond to real effects in the data, and lead to artificially large left tails in the templates. We identify these regions (separately for each year's MC, as the hole locations are different) and remove any jet pairs that contain a reco jet in one of these regions. Fig. 3.4 shows these cells highlighted for 2017 simulation on the left, and the effect on an example template on the right.

Step 7 rejects pairs in which the reco jet fails tight jet ID. In the main analysis, we reject the entire event if any jet with $p_{\mathrm{T}} > 30$ GeV fails this ID. We apply the same ID here to avoid including these jets in the templates. It is found to have minimal effect, except for very high eta ($|\eta| > 4.0$).

Step 8 rejects jets that fail loose pileup jet ID [61]. Jets that fail this pileup ID are

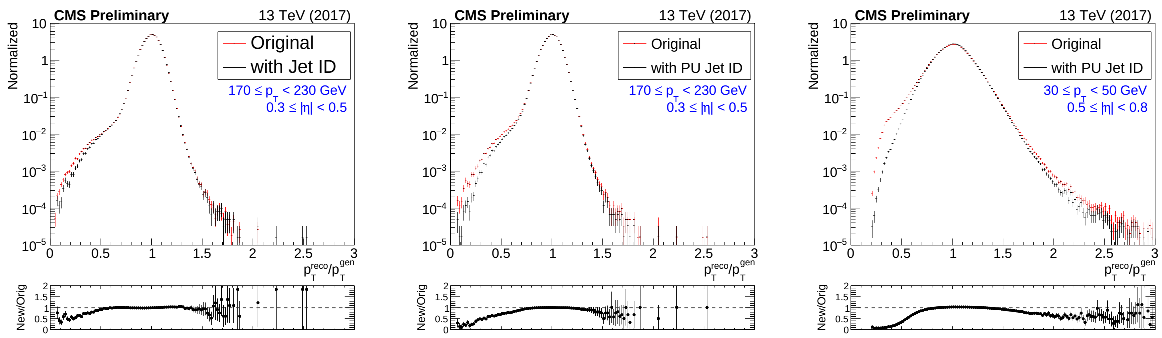

Figure 3.5: (left) and (center) show the effect of applying tight jet ID and loose pileup jet ID, respectively, to the same template as in Fig. 3.3. Jet ID has a relatively small effect, while the pileup jet ID leads to a fairly large reduction in the left tail. (right) The effect of pileup jet ID on a lower p_{T} template ($30 \leq p_{\mathrm{T}} < 50$ GeV). Here we observe a reduction in both tails.

not included in the Rebalance and Smear method described in Chapter 8. Since we do not use them there, we do not want to include them in the templates. Vetoing jets that fail pileup ID has a fairly significant effect on the tails. At higher p_{T}, we see a reduction in the left tail from gen jets that were mis-matched to a low-p_{T} pileup jet. At lower p_{T}, we see a reduction in both tails (see Fig. 3.5).

In the final step, we have identified jet pairs and need to fill histograms. The only thing left to decide is how to fill the b and non-b jet templates. There are three options that have all been tried:

1. Fill only a single histogram, based on the gen jet flavor (found by identifying flavor of hadrons within jet).

2. Fill only a single histogram, based on whether the reco jet is medium b-tagged.

3. Fill both histograms, with weights given by the probability of tagging that jet as a b jet (corrected with data/MC scale factors).

We evaluate each method empirically by observing the agreement in N_{b} in QCD-enriched control regions after applying the full Rebalance and Smear method described

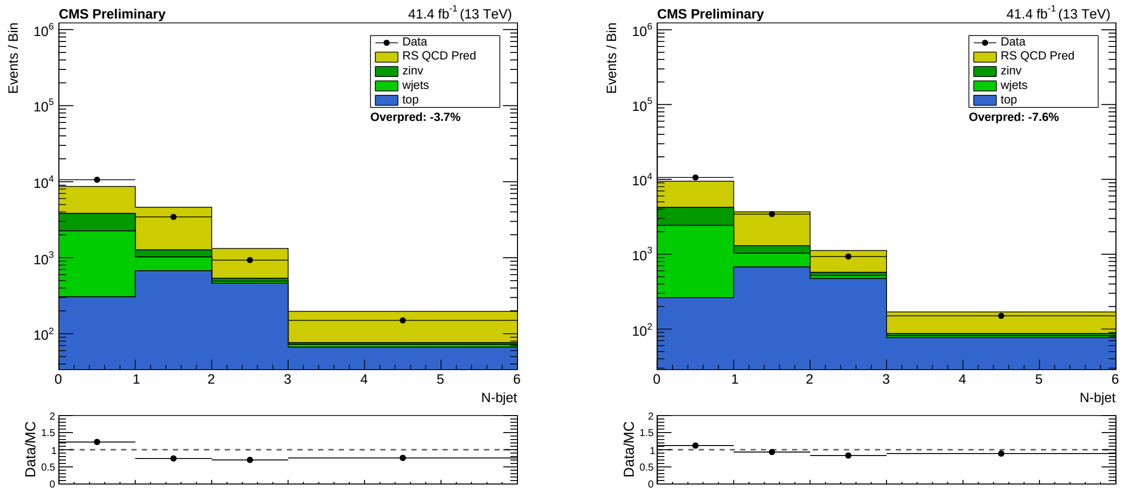

Figure 3.6: The N_{b} distributions in the inverted-$\Delta\phi_{\mathrm{min}}$ + M_{T2} sideband control region for $450 \leq H_{\mathrm{T}} < 1200$ GeV (described in Chapter 8), when (left) filling a single template based on reco jet medium b tagging, and (right) filling both histograms with weights given by the corrected probability of tagging a given jet as a medium b jet. Observed data are shown as black points, and predictions from the full Rebalance and Smear method (Chapter 8) utilizing the templates are shown in yellow. Significant improvement is seen when using the second method.

in Chapter 8. It is found that (2) slightly improves on (1), and (3) is quite a bit better than both. An example of the effect is shown in Fig. 3.6.

3.4.2 Fits and jet energy resolution corrections

Once the templates are derived, it is useful to split each template into a Gaussian "core" and non-Gaussian "tails" The reason for this is twofold. First, it is known that jet resolution is slightly different in data than in MC. This only applies to the standard Gaussian smearing of the jets, so in order to correct for the differences the Gaussian core must be isolated. Second, is is useful for studies on systematic uncertainties to be able to alter the core/tails of the templates individually. For example, to study the effect of mis-modeling in the tails, one should be able to modify the size of the tails without affecting the size/shape of the Gaussian core.

To fit the core of a template, a Gaussian is fitted in the range (mean ± RMS), where mean and RMS are the mean and standard deviation of the template. When measuring jet energy resolution and deriving scale factors, CMS only defines the core of the jet response function as extending out to $\pm 2\sigma$ of the fitted Gaussian [62]. To be consistent, we use the same definition here, and in order to avoid discontinuities, we linearly scale the Gaussian to 0 between ±1 and 2 σ. More precisely, if $g(x)$ is the full fitted Gaussian, and μ and σ are its mean and standard deviation, our defined core function is

$$\text{Core}(x) = \begin{cases} g(x) & \text{if } |x-\mu| \leq 1\sigma \\ g(x)\left(2 - \frac{|x-\mu|}{\sigma}\right) & \text{if } 1\sigma < |x-\mu| \leq 2\sigma \\ 0 & \text{if } |x-\mu| > 2\sigma \end{cases} \cdot$$

The tails are then simply defined as the full template function minus Core(x). Examples of this fitting procedure are shown in Fig. 3.1. The truncated-Gaussian cores are shown in red, and the tails in green.

The first way that these fits are used is in the correction of the templates for jet energy resolution differences between data and MC. The JetMET group provides year-dependent scale factors binned in η, shown in Table 3.1. When using MC-derived templates for real data, the core of the template for a given jet is first widened by this scale factor. This is done in a way that preserves the relative core/tail normalization. Specifically, if α is the scale factor by which we want to widen the core, the modified template is given by

$$f_\alpha(x) = \frac{1}{\alpha} \cdot \text{Core}((x-1)/\alpha + 1) + \text{Tail}(x).$$

The fitted tails are used in a similar way to study systematic uncertainties in the tail modeling. The tails are not stretched along the x direction, but instead simply scaled

Table 3.1: Jet energy resolution scale factors (data resolution divided by MC resolution) provided by the JetMET group. Uncertainties are statistical plus systematic.

	2016	2017	2018
$0.0 \leq \lvert\eta\rvert < 0.5$	1.160 ± 0.065	1.143 ± 0.022	1.150 ± 0.042
$0.5 \leq \lvert\eta\rvert < 0.8$	1.195 ± 0.065	1.182 ± 0.048	1.134 ± 0.080
$0.8 \leq \lvert\eta\rvert < 1.1$	1.146 ± 0.063	1.099 ± 0.046	1.102 ± 0.052
$1.1 \leq \lvert\eta\rvert < 1.3$	1.161 ± 0.103	1.114 ± 0.140	1.134 ± 0.112
$1.3 \leq \lvert\eta\rvert < 1.7$	1.128 ± 0.099	1.131 ± 0.147	1.104 ± 0.211
$1.7 \leq \lvert\eta\rvert < 1.9$	1.100 ± 0.108	1.160 ± 0.097	1.149 ± 0.159
$1.9 \leq \lvert\eta\rvert < 2.1$	1.143 ± 0.121	1.239 ± 0.191	1.148 ± 0.209
$2.1 \leq \lvert\eta\rvert < 2.3$	1.151 ± 0.114	1.260 ± 0.150	1.114 ± 0.191
$2.3 \leq \lvert\eta\rvert < 2.5$	1.296 ± 0.237	1.409 ± 0.202	1.347 ± 0.274
$2.5 \leq \lvert\eta\rvert < 2.8$	1.342 ± 0.209	1.991 ± 0.568	2.137 ± 0.524
$2.8 \leq \lvert\eta\rvert < 3.0$	1.779 ± 0.201	2.292 ± 0.374	1.650 ± 0.941
$3.0 \leq \lvert\eta\rvert < 3.2$	1.187 ± 0.124	1.270 ± 0.109	1.225 ± 0.194
$\lvert\eta\rvert \geq 3.2$	1.192 ± 0.149	1.154 ± 0.152	1.082 ± 0.198

up/down (i.e., their normalization relative to the core is increased/decreased). Formally, to increase the tail size by a factor β, the modified template is given by

$$f_\beta(x) = N_\beta(\text{Core}(x) + \beta \cdot \text{Tail}(x)),$$

where N_β is a normalization factor to ensure that the template integral remains equal to one. The effect of varying the tail normalization on the QCD background estimate is studied in Sec. 8.7.

Chapter 4

Overview of the M_{T2} Analysis

Physics beyond the Standard Model at the LHC may manifest itself in a wide variety of potential final state topologies. Phase space covering a wide range of the relevant kinematic variables, such as hadronic energy, missing transverse momentum, and number of leptons, should be carefully searched for signs of new physics. In this dissertation, we focus on the "all-hadronic" final state, which encompasses events with no prompt leptons and large amounts of hadronic energy and missing transverse momentum. Section 4.1 discusses the theoretical motivation for an all-hadronic search, Sec. 4.2 describes the various backgrounds for such a search, and Sec. 4.3 introduces M_{T2}, a $p_{\mathrm{T}}^{\mathrm{miss}}$-based kinematic variable that is used to discriminate signal from background.

4.1 Motivation for an all-hadronic search

Many theories of BSM physics contain heavy particles (masses on the order of hundreds of GeV to TeV), whose decays produce a multitude of highly energetic particles. Thus, the events tend to have lots of visible energy (in the form of highly energetic jets and/or leptons) as well as *invisible* energy (in the form of $p_{\mathrm{T}}^{\mathrm{miss}}$), if the decay products are not detectable.

Most SM processes that produce significant true $p_{\mathrm{T}}^{\mathrm{miss}}$ from neutrinos also have leptons,

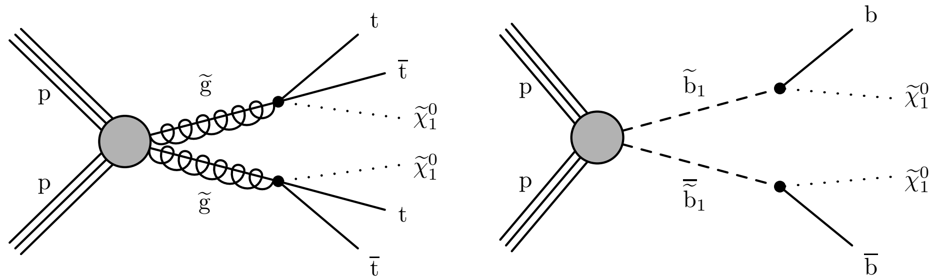

Figure 4.1: Example Feynman diagrams of strongly-produced SUSY particles decaying hadronically. (left) Gluino pair production, where each gluino decays into a neutralino and two top quarks. (right) Bottom squark pair production, where each squark decays into a neutralino and bottom quark.

since the neutrinos often come from $W \to \ell\nu$ decay. This is especially true for events with many jets and b-tagged jets, which often contain a $t\bar{t}$ pair. So by vetoing events with prompt leptons, it is possible to significantly reduce a large class of SM backgrounds. The flip side of this is that a zero-lepton selection will be contaminated by events from QCD multijet production, with significant fake $p_{\mathrm{T}}^{\mathrm{miss}}$ from jet mis-measurement. However, this can be managed with the use of a variety of discriminating kinematic variables, discussed in Sec. 4.3 and Chapter 5.

Based on general concerns, then, there is motivation for a search with lots of hadronic energy and missing transverse momentum, but no leptons. There are also many specific examples of models of BSM physics that produce such final states. Decay chains in simplified models of supersymmetry, discussed in Sec. 1.3.1, are often completely hadronic. Two examples are shown in Fig. 4.1. The left diagram shows the pair production of two gluinos, each decaying to a $t\bar{t}$ pair and an invisible $\tilde{\chi}_1^0$. Each top quark then has a 68% probability of decaying purely hadronically. The right diagram shows the pair production of bottom squarks, each of which decay to a bottom quark and $\tilde{\chi}_1^0$.

Additionally, the pair production of leptoquarks, discussed in Sec. 1.3.2, may produce purely hadronic final states if each leptoquark decays into a quark and neutrino (e.g., LQ $\rightarrow t\bar{\nu}_\tau$). In the case that the leptoquarks are scalars (spin 0), this process is indistinguishable from the pair production of squarks that decay to a quark and $\widetilde{\chi}^0_1$ when $m_{\widetilde{\chi}^0_1} = 0$, so any analysis optimized to search for hadronically-decaying squark pairs is naturally optimized to search for leptoquark pairs.

4.2 Sources of backgrounds

The M_{T2} analysis targets new physics signatures in all-hadronic events with large amounts of missing transverse momentum. This selection leads to backgrounds that fall into three main categories:

1. **Invisible Z**: $Z \rightarrow \nu\bar{\nu}$ events produced in association with jets contain genuine $p_{\mathrm{T}}^{\mathrm{miss}}$, and represent an *irreducible* background to the analysis. While there is no real handle to eliminate these events, as they look exactly like signal, the events contain no inherent b jets so the background becomes less important in signal regions with large numbers of b-tagged jets. This background is estimated using $Z \rightarrow \ell^+\ell^-$ events, as described in Chapter 6.

2. **Lost lepton**: events with leptonically-decaying W bosons contain genuine $p_{\mathrm{T}}^{\mathrm{miss}}$ from the neutrino, but are largely rejected based on the presence of the charged lepton. However, they can populate the signal regions if the lepton is in some way "lost" (usually, if it is outside of detector acceptance, or it is not isolated). The primary processes making up this background are $t\bar{t}$ and W+jets ($t\bar{t}$ is more important in regions with b-tagged jets), but there are also minor contributions from rarer processes such as single top, $t\bar{t}Z$, $t\bar{t}W$, $t\bar{t}H$, and $tt\bar{t}\bar{t}$. These are combined with $t\bar{t}$ and referred to as "top" in the plots and tables in this dissertation. The lost lepton background is primarily estimated with a single-lepton control sample,

as described in Chapter 7.

3. **QCD multijet**: the process with by far the largest cross section at the LHC is the QCD production of multijet events. These events have no genuine $p_{\mathrm{T}}^{\mathrm{miss}}$, but can enter high-$p_{\mathrm{T}}^{\mathrm{miss}}$ signal regions if one or more jets is mis-measured. There are a number of handles that can be used to reject these events, such as the $\Delta\phi$ between the jets and $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ vector, and the M_{T2} variable described in the next section. The remaining background is estimated through a data-driven procedure known as "Rebalance and Smear", detailed in Chapter 8.

4.3 The M_{T2} variable

Defining the *transverse energy* of a particle as $E_{\mathrm{T}}^2 \equiv m^2 + p_{\mathrm{T}}^2$, we can define the *transverse mass* of a two-particle system as

$$\begin{aligned} M_{\mathrm{T}}^2 &\equiv (E_{\mathrm{T},1} + E_{\mathrm{T},2})^2 - (\vec{p}_{\mathrm{T},1} + \vec{p}_{\mathrm{T},2})^2 \\ &= m_1^2 + m_2^2 + 2(E_{\mathrm{T},1}E_{\mathrm{T},2} - \vec{p}_{\mathrm{T},1} \cdot \vec{p}_{\mathrm{T},2}). \end{aligned} \tag{4.1}$$

Assuming massless particles (or, equivalently, particles that are highly relativistic such that $E \gg m$ and $E_{\mathrm{T}} \approx p_{\mathrm{T}}$), this simplifies to

$$M_{\mathrm{T}}^2 = 2p_{\mathrm{T},1}p_{\mathrm{T},2}(1 - \cos\theta), \tag{4.2}$$

where θ is the angle between the particle transverse momentum vectors.

This variable is frequently useful when a particle decays to something visible and something invisible (e.g. $W \to e\nu$). Assuming that the invisible particle is the dominant source of missing energy in the event, the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ vector is approximately the transverse momentum vector of the invisible particle and can be plugged into Eq. 4.2 to compute

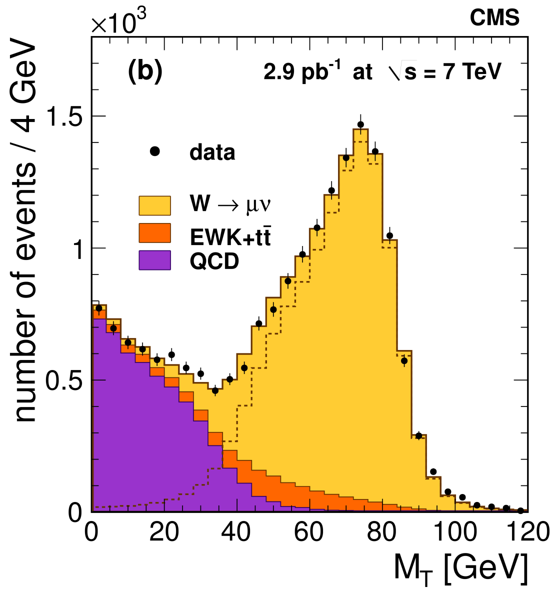

Figure 4.2: Transverse mass of the muon-$\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ system from a CMS measurement of the W production cross section [63]. The transverse mass has a rough upper bound at $M_W = 80$ GeV, limited by experimental resolution and the inherent decay width of the W of $\Gamma(W) = 2.1$ GeV.

the transverse mass of the system. As the transverse mass is just the invariant mass computed only with the transverse components of the particle momenta, it naturally has an upper bound equal to the parent particle mass. An example from a CMS measurement of the W production cross section [63] is shown in Fig. 4.2. In this case, the two-particle system is $W \to \mu\nu$ decay, and $M_{\mathrm{T}} \lesssim M_W = 80$ GeV. The location of the M_{T} "cliff" can be used to extract a measurement of the W mass. The events with transverse mass larger than the W mass are due to the decay width of the W of $\Gamma(W) = 2.1$ GeV, experimental resolution on the muon momentum and $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ vectors, and non-neutrino contributions to $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$.

Computing the transverse mass of a decaying SUSY particle would be useful in distinguishing signal from background, since they tend to be much heavier than SM particles that make up the background. However, the SUSY particles are *pair-produced*, and computing the transverse mass of each system is not possible because it is not possible to resolve the $\vec{p}_{\rm T}^{\rm miss}$ vector into components from each invisible particle. An example of this is the SUSY bottom squark pair production shown in Fig. 4.1 (right). Each bottom squark decays into a b quark, which appears as a jet in the detector (potentially b-tagged), and a neutralino, which is invisible.

Absent this information, the best we can do is to try *all possible partitions* of the $\vec{p}_{\rm T}^{\rm miss}$ vector and choose the one that gives the weakest $M_{\rm T}$ bound. This minimized quantity should be bounded above by the mass of the pair-produced parent particles. So, following [64], we define

$$M_{\rm T2} = \min_{\vec{p}_{\rm T}^{X(1)}+\vec{p}_{\rm T}^{X(2)}=\vec{p}_{\rm T}^{\rm miss}} \left[\max\left(M_{\rm T}^{(1)}, M_{\rm T}^{(2)}\right)\right], \tag{4.3}$$

where (from Eq. 4.1)

$$M_{\rm T}^{(i)2} = m_{{\rm vis}(i)}^2 + m_X^2 + 2\left(E_{\rm T}^{{\rm vis}(i)} E_{\rm T}^{X(i)} - \vec{p}_{\rm T}^{{\rm vis}(i)} \cdot \vec{p}_{\rm T}^{X(i)}\right) \quad (i = 1, 2), \tag{4.4}$$

vis(i) represents the $i^{\rm th}$ "visible" system, $E_{\rm T}^{X(i)}$ and $\vec{p}_{\rm T}^{X(i)}$ are the assigned transverse energy and momentum of each invisible particle, and m_X is the mass of the invisible particle. The maximum of $M_{\rm T}^{(1)}$ and $M_{\rm T}^{(2)}$ is used, since if the correct momenta are chosen then *both* should be bounded above by the parent mass.

A couple of complications arise when trying to practically implement this. First, there are typically more than two visible objects in an event (either the event is accompanied by ISR or pileup jets, or the decay cascade naturally produces more than one visible object. For example, gluino decay into top quarks, illustrated in Fig. 4.1 (left), produces multiple jets per decaying gluino). If the original pair-produced particles are produced back-to back, the visible systems tend to be in opposite hemispheres. So before computing $M_{\rm T2}$, all jets in the event are clustered into two *pseudojets* following the algorithm described in Ref. [49], Section 13.4, and these pseudojets are used as the two visible systems. First, two initial seed axes are chosen. In this analysis they are chosen by identifying the two jets that have the largest dijet invariant mass. Next, other jets are associated to one of these axes according to a clustering criterion. Here we use the minimal Lund distance, meaning that jet k is associated to hemisphere i rather than hemisphere j if

$$(E_i - p_i \cos\theta_{ik})\frac{E_i}{(E_i + E_k)^2} < (E_j - p_j \cos\theta_{jk})\frac{E_j}{(E_j + E_k)^2}, \tag{4.5}$$

where E_i and p_i are the energy and momentum of pseudojet i, and θ_{ik} is the angle between pseudojet i and jet k. After all jets are associated to one or the other axis, the axes are recalculated as the sum of the momenta of all jets connected to a hemisphere. The association is iterated using these new axes until no jets switch from one group to the other.

A second complication arises in choosing the mass terms in Eq. 4.4. Computing the masses of the two pseudojets ($m_{\mathrm{vis}(i)}$) is possible, but mis-measured multijet events with large pseudojet masses may give rise to large M_{T2}, eliminating some of the discriminating power of the variable. Setting these pseudojet masses to 0 is found to further suppress multijet events while maintaining signal sensitivity. The invisible particle mass, m_X, on the contrary is an unknown parameter and could not be set to its proper value even if desired. It is found that setting $m_X = 0$ is sufficient in maintaining discriminating power. Hence, in this analysis both mass terms in Eq. 4.4 are set to 0 when computing M_{T2}.

As discussed in Sec. 4.2, a big challenge in hadronic searches is the huge cross section for QCD production of multijet events at the LHC. While requiring large $p_{\mathrm{T}}^{\mathrm{miss}}$ suppresses much of this, there is still a significant contribution from events with $p_{\mathrm{T}}^{\mathrm{miss}}$ from mis-measured jets. The M_{T2} variable allows further suppression of these events, by taking advantage of the different event topologies between signal and mis-measured multijet events.

This is illustrated in Figs. 4.3 and 4.4. Fig. 4.3 shows distributions of M_{T2} for signal and background after an inclusive selection, while Fig. 4.4 shows 2D distributions of M_{T2} vs. $p_{\mathrm{T}}^{\mathrm{miss}}$ separately for signal and QCD multijet events. In SUSY events, the large mass scale of the produced sparticles and acoplanarity of the visible objects tends to concentrate the events in the high-M_{T2} region (M_{T2} follows $p_{\mathrm{T}}^{\mathrm{miss}}$ quite closely). On the contrary, mis-measured multijet events populate the low-M_{T2} region regardless of $p_{\mathrm{T}}^{\mathrm{miss}}$ or jet p_{T}, since they tend to be “back-to-back”. Electroweak backgrounds have larger M_{T2} tails than does QCD, since they contain real missing energy from neutrinos. Because of this, after requiring high M_{T2}, QCD multijet events become a sub-dominant background in all signal regions.

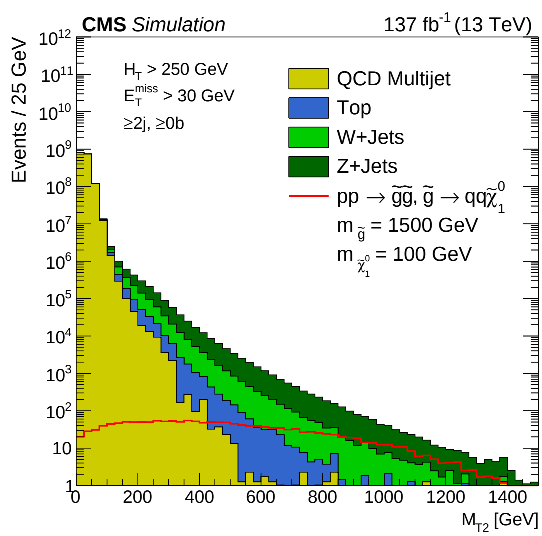

Figure 4.3: M_{T2} distributions for an example signal and backgrounds after an inclusive selection of $H_{\mathrm{T}} > 250$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV, and $N_{\mathrm{j}} \geq 2$. The signal shown is gluino pair production where each gluino decays to two light-flavor quarks and a neutralino, with $m_{\widetilde{g}} = 1500$ GeV and $m_{\widetilde{\chi}_1^0} = 100$ GeV. The QCD multijet background falls rapidly with M_{T2}, while signal extends out to large values. The electroweak backgrounds have larger tails, since they contain true $p_{\mathrm{T}}^{\mathrm{miss}}$, but some discriminating power remains.

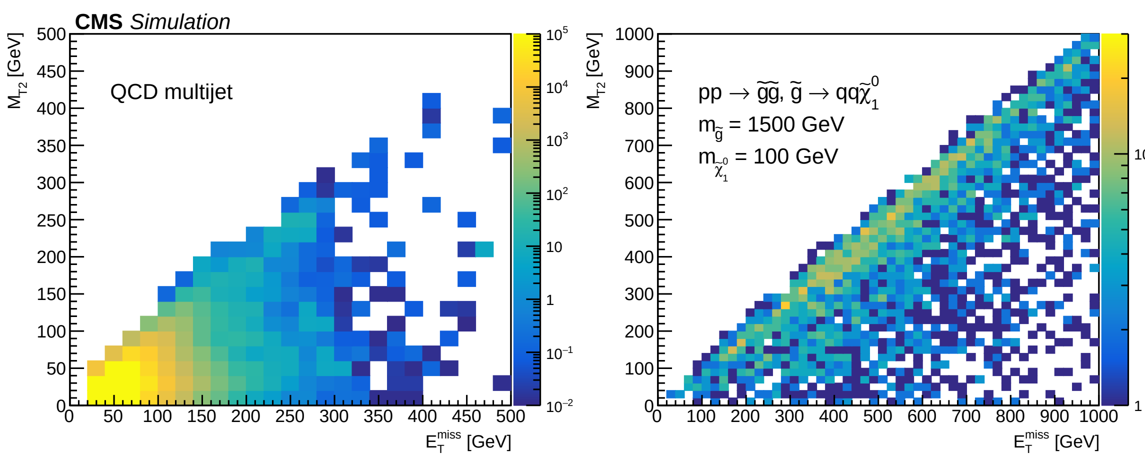

Figure 4.4: Distributions of M_{T2} vs. $p_{\mathrm{T}}^{\mathrm{miss}}$, for QCD multijet events (left), and gluino pair production (right). For signal, $M_{\mathrm{T2}} \sim p_{\mathrm{T}}^{\mathrm{miss}}$, while for background, M_{T2} tends to be much smaller than $p_{\mathrm{T}}^{\mathrm{miss}}$.

Chapter 5

Event Selection and Triggering

The general strategy of the M_{T2} analysis is to apply a baseline selection motivated by the available triggers and by the desire to reduce QCD multijet background to manageable levels, and then to categorize the selected events into bins of differing levels of hadronic activity, number of b-tagged jets, and missing transverse energy. The "namesake" variable M_{T2}, used to reduce QCD background, was described in Sec. 4.3, but a number of other variables are also used to constrain backgrounds and categorize events. Sec. 5.1 defines the relevant physics objects (jets, leptons, etc.) and kinematic variables, Sec. 5.2 describes the triggers used in the analysis, Sec. 5.3 outlines the baseline selection, and Sec. 5.4 gives the precise definitions of the signal regions.

5.1 Object and variable definitions

In CMS, individual particles are identified by combining information from the tracker, calorimeters, and muon system using the "particle-flow" algorithm, described in Sec. 3.1.1. This particle-level information is then used to cluster jets, reconstruct vertices, and compute the missing transverse energy. Here we list the various physics objects used in this analysis and the selections applied on them.

5.1.1 Vertices

"Vertex finding" is a process of finding points in space from which groups of reconstructed particle tracks that, loosely, come from the same "interaction", emanate. There is generally a single "primary vertex", where the hard interaction took place, pileup vertices from pileup interactions, and, potentially, secondary vertices from longer-lived decaying particles such as b hadrons. Algorithms for reconstructing vertices are described in [65]. For this analysis, we consider a reconstructed vertex as good if it satisfies:

- not "fake" - if no vertices are reconstructed from tracks (rare), CMS software assigns a default vertex based on the beam-spot (luminous region produced by proton beam collision), and labels it as "fake". We do not consider such vertices.
- $N_{\mathrm{dof}} > 4$ - N_{dof} is the number of degrees of freedom in the fit of the position of the vertex, strongly correlated with the number of tracks consistent with originating from the vertex.
- $|z| < 25$ cm - the longitudinal distance from the beam-spot.
- $\rho < 2$ cm - the distance from the beam axis.

When more than one good reconstructed vertex is found in the event, the reconstructed vertex with the largest value of summed physics-object p_{T}^2 is taken to be the primary interaction vertex.

5.1.2 Jets

Charged hadrons from pileup interactions are identified and removed based on the "charged hadron subtraction" algorithm [66]. Jets are then clustered using the anti-k_{T} algorithm with distance parameter $R = 0.4$. Jet energies are corrected for pileup contamination and detector response with CMS-derived era-dependent jet energy corrections.

We select jets that satisfy $p_{\mathrm{T}} > 30$ GeV and $|\eta| < 2.4$, and pass PF jet loose ID (2016) or tight ID (2017+18). For events with only one jet, we require tighter ID requirements to reject noisy jets. All jet ID cuts are summarized in Table 5.1.

Table 5.1: Jet ID definitions

	2016	2017-18	Monojet (all years)
Neutral hadron fraction	< 0.99	< 0.90	< 0.80
Neutral EM fraction	< 0.99	< 0.90	< 0.70
Number of constituents	> 1	> 1	> 1
Charged hadron fraction	> 0	> 0	> 0.05
Charged multiplicity	> 0	> 0	> 0
Charged EM fraction	< 0.99	–	< 0.99

We define N_{j} as the number of jets passing the above selections, and H_{T} as the scalar sum of p_{T} values for all such jets.

Jets originating from b quarks are tagged using the DeepCSV algorithm [59], at the medium working point. For the purposes of counting the number of b tags (N_{b}), we loosen the p_{T} threshold for b-tagged jets to 20 GeV. This helps to add sensitivity to compressed-spectrum signals with jets from b quarks.

5.1.3 $p_{\mathrm{T}}^{\mathrm{miss}}$

We use type 1-corrected PFMET, defined in Sec. 3.2, using the same jet energy corrections as applied to the jets. We additionally define $\vec{H}_{\mathrm{T}}^{\mathrm{miss}}$ as

$$\vec{H}_{\mathrm{T}}^{\mathrm{miss}} = -\sum_{\mathrm{jets}} \vec{p}_{\mathrm{T},i}, \tag{5.1}$$

where the sum is taken over all jets passing the above requirements. The difference with respect to $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ is that this excludes forward or low-p_{T} jets and unclustered energy.

5.1.4 $p_{\mathrm{T}}^{\mathrm{miss}}$ filters

In addition to real missing energy due to invisible particles, events may have some amount of "fake $p_{\mathrm{T}}^{\mathrm{miss}}$" due to either detector effects or external sources (e.g. cosmic rays or beam-halo particles). We have discussed fake $p_{\mathrm{T}}^{\mathrm{miss}}$ from standard jet mis-measurement due to stochastic smearing in the calorimeters, but more pathological effects are also possible, such as noisy calorimeter cells or bad track reconstructions. To eliminate as best as possible events containing such sources of fake $p_{\mathrm{T}}^{\mathrm{miss}}$, the JetMET group at CMS recommends a set of "$p_{\mathrm{T}}^{\mathrm{miss}}$ filters" that use features of the reconstructed event to identify certain classes of bad events. We apply all standard recommended filters, listed here:

- primary vertex filter
- CSC super-tight beam halo 2016 filter (despite name, used in all 3 years)
- HBHE noise filter
- HBHE iso noise filter
- EE badSC noise filter
- ECAL dead cell trigger primitive filter
- bad muon filter
- ECAL bad calibration filter (2017+18 only)

There are also a few custom $p_{\mathrm{T}}^{\mathrm{miss}}$ filters developed by this analysis or the CMS SUSY group that are applied to protect against other observed sources of fake $p_{\mathrm{T}}^{\mathrm{miss}}$. First, we reject any event containing a jet with $p_{\mathrm{T}} > 30$ GeV and $|\eta| < 4.7$ which fails the PF jet loose/tight ID as described above. Since this jet would not enter the collection used to compute M_{T2}, the pseudojets would likely be imbalanced and the resulting M_{T2} biased. This is not applied to fast simulation MC samples since the input variables are not correctly modeled, and M_{T2} is expected to already be high anyway.

Next, we require that the ratio of PFMET over caloMET ($p_{\mathrm{T}}^{\mathrm{miss}}$ computed only with calorimeter deposits) is less than 5. This was found to remove events with a bad high-p_{T} muon track inside a jet, which are not removed by either the lepton vetoes or the bad muon track filter. Also to reduce the effect of mis-measured muons, we veto events that contain a jet with $p_{\mathrm{T}} > 200$ GeV and a muon fraction larger than 50%, and that satisfies $|\Delta\phi(\mathrm{jet}, p_{\mathrm{T}}^{\mathrm{miss}})| > \pi - 0.4$.

To remove certain known pathological events in fast simulation MC, we remove events in such MC containing a jet satisfying $p_{\mathrm{T}} > 20$ GeV, $|\eta| < 2.5$, charged hadron fraction < 0.1, and no matching generator-level jet within $\Delta R < 0.3$.

Finally, an issue with two HCAL endcap modules during 2018 data taking required the addition of a special filter to reject events containing jets or electrons in the affected region. Details are given in Sec. 5.6.

5.1.5 Electrons

While the analysis signal regions are purely hadronic, it is still necessary to define lepton candidates, both in order to define a lepton veto and to select events for the leptonic control regions. Electron candidates are required to satisfy $p_{\mathrm{T}} > 10$ GeV and $|\eta| < 2.4$. Good electrons are identified using cut-based ID working points developed by the EGamma group at CMS: the "veto" working point is used for the signal region veto and the single lepton control region, and the "loose" working point is used for the dileptonic $Z \to \ell^+\ell^-$ control region. The cuts are summarized in Table 5.2, and definitions for the various variables are listed here.

- $\sigma_{i\eta i\eta}$ - a variable describing the width of the shower in the ECAL; computed using the reconstructed hits in a 5x5 seed cluster
- $|\Delta\eta_{\mathrm{Seed}}|$ - difference in η between a ECAL cluster position and track direction at

vertex extrapolated to ECAL

- $|\Delta\phi_{In}|$ - difference in ϕ between a ECAL cluster position and track direction at vertex extrapolated to ECAL
- H/E - ratio of energy in HCAL behind ECAL cluster to the energy in the ECAL cluster
- $|1/E - 1/p|$ - tests consistency of ECAL cluster energy and track momentum
- $|d_{\mathrm{xy}}|$ - transverse impact parameter of the track with respect to the primary vertex
- $|d_{\mathrm{z}}|$ - longitudinal impact parameter of the track with respect to the primary vertex
- conversion veto - reject electron candidates that look like photon conversions to e^+e^- pairs

Table 5.2: Cut-based electron ID for the Veto and Loose ID working points, for electrons in the barrel (endcap).

	Veto ID	Loose ID
$\sigma_{i\eta i\eta}$ (RecHits in 5x5 seed cluster)	< 0.0126 (0.0457)	< 0.0112 (0.0425)
$\lvert\Delta\eta_{\mathrm{Seed}}\rvert$	< 0.00463 (0.00814)	< 0.00377 (0.00674)
$\lvert\Delta\phi_{In}\rvert$	< 0.148 (0.19)	< 0.0884 (0.169)
H/E	$< 0.05 + 1.16/E_{\mathrm{SC}} + 0.0324\ \rho/E_{\mathrm{SC}}$ $(< 0.05 + 2.54/E_{\mathrm{SC}} + 0.183\ \rho/E_{\mathrm{SC}})$	$< 0.05 + 1.16/E_{\mathrm{SC}} + 0.0324\ \rho/E_{\mathrm{SC}}$ $(0.0441 + 2.54/E_{\mathrm{SC}} + 0.183\ \rho/E_{\mathrm{SC}})$
$\lvert\frac{1}{E} - \frac{1}{p}\rvert$	< 0.209 (0.132)	< 0.193 (0.169)
$\lvert d_{\mathrm{xy}}\rvert$ (w.r.t. primary vertex)	< 0.2 (0.2) cm	< 0.2 (0.2) cm
$\lvert d_{\mathrm{z}}\rvert$(w.r.t. primary vertex)	< 0.5 (0.5) cm	< 0.5 (0.5) cm
# of expected missing inner hits	≤ 2 (3)	≤ 1 (1)
conversion veto	yes	yes

In addition to the ID described above, electrons are required to be isolated. This is defined using relative mini-PF isolation, as miniPFIso/$p_{\mathrm{T}} < 0.1$. "PF isolation" is just the sum of the p_{T} values of all particle-flow candidates within a cone around the electron candidate, and the "mini" part means that this cone size gets smaller with higher electron

p_{T}. Precisely, the cone size used is

$$\Delta R = \begin{cases} 0.2 & \text{if } p_{\mathrm{T}} < 50 \text{ GeV} \\ 10 \text{ GeV}/p_{\mathrm{T}} & \text{if } 50 < p_{\mathrm{T}} < 200 \text{ GeV} \\ 0.05 & \text{if } p_{\mathrm{T}} > 200 \text{ GeV} \end{cases} \tag{5.2}$$

A correction to the isolation to account for pileup contamination is applied, based on the event-level energy density and the effective area of the electron cone.

5.1.6 Muons

Muon candidates are required to pass $p_{\mathrm{T}} > 10$ GeV and $|\eta| < 2.4$, and a loose ID selection defined as:

- matched to a particle-flow muon
- either a global muon (tracker+muon system) or a tracker-only muon
- $|d_{\mathrm{xy}}| < 0.2$ cm (transverse impact parameter with respect to the primary vertex)
- $|d_{\mathrm{z}}| < 0.5$ cm (longitudinal impact parameter with respect to the primary vertex)

We require the muons to be isolated using the same relative mini PF isolation as used for the electrons, this time requiring miniPFIso/$p_{\mathrm{T}} < 0.2$.

5.1.7 Isolated tracks

In addition to vetoing events with the reconstructed leptons described above, we further add a veto for events with “isolated tracks” that were not reconstructed as leptons, either because they are charged hadrons or they failed some criteria to be promoted to a full reconstructed lepton. This allows better rejection of backgrounds with hadronically decaying τ leptons (these frequently produce isolated pions) or with isolated leptons that were not caught be the lepton veto, without appreciably affecting signal efficiency. We select charged particle-flow candidates with different requirements depending on the type of candidate.

For particle-flow electrons and muons, we require them to pass $p_{\mathrm{T}} > 5$ GeV, $|\eta| < 2.4$, $|d_{\mathrm{z}}| < 0.1$ cm, $|d_{\mathrm{xy}}| < 0.2$ cm, and a track isolation cut of $\mathrm{iso}/p_{\mathrm{T}} < 0.2$. The track isolation is computed as the sum of all charged hadron particle-flow candidates with in a cone of $\Delta R < 0.3$, and that satisfy $|d_{\mathrm{z}}| < 0.1$ cm with respect to the primary vertex. For lepton counting, particle-flow leptons within $\Delta R < 0.01$ of selected reconstructed leptons are removed.

Charged particle-flow hadrons are required to pass $p_{\mathrm{T}} > 10$ GeV, $|\eta| < 2.4$, $|d_{\mathrm{z}}| < 0.1$ cm, $|d_{\mathrm{xy}}| < 0.2$ cm, and a track isolation cut of $\mathrm{iso}/p_{\mathrm{T}} < 0.1$, computed in the same way as above.

5.1.8 $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$

The variable $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ (referred to as $\Delta\phi_{\mathrm{min}}$ in the following) is defined as the minimum $\Delta\phi$ between $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ and any of the four highest p_{T} jets in the event. For this variable only, we consider jets with $p_{\mathrm{T}} > 30$ GeV and $|\eta| < 4.7$.

5.2 Triggers

5.2.1 Signal region triggers

The baseline H_{T} and $p_{\mathrm{T}}^{\mathrm{miss}}$ cuts used for the signal regions are constrained by the available triggers utilized in the data-taking, which can be different from year to year. The signal regions use an OR of various H_{T}- and $p_{\mathrm{T}}^{\mathrm{miss}}$-based triggers, that cover different (but overlapping) regions of phase space.

The exact triggers used are an OR of the following:

2016 data

- `HLT_PFHT900` - pure-H_T trigger; nominally turns on at 900 GeV but observed plateau near 1000 GeV.
- `HLT_PFJet450` - triggers on any jet with $p_T > 450$ GeV. Not strictly necessary, but used anyway to cover for any potential small inefficiency.
- `HLT_PFHT300_PFMET100` - triggers on a combination of $H_T > 300$ GeV and $p_T^{miss} >$ 100 GeV. Again not strictly necessary, but covers for any potential inefficiency of the p_T miss triggers in the intermediate H_T regions.
- `HLT_PFMET[NoMu]120_PFMHT[NoMu]120_IDTight` - a pure p_T^{miss} trigger, where the p_T^{miss} is computed either with or without muons. Observed plateau around $p_T^{miss} =$ 250 GeV.

2017+18 data (similar triggers to 2016, but with raised thresholds)

- `HLT_PFHT1050`
- `HLT_PFJet500`
- `HLT_PFHT800_PFMET75_PFMHT75` OR `HLT_PFHT500_PFMET100_PFMHT100`
- `HLT_PFMET[NoMu]120_PFMHT[NoMu]120_PFHT60_IDTight`

An illustration of the trigger coverage in the baseline signal region (defined in Sec. 5.3) is shown in Fig. 5.1. At low H_T, the only available trigger is the pure-p_T^{miss} trigger, which necessitates the p_T^{miss} cut of 250 GeV. At higher H_T (>1200 GeV), the pure-H_T trigger can be used and allows for the relaxing of the p_T^{miss} cut to 30 GeV. The H_T +p_T^{miss} triggers (and jet p_T trigger) provide redundancy in the intermediate to high H_T regions.

Measuring the efficiency of the triggers is important to ensure the signal regions are fully covered and trigger at near 100%. To make this measurement, a selection in the

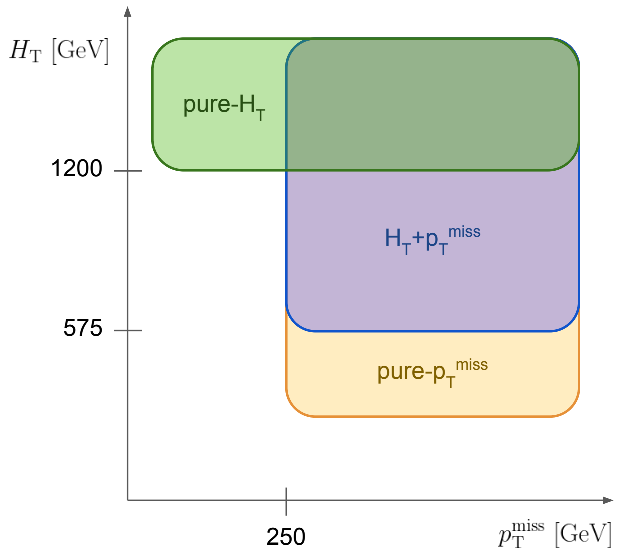

Figure 5.1: Illustration of the trigger coverage of the baseline signal region (defined in Sec. 5.3). The pure-p_T^{miss} triggers cover the low-H_T regions, starting at p_T^{miss} = 250 GeV. The pure-H_T trigger covers the high-H_T regions, starting at H_T = 1200 GeV. The H_T +p_T^{miss} triggers provide redundancy and cover for any small inefficiencies in the intermediate H_T regions.

plateau of an orthogonal "reference trigger" is made, and then the efficiency of the desired trigger is plotted as a function of the relevant kinematic variable.

The efficiency of the pure-H_T triggers is measured in two ways: one using an electron trigger reference, and one using a p_T^{miss} trigger reference. Both methods start by selecting events that pass all p_T^{miss} filters, and that have at least 2 jets with $p_T > 30$ GeV and no leptons (other than the potential reference electron). The electron method further requires that the event passes a single electron trigger, and contains an electron with $p_T > 35$ GeV that is isolated and passes a tight cut-based ID. The p_T^{miss} method requires that the event passes the pure-p_T^{miss} trigger and has $p_T^{miss} > 300$ GeV. Both methods give consistent results, with plateau efficiencies of >98%. After combining with the `HLT_PFJetXXX` trigger, the efficiency is >99% in all years. Example measurements are shown in Fig. 5.2.

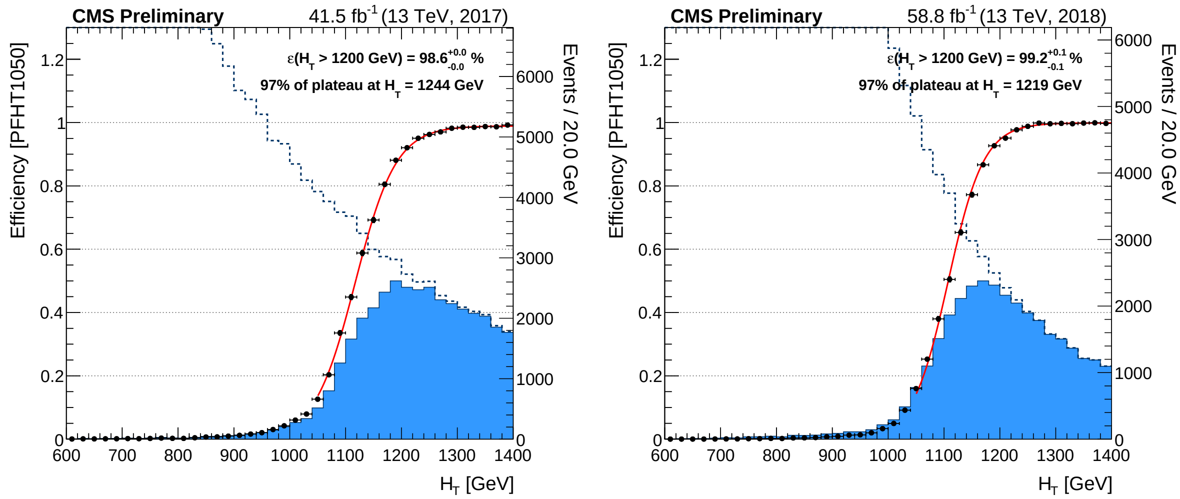

Figure 5.2: Trigger efficiency measurements for the pure-H_T `HLT_PFHT1050` trigger, (left) for 2017 data using a $p_\mathrm{T}^\mathrm{miss}$-based reference trigger, and (right) for 2018 data using an electron-based reference trigger. The dashed (solid) blue histograms are the denominator (numerator) event counts, and the red lines are logistic function fits to the efficiency curves.

The pure-$p_\mathrm{T}^\mathrm{miss}$ trigger efficiencies are measured using the electron-based method described above. The H_T and $p_\mathrm{T}^\mathrm{miss}$ legs of the H_T $+p_\mathrm{T}^\mathrm{miss}$ triggers must be measured separately. The H_T leg is measured using a $p_\mathrm{T}^\mathrm{miss}$ trigger reference, and the $p_\mathrm{T}^\mathrm{miss}$ leg is measured using an electron trigger reference (with the additional requirement of $H_\mathrm{T} > 700$ GeV to ensure that we are in the plateau of the H_T leg). Efficiencies of close to 100% are observed in all cases. Example measurements of the $p_\mathrm{T}^\mathrm{miss}$ and H_T $+p_\mathrm{T}^\mathrm{miss}$ trigger efficiencies are shown in Fig. 5.3.

5.2.2 Control region triggers

Control region definitions are given in Sec. 5.5. Here we only summarize the triggers used, and it is sufficient to know that there are both single lepton and dilepton control regions used for predicting the electroweak backgrounds, and an inclusive low-H_T data sample used in the estimate of the QCD multijet background.

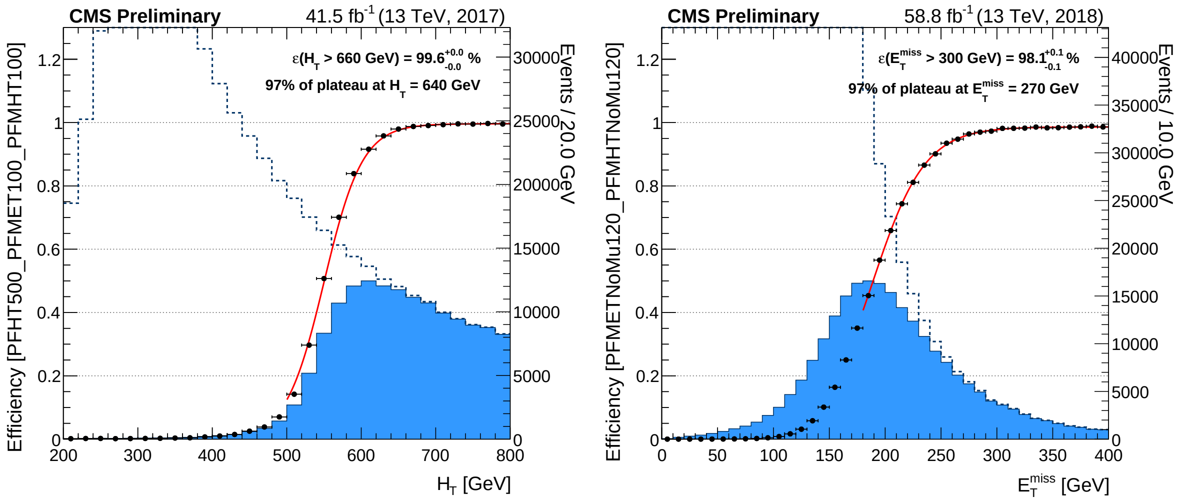

Figure 5.3: Trigger efficiency measurements (left) for the H_{T}-leg of the HLT_PFHT500_PFMET100 trigger in 2017 data, measured with a $p_{\mathrm{T}}^{\mathrm{miss}}$-based reference trigger, and (right) for the pure-$p_{\mathrm{T}}^{\mathrm{miss}}$ HLT_PFMETNoMu120_PFMHTNoMu120 trigger in 2018 data, measured using an electron-based reference trigger. The dashed (solid) blue histograms are the denominator (numerator) event counts, and the red lines are logistic function fits to the efficiency curves.

The single lepton control region is triggered with the same triggers as the signal region, since the cuts on the relevant kinematic variables are the same (the only difference is the requirement of exactly one lepton, instead of zero).

This is not true for the dilepton control region, since as we will see the leptons are removed from the H_{T} and added to the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ vector before cutting on these variables. So the standard H_{T} and $p_{\mathrm{T}}^{\mathrm{miss}}$ triggers will not work, and special dilepton triggers are used instead.

The dimuon selection is triggered with a combination of isolated and non-isolated dimuon triggers and non-isolated single-muon triggers (non-isolated and single-muon paths are to recover inefficiencies at high muon p_{T}).

Similarly, the dielectron selection is triggered with a combination of isolated and non-isolated dielectron triggers as well as a single-photon trigger.

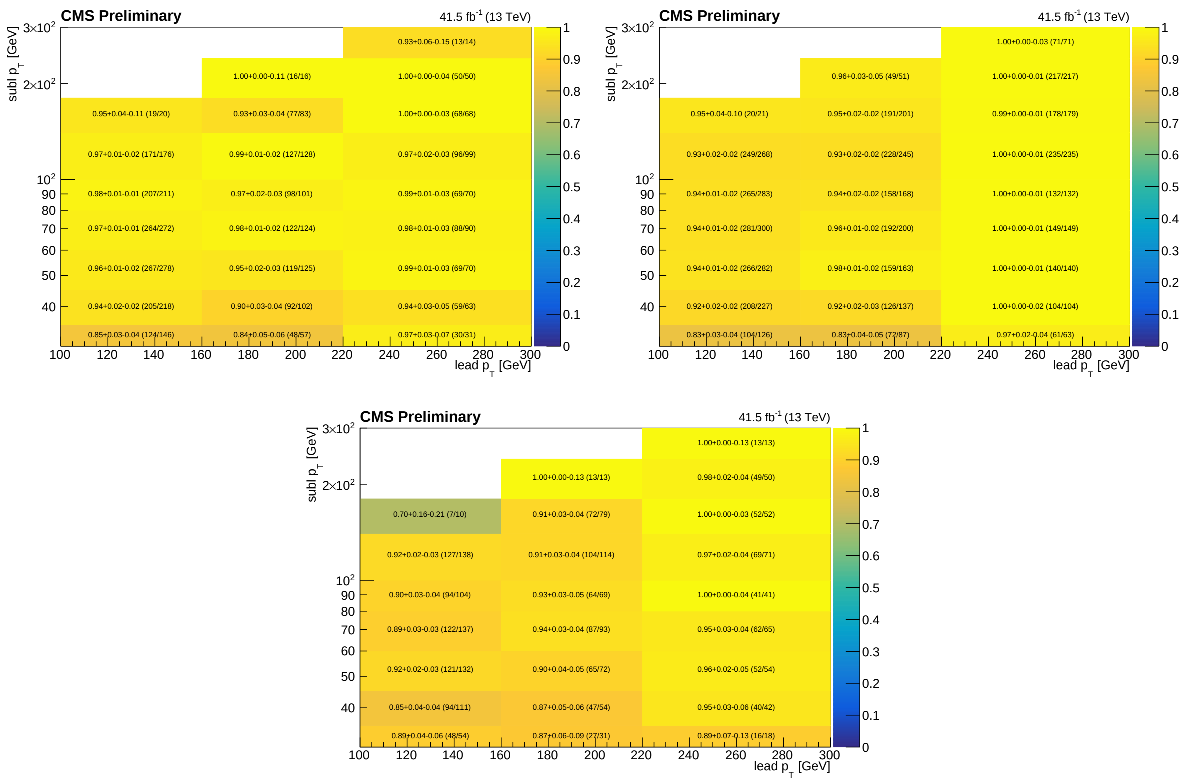

Figure 5.4: Measured dilepton trigger efficiencies as a function of leading and subleading lepton p_{T} in 2017 data, for (top left) dimuon, (top right) dielectron, and (bottom) $e\mu$ selections. Triggers used are a logical OR of the various triggers described in Sec. 5.2.2. Measurements in 2016 and 2018 are similar.

Finally, there is an different-flavor (i.e. $e\mu$) control region used for estimating $t\bar{t}$ contamination in the same-flavor region, which is triggered with a combination of isolated and non-isolated $e\mu$ triggers, non-isolated single-muon triggers, and a single-photon trigger.

Efficiencies of these dilepton can be measured using pure-H_{T} triggers as a reference (both non-prescaled and lower threshold prescaled triggers). Plots of measured trigger efficiency as a function of leading and subleading lepton p_{T} in the relevant section of phase space are shown in Fig. 5.4, for 2017 data. Measurements for 2016 and 2018 are similar. Inefficiencies up to $\sim$10% are present; these are accounted for by applying them as weights to MC on a year-by-year basis as a function of leading and subleading p_{T}.

Though not technically a control region, an inclusive selection of low-H_{T} events is used in the Rebalance and Smear method for computing QCD multijet background, described in Chapter 8. This selection makes use of low-threshold pure-H_{T} triggers, that are necessarily prescaled due to their very high rates. Though event-by-event prescale values are stored in the data (the exact prescale rate depends on the instantaneous luminosity and the specific trigger menu being used, so vary with time), yearly "effective prescales" are measured as a consistency check. This is done by measuring the ratio of rates of different triggers as a function of H_{T}, and fitting the ratio in the plateau. Comparing to the high threshold unprescaled trigger, the effective prescales of all others can be determined.

A plot measuring the effective prescales in 2017 data is shown in Fig. 5.5 (left). On the right is the observed H_{T} spectrum using these prescale values, which is smooth as is expected if the correct prescales are measured.

5.3 Baseline selection

Using the object and variable definitions detailed in Sec. 5.1, the baseline selection used for all analysis signal regions is as follows:

- at least one good vertex, as defined in Sec. 5.1.1
- pass all $p_{\mathrm{T}}^{\mathrm{miss}}$ filters, as defined in Sec. 5.1.4
- pass the logical OR of all signal region triggers listed in Sec. 5.2.1
- $H_{\mathrm{T}} > 1200$ GeV and $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV, or $H_{\mathrm{T}} > 250$ GeV and $p_{\mathrm{T}}^{\mathrm{miss}} > 250$ GeV. These cuts are based on the available trigger thresholds, and illustrated in Fig. 5.1. For $N_{\mathrm{j}} = 1$ regions, $p_{\mathrm{T}}^{\mathrm{miss}} > 250$ GeV is always required.

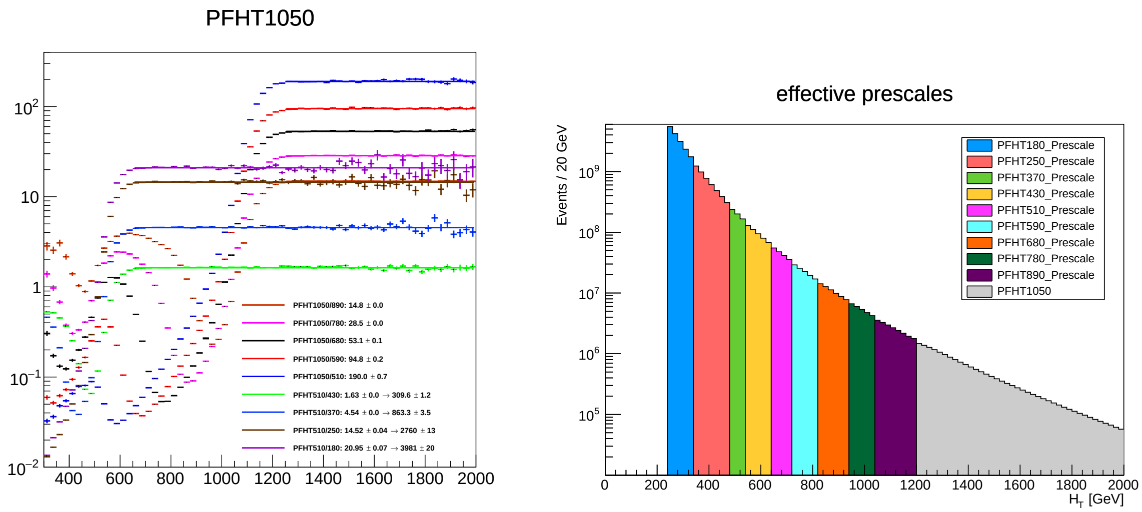

Figure 5.5: (left) A plot of the ratio of rates of various pure-H_T triggers in 2017 data. The value at the plateau gives the relative prescale between two triggers, and comparing to the unprescaled `HLT_PFHT1050` gives the absolute prescale. (right) The observed H_T spectrum using these prescales, which is smooth as is expected if the correct values are measured.

- $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}}) > 0.3$. This protects against large $p_{\mathrm{T}}^{\mathrm{miss}}$ from jet mis-measurement in multijet events, and rejects a large fraction of QCD background.
- $|\vec{H}_{\mathrm{T}}^{\mathrm{miss}} - \vec{p}_{\mathrm{T}}^{\mathrm{miss}}|/p_{\mathrm{T}}^{\mathrm{miss}} < 0.5$. This requires that $\vec{H}_{\mathrm{T}}^{\mathrm{miss}}$ be similar to $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$, and in doing so protects against bias in the shape of M_{T2}. $p_{\mathrm{T}}^{\mathrm{miss}}$ is sensitive to reconstructed objects with $p_{\mathrm{T}} < 30$ GeV or $|\eta| > 2.4$, whereas these are not used in $\vec{H}_{\mathrm{T}}^{\mathrm{miss}}$ or in the construction of pseudojets for M_{T2}, so a large contribution from these objects to the missing transverse energy can bias the M_{T2} distribution.
- lepton veto: to reduce the background from events with a W boson decay, we reject events if they contain
  - a reconstructed electron or muon as defined in Secs. 5.1.5 and 5.1.6
  - a particle-flow electron, muon, or hadron candidate as defined in Sec. 5.1.7, with the additional requirement of $M_{\mathrm{T}}(\mathrm{cand}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}}) < 100$ GeV (to avoid vetoing signal events, which contain $p_{\mathrm{T}}^{\mathrm{miss}}$ not from the leptonic W decay and

hence have high M_{T}).

- $M_{\mathrm{T2}} > 200$ GeV (only for events with at least 2 jets). This provides a large rejection of QCD multijet events that have large $p_{\mathrm{T}}^{\mathrm{miss}}$ from mis-measured jets, as discussed in Sec. 4.3. The threshold is raised to $M_{\mathrm{T2}} > 400$ GeV for events with $H_{\mathrm{T}} > 1500$ GeV, in order to ensure that QCD multijet events remain a sub-dominant background.

5.4 Signal region definitions

After the baseline selection defined in Sec. 5.3, signal regions are further categorized into different bins of H_{T}, N_{j}, N_{b}, and M_{T2}.

First, we categorize $N_{\mathrm{j}} \geq 2$ (multijet) events by H_{T}, N_{j}, and N_{b}. Each $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ bin is referred to as a *topological region*.

- Five H_{T} regions: $[250, 450]$, $[450, 575]$, $[575, 1200]$, $[1200, 1500]$, $[1500, \infty]$ GeV. These bins are referred to here and throughout this dissertation as Very Low, Low, Medium, High, and Extreme H_{T}, respectively.
- For the first three H_{T} regions, we use 11 bins in N_{j} and N_{b}:
  - 2–3 jets; 0, 1, 2 b tags
  - 4–6 jets; 0, 1, 2 b tags
  - 2–6 jets; ≥3 b tags
  - ≥7 jets; 0, 1, 2, ≥3 b tags
- For the highest two H_{T} regions, we further subdivide the $N_{\mathrm{j}} \geq 7$ regions for a total of 17 bins in N_{j} and N_{b}:
  - 2–3 jets; 0, 1, 2 b tags
  - 4–6 jets; 0, 1, 2 b tags
  - 2–6 jets; ≥3 b tags

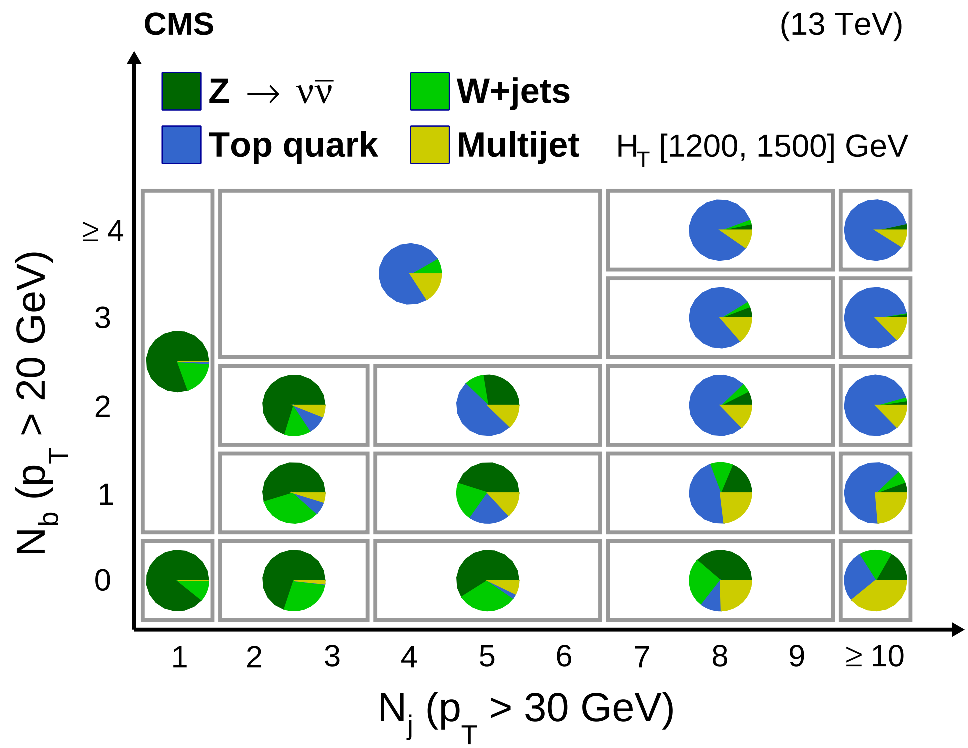

Figure 5.6: Topological region binning (defined by the solid gray lines) for the High H_{T} region ($1200 \leq H_{\mathrm{T}} < 1500$ GeV). The proportion of each background type is shown as a pie chart in each topological region. QCD multijet becomes more important with increasing N_{j}, and top backgrounds become more dominant at high N_{b}. Background composition is estimated using the data-driven techniques described in the following chapters.

- 7–9 jets; 0, 1, 2, 3, ≥4 b tags
- ≥10 jets; 0, 1, 2, 3, ≥4 b tags

The division into topological regions for the High H_{T} region is shown in Fig.5.6, along with background composition in each region. We see that QCD multijet background becomes more important with increasing N_{j}, and top background becomes more dominant with increasing N_{b}.

Finally, we further divide each topological region into bins of M_{T2}. The bin thresholds are chosen according to the following criteria:

- The lower edge of the first M_{T2} bin is 200 GeV, except for $H_{\mathrm{T}} > 1500$ GeV (Extreme H_{T}) where it is 400 GeV.
- In each topological region, we select the lower threshold of the last M_{T2} bin such that this bin is expected, from simulation, to contain approximately one background event. Moreover, the upper limit on H_{T} effectively places an upper limit on M_{T2}. Therefore, this lower M_{T2} threshold should not be larger than the upper limit on H_{T}, in each H_{T} region.
- Bin widths are nominally chosen to be 100 GeV. In each topological region, we merge M_{T2} bins which are expected to contain less than one background event. In a few cases, we merge intermediate M_{T2} bins to minimize the number of signal region bins.

The final M_{T2} binning for the multijet search regions can be found in Tables A.2 through A.12 in Appendix A.

In addition to the $N_{\mathrm{j}} \geq 2$ bins described above, we also select events with $N_{\mathrm{j}} = 1$ (monojet), with the jet required to have $p_{\mathrm{T}}(\mathrm{jet}) > 250$ GeV. Events are then categorized as follows:

- Two N_{b} regions: $N_{\mathrm{b}} = 0$ and $N_{\mathrm{b}} \geq 1$ (since the p_{T} threshold for N_{b} is lower than that for N_{j}, it is possible that $N_{\mathrm{b}} > 1$ even when $N_{\mathrm{j}} = 1$)
- For $N_{\mathrm{b}} = 0$, H_{T} bin edges of $[250, 350, 450, 575, 700, 1000, 1200, \infty]$ GeV
- For $N_{\mathrm{b}} \geq 1$, H_{T} bin edges of $[250, 350, 450, 575, 700, \infty]$ GeV

In total, there are 270 multijet and 12 monojet regions, for a grand total of 282 signal region search bins.

5.5 Control regions

To estimate the various backgrounds, we make use of a number of control regions orthogonal to the signal region. These include a single lepton selection, a dileptonic selection (mostly $Z \to \ell^+\ell^-$ events), and a selection enriched in QCD multijet events.

5.5.1 Single lepton control region

A single lepton selection is used to estimate backgrounds from $t\bar{t}$ and W+jets ("lost lepton" background). The selection is the same as the baseline selection described in Sec. 5.3, with the exception of the lepton veto. Instead, we require exactly one candidate passing the reconstructed lepton or particle-flow lepton selections (e or μ only). We further require this lepton candidate to satisfy $M_{\mathrm{T}}(\mathrm{cand}, p_{\mathrm{T}}^{\mathrm{miss}}) < 100$ GeV to reduce signal contamination. Events in which both the lepton and missing energy are from leptonic W decay tend to have $M_{\mathrm{T}} \lesssim M_W = 80$ GeV, while signal events (in which the predominant source of $p_{\mathrm{T}}^{\mathrm{miss}}$ is *not* W decay) tend to have much larger M_{T}. As the H_{T} and $p_{\mathrm{T}}^{\mathrm{miss}}$ selections are the same as in the baseline signal region selection, the same triggers as used in the signal region are used for this single lepton control region.

When a true lepton is within detector acceptance, it is generally reconstructed in some form, even if not classified as an isolated lepton candidate. We therefore remove the closest jet within $\Delta R < 0.4$ of the lepton candidate and count the lepton as a visible object for the purpose of computing the variables H_{T}, $H_{\mathrm{T}}^{\mathrm{miss}}$, $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$, $|\vec{H}_{\mathrm{T}}^{\mathrm{miss}} - \vec{p}_{\mathrm{T}}^{\mathrm{miss}}|/p_{\mathrm{T}}^{\mathrm{miss}}$, and M_{T2}. The lepton is *not* however used in computing N_{j}, as most leptons have low p_{T}.

The baseline single lepton control region is separated into $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ topological regions in the same way as the signal region, with the exception of events with $N_{\mathrm{j}} \geq 7$. The regions with $\geq$7 jets and $\geq$1 b tag are all predicted using control region bins with

7–9 or ≥10 jets and 1–2 b tags (but with the same H_{T} binning as the signal regions). This is motivated by the low control region statistics in bins with ≥10 jets or ≥7 jets and ≥2 b tags, as well as potential signal contamination in bins with ≥7 jets and ≥3 b tags. Similarly, regions with either 7–9 or ≥10 jets and 0 b tags are all predicted using control region bins with ≥ 7 jets and 0 b tags, due to low control region statistics in bins with ≥10 jets.

The single lepton control region is not binned in M_{T2} like the signal regions. Instead, the lost lepton prediction along the M_{T2} dimension is estimated using a hybrid technique described in Chapter 7.

5.5.2 Dilepton control regions

We make use of two dilepton control regions: a same-flavor $e^{\pm}e^{\mp}/\mu^{\pm}\mu^{\mp}$ control region, consisting mostly of $Z \rightarrow \ell^{+}\ell^{-}$ events and used to estimate the $Z \rightarrow \nu\bar{\nu}$ background, and an different-flavor $e^{\pm}\mu^{\mp}$ control region, used to estimate contamination from flavor-symmetric processes (mainly $t\bar{t}$) in the same-flavor region.

In either case, we require two opposite-charge leptons passing the reconstructed lepton selections given in Secs. 5.1.5 and 5.1.6. Dileptonic triggers are used, and to improve trigger efficiency we further require that electrons pass a slightly tighter cut-based ID requirement ("loose" instead of "veto"). The leptons must satisfy $p_{\mathrm{T}}(\ell\ell) > 200$ GeV (to ensure similar kinematics to high-p_{T} $Z \rightarrow \nu\bar{\nu}$ events that populate the signal regions), and leading/subleading lepton $p_{\mathrm{T}} > 100/35$ GeV. The invariant mass $m_{\ell\ell}$ must satisfy $|m_{\ell\ell} - m_Z| < 20$ GeV in order to ensure we select mainly Z events.

Since the leptons take the place of the neutrinos in $Z \rightarrow \nu\bar{\nu}$ events, to emulate the kinematics the lepton $\vec{p}_T$ vectors are added to the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ vector for the computation of all kinematic variables. Further, since leptons are usually reconstructed as (at least components of) jets, the closest jet within $\Delta R < 0.4$ of each lepton is removed from the event before computing kinematic quantities. The variables N_{j}, N_{b}, H_{T}, $H_{\mathrm{T}}^{\mathrm{miss}}$, $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$, $|\vec{H}_{\mathrm{T}}^{\mathrm{miss}} - \vec{p}_{\mathrm{T}}^{\mathrm{miss}}|/p_{\mathrm{T}}^{\mathrm{miss}}$, and M_{T2} are all potentially modified by these changes.

In addition to the main same-flavor and different-flavor control regions described above, there is additionally an auxiliary different-flavor region enriched in $t\bar{t}$, used to measure the ratio of same-flavor to different-flavor events from flavor-symmetric processes. This is defined by inverting the $p_{\mathrm{T}}(\ell\ell)$ and $m_{\ell\ell}$ selections, so that $p_{\mathrm{T}}(\ell\ell) < 200$ GeV and $|m_{\ell\ell} - m_Z| > 20$ GeV.

Similarly as for the single lepton control region, the dilepton control region is binned into topological in the same way as the signal region, with the exception of bins with $\geq$7 jets. This time, signal regions with 7–9 or $\geq$10 jets and 0 b tags are predicted with control region bins with $\geq$7 jets and 0 b tags, and signal regions with 7–9 or $\geq$10 jets and $\geq$ 1 b tag are predicted using control regions with $\geq$ 7 jets and $\geq$ 1 b tag (still with the same H_{T} binning as the signal regions).

5.5.3 QCD-enriched control regions

We define a few control regions orthogonal to the signal region that are enriched in QCD multijet events, in order to validate the estimate of the QCD multijet background. This is done by inverting cuts on the two kinematic variables mainly responsible for rejecting QCD background. The baseline signal region selection is applied, with the exception of one of the following:

- the $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ requirement is inverted to $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}}) < 0.3$.
- the M_{T2} requirement is shifted to $100 < M_{\mathrm{T2}} < 200$ GeV
- *both* the $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ and M_{T2} cuts are changed as above

The same triggers are used as in the signal regions.

5.6 2018 HEM-15/16 failure

During 2018 data taking, two modules in the $z < 0$ side HCAL endcap (HEM), representing about a 45° angle, were lost. This means that any energy deposited in these modules was not recorded. The effects of this are (1) under-measured jets in the region, as a fraction of the energy is lost, and (2) increased electron and photon fake rates, as jets in the region have inflated EM-to-hadronic energy ratios.

Studies showed this has negligible impact on backgrounds with real $p_{\mathrm{T}}^{\mathrm{miss}}$ (i.e. invisible Z and lost lepton), but a significant effect on the fake-$p_{\mathrm{T}}^{\mathrm{miss}}$ QCD multijet background. Fig. 5.7 shows on top the actual observed effect on an imbalanced dijet control region, and on bottom the estimated effect on QCD yields in the signal regions, found by modifying the Rebalance & Smear method with special HEM response functions measured in HEM-emulated data. Up to an 8-fold increase in QCD background is observed. Additionally, we observed an excess of events in the single lepton control region with the single lepton being a fake electron in the HEM region.

To avoid contaminating our signal and control regions with these HEM-affected events, we apply a special veto to the affected portion of 2018 data, as well as a fraction of the 2018 MC events corresponding to the fraction of affected data (about 39 out of 58 fb^{-1}, or 66%). The filter rejects events containing certain objects in the HEM region, defined as $\eta \in [-4.7, -1.4]$ and $\phi \in [-1.6, -0.8]$. The objects that trigger the veto are as follows:

- For events with $N_{\mathrm{j}} = 1$ or $H_{\mathrm{T}} < 575$ GeV, any jet with $p_{\mathrm{T}} > 20$ GeV.
- Also for events with $N_{\mathrm{j}} = 1$ or $H_{\mathrm{T}} < 575$ GeV, any lost track with $d_{\mathrm{z}} < 0.2$ cm and $d_{\mathrm{xy}} < 0.1$ cm (a “lost track” is any track not reconstructed as a PF candidate).
- For events with $N_{\mathrm{j}} \geq 2$ and $H_{\mathrm{T}} \geq 575$ GeV, any jet with $p_{\mathrm{T}} > 30$ GeV.
- The electron for any events in the single lepton control region.

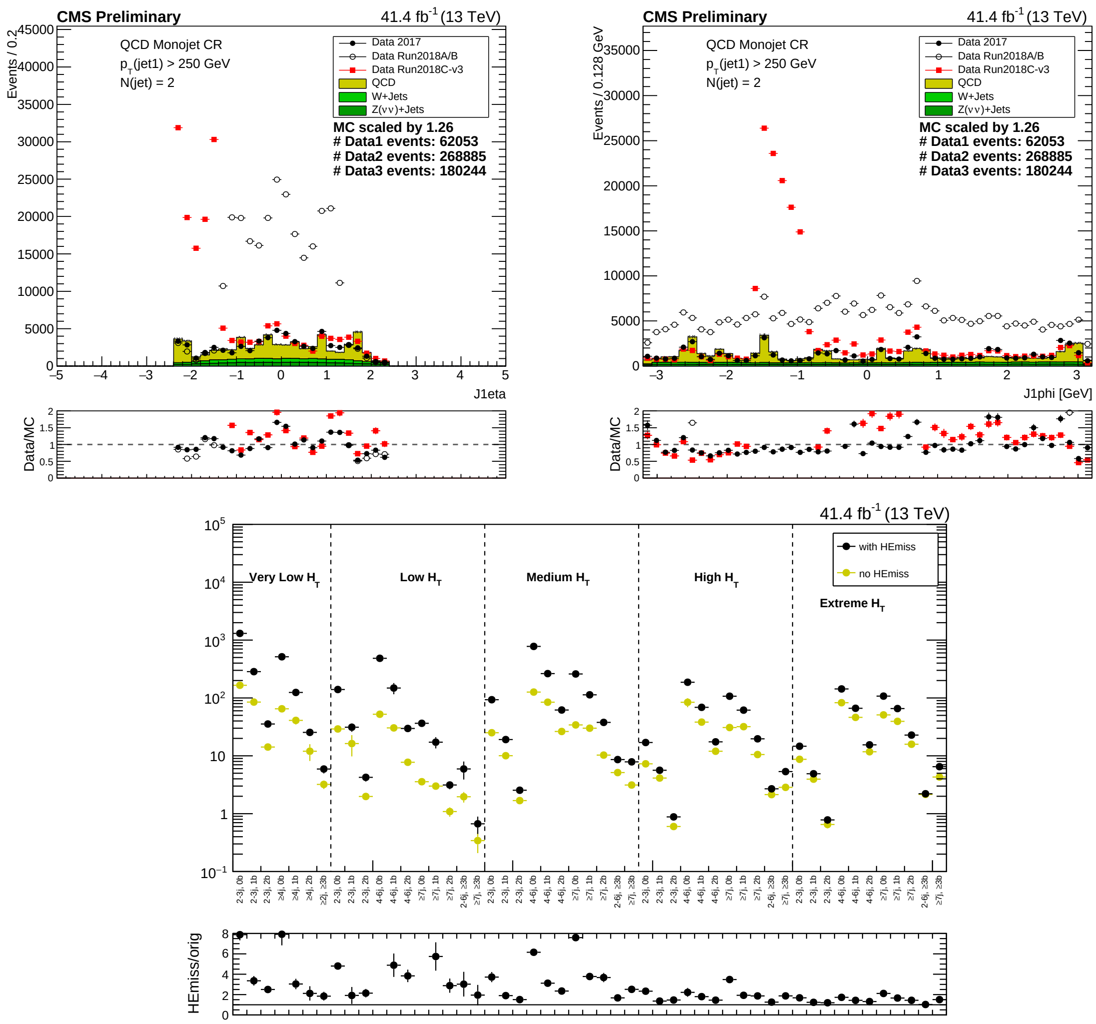

Figure 5.7: (top) Sub-leading jet η and ϕ before and after the HEM issue in an imbalanced dijet control region ($N_\mathrm{j} = 2$, $\Delta\phi_\mathrm{min} < 0.3$). Black points are 2017, and are flat in ϕ and agree with MC. White points are early pre-HEM issue 2018 data, flat in ϕ but with a uniform barrel-only excess due to an unrelated filter issue that was since fixed. Red points are post-HEM issue 2018 data, showing a clear excess in the HEM region $\eta \in [-4.7, -1.4]$, $\phi \in [-1.6, 0.8]$. Dijet events that are otherwise balanced but that have a jet in the HEM region become imbalanced due to the lost energy. (bottom) Estimated effect of HEM issue on QCD signal region yields, found by modifying Rebalance & Smear templates with special response functions measured in HEM-emulated data. The issue causes up to an 8-fold increase in estimated background, and the effect is worst at low H_T.

Chapter 6

Invisible Z Background

One of the main sources of background to the M_{T2} analysis comes from the production of a Z boson in association with one or more jets, where the Z decays "invisibly" to two neutrinos. This is an *irreducible* background, since the signature is exactly like that of signal: no prompt charged leptons, real $p_{\mathrm{T}}^{\mathrm{miss}}$, and multiple jets. For signal models that produce b quarks, however, the number of b-tagged jets becomes a discriminating variable as there are no inherent b jets in Z+jets production.

To estimate this background in a data-driven way, we make use of the fact that Z bosons may also decay into two charged leptons with known branching fraction. The rate of $Z \to \ell^+\ell^-$ production is measured in data, and this is translated into an estimate on $Z \to \nu\bar{\nu}$ by multiplying by a transfer factor derived from Monte Carlo. Section 6.1 describes how this is done in each $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ topological region, Sec. 6.2 details the "hybrid template" method for extrapolating the estimate along the M_{T2} dimension, and Sec. 6.3 describes the systematic uncertainties assessed on the final estimate.

6.1 Estimating $Z \to \nu\bar{\nu}$ from $Z \to \ell^+\ell^-$

In previous iterations of this analysis [67, 68], the primary $Z \to \nu\bar{\nu}$ estimate was performed using a control sample of γ+jets events rather than $Z \to \ell^+\ell^-$ events. The cross section for γ+jets production is much larger than that for Z+jets production, and so the number of events in γ+jets control regions is much higher than for the corresponding $Z \to \ell^+\ell^-$ control regions. However, systematic uncertainties associated with photon reconstruction efficiencies and fake rates are significantly larger than those for an estimate based on $Z \to \ell^+\ell^-$ events. As the integrated luminosity has increased four-fold from the 2016 analysis, the loss in statistical precision is less important and so the $Z \to \ell^+\ell^-$ method is used instead.

The exact definition of the dilepton control region used in the $Z \to \nu\bar{\nu}$ estimate is described in Sec. 5.5.2. It is constructed to be enriched in $Z \to \ell^+\ell^-$ events (as opposed to other dilepton-producing processes, such as $t\bar{t}$), and the lepton $\vec{p}_T$ vectors are added to the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ in order to mimic the kinematics of $Z \to \nu\bar{\nu}$ events.

Data vs. MC comparisons of the main analysis kinematic variables in the baseline $Z \to \ell^+\ell^-$ control region are shown in Fig. 6.1. There is some level of disagreement between data and MC, but this is to be expected. The estimate is primarily data-driven, and MC is used only at leading order, so these disagreements generally do not affect the final estimate. Where MC is used, primarily in the far M_{T2} tails, appropriate systematic uncertainties are assigned.

The $Z \to \nu\bar{\nu}$ estimate is first performed in each $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ topological region, integrated over M_{T2} (for the monojet region, the H_{T} dimension is equivalent to $\vec{p}_T^{\mathrm{jet1}}$, so there is no integration and the estimate is performed in each analysis bin). For all regions with $\geq$7 jets and $\geq$1 b tag, an inclusive control region with $\geq$7 jets and $\geq$1 b tag is used, to avoid statistical fluctuations in these regions where $Z \to \nu\bar{\nu}$ is a subleading background. Similarly, for regions with 7–9 jets or $\geq$10 jets and 0 b tags, an inclusive control region with $\geq$7 jets and 0 b tags is used. Since MC is used directly to predict the N_{j} and N_{b} shapes here, agreement between data and MC is explicitly checked in these high-N_{j} regions. Fig. 6.2 shows data vs. MC comparisons of N_{j} and N_{b} in the $\geq$7 jet

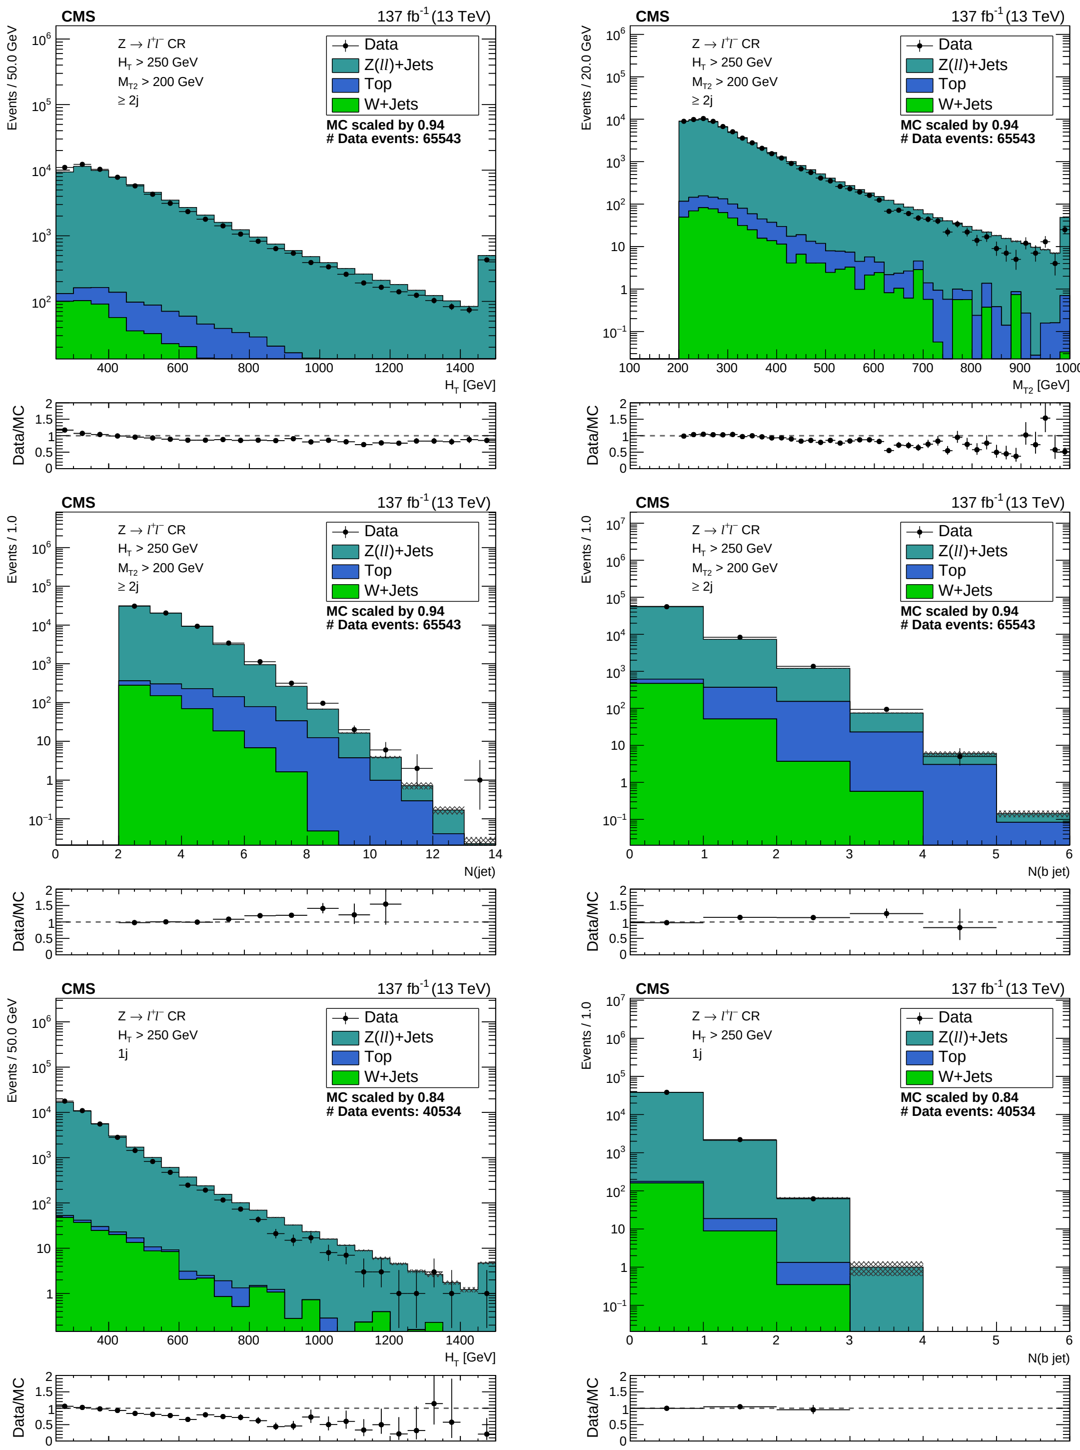

Figure 6.1: Data vs. MC comparisons in the baseline $Z \to \ell^+\ell^-$ control region, for $N_\mathrm{j} \geq 2$ (top two rows), and $N_\mathrm{j} = 1$ (bottom row). From left to right, top to bottom, the variables plotted are H_T, $M_{\mathrm{T}2}$, N_j, N_b, H_T, and N_b.

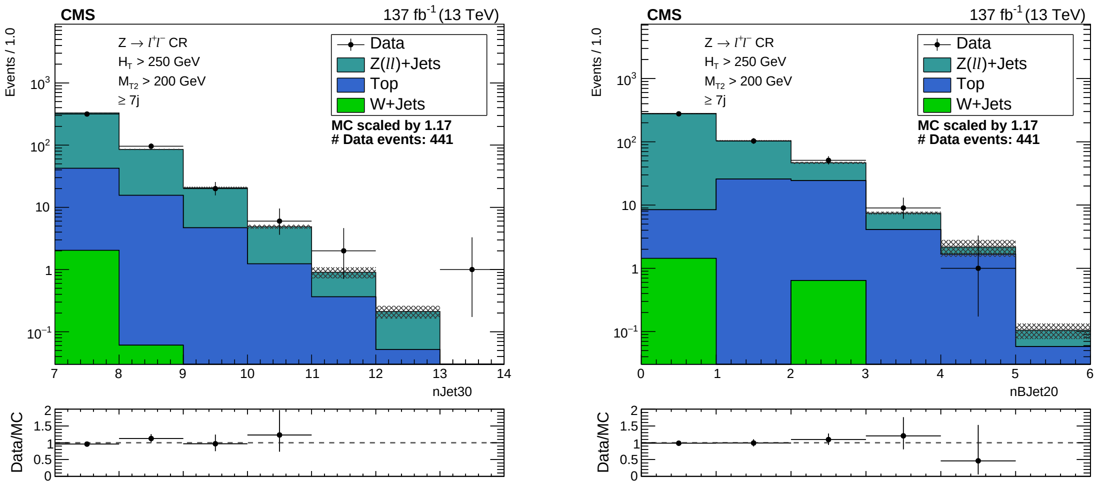

Figure 6.2: Comparisons of N_j and N_b in the $N_\mathrm{j} \geq 7$ region. In these high-N_j regions, shapes of N_j and N_b are taken directly from MC as statistics in data are insufficient. Monte Carlo is seen to agree well with the observed data in the control region, validating its use to predict signal region yields.

region; sufficient agreement is seen to justify the use of MC in predicting the N_j and N_b shapes.

For regions with $H_\mathrm{T} > 1500$ GeV, all events with $M_{\mathrm{T2}} > 200$ GeV are used in the control region even though the signal region starts at $M_{\mathrm{T2}} > 400$ GeV. Once the per-topological region estimate is done, a hybrid approach using both data and MC is used to extrapolate along the M_{T2} dimension, as described in Sec. 6.2.

The final estimate in each $(H_\mathrm{T}, N_\mathrm{j}, N_\mathrm{b}, M_{\mathrm{T2}})$ signal region can then be summarized as

$$N^{\mathrm{SR}}_{Z\to\nu\bar{\nu}} = \left[N^{\mathrm{CRSF}}_{2-\mathrm{lep}}\; P_{Z\to\ell^+\ell^-}\; R^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}_{\mathrm{MC}} \right](\Omega) \times k_{\mathrm{hybrid}}(M_{\mathrm{T2}}|\Omega), \tag{6.1}$$

where

- $N^{\mathrm{CRSF}}_{2-\mathrm{lep}}$, $P_{Z\to\ell^+\ell^-}$, and $R^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}_{\mathrm{MC}}$ are measured in each topological region (referred to here as $\Omega \equiv (H_\mathrm{T}, N_\mathrm{j}, N_\mathrm{b})$), integrated over M_{T2}.
- $N^{\mathrm{CRSF}}_{2-\mathrm{lep}}$ is the number of observed events in data in the same-flavor dilepton control

region.

- $P_{Z\to\ell^+\ell^-}$ is the *purity*, or fraction of $Z \to \ell^+\ell^-$ events in the dilepton control region (see Sec. 6.1.1).
- $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ is the ratio between $Z \to \nu\bar{\nu}$ and $Z \to \ell^+\ell^-$ MC yields in this region.
- $k_{\mathrm{hybrid}}(M_{\mathrm{T2}}|\Omega)$ is a normalized template used to distribute events as a function of M_{T2} in each topological region (see Sec. 6.2).

The values of $N_{2-\mathrm{lep}}^{\mathrm{CRSF}}$, $P_{Z\to\ell^+\ell^-}$, and $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ for each topological region are given in Tables 6.1 and 6.2.

6.1.1 $Z \to \ell^+\ell^-$ purity in the dilepton control region

While $Z \to \ell^+\ell^-$ is the process of interest in the dilepton control region, other SM processes can produce similar same-flavor dilepton signatures. Most prominent is $t\bar{t}$ production, but other processes contribute in more minor ways, such as $t\bar{t}V$, diboson, and W+jets (with a fake lepton) production.

Certain kinematic cuts placed in the dilepton control region ensure that the contributions from these other processes are as small as possible. The $|m_{\ell\ell} - m_Z| < 20$ GeV cut requires that the dilepton mass is consistent with Z boson decay. Furthermore, the $p_{\mathrm{T}}(\ell\ell) > 200$ GeV cut selects events in which the leptons are the primary source of $p_{\mathrm{T}}^{\mathrm{miss}}$ in the event (once their $\vec{p}_T$ vectors are added to the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$). When this is the case, the $p_{\mathrm{T}}^{\mathrm{miss}} > 250$ GeV cut means that $p_{\mathrm{T}}(\ell\ell)$ must be large. If, however, there are contributions to the $p_{\mathrm{T}}^{\mathrm{miss}}$ from actual invisible particles (such as the neutrinos present in all of the secondary production modes listed above), $p_{\mathrm{T}}(\ell\ell)$ is generally smaller. This is illustrated in Fig. 6.3, which shows a distribution of $p_{\mathrm{T}}(\ell\ell)$ in the baseline dilepton control region (with the $m_{\ell\ell}$ cut inverted for $p_{\mathrm{T}}(\ell\ell) < 200$ GeV). We see that the $p_{\mathrm{T}}(\ell\ell) > 200$ GeV

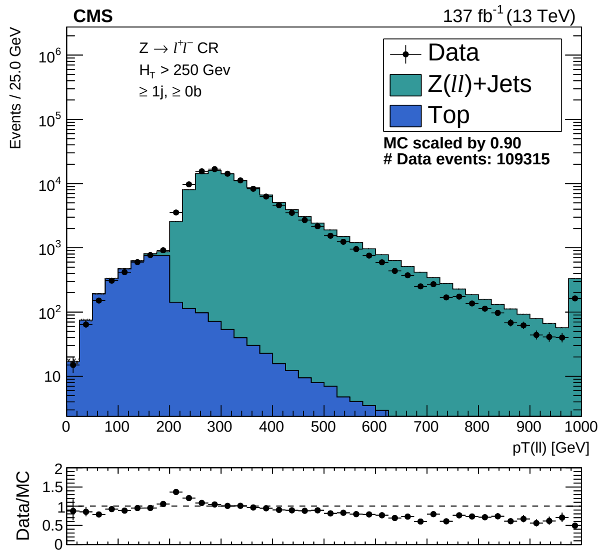

Figure 6.3: $p_{\mathrm{T}}(\ell\ell)$ in the baseline dilepton control region, with the exception of the $|m_{\ell\ell} - m_Z| < 20$ GeV cut which is inverted for the $p_{\mathrm{T}}(\ell\ell) < 200$ GeV portion of the plot. We see that the $m_{\ell\ell}$ and $p_{\mathrm{T}}(\ell\ell)$ requirements ensure that the control region is predominantly from $Z \to \ell^+\ell^-$, and inverting these requirements produces a region enriched in top processes (used to measure $R^{\mathrm{SF/OF}}$).

region is almost entirely from $Z \to \ell^+\ell^-$ events, and the $p_{\mathrm{T}}(\ell\ell) < 200$ GeV region comes predominantly from top ($t\bar{t}$ and $t\bar{t}V$) processes.

Despite the attempts made to purify the control region with $Z \to \ell^+\ell^-$ events, there is nevertheless some contamination from non-Z processes. To estimate and subtract off this contamination, we make use of the fact that all of the secondary processes are flavor-symmetric, in that they produce opposite-flavor ($e\mu$) events at the same rate as same-flavor (ee or $\mu\mu$) events. Hence, we can measure their rate in an opposite-flavor (OF) control region (same cuts as the nominal same-flavor (SF) dilepton control region, but requiring one e and one μ), and then convert this to an estimated same-flavor rate by multiplying by a transfer factor $R^{\mathrm{SF/OF}}$.

While this SF/OF ratio is in theory equal to 1 at the production level (since branching fractions to electrons and muons are the same), different reconstruction efficiencies for

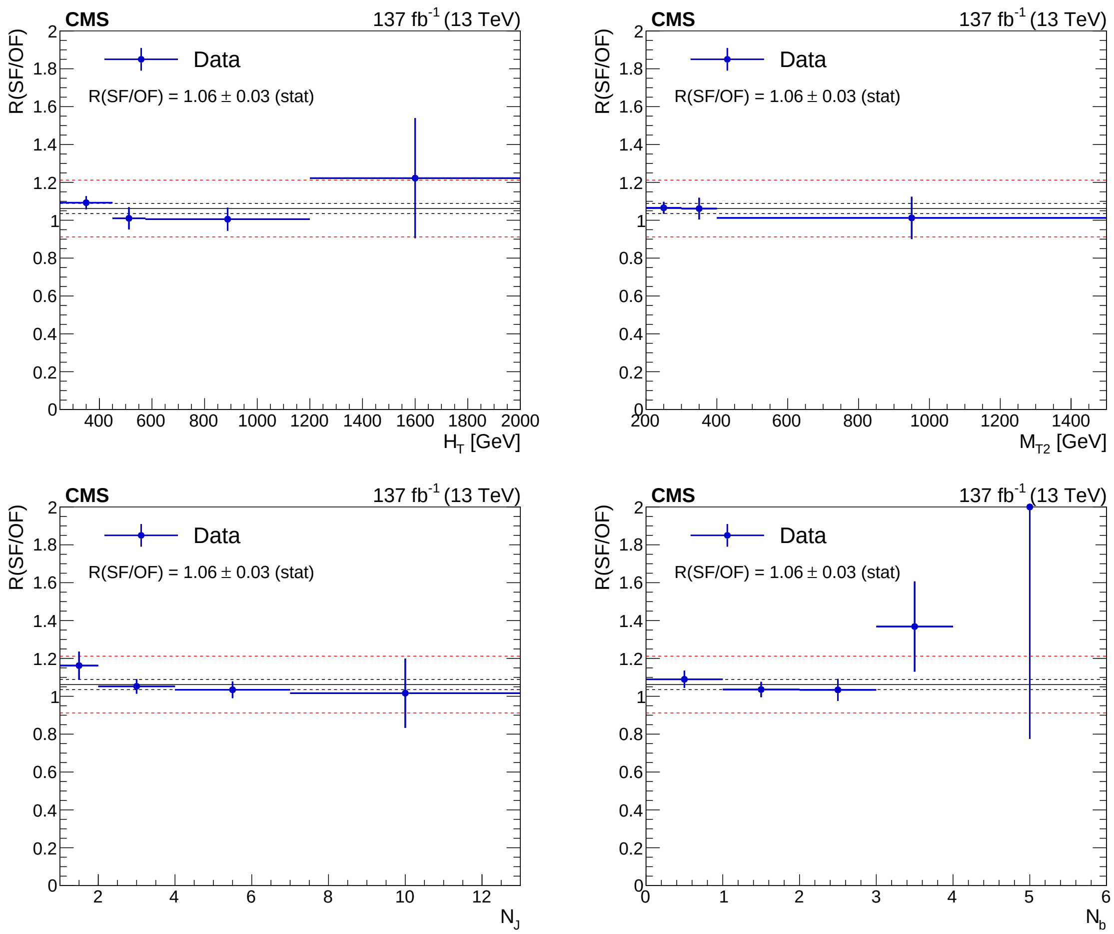

Figure 6.4: Measured $R^{\mathrm{SF/OF}}$ as a function of H_{T}, M_{T2}, N_{j}, and N_{b}. The black lines show the measured value of 1.06 ± 0.03. The dotted red lines show the 15% systematic uncertainty assessed to cover for any kinematic variation in the ratio.

electrons and muons cause it to deviate from 1 at the reconstructed level. We measure this ratio in data in a top-enriched dilepton control region, which is the same as the baseline dilepton control region but with $|m_{\ell\ell} - m_Z| > 20$ GeV and $p_{\mathrm{T}}(\ell\ell) < 200$ GeV (the portion of Fig. 6.3 with $p_{\mathrm{T}}(\ell\ell)$ below 200 GeV). The measurement is done inclusively in all analysis variables, to get a single number that is applied to every topological region. The measured ratio is plotted as a function of the analysis variables to ensure that it is actually constant, and a systematic uncertainty on $R^{\mathrm{SF/OF}}$ is assessed to cover for any observed variation.

This is illustrated in Fig. 6.4, in which $R^{\mathrm{SF/OF}}$ as measured in data is plotted as a function of H_{T}, M_{T2}, N_{j}, and N_{b}. The ratio is measured to be 1.06 ± 0.03, and is seen to be approximately independent of the kinematic variables. To cover for the small observed variation, a systematic uncertainty of 0.15 is placed on $R^{\mathrm{SF/OF}}$ and propagated to the final estimate.

Once $R^{\mathrm{SF/OF}}$ is measured, we can compute the purity $P_{Z\to\ell^+\ell^-}$ used in Eq. 6.1 (measured independently in each topological region as) as

$$P_{Z\to\ell^+\ell^-} = 1 - R^{\mathrm{SF/OF}} N^{\mathrm{CROF}}_{2-\mathrm{lep}} / N^{\mathrm{CRSF}}_{2-\mathrm{lep}}, \tag{6.2}$$

where $N^{\mathrm{CROF}}_{2-\mathrm{lep}}$ is the number of observed events in data in the opposite-flavor dilepton control region. The values of $N^{\mathrm{CROF}}_{2-\mathrm{lep}}$ and $P_{Z\to\ell^+\ell^-}$ in each topological region are given in Tables 6.1 and 6.2.

6.2 M_{T2} extrapolation

To project the per-topological region estimate described in the previous section along the M_{T2} dimension, we use a "hybrid" M_{T2} template derived using both data and MC ($k_{\mathrm{hybrid}}(M_{\mathrm{T2}}|\Omega)$ in Eq. 6.1).

In every H_{T} region, the M_{T2} shape for both $Z \to \nu\bar{\nu}$ and $Z \to \ell^+\ell^-$ MC events is observed to be independent of N_{b}. An example of this is shown in Fig. 6.5, for the Medium H_{T} region. Furthermore, in the Extreme H_{T} region ($[1500, \infty]$ GeV), the M_{T2} shape is observed to also be independent of N_{j} up to M_{T2} values of about 1 TeV.

Thus, we build M_{T2} shape templates from data and MC in each $(H_{\mathrm{T}}, N_{\mathrm{j}})$ bin for regions with $H_{\mathrm{T}} < 1500$ GeV, integrating over N_{b}, while one single hybrid template is

Table 6.1: The control region Drell-Yan (DY) MC yield, same-flavor (SF) and opposite-flavor (OF) data yields, purity, and $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ for the $Z \to \nu\bar{\nu}$ estimate as a function of H_T, N_j, and N_b, in the monojet, Very Low H_T, and Low H_T regions. Note that the the regions marked with a * or † share the same control region (with the same yield and purity), and the distribution of events in the N_b and/or N_j dimensions are folded into the $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ value.

Invisible Z							
Region			DY yield	SF yield	OF yield	Purity	$R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$
H_T [GeV]	N_j	N_b					
$[250, 350]$	1	0	30647.5	27100	6	$0.998^{+0.001}_{-0.001}$	6.15 ± 0.02
$[350, 450]$	1	0	9680.6	7891	5	$0.996^{+0.002}_{-0.003}$	5.11 ± 0.02
$[450, 575]$	1	0	3414.1	2375	3	$0.996^{+0.002}_{-0.004}$	4.59 ± 0.02
$[575, 700]$	1	0	985.6	610	2	$0.988^{+0.008}_{-0.016}$	4.3 ± 0.02
$[700, 1000]$	1	0	480.3	265	1	$0.992^{+0.007}_{-0.019}$	4.09 ± 0.03
$[1000, 1200]$	1	0	47.5	21	0	$1.000^{+0.000}_{-0.094}$	4.11 ± 0.1
>1200	1	0	16.0	7	0	$1.000^{+0.000}_{-0.281}$	6.53 ± 0.38
$[250, 350]$	1	≥ 1	1771.9	1593	2	$0.997^{+0.002}_{-0.004}$	5.69 ± 0.08
$[350, 450]$	1	≥ 1	531.3	460	1	$0.993^{+0.006}_{-0.016}$	4.76 ± 0.08
$[450, 575]$	1	≥ 1	178.6	152	1	$0.986^{+0.012}_{-0.032}$	4.35 ± 0.08
$[575, 700]$	1	≥ 1	48.0	40	0	$1.000^{+0.000}_{-0.049}$	4.11 ± 0.1
>700	1	≥ 1	23.0	20	1	$0.947^{+0.044}_{-0.123}$	7.38 ± 1.73
$[250, 450]$	2-3	0	32020.2	31566	11	$0.996^{+0.001}_{-0.002}$	6.01 ± 0.04
$[250, 450]$	2-3	1	3461.3	4038	8	$0.984^{+0.005}_{-0.008}$	5.7 ± 0.05
$[250, 450]$	2-3	2	471.8	523	3	$0.980^{+0.011}_{-0.020}$	5.54 ± 0.11
$[250, 450]$	4-6	0	3892.5	4181	5	$0.996^{+0.002}_{-0.003}$	6.35 ± 0.11
$[250, 450]$	4-6	1	662.4	875	5	$0.971^{+0.013}_{-0.020}$	6.2 ± 0.13
$[250, 450]$	4-6	2	150.5	212	3	$0.934^{+0.036}_{-0.064}$	5.68 ± 0.2
$[250, 450]$	≥ 7	0	6.5	12	2	$0.911^{+0.058}_{-0.118}$	6.32 ± 2.15
$[250, 450]$	≥ 7	1	1.6	3	0	$1.000^{+0.000}_{-0.657}$	7.31 ± 2.92 (*)
$[250, 450]$	≥ 7	2	1.6	3	0	$1.000^{+0.000}_{-0.657}$	2.23 ± 0.99 (*)
$[250, 450]$	2-6	≥ 3	23.5	36	2	$0.941^{+0.038}_{-0.078}$	6.51 ± 0.5
$[250, 450]$	≥ 7	≥ 3	1.6	3	0	$1.000^{+0.000}_{-0.657}$	0.38 ± 0.26 (*)
$[450, 575]$	2-3	0	8588.5	7288	5	$0.996^{+0.002}_{-0.003}$	5.16 ± 0.06
$[450, 575]$	2-3	1	951.2	943	5	$0.976^{+0.010}_{-0.016}$	5.07 ± 0.06
$[450, 575]$	2-3	2	113.0	128	3	$0.941^{+0.032}_{-0.057}$	4.97 ± 0.09
$[450, 575]$	4-6	0	2789.8	2590	5	$0.993^{+0.003}_{-0.005}$	5.7 ± 0.11
$[450, 575]$	4-6	1	499.2	599	5	$0.970^{+0.013}_{-0.021}$	5.55 ± 0.11
$[450, 575]$	4-6	2	104.3	136	3	$0.937^{+0.034}_{-0.061}$	5.56 ± 0.13
$[450, 575]$	≥ 7	0	26.2	40	0	$1.000^{+0.000}_{-0.049}$	6.05 ± 0.7
$[450, 575]$	≥ 7	1	8.5	19	0	$1.000^{+0.000}_{-0.104}$	5.02 ± 0.64 (*)
$[450, 575]$	≥ 7	2	8.5	19	0	$1.000^{+0.000}_{-0.104}$	1.5 ± 0.21 (*)
$[450, 575]$	2-6	≥ 3	12.3	18	2	$0.941^{+0.038}_{-0.078}$	6.07 ± 0.29
$[450, 575]$	≥ 7	≥ 3	8.5	19	0	$1.000^{+0.000}_{-0.104}$	0.17 ± 0.04 (*)

built for the Extreme $H_\mathrm{T} > 1500$ GeV region, integrating also over N_j. The one exception is for regions with 2–6 jets and ≥3 b tags. To avoid sculpting the N_j distribution by requiring $N_\mathrm{b} \geq 3$, we use regions with $N_\mathrm{j} \geq 3$ to obtain M_{T2} shape templates in these

Table 6.2: Same as Table 6.1, but for the Medium, High, and Extreme H_{T} regions

Invisible Z							
Region			DY yield	SF yield	OF yield	Purity	$R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$
H_{T} [GeV]	N_{j}	N_{b}					
$[575, 1200]$	2-3	0	7454.8	5509	6	$0.993^{+0.003}_{-0.004}$	4.87 ± 0.06
$[575, 1200]$	2-3	1	832.1	738	5	$0.965^{+0.015}_{-0.024}$	4.79 ± 0.06
$[575, 1200]$	2-3	2	85.2	97	3	$0.890^{+0.060}_{-0.107}$	4.85 ± 0.08
$[575, 1200]$	4-6	0	4218.6	3619	5	$0.993^{+0.003}_{-0.004}$	5.32 ± 0.08
$[575, 1200]$	4-6	1	799.8	893	6	$0.960^{+0.016}_{-0.024}$	5.23 ± 0.08
$[575, 1200]$	4-6	2	159.2	200	5	$0.909^{+0.039}_{-0.062}$	5.26 ± 0.09
$[575, 1200]$	2-6	≥ 3	17.5	34	3	$0.843^{+0.086}_{-0.153}$	5.38 ± 0.14
$[575, 1200]$	7-9	0	148.4	173	2	$0.994^{+0.004}_{-0.008}$	6.03 ± 0.35 (*)
$[575, 1200]$	7-9	1	58.0	120	3	$0.929^{+0.039}_{-0.069}$	4.51 ± 0.27 (†)
$[575, 1200]$	7-9	2	58.0	120	3	$0.929^{+0.039}_{-0.069}$	1.33 ± 0.08 (†)
$[575, 1200]$	7-9	3	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.18 ± 0.01 (†)
$[575, 1200]$	7-9	≥ 4	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.033 ± 0.005 (†)
$[575, 1200]$	≥ 10	0	148.4	173	2	$0.994^{+0.004}_{-0.008}$	0.046 ± 0.015 (*)
$[575, 1200]$	≥ 10	1	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.056 ± 0.018 (†)
$[575, 1200]$	≥ 10	2	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.017 ± 0.006 (†)
$[575, 1200]$	≥ 10	3	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.0039 ± 0.0016 (†)
$[575, 1200]$	≥ 10	≥ 4	58.0	120	3	$0.929^{+0.039}_{-0.069}$	0.0 ± 0.0 (†)
$[1200, 1500]$	2-3	0	300.3	194	2	$0.989^{+0.007}_{-0.015}$	5.01 ± 0.34
$[1200, 1500]$	2-3	1	33.3	24	2	$0.955^{+0.029}_{-0.059}$	4.96 ± 0.35
$[1200, 1500]$	2-3	2	2.3	3	0	$1.000^{+0.000}_{-0.657}$	5.58 ± 0.59
$[1200, 1500]$	4-6	0	288.8	236	2	$0.986^{+0.009}_{-0.018}$	5.46 ± 0.32
$[1200, 1500]$	4-6	1	59.1	73	2	$0.941^{+0.038}_{-0.077}$	5.17 ± 0.31
$[1200, 1500]$	4-6	2	10.0	11	0	$1.000^{+0.000}_{-0.179}$	5.66 ± 0.4
$[1200, 1500]$	2-6	≥ 3	1.3	0	0	$1.000^{+0.000}_{-0.000}$	4.95 ± 0.67
$[1200, 1500]$	7-9	0	28.1	25	0	$1.000^{+0.000}_{-0.079}$	5.84 ± 0.78 (*)
$[1200, 1500]$	7-9	1	11.5	13	2	$0.918^{+0.053}_{-0.109}$	4.12 ± 0.57 (†)
$[1200, 1500]$	7-9	2	11.5	13	2	$0.918^{+0.053}_{-0.109}$	1.13 ± 0.16 (†)
$[1200, 1500]$	7-9	3	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.22 ± 0.04 (†)
$[1200, 1500]$	7-9	≥ 4	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.03 ± 0.014 (†)
$[1200, 1500]$	≥ 10	0	28.1	25	0	$1.000^{+0.000}_{-0.079}$	0.13 ± 0.11 (*)
$[1200, 1500]$	≥ 10	1	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.12 ± 0.1 (†)
$[1200, 1500]$	≥ 10	2	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.047 ± 0.04 (†)
$[1200, 1500]$	≥ 10	3	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.013 ± 0.012 (†)
$[1200, 1500]$	≥ 10	≥ 4	11.5	13	2	$0.918^{+0.053}_{-0.109}$	0.0085 ± 0.0106 (†)
>1500	2-3	0	169.4	135	0	$1.000^{+0.000}_{-0.015}$	5.24 ± 0.28
>1500	2-3	1	17.8	13	0	$1.000^{+0.000}_{-0.152}$	4.87 ± 0.29
>1500	2-3	2	1.2	0	0	$1.000^{+0.000}_{-0.000}$	5.11 ± 0.66
>1500	4-6	0	184.9	153	2	$0.993^{+0.005}_{-0.009}$	5.49 ± 0.29
>1500	4-6	1	35.7	31	2	$0.896^{+0.067}_{-0.137}$	5.52 ± 0.31
>1500	4-6	2	5.9	5	2	$0.786^{+0.138}_{-0.282}$	5.17 ± 0.38
>1500	2-6	≥ 3	0.7	1	0	$1.000^{+0.000}_{-1.970}$	4.93 ± 0.81
>1500	7-9	0	22.8	27	0	$1.000^{+0.000}_{-0.073}$	5.59 ± 0.32 (*)
>1500	7-9	1	9.1	9	0	$1.000^{+0.000}_{-0.219}$	4.78 ± 0.32 (†)
>1500	7-9	2	9.1	9	0	$1.000^{+0.000}_{-0.219}$	1.24 ± 0.1 (†)
>1500	7-9	3	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.14 ± 0.02 (†)
>1500	7-9	≥ 4	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.024 ± 0.009 (†)
>1500	≥ 10	0	22.8	27	0	$1.000^{+0.000}_{-0.073}$	0.22 ± 0.02 (*)
>1500	≥ 10	1	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.24 ± 0.03 (†)
>1500	≥ 10	2	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.1 ± 0.02 (†)
>1500	≥ 10	3	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.00032 ± 0.00012 (†)
>1500	≥ 10	≥ 4	9.1	9	0	$1.000^{+0.000}_{-0.219}$	0.0014 ± 0.001 (†)

regions.

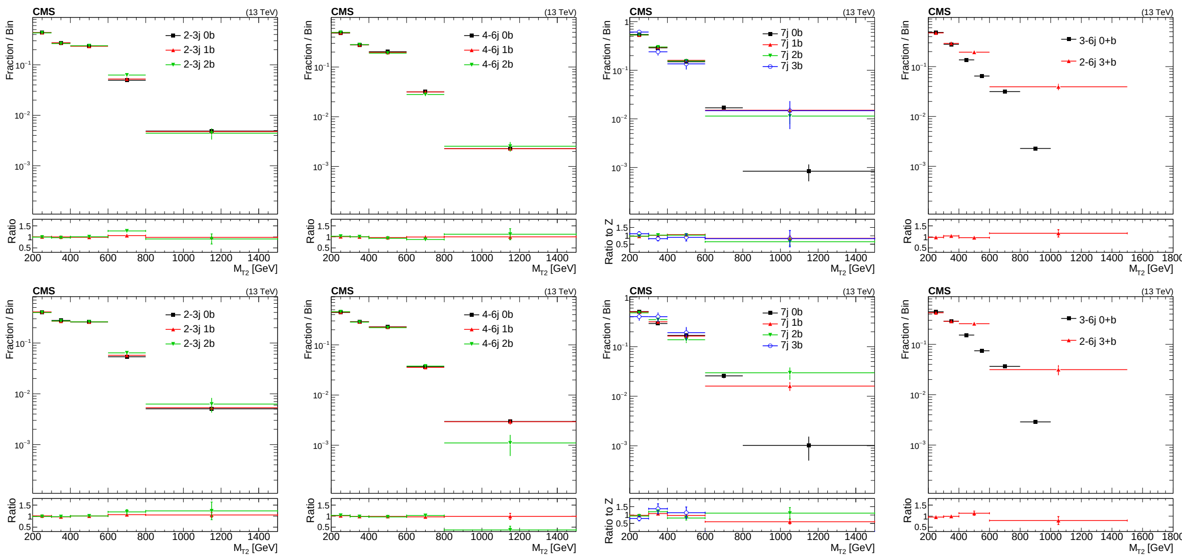

Figure 6.5: Example M_{T2} distributions in various bins of N_{j} and N_{b} in the H_{T} region $[575, 1000]$ GeV, for $Z \to \ell^+\ell^-$ (top) and $Z \to \nu\bar{\nu}$ (bottom) MC. We see that the M_{T2} shape is independent of N_{b} in all N_{j} regions.

Starting from the highest M_{T2} bin in each topological region in the dilepton control region, we merge bins until the sum of expected yields in the merged bins is at least 50 events, as predicted by MC for the full integrated luminosity of 137 fb^{-1}. For the lower non-merged M_{T2} bins, which have larger statistics, the M_{T2} shapes are built directly from $Z \to \ell^+\ell^-$ data, corrected by the $Z \to \nu\bar{\nu}/Z \to \ell^+\ell^-$ MC ratio in order to account for lepton reconstruction effects. The $Z \to \nu\bar{\nu}$ MC M_{T2} shape is instead used to distribute events across the merged M_{T2} bins, after renormalizing the MC to the total data yield in the same bins. For the $H_{\mathrm{T}} > 1500$ GeV region, we use N_{j}-binned $Z \to \nu\bar{\nu}$ MC shapes for the extrapolation, to avoid the mild N_{j} dependence observed at high M_{T2}.

For the lower M_{T2} bins, where the M_{T2} shapes are built directly from $Z \to \ell^+\ell^-$ data, an uncertainty corresponding to the statistical uncertainty in data in each M_{T2} bin is accounted for. For the bins where $Z \to \nu\bar{\nu}$ MC is used, an additional systematic

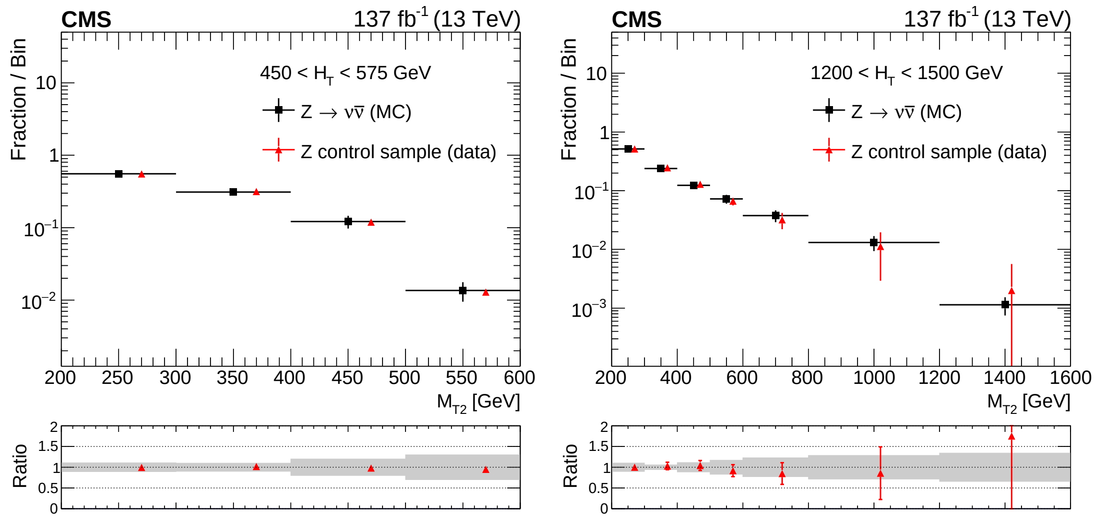

Figure 6.6: Comparisons of M_{T2} shape from $Z \rightarrow \nu\bar{\nu}$ MC (black) and $Z \rightarrow \ell^+\ell^-$ data from the dilepton control region (red), for the Low (left) and High (right) H_{T} regions. The grey band in the ratio plot corresponds to the systematic assigned to the $Z \rightarrow \nu\bar{\nu}$ shape, as described in Sec. 6.3.

uncertainty, as large as 40% in the last M_{T2} bin, is assessed, as described in Sec. 6.3.

This M_{T2} shape procedure is validated by comparing M_{T2} shapes of $Z \rightarrow \nu\bar{\nu}$ MC and $Z \rightarrow \ell^+\ell^-$ data from the dilepton control region. Examples of this for the Low and High H_{T} regions are shown in Fig. 6.6. M_{T2} shape obtained from dilepton data is shown in red markers, and that from $Z \rightarrow \nu\bar{\nu}$ MC is shown in black. The grey band corresponds to the systematic uncertainty assigned to the $Z \rightarrow \nu\bar{\nu}$ shape, as described in Sec. 6.3. It is seen than any discrepancies in shape are covered by the assigned systematic.

6.3 Systematic uncertainties

The following systematics are assessed on the $Z \rightarrow \nu\bar{\nu}$ background prediction:

- Control region statistical error: the Poisson error on the observed data count in each $Z \rightarrow \ell^+\ell^-$ control region. This is correlated among all bins that share a common control region (including the M_{T2} bins that utilize a merged control region), but is otherwise uncorrelated.

- $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ (stat): from MC statistical uncertainty
- $R_{\mathrm{MC}}^{Z\to\nu\bar{\nu}/Z\to\ell^+\ell^-}$ (syst): $\mathcal{O}$(5-10%) uncertainty, mainly from lepton efficiency uncertainties in the dilepton control region. Jet energy scale uncertainties also contribute.
- Purity (stat): from the 3% statistical uncertainty on $R^{\mathrm{SF/OF}}$
- Purity (syst): from the assigned 15% systematic uncertainty on $R^{\mathrm{SF/OF}}$ to cover for any kinematic variation
- M_{T2} shape uncertainty: For the data-driven component of the hybrid templates (low-M_{T2}, high stats bins), this is covered by the CR statistical error listed above. For the MC-driven component (higher M_{T2}), this is an extra assigned uncertainty based on MC variations in the M_{T2} shape, accounting for theoretical (renormalization and factorization scales, parton distribution functions) and experimental (jet energy scale, $p_{\mathrm{T}}^{\mathrm{miss}}$ modeling) effects. These effects give at most a 20% variation in the last M_{T2} bin in each topological region when these parameters are varied in MC. The uncertainty in the last bin is increased to 40% to account for possible mis-modeling of M_{T2} in MC. The uncertainty is implemented as a correlated linear morphing of the M_{T2} shape in the higher M_{T2} bins, with a maximum amplitude of 40% in the last bin, done in such a way as to preserve normalization and only affect the shape. These shape uncertainties are not correlated between topological regions.

Chapter 7

Lost Lepton Background

A second source of background comes from events where a genuine prompt lepton is produced from the decay of a W boson. While the tight lepton veto employed in the analysis aims to eliminate as many of these events as possible, a significant number still enter the signal regions because the lepton gets "lost". This can happen for a number of reasons, such as the lepton being outside of the acceptance window (i.e. it is too forward or the p_{T} is too small) or not being isolated, or the reconstruction algorithms failing to identify the candidate as a lepton.

Generally the energy of the lepton is still accounted for (making "lost" a bit of a misnomer), so there is no fake $p_{\mathrm{T}}^{\mathrm{miss}}$ from the lepton. However, there is real $p_{\mathrm{T}}^{\mathrm{miss}}$ from the neutrino from the W decay, and this is often enough to allow the event to enter the signal region. The dominant production mechanisms for the lost lepton background are W+jets and $t\bar{t}$ production, but there are also smaller contributions from rarer processes such as single top, $t\bar{t}W$, $t\bar{t}Z$, $t\bar{t}H$, and $tt\bar{t}\bar{t}$.

The lost lepton background is estimated in a data-driven way using a control sample of events where there *is* a reconstructed lepton. Section 7.1 describes how this is done in each $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ topological region, Sec. 7.2 explains the method for extrapolating along the M_{T2} dimension, and Sec. 7.3 lists the systematic uncertainties assessed on the final estimate.

7.1 Prediction from single lepton control regions

The lost lepton estimate is performed with single lepton control region, described in detail in Sec. 5.5.1. It mainly inverts the lepton veto employed in the signal regions, but with the additional requirement that $M_{\mathrm{T}}(\mathrm{cand}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}}) < 100$ GeV, in order to reduce signal contamination. All individual processes (W+jets, $t\bar{t}$, etc.) are summed for the purposes of computing the estimate (i.e. there is only a single transfer factor for all processes combined).

Data vs. MC comparisons of the main analysis kinematic variables in the baseline single lepton control region are shown in Fig. 7.1. In these plots and throughout this chapter, all processes with a top quark ($t\bar{t}$, single top, $t\bar{t}$V) are summed together and shown as a single histogram labeled "Top". As with the dilepton control region, there is some level of disagreement between data and MC, but this is to be expected. The estimate is again primarily data driven, and these disagreements generally do not affect the final estimate. As described below, in high N_{j} regions MC is used to extrapolate along the N_{b} dimension. To account for known $t\bar{t}$ + heavy flavor mis-modeling in MC, we correct the MC with a reweighting procedure described in Sec. 7.1.1.

As with the $Z \to \nu\bar{\nu}$ estimate, the lost lepton estimate is first performed in each $(H_{\mathrm{T}}, N_{\mathrm{j}}, N_{\mathrm{b}})$ topological region, integrated over M_{T2} (for the monojet region, the H_{T} dimension is equivalent to $\vec{p}_{T}^{\,\mathrm{jet1}}$, so there is no integration and the estimate is performed in each analysis bin). For all regions with 7–9 or ≥10 jets and ≥1 b tag, an inclusive control region with ≥7 jets and 1–2 b tags is used, to avoid low statistics and higher signal contamination in regions with high jet and b-jet multiplicity. Similarly, regions with either 7–9 or ≥10 jets and 0 b tags are all predicted using control region bins with ≥7 jets and 0 b tags, due to low control regions statistics in regions with ≥10 jets.

Fig. 7.2 shows N_{j} distributions for data and MC in the ≥7 jet region, for both $N_{\mathrm{b}} = 0$

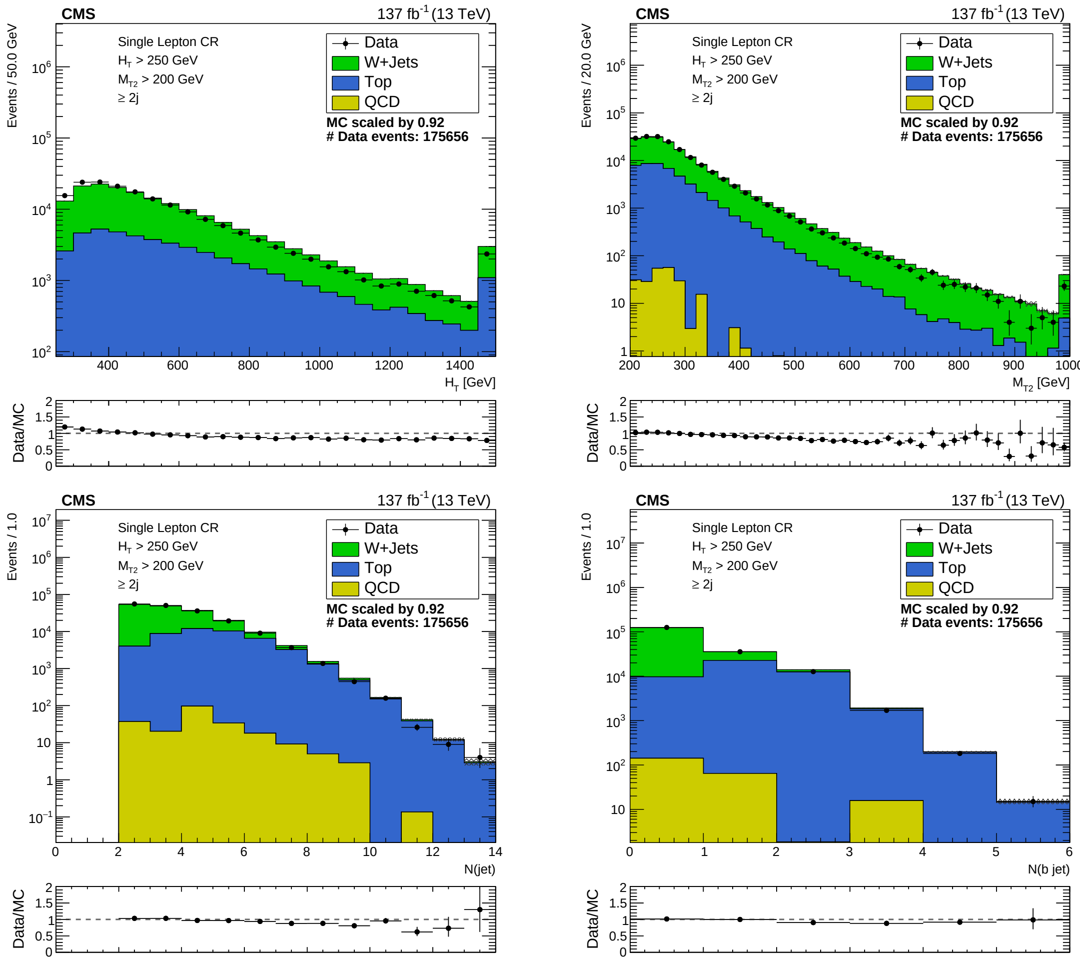

Figure 7.1: Data vs. MC comparisons in the baseline single lepton control region, for $N_{\mathrm{j}} \geq 2$. From left to right, top to bottom, the variables plotted are H_{T}, M_{T2}, N_{j}, and N_{b}.

and $N_{\mathrm{b}} \geq 1$. Good agreement is observed, justifying the use of MC to extrapolate into the $\geq$10 jet regions as described above. For the N_{b} distributions, agreement is not as good, so a reweighting is performed to correct MC as described in the following section.

For regions with $H_{\mathrm{T}} > 1500$ GeV, all events with $M_{\mathrm{T2}} > 200$ GeV are used in the control region even though the signal region starts at $M_{\mathrm{T2}} > 400$ GeV. Once the per-topological region estimate is done, a hybrid approach using both data and MC is used to extrapolate along the M_{T2} dimension, as described in Sec. 6.2.

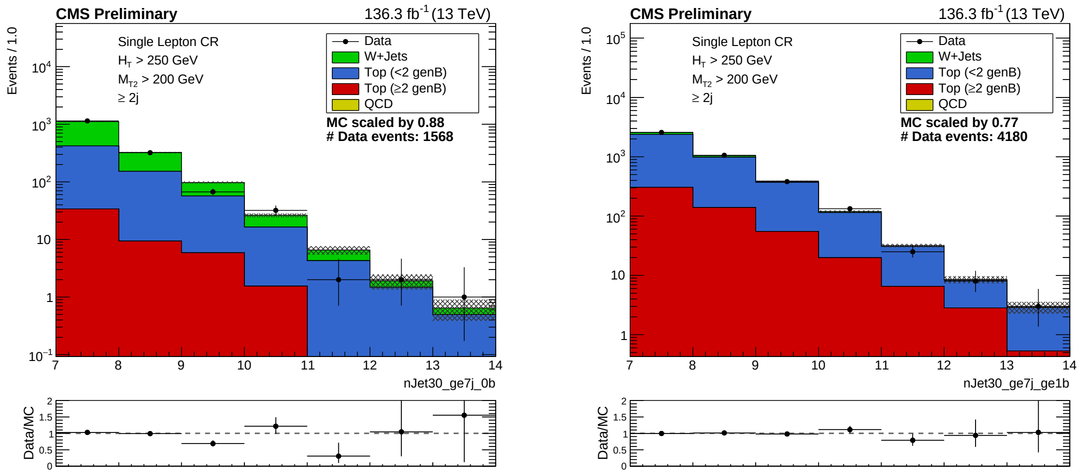

Figure 7.2: Data vs. MC N_j comparisons in the baseline single lepton control region with $N_\mathrm{j} \geq 7$, for $N_\mathrm{b} = 0$ (left) and $N_\mathrm{b} \geq 1$ (right). Agreement is sufficient to justify using MC to extrapolate along the N_j dimension for high-N_j regions.

The final estimate in each $(H_\mathrm{T}, N_\mathrm{j}, N_\mathrm{b}, M_\mathrm{T2})$ signal region can then be summarized as

$$N_\mathrm{LL}^\mathrm{SR} = N_{1\ell}^\mathrm{CR}(M_\mathrm{T2}, \Omega)\, R_\mathrm{MC}^{0\ell/1\ell}(M_\mathrm{T2}, \Omega)\, k_\mathrm{LL}(M_\mathrm{T2}|\Omega), \tag{7.1}$$

where

- $N_{1\ell}^\mathrm{CR}$ is the number of observed events in data in the single lepton control region.
- $R_\mathrm{MC}^{0\ell/1\ell}$ is the ratio between zero lepton and single lepton MC yields in this region.
- $k_\mathrm{LL}(M_\mathrm{T2}|\Omega)$ is a normalized template used to distribute events as a function of M_T2 in each topological region (see Sec. 7.2), only necessary in regions with ≥2 jets.

7.1.1 $t\bar{t}$ + heavy flavor modeling

As described previously, regions with ≥7 jets and ≥1 b tag, are predicted using a single control region with ≥7 jets and 1–2 b tags. MC is then used to extrapolate both into the ≥10 jet region and into the high-N_b regions. While the MC modeling of N_j

is sufficient out-of-the-box (Fig. 7.2), this is not the case for N_b modeling. Fig. 7.3 (left) shows a comparison of N_b in data and MC in the ≥7 jet portion of the single lepton control region. At high N_b, MC under-predicts the number of events in data by a significant amount.

This disagreement can be attributed to known MC mis-modeling of $t\bar{t}$ + heavy flavor events (i.e. $t\bar{t}b\bar{b} + X$). A CMS measurement of the ratio $\sigma(t\bar{t}b\bar{b})/\sigma(t\bar{t}jj)$ finds that this ratio is a factor of 1.7 ± 0.5 higher in data than in MC [69]. Hence, to correct for this we identify $t\bar{t}$ and $t\bar{t}$V MC events with two additional generator-level b jets not from top decay, and weight them by an additional factor of 1.7 ± 0.5. The uncertainty is propagated to the final estimate as a systematic uncertainty. Fig. 7.3 (right) shows the N_b distribution after this reweighting procedure, and one sees that agreement is significantly improved.

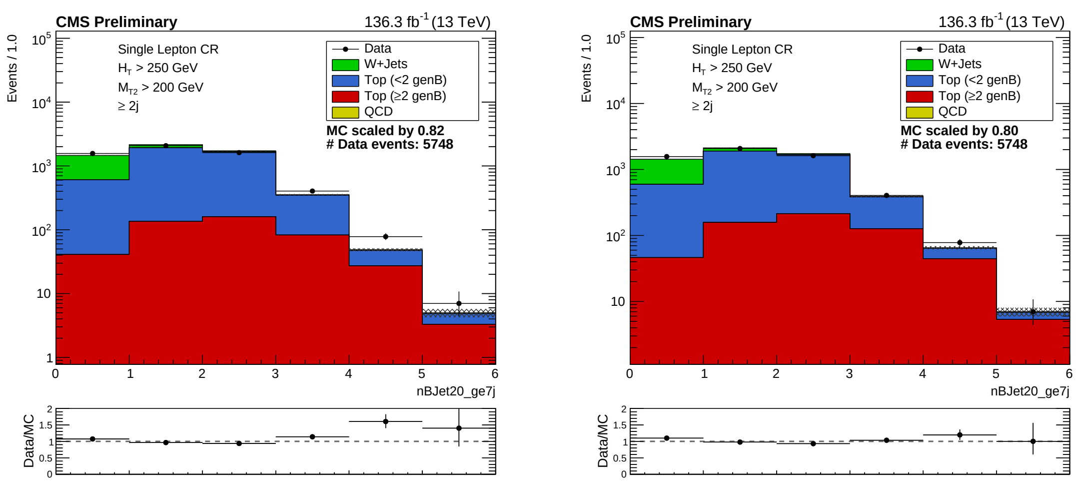

Figure 7.3: Data vs. MC N_b comparisons in the baseline single lepton control region with $N_\mathrm{j} \geq 7$, with no extra weights (left) and with the $t\bar{t}b\bar{b}$ component (red) scaled by 1.7 (right), as described in Sec. 7.1.1. The reweighting significantly improves agreement in the high-N_b bins.

7.1.2 Signal contamination

Despite the selections intended to reduce contributions from signal to the single lepton control regions, signal contamination can be non-negligible in some regions of phase space where the signal is kinematically similar to the background. A contribution from signal to the control region would result in an overestimation of the lost lepton background.

The only signals considered in this analysis which show potential contamination issues are those with prompt lepton decays. Namely, gluino pair production where the gluinos decay to top quarks, and direct top squark pair production where the squarks decay to top quarks. The contributions from other signals are found to be negligible in the control regions. For points near the expected exclusion limits at high masses, the signal contamination is maximally 5% of the expected background yields in the control regions, including the hybrid M_{T2} binning described in the following section. The main place where this effect becomes significant is for top squark pair production points in which the mass difference between the top squark and neutralino is near the top quark mass, such that the signal looks very similar to SM $t\bar{t}$ production.

To account for this in our interpretations, we treat the amount by which the lost lepton background would be overestimated as a reduction in signal efficiency. Specifically, in each analysis bin, we define

$$N_{\mathrm{sig}}^{\mathrm{SR}'} = N_{\mathrm{sig}}^{\mathrm{SR}} - TF \cdot N_{\mathrm{sig}}^{\mathrm{CR}}, \tag{7.2}$$

where $N_{\mathrm{sig}}^{\mathrm{SR}}$ and $N_{\mathrm{sig}}^{\mathrm{CR}}$ are the predicted signal in the signal region and control region bins, respectively, and TF is the transfer factor from control to signal region used in the lost lepton estimate (in the notation of Eq. 7.1, $TF = R_{\mathrm{MC}}^{0\ell/1\ell} \cdot k_{\mathrm{LL}}(M_{\mathrm{T2}}|\Omega)$). Then the quantity $N_{\mathrm{sig}}^{\mathrm{SR}'} \leq N_{\mathrm{sig}}^{\mathrm{SR}}$ is used in calculating the limit on the signal cross section.

This treatment has been used in several other CMS SUSY analyses (e.g. [70]), and has the useful property that $N_{\text{sig}}^{\text{SR}'}$ depends linearly on the signal cross section. This can be seen by rewriting as

$$N_{\text{sig}}^{\text{SR}'} = \sigma_{\text{sig}} \cdot \mathcal{L} \cdot (\varepsilon_{\text{sig}}^{\text{SR}} - TF \cdot \varepsilon_{\text{sig}}^{\text{CR}}), \tag{7.3}$$

where σ_{sig} is the signal cross section, $\mathcal{L}$ is the integrated luminosity, and $\varepsilon_{\text{sig}}^{\text{SR}}$ and $\varepsilon_{\text{sig}}^{\text{CR}}$ are the efficiencies for the signal to populate the signal and control regions, respectively.

7.2 M_{T2} extrapolation

In a similar way as the $Z \to \nu\bar{\nu}$ background estimate, the control region prediction in each topological region is extrapolated along the M_{T2} dimension using a data/MC "hybrid" shape template. Starting from the highest M_{T2} bin in each (H_{T},N_{j},N_{b}) control region, we merge with lower bins until there are at least 50 expected MC events in the merged bin.

Lower M_{T2} bins with large statistics are used directly to predict the backgrounds, so $N_{1\ell}^{\text{CR}}$ and $R_{\text{MC}}^{0\ell/1\ell}$ are measured in the specific M_{T2} bin and $k_{\text{LL}}(M_{\text{T2}}|\Omega)$ is equal to 1. For the higher M_{T2} bins past the merging threshold, the control region is integrated over the merged bin for the purposes of computing $N_{1\ell}^{\text{CR}}$ and $R_{\text{MC}}^{0\ell/1\ell}$. Then k_{LL} is the ratio of the MC signal region yield in an individual M_{T2} bin to that in the integrated bin (i.e., the M_{T2} shape within the merged bin is taken directly from signal region MC).

The MC modeling of M_{T2} is checked in data, in single lepton events with either $N_{\text{b}} = 0$ or $N_{\text{b}} \geq 1$, as shown in the left and right panels of Fig. 7.4. The predicted distributions in the comparison are obtained by summing all the relevant topological regions, after normalizing MC event yields to data and distributing among the M_{T2} bins using the procedure described above. The gray band in the ratio plot represents the systematic

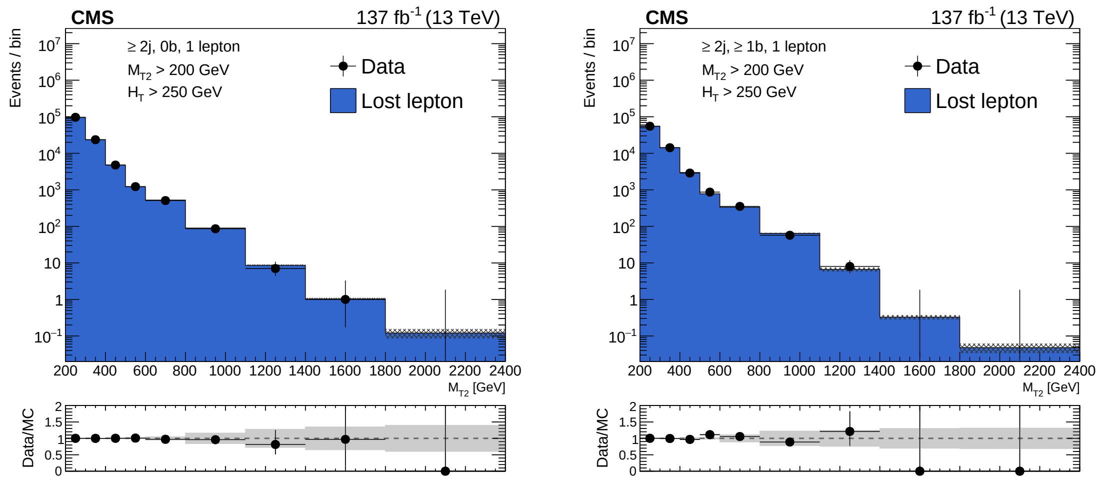

Figure 7.4: Distributions of the M_{T2} variable in data and MC for the single lepton control region, after normalizing the simulation to data in each topological region and distributing events among the M_{T2} bins using the procedure described in Sec. 7.2. On the left is a selection with $N_{\mathrm{b}} = 0$, which probes mostly W+jets events, and on the right is a selection with $N_{\mathrm{b}} \geq 1$, which probes primarily top quark processes. The solid gray band in the ratio plot represents the systematic uncertainty assessed to the predicted M_{T2} shape, as described in Sec. 7.3.

uncertainty assessed to the predicted M_{T2} shape, as described in the following section. Fig. 7.5 shows the same, but separately for each H_{T} region.

7.3 Systematic uncertainties

The following systematics are assessed on the lost lepton background prediction:

- Control region statistical error: the Poisson error on the observed data count in each single lepton control region. This is correlated among all bins that share a common control region (including the M_{T2} bins that utilize a merged control region), but is otherwise uncorrelated.
- MC statistical error: the statistical uncertainty due to the limited size of the MC samples on the $0\ell/1\ell$ transfer factor (and on $k_{\mathrm{LL}}(M_{\mathrm{T2}})$ for bins in which MC is used

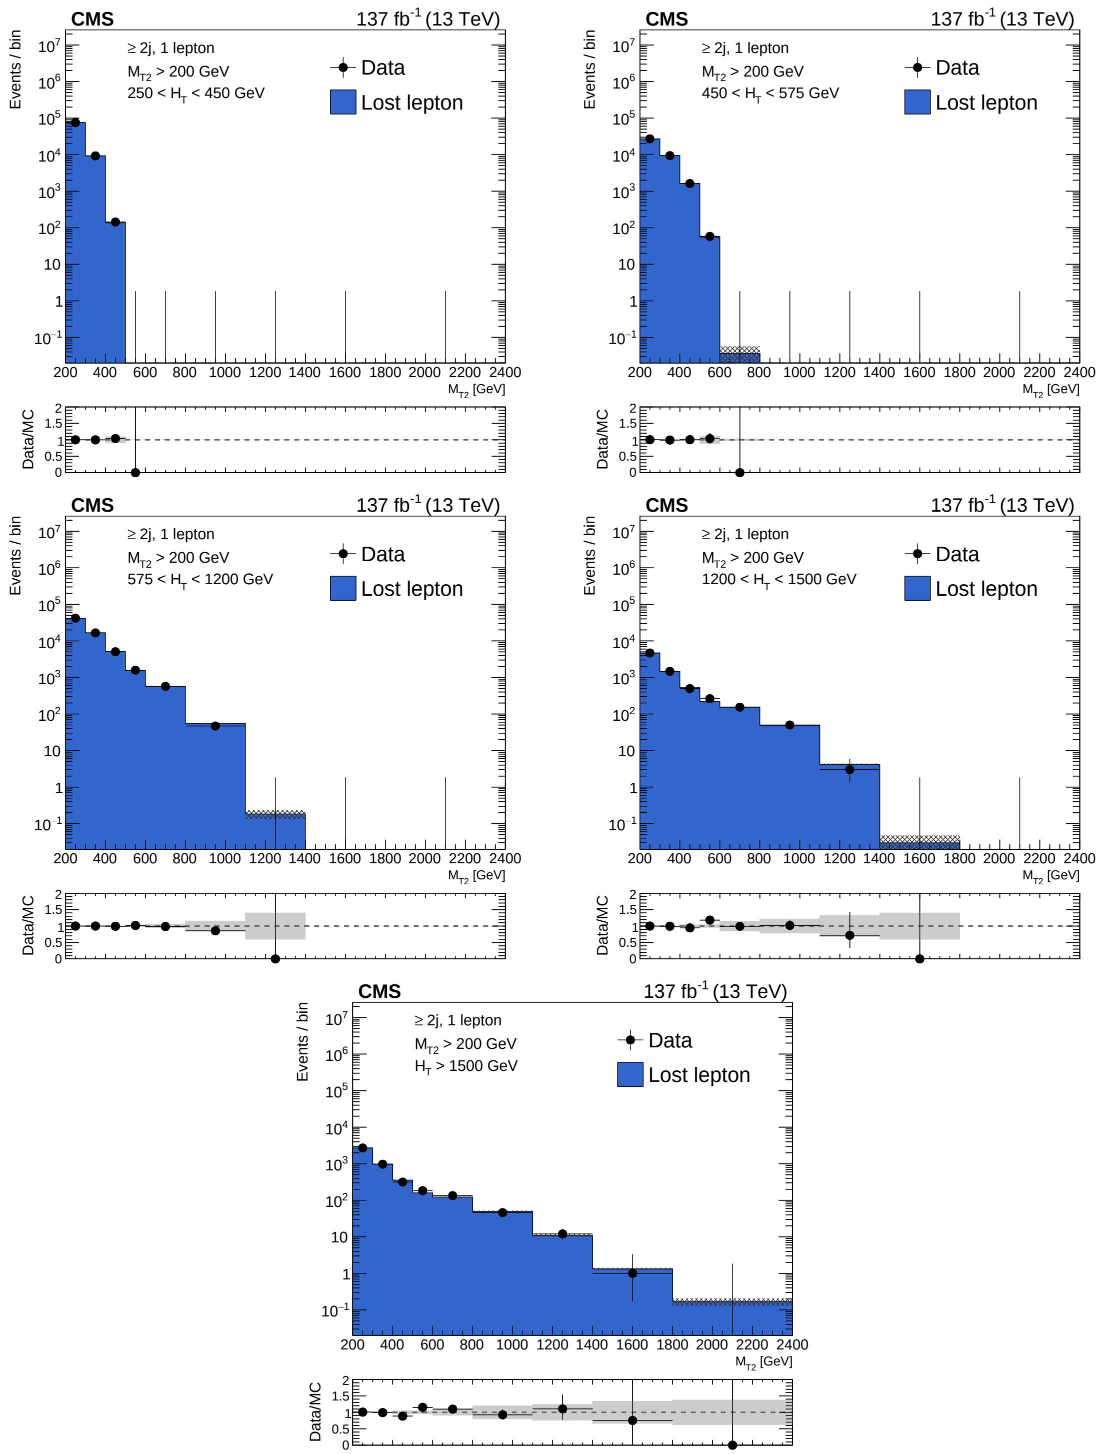

Figure 7.5: The same as Fig. 7.4, but plotted separately for the Very Low, Low, Medium, High, and Extreme H_T regions.

to predict the M_{T2} shape)

- e/μ selection efficiency: to account for differences in lepton efficiency between data and MC, scale factors are derived in bins of p_{T} and η for electrons and muons by the CMS SUSY group, along with corresponding uncertainties. These uncertainties are propagated on a per-event basis to the final predicted yield, and the variation is taken as a systematic uncertainty. This is as large as 7% and correlated among all bins.

- τ selection efficiency: from MC studies, we take a 10% relative uncertainty on the selection efficiency for 1-prong τ's, and a 100% relative uncertainty on the selection efficiency for 3-prong τ's. This is propagated to the final predicted yield and the variation is taken as a systematic uncertainty. This is as large as 3% and correlated among all bins.

- M_{T} cut efficiency: the differences between MC and data of the M_{T} cut used for the control region has been studied in a $Z \to \ell^+\ell^-$ control region in which one lepton is treated as missing (to mimic $W \to \ell\nu$ decay) has been studied. Based on this, we apply a 3% systematic uncertainty, correlated everywhere.

- b tagging efficiency: the CMS b tagging group provides uncertainties on the scale factors used to correct MC b tagging efficiencies. These are propagated to the final estimate, and the variation is taken as a systematic. This is as large as 4%, and correlated among all bins.

- $t\bar{t}b\bar{b}$ reweighting: as described in Sec. 7.1.1, MC events with both $t\bar{t}$ and $b\bar{b}$ pairs are weighted by an additional factor of 1.7 ± 0.5. This uncertainty is propagated to the final estimate, and the variation is taken as a systematic uncertainty. The magnitude of this uncertainty is from 0–25% (largest in bins with very high N_{j} and N_{b}), and correlated everywhere.

- jet energy scale: from MC studies using the jet energy correction uncertainties provided by the CMS JetMET group, we apply a 5% uncertainty, correlated among all bins.
- MC renormalization/factorization scales: derived by varying the generator weight of the MC. This is typically a few percent but up to 10% in some bins, and correlated everywhere.
- M_{T2} shape uncertainty: much like for the $Z \to \nu\bar{\nu}$ estimate, this is an extra uncertainty assigned to the M_{T2} shape, accounting for both theoretical and experimental effects. It is based on MC variation studies and validated using the M_{T2} shape comparisons shown in Figs. 7.4 and 7.5 (solid gray band in ratio plot). The uncertainty is implemented as a linear morphing of the M_{T2} shape, starting from the first M_{T2} bin in each topological region where the MC shape is used, and growing to a maximum of 40% in the last M_{T2} bin. This uncertainty is correlated within a given topological region, but not between different topological regions.

Chapter 8

QCD Multijet Background: The Rebalance and Smear Method

The third and final background of the M_{T2} analysis arises from mis-measured jets in QCD multijet events (and, to a much lesser degree, in events with hadronically-decaying top quarks or vector bosons). This background is greatly suppressed by the M_{T2} and $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ cuts and hence is the smallest of the three backgrounds. However, it is also the most difficult to model and estimate since it depends strongly on the peculiarities of the CMS detector and its imperfect response to jets.

QCD Monte Carlo cannot be relied upon to model this correctly (and statistics are too poor anyway, due to the high cross section and low acceptance), so a data-driven technique is required. This iteration of the analysis employs a new "Rebalance and Smear" method to estimate the multijet background. We briefly describe the old method and reasons for switching, then explain in detail the new technique.

8.1 The $\Delta\phi$-ratio method

Previous iterations of this analysis [67, 68] used the "$\Delta\phi$-ratio" method to estimate QCD background. The method utilizes the variable $\Delta\phi_{\mathrm{min}} \equiv \Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ defined in Sec. 5.1. Events with a badly measured jet tend to have small $\Delta\phi_{\mathrm{min}}$, as a single mis-measured jet drives the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$. Consequently, inverting the $\Delta\phi_{\mathrm{min}} > 0.3$ requirement in the signal regions gives a control region enriched in QCD multijet events.

The $\Delta\phi$-ratio method estimates multijet contribution to the signal region by scaling events in this low-$\Delta\phi_{\mathrm{min}}$ control region by a transfer factor $r_\phi(M_{\mathrm{T2}}) = N(\Delta\phi_{\mathrm{min}} > 0.3)/N(\Delta\phi_{\mathrm{min}} < 0.3)$. Note that the numerator is "signal region-like" ($\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ far from any jets), while the denominator is "background-like" ($\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ close to a jet). From simulation, the functional form of this ratio as a function of M_{T2} is found to be well-described by a power law,

$$r_\phi(M_{\mathrm{T2}}) = \frac{N(\Delta\phi_{\mathrm{min}} > 0.3)}{N(\Delta\phi_{\mathrm{min}} < 0.3)} = a \cdot M_{\mathrm{T2}}^b, \tag{8.1}$$

for sufficiently high M_{T2}. Below $M_{\mathrm{T2}} \approx 60$ GeV, the power law form breaks down as the dominant source of $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ is not from jet mis-measurement (e.g. energy from pileup may contribute more significantly).

This ratio is measured in data in a low-M_{T2} sideband, with an upper bound of $M_{\mathrm{T2}} = 100$ GeV. Above this, the contribution from electroweak processes ($t\bar{t}$+jets and V+jets) is too high relative to QCD to allow an accurate measurement. The lower bound is 60 GeV for $H_{\mathrm{T}} < 1200$ GeV, and 70 GeV for $H_{\mathrm{T}} \geq 1200$ GeV. Data for the measurement is taken from pure-H_{T} triggers (prescaled for $H_{\mathrm{T}} < 1200$ GeV). The fit is done inclusively in N_{j} and N_{b}, and in the same H_{T} bins as the main analysis. Fig. 8.1 shows example fits for 2017 data in the medium and high H_{T} regions.

Each event in the low-$\Delta\phi_{\mathrm{min}}$ control region (integrated across N_{j} and N_{b}) is weighted by $r_\phi(M_{\mathrm{T2}})$ to get its contribution to the signal region. It remains to distribute events among the N_{j} and N_{b} bins. This is done through transfer factors f_j and r_b. The first, f_j is the fraction of events falling into a particular N_{j} bin, and the second, r_b, is the fraction of events in a given N_{j} bin falling into a particular N_{b} bin. It is found through simulation that both f_j and r_b are invariant with respect to M_{T2}, and r_b is invariant with respect to

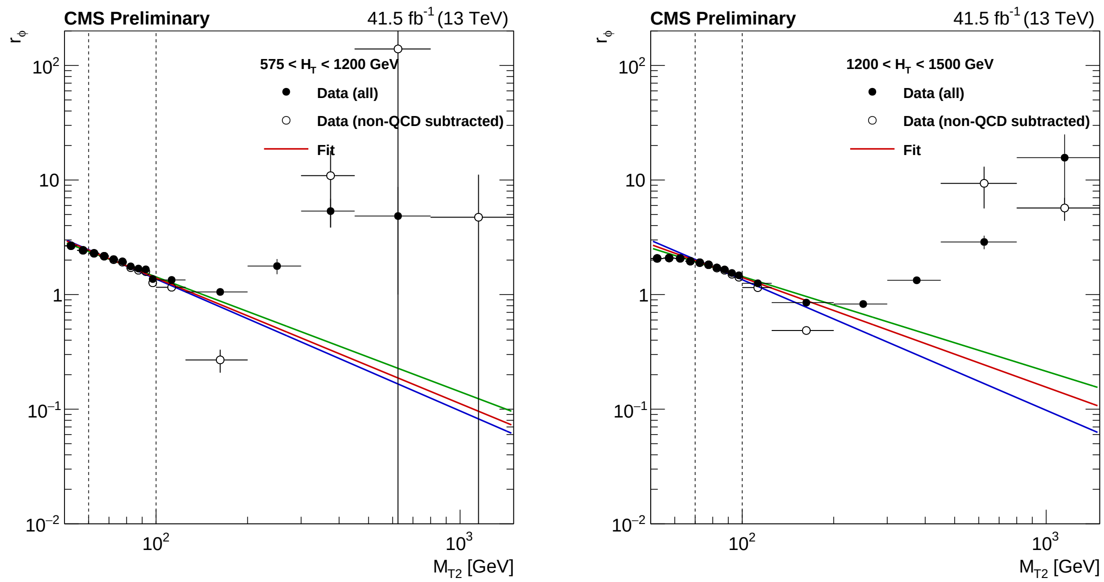

Figure 8.1: Example measurements of $r_\phi(M_{\mathrm{T2}})$ in 2017 data, for the medium and high H_{T} regions. Black points are straight from data; white points have contribution from electroweak MC subtracted off from both the numerator and denominator. Vertical dashed lines show the fit region. The red line is the central fit; the green and blue lines show the variations from extending the fit window one bin on the low or high edge, respectively.

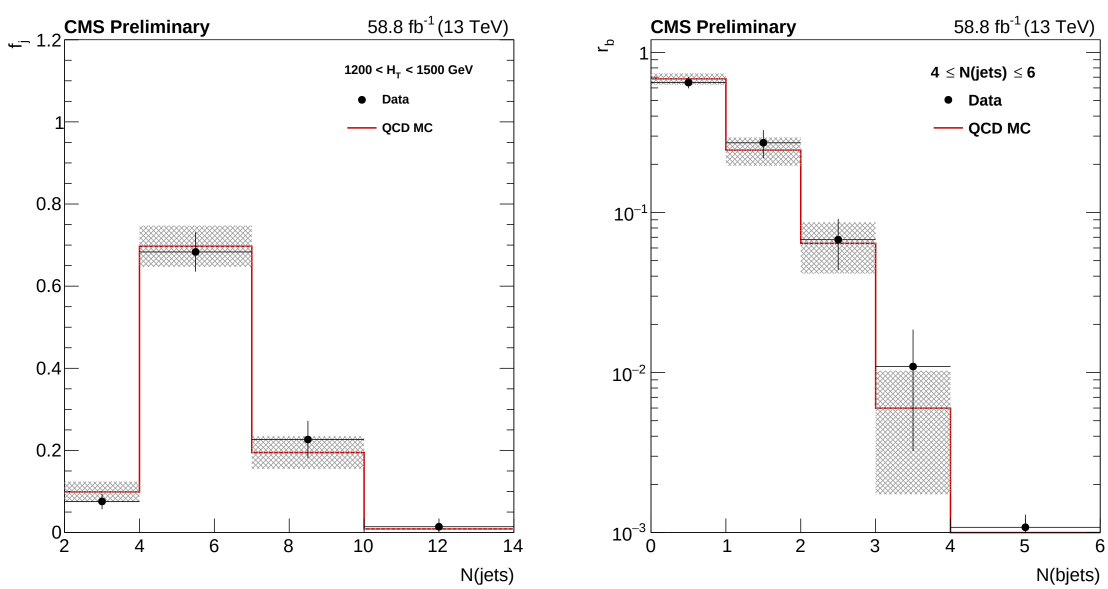

Figure 8.2: Transfer factors f_j (left) and r_b (right) measured in 2018 data. f_j is shown for $1200 \leq H_{\mathrm{T}} < 1500$ GeV, and r_b is shown for $4 \leq N_{\mathrm{j}} \leq 6$. Data agrees with simulation within the uncertainties.

H_T. It is further found that the shapes are equivalent at $\Delta\phi_{\min} < 0.3$ and $\Delta\phi_{\min} > 0.3$. These facts allow us to measure f_j and r_b in a QCD-enriched control region ($\Delta\phi_{\min} < 0.3$, $100 < M_{T2} < 200$ GeV), f_j in bins of H_T and r_b in bins of N_j. Example measurements of f_j and r_b in 2018 data are shown in Fig. 8.2.

Once all three of r_ϕ, f_j, and r_b are measured, we can get a final signal region estimate as

$$N_{SR}(H_T, N_j, N_b, M_{T2}) = \left(\sum_{\Delta\phi_{\min}<0.3} r_\phi(M_{T2}) \right) \cdot f_j(H_T) \cdot r_b(N_j), \tag{8.2}$$

where the sum is taken over all events in a given H_T region, inclusively in N_j and N_b.

8.2 Overview of Rebalance and Smear

While the $\Delta\phi$-ratio method has been used as the primary multijet estimation technique in the past, there are a number of motivations to look for a better, more robust method. Most importantly, the $\Delta\phi$-ratio method relies on a fairly severe extrapolation (r_ϕ is measured in $60 < M_{T2} < 100$ GeV, and is used to predict signal region yields at $M_{T2} > 200$ GeV), with no way to explicitly check its validity in the region of interest. Moreover, the power law fit function itself is empirically derived from simulation, with no underlying theoretical motivation that would give one confidence in its applicability.

An alternate method that is used as the primary multijet estimation technique in this iteration of the analysis is known as Rebalance and Smear (R&S). The method consists of two distinct steps. The first, "rebalancing", seeks to adjust the p_T of jets in multijet events such that the resulting p_T^{miss} is approximately zero, with the aim of reproducing the true hard-scatter event which has no p_T^{miss}. This is performed through a likelihood maximization, accounting for jet energy resolution. The output of the rebalancing is an inclusive sample of multijet events with approximately zero p_T^{miss} that are used as a seed for the second step, the "smearing". In this step, the p_T values of the rebalanced jets are smeared according to jet response functions, in order to model the instrumental effects that lead to nonzero p_T^{miss}. The smearing step is repeated many times for each rebalanced event, which allows the accumulation of events in the tails of kinematic distributions such as p_T^{miss} and M_{T2} and for a more precise estimate of the multijet background in the signal regions.

Both the rebalancing and smearing steps make use of "jet response templates", which are distributions of the ratio of reconstructed jet p_T to generator-level jet p_T. The templates are derived from simulation in bins of jet p_T and η, separately for b-tagged and non-b-tagged jets. Details on the derivation of the templates are given in Sec. 3.4.

For R&S in data, events from pure-H_T triggers are used. Events with $H_T < 1200$ GeV come from prescaled triggers, and get weighted by the corresponding event-level prescale value in the final prediction.

In addition to the trigger selections, events must contain at least one good vertex, two jets with $p_T > 10$ GeV, and pass the standard event cleaning filters in order to be used in the R&S. No other selections are applied.

8.2.1 Rebalancing

The rebalancing procedure adjusts the p_T of jets in an event with the aim of reproducing the true hard-scatter event which has no p_T^{miss}. Note that only the magnitude of the jet p_T is modified, while the jet direction remains unchanged.

Of all jets in the event, a jet qualifies for use in the rebalancing and smearing procedure if it has $p_T > 10$ GeV, and if it is not identified as a jet from pileup in the case that $p_T < 100$ GeV. All other jets are left unchanged but are still used in the calculation of $\vec{p}_T^{miss}$ and other jet-related quantities. An event with n qualifying jets is rebalanced by varying the $p_{\mathrm{T}}^{\mathrm{reb}}$ of each jet to maximize the likelihood function

$$L = \prod_{i=1}^{n} \mathrm{P}(p_{\mathrm{T},i}^{\mathrm{reco}} | p_{\mathrm{T},i}^{\mathrm{reb}}) \times G\left(\frac{p_{\mathrm{T,reb,x}}^{\mathrm{miss}}}{\sigma_{\mathrm{T}}^{\mathrm{soft}}}\right) \times G\left(\frac{p_{\mathrm{T,reb,y}}^{\mathrm{miss}}}{\sigma_{\mathrm{T}}^{\mathrm{soft}}}\right), \tag{8.3}$$

where

$$G(x) \equiv e^{-x^2/2} \tag{8.4}$$

and

$$\vec{p}_{\mathrm{T,reb}}^{\mathrm{miss}} \equiv \vec{p}_{\mathrm{T}}^{\mathrm{miss}} - \sum_{i=1}^{n} \left(\vec{p}_{\mathrm{T},i}^{\mathrm{reb}} - \vec{p}_{\mathrm{T},i}^{\mathrm{reco}}\right). \tag{8.5}$$

The term $\mathrm{P}(p_{\mathrm{T},i}^{\mathrm{reco}} | p_{\mathrm{T},i}^{\mathrm{reb}})$ in Eq. 8.3 is the probability for a jet with p_{T} of $p_{\mathrm{T},i}^{\mathrm{reb}}$ to be assigned a p_{T} of $p_{\mathrm{T},i}^{\mathrm{reco}}$ after reconstruction. This probability is taken directly from the jet response templates. The two $G(x)$ terms in Eq. 8.3 enforce an approximate balancing condition. The $\vec{p}_{\mathrm{T,reb}}^{\mathrm{miss}}$ terms in Eq. 8.3 represent the missing transverse momentum after rebalancing, and are obtained by simply propagating to $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ the changes in jet p_{T} from rebalancing. For the balancing of the x and y components of the missing transverse momentum, we use $\sigma_{\mathrm{T}}^{\mathrm{soft}} = 20$ GeV, which is approximately the width of the distributions of the x and y components of $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ in minimum bias events. This parameter represents the inherent missing energy due to low-p_{T} jets, unclustered energy, and jets from pileup that cannot be eliminated by rebalancing. A systematic uncertainty is assessed to cover for the effects of the variation of $\sigma_{\mathrm{T}}^{\mathrm{soft}}$.

In practice, the likelihood maximization is done by minimizing $-\log L$ using MINUIT [71] via ROOT. The minimization is done by finding the n parameters $c_1, \ldots, c_n$ such that $p_{\mathrm{T},i}^{\mathrm{reb}} \equiv \frac{1}{c_i} p_{\mathrm{T},i}^{\mathrm{reco}}$ minimize $-\log L$. To calculate $P(p_{\mathrm{T},i}^{\mathrm{reco}} | p_{\mathrm{T},i}^{\mathrm{reb}})$ we look at the response template for jets with $p_{\mathrm{T}} = \frac{1}{c_i} p_{\mathrm{T},i}^{\mathrm{reco}}$ and $\eta = \eta_i$ and find the probability of c_i. The rebalanced event will have jets with p_{T} scaled by their corresponding $\frac{1}{c_i}$.

8.2.2 Smearing

Once a sample of rebalanced events has been obtained the next step is to smear the jets in these events many times. Each rebalanced event is smeared ($100 \times$ prescale) times for data, or just 100 times for MC. The number of smears is capped at 5000, and a corresponding extra weight is applied to events where $100 \times$ prescale > 5000. For each smearing, the p_{T} of each jet in the rebalanced event is scaled by a random factor drawn from the corresponding jet response template. If an event contains jets that were not considered in the rebalancing procedure (i.e. they have $p_{\mathrm{T}} < 10$ GeV, or have $p_{\mathrm{T}} < 100$ GeV and fail the pileup jet ID) then those jets remain in the event without any smearing.

After the smearing has been done, all jet-related quantities (H_{T}, $p_{\mathrm{T}}^{\mathrm{miss}}$, M_{T2}, $\Delta\phi_{\mathrm{min}}$, etc.) are recalculated and analysis selections are applied. Histograms are filled for each smeared event that passes the analysis selections with a weight of $1/N_{\mathrm{smears}}$. The R&S predictions for kinematic distributions and event yields are taken from these histograms. For the purposes of calculating statistical uncertainty, multiple smears of the same input event are taken as fully correlated when filling histograms. For example, if an event is smeared 100 times, and 3 of those smeared events populate a given histogram bin, then that bin is filled once with a weight of 3/100, rather than three separate times with a weight of 1/100.

8.3 Performance in Monte Carlo

Figures 8.3–8.5 show kinematic distributions from QCD Monte Carlo and R&S based on the same QCD Monte Carlo after a loose selection of $H_{\mathrm{T}} > 1200$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV, and $M_{\mathrm{T2}} > 50$ GeV. Figures 8.6–8.8 show the same distributions but with $450 < H_{\mathrm{T}} < 1200$ GeV. The R&S method models the shapes of these distributions quite well. There is an overall normalization difference of $\mathcal{O}$(few percent) introduced by the $p_{\mathrm{T}}^{\mathrm{miss}}$ and M_{T2} cuts due to differences in modeling of very low $p_{\mathrm{T}}^{\mathrm{miss}}$ events.

Figure 8.9 shows Monte Carlo closure in the topological regions after the baseline selection. Statistics are quite low in the straight MC, but yields agree reasonably well where comparison is possible.

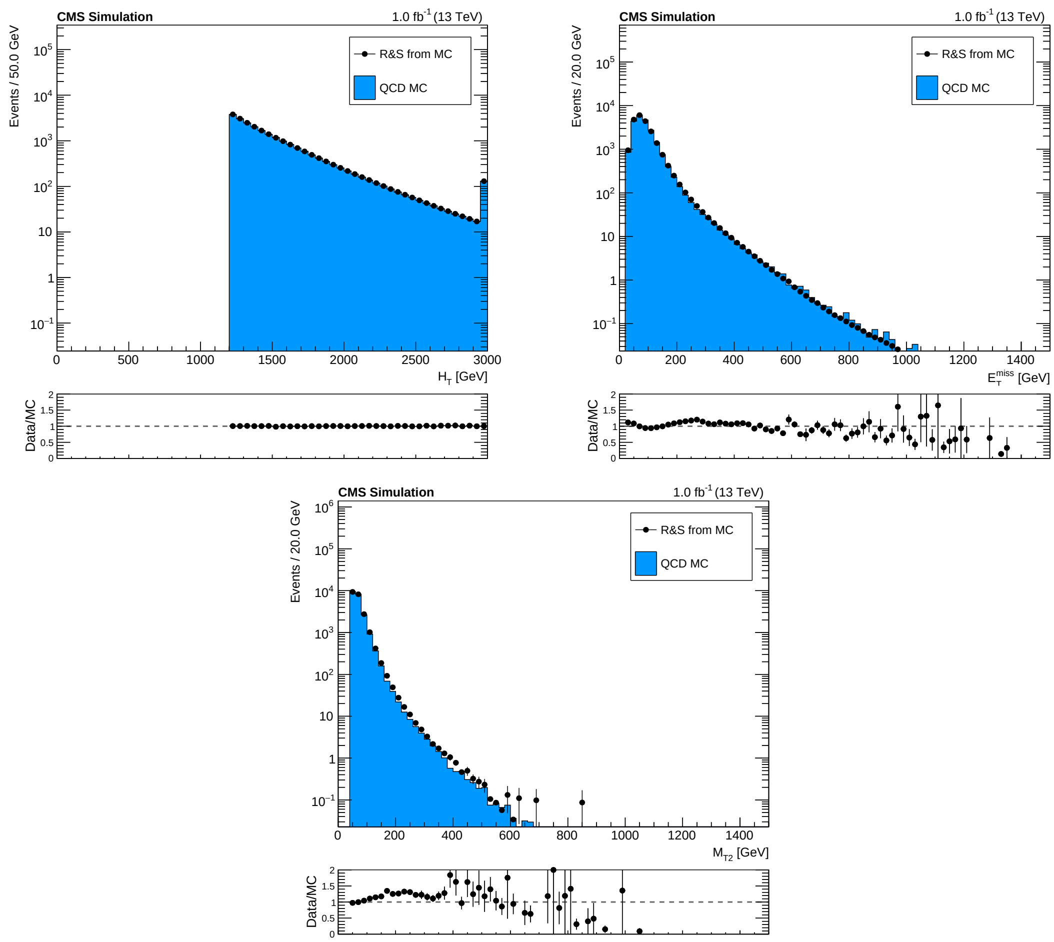

Figure 8.3: H_{T}, $p_{\mathrm{T}}^{\mathrm{miss}}$, and M_{T2} distributions for Monte Carlo and R&S based on MC. The selection is $H_{\mathrm{T}} > 1200$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV and $M_{\mathrm{T2}} > 50$ GeV.

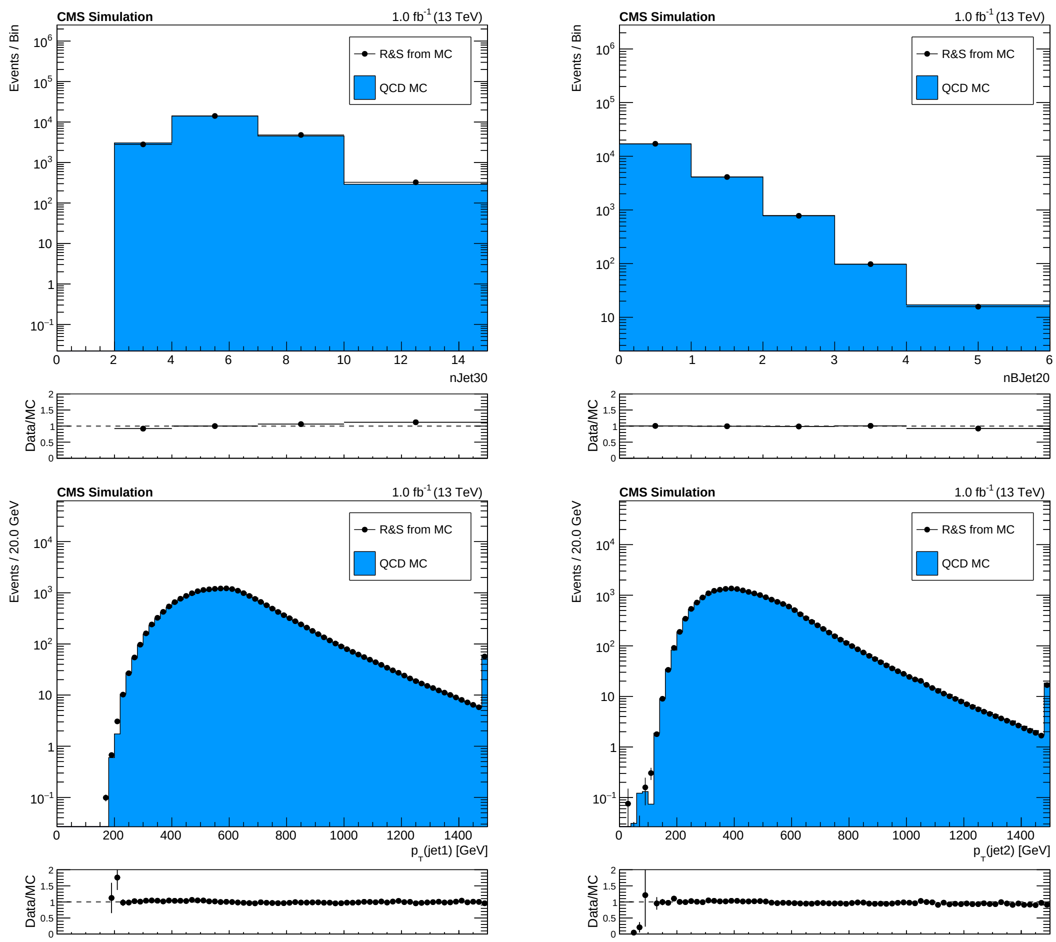

Figure 8.4: N_j, N_b and leading- and subleading-jet p_T distributions for Monte Carlo and R&S based on MC. The selection is $H_\mathrm{T} > 1200$ GeV, $p_\mathrm{T}^\mathrm{miss} > 30$ GeV and $M_\mathrm{T2} > 50$ GeV.

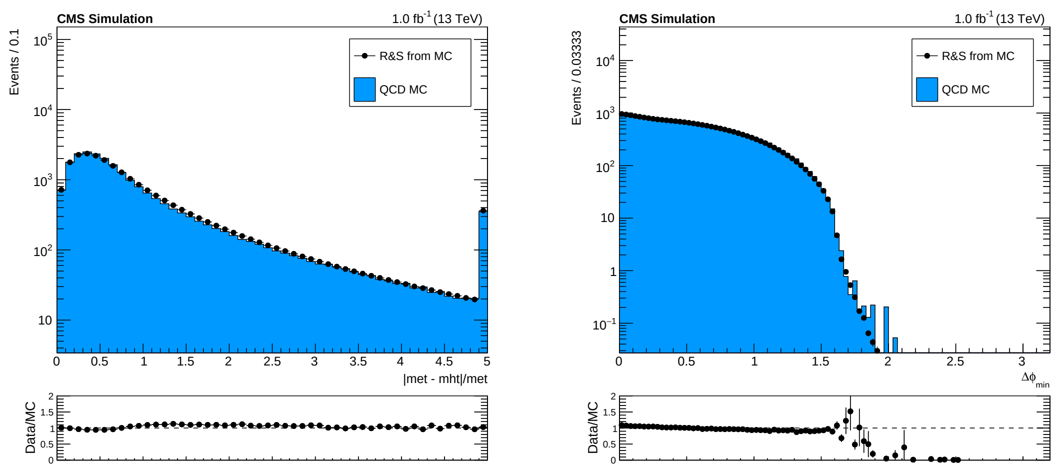

Figure 8.5: $|\vec{H}_{\mathrm{T}}^{\mathrm{miss}} - \vec{p}_{\mathrm{T}}^{\mathrm{miss}}|/p_{\mathrm{T}}^{\mathrm{miss}}$ and $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ distributions for Monte Carlo and R&S based on MC. The selection is $H_{\mathrm{T}} > 1200$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV and $M_{\mathrm{T2}} > 50$ GeV.

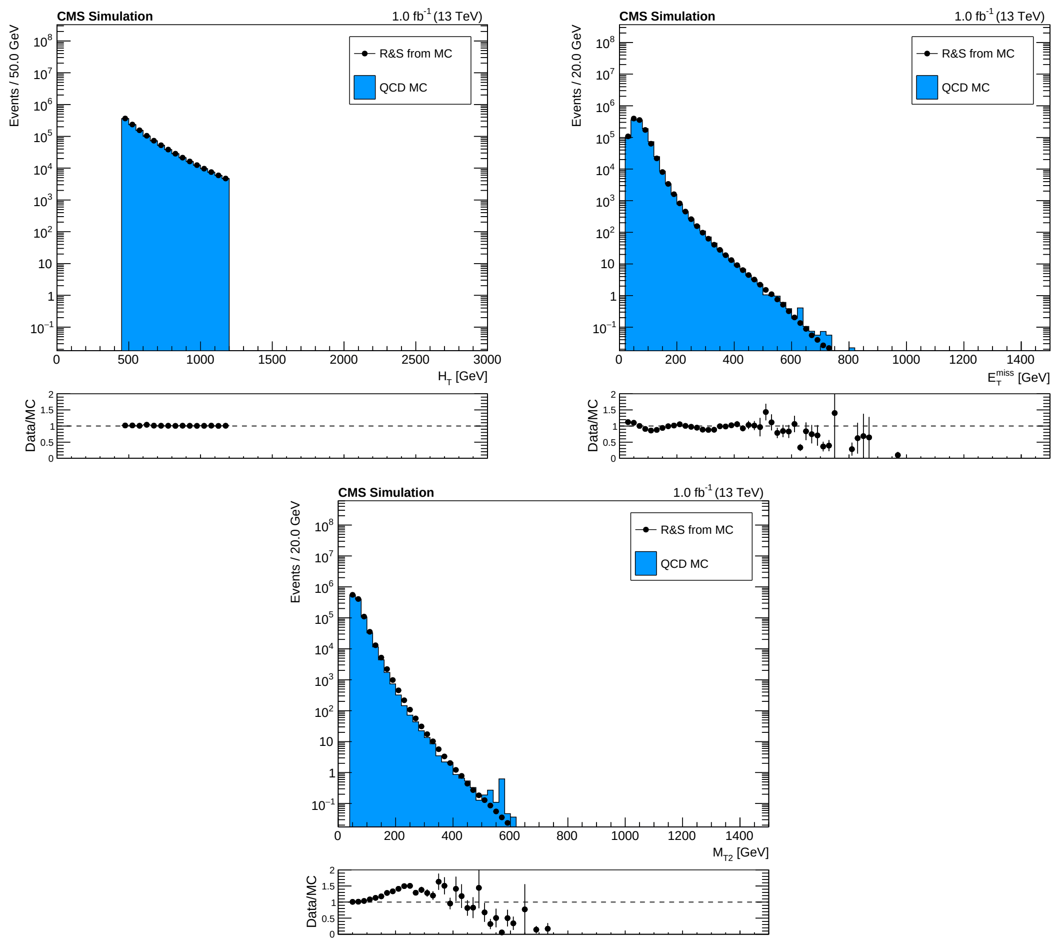

Figure 8.6: H_{T}, $p_{\mathrm{T}}^{\mathrm{miss}}$, and M_{T2} distributions for Monte Carlo and R&S based on MC. The selection is $450 < H_{\mathrm{T}} < 1200$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV and $M_{\mathrm{T2}} > 50$ GeV.

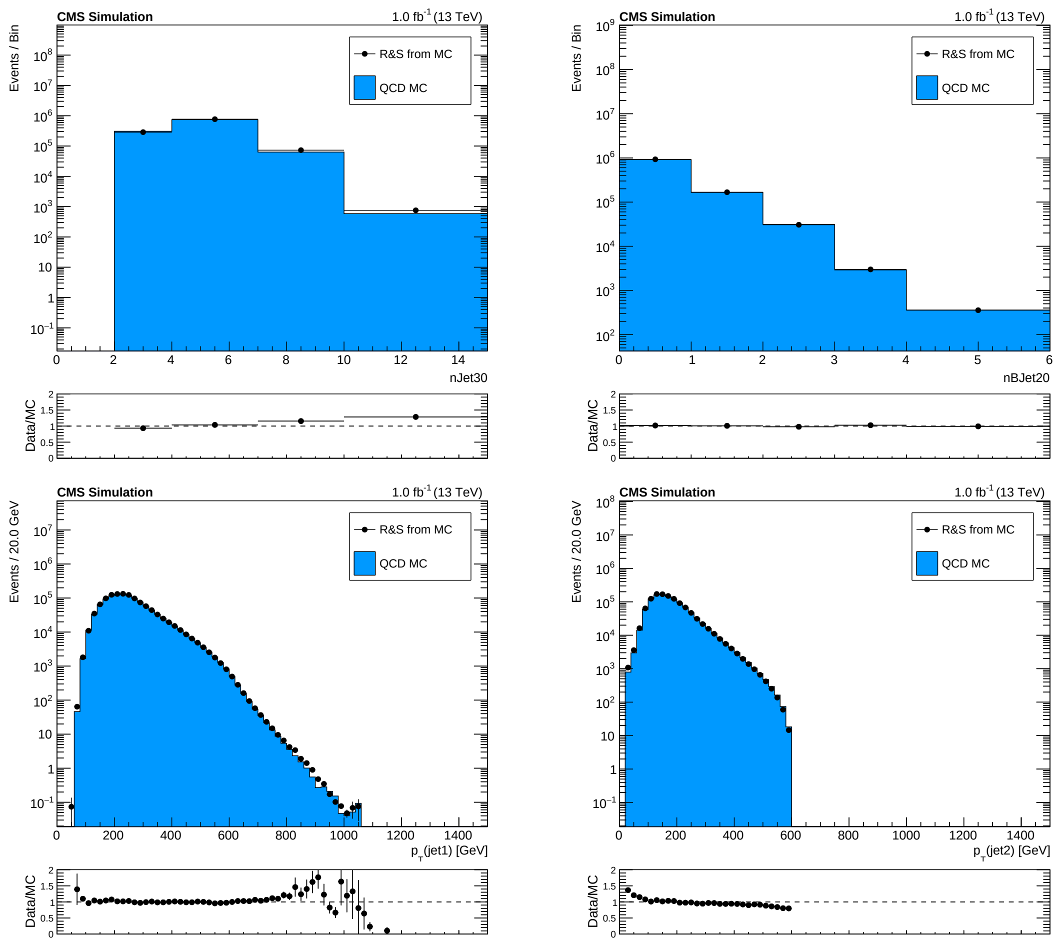

Figure 8.7: N_j, N_b and leading- and subleading-jet p_T distributions for Monte Carlo and R&S based on MC. The selection is $450 < H_\mathrm{T} < 1200$ GeV, $p_\mathrm{T}^\mathrm{miss} > 30$ GeV and $M_\mathrm{T2} > 50$ GeV.

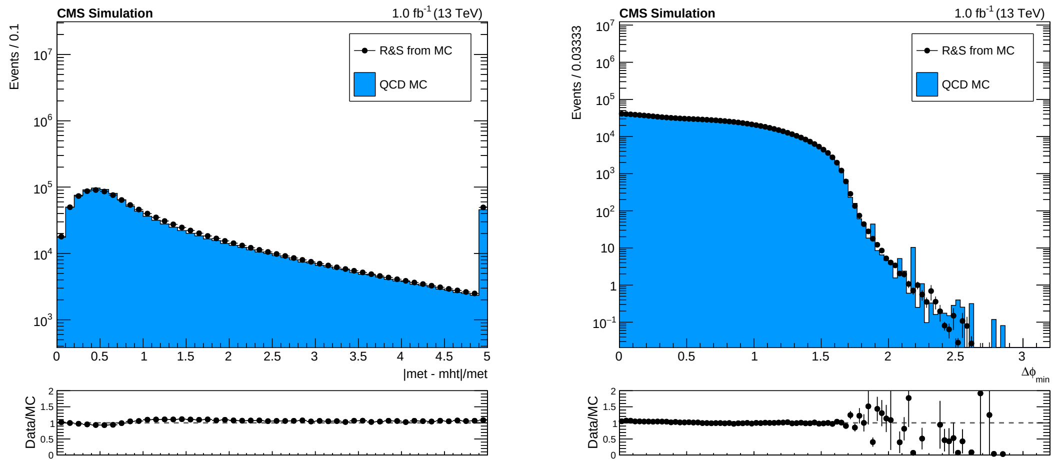

Figure 8.8: $|\vec{H}_{\mathrm{T}}^{\mathrm{miss}} - \vec{p}_{\mathrm{T}}^{\mathrm{miss}}|/p_{\mathrm{T}}^{\mathrm{miss}}$ and $\Delta\phi(\mathrm{j}_{1234}, \vec{p}_{\mathrm{T}}^{\mathrm{miss}})$ distributions for Monte Carlo and R&S based on MC. The selection is $450 < H_{\mathrm{T}} < 1200$ GeV, $p_{\mathrm{T}}^{\mathrm{miss}} > 30$ GeV and $M_{\mathrm{T2}} > 50$ GeV.

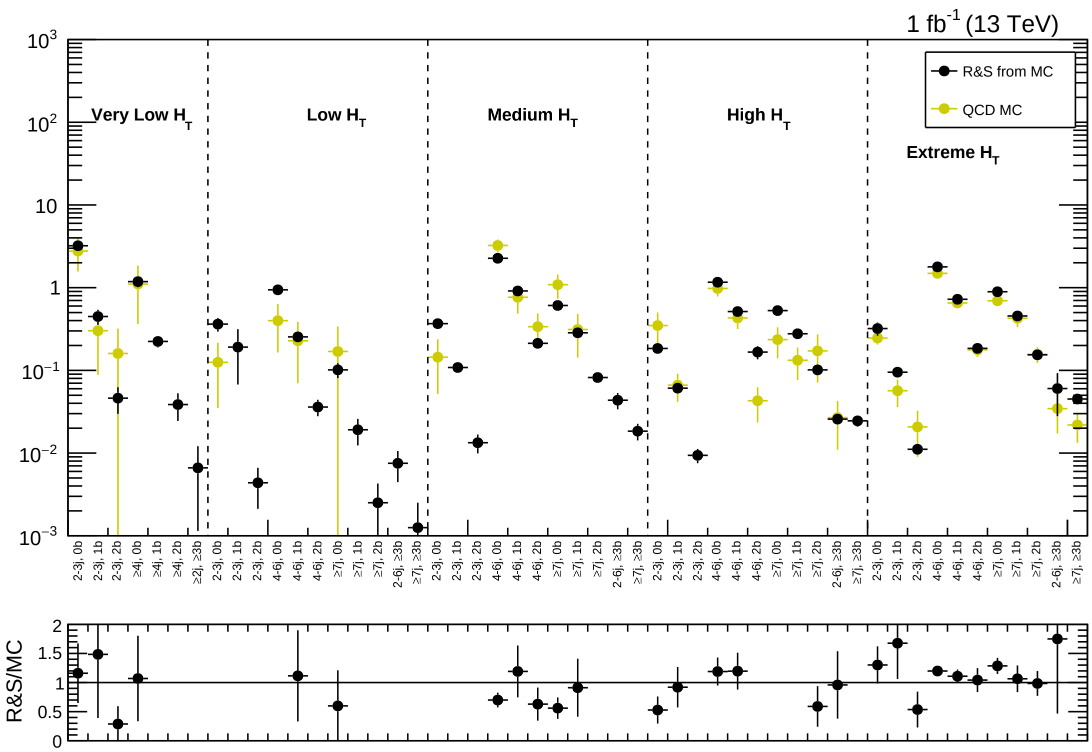

Figure 8.9: R&S Monte Carlo closure in topological regions after the baseline signal region selection. The bottom histogram shows the ratio of yields in each region.

8.4 Electroweak contamination

The input to the rebalancing step in data comes from pure-H_{T} triggers with no attempt to remove any possible contamination from non-QCD processes. Most electroweak events are rebalanced to have $p_{\mathrm{T}}^{\mathrm{miss}}$ close to zero just like actual QCD events, and contribute an extremely small amount to the final prediction since the cross section for electroweak processes is much smaller than the QCD cross section. However, some configurations of electroweak events prove difficult to rebalance, such as events with $p_{\mathrm{T}}^{\mathrm{miss}}$ in one hemisphere and all jets in the other hemisphere. An example of one such Monte Carlo event is shown in Figure 8.10 The $p_{\mathrm{T}}^{\mathrm{miss}}$ in these events is reduced in the rebalancing step but can still be rather large. When the $p_{\mathrm{T}}^{\mathrm{miss}}$ after rebalancing is large, almost all smeared events will also have large $p_{\mathrm{T}}^{\mathrm{miss}}$ and will therefore contribute to the final prediction much more than if the smeared $p_{\mathrm{T}}^{\mathrm{miss}}$ was actually a product of sampling the tails of the jet response templates.

In order to remove contamination to the R&S prediction from electroweak events that are difficult to rebalance, we require the $p_{\mathrm{T}}^{\mathrm{miss}}$ after rebalancing to be less than 100 GeV. Figures 8.11 and 8.12 show the rebalanced $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution for smeared QCD and electroweak MC events that enter the signal regions. The electroweak events are the sum of events from $Z \to \nu\bar{\nu}$, W+jets, and $t\bar{t}$ Monte Carlo. From these distributions we can determine the effect of requiring $p_{\mathrm{T}}^{\mathrm{miss}} < 100$ GeV after rebalancing. We compare the QCD yield integrated over all rebalanced $p_{\mathrm{T}}^{\mathrm{miss}}$ to the sum of the QCD and electroweak yields with rebalanced $p_{\mathrm{T}}^{\mathrm{miss}} < 100$ GeV and take a scale factor to correct for the difference. This scale factor is found to be 0.98 or 0.99 in all H_{T} regions, so the effect is tiny. Table 8.1 summarizes the computation of this scale factor for each H_{T} region.

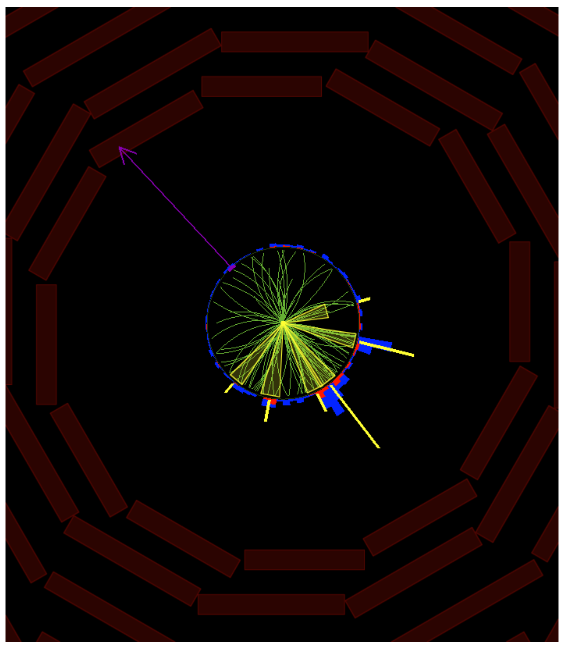

Figure 8.10: Example of a $Z \to \nu\bar{\nu}$ event in Monte Carlo with a configuration that leads to large $p_{\mathrm{T}}^{\mathrm{miss}}$ after rebalancing. This event has 517 GeV of $p_{\mathrm{T}}^{\mathrm{miss}}$ before rebalancing and 311 GeV of $p_{\mathrm{T}}^{\mathrm{miss}}$ after rebalancing.

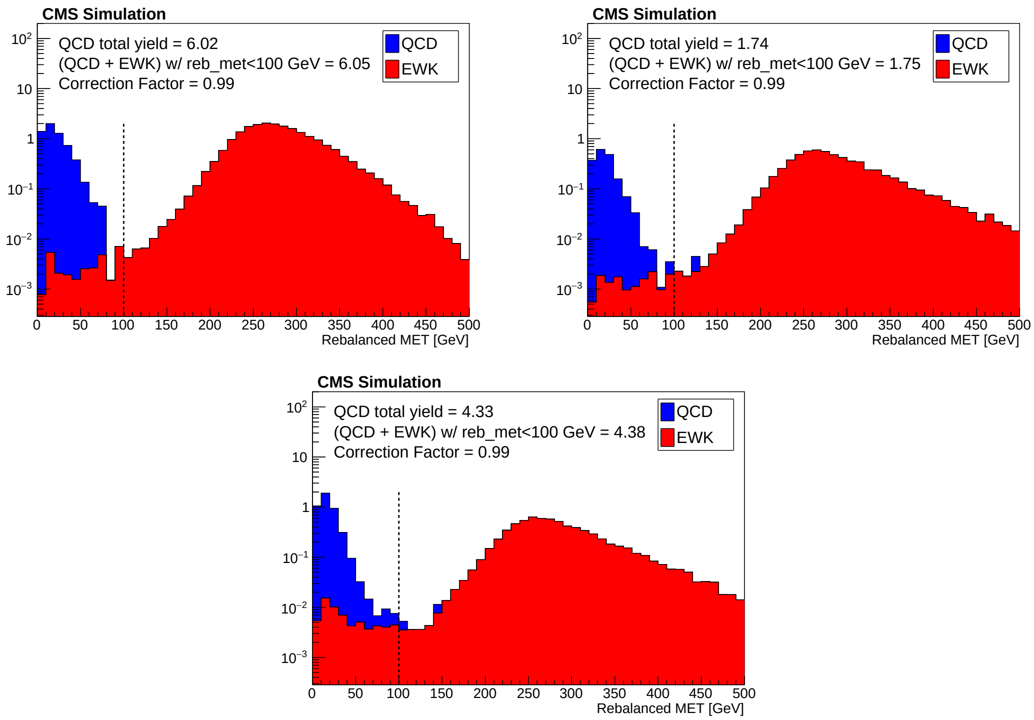

Figure 8.11: Rebalanced $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution for QCD and electroweak smeared events in the very low, low, and medium H_{T} regions after the baseline selection.

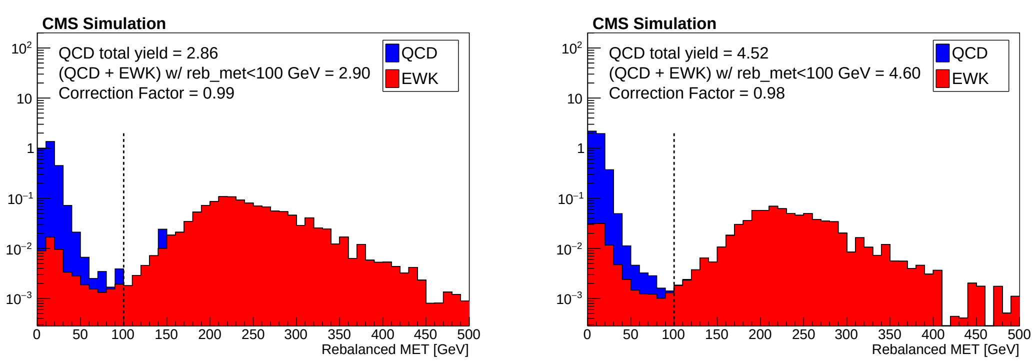

Figure 8.12: Rebalanced $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution for QCD and electroweak smeared events in the high and extreme H_{T} regions after the baseline selection.

Table 8.1: Derivation of scale factors to correct for loss of QCD events due to the rebalanced $p_{\mathrm{T}}^{\mathrm{miss}} < 100$ GeV requirement. The scale factor is found to be 1.00 in all H_{T} regions.

	Very Low H_{T}	Low H_{T}	Med H_{T}	High H_{T}	Ext H_{T}
QCD total yield	6.02	1.74	4.33	2.86	4.52
QCD + EWK, reb $p_{\mathrm{T}}^{\mathrm{miss}} < 100$ GeV	6.05	1.75	4.38	2.90	4.60
Correction Factor	0.99	0.99	0.99	0.99	0.98

8.5 Performance in data control regions

In order to gauge the performance of the R&S method in data we define three control regions that are orthogonal to the search regions and enriched in QCD events. The first control region is obtained from the baseline selection by inverting the $\Delta\phi_{\min}$ cut, requiring $\Delta\phi_{\min} < 0.3$. The second control region is the M_{T2} sideband $100 < M_{\mathrm{T2}} < 200$ GeV. The third control region is defined by both inverting the $\Delta\phi_{\min}$ selection and selecting the M_{T2} sideband. Figures 8.13–8.16 show several kinematic distributions in the control regions for $450 < H_{\mathrm{T}} < 1200$ GeV and $H_{\mathrm{T}} > 1200$ GeV separately. Non-QCD background contributions in these control regions are taken from Monte Carlo. Selecting just the M_{T2} sideband or just inverting $\Delta\phi_{\min}$ for $450 < H_{\mathrm{T}} < 1200$ GeV is not enough to make QCD a significant fraction of the total background in this H_{T} region, so these plots are not shown.

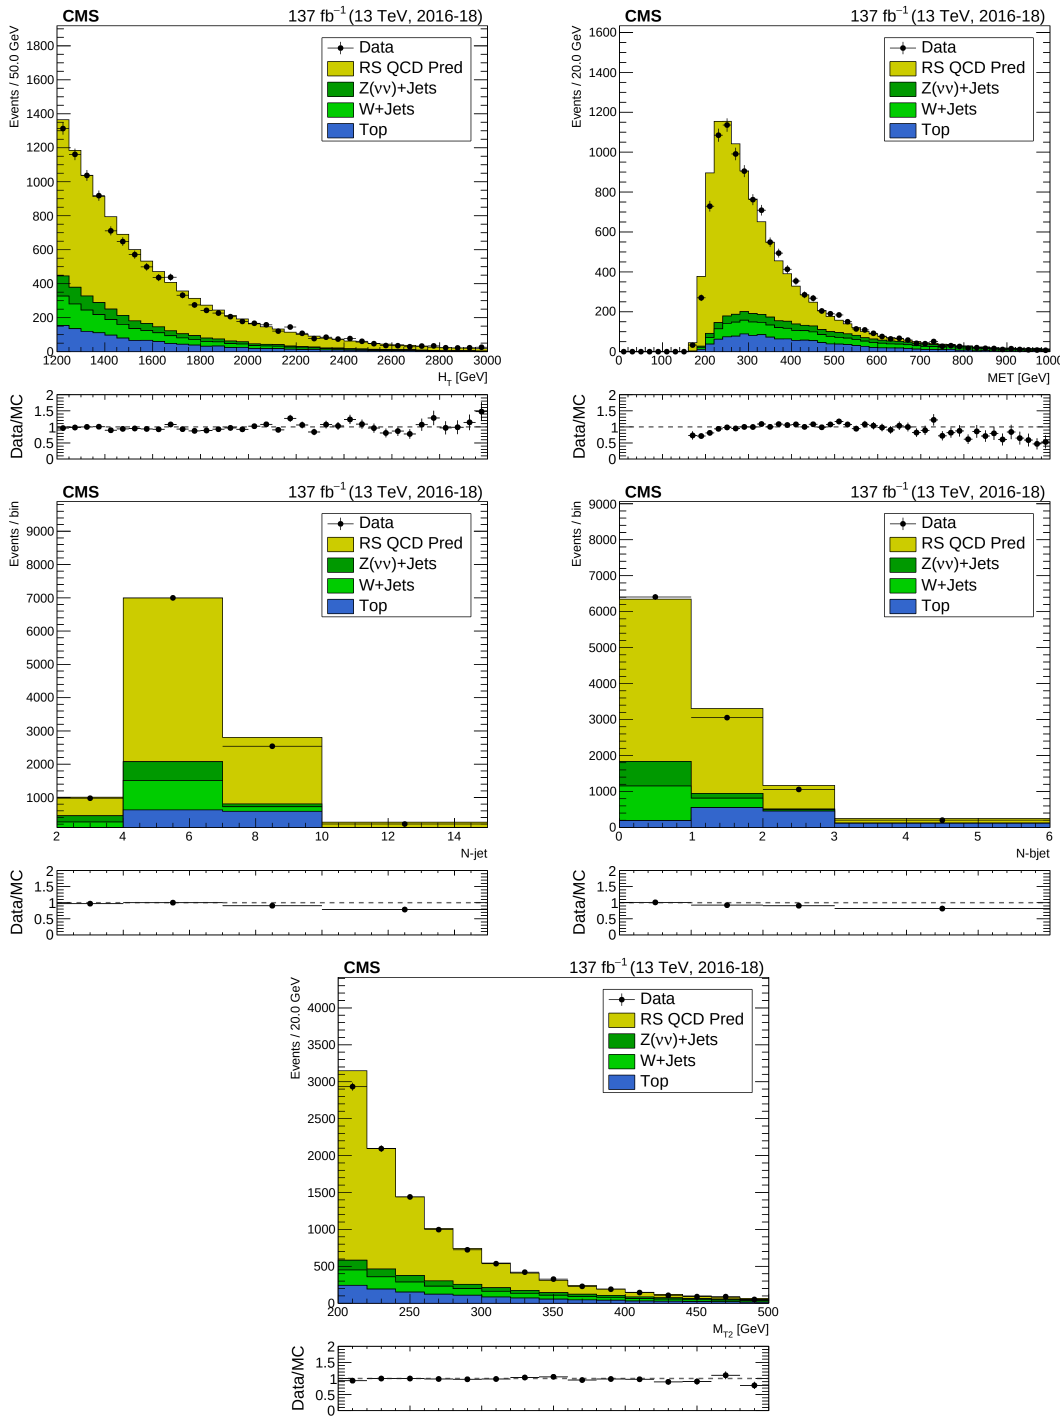

Figure 8.13: Comparison of kinematic distributions for data and background in the inverted $\Delta\phi_{\min}$ control region for $H_T > 1200$ GeV. The QCD background is from the rebalance and smear data-driven prediction. Non-QCD backgrounds are taken from Monte Carlo.

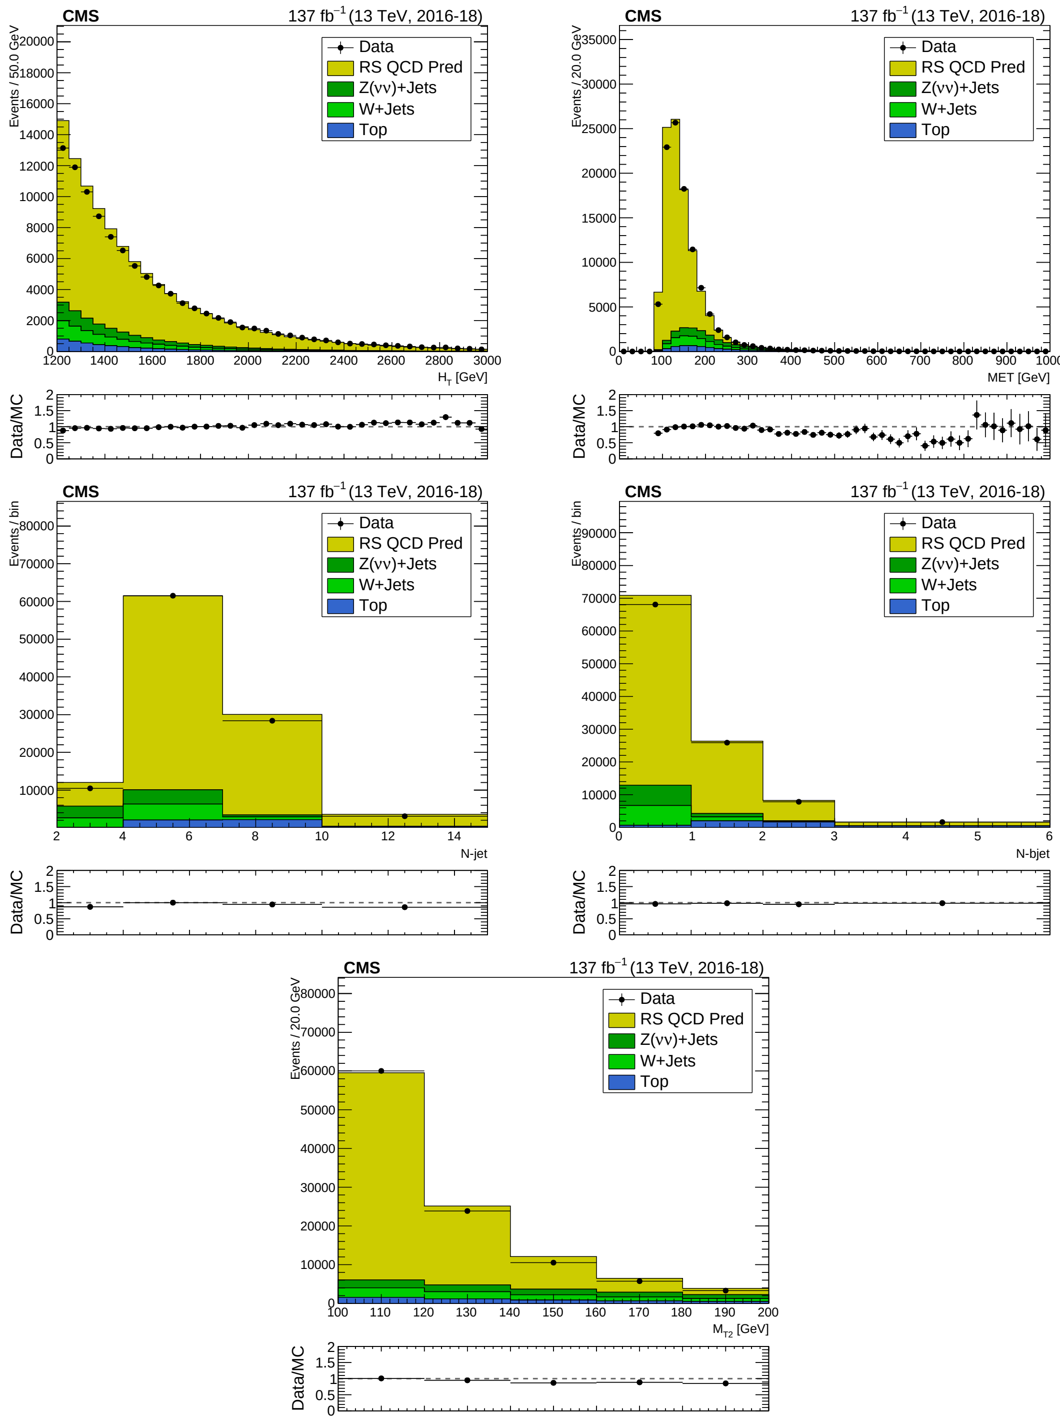

Figure 8.14: Comparison of kinematic distributions for data and background in the M_{T2} sideband control region ($100 < M_{\mathrm{T2}} < 200$ GeV) for $H_{\mathrm{T}} > 1200$ GeV. The QCD background is from the rebalance and smear data-driven prediction. Non-QCD backgrounds are taken from Monte Carlo.

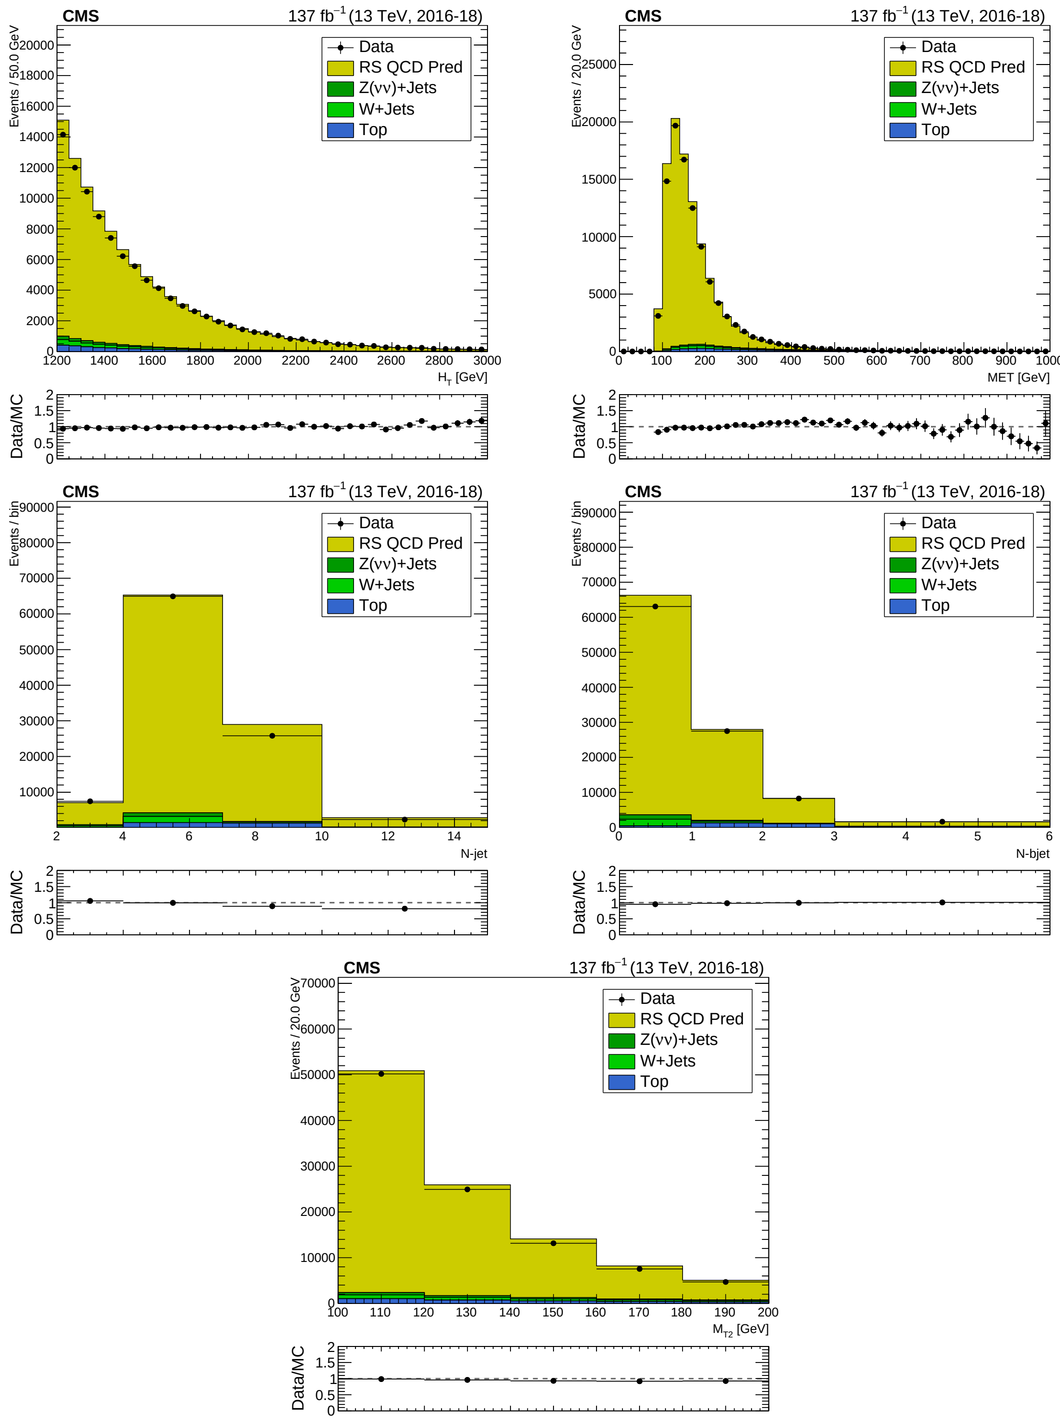

Figure 8.15: Comparison of kinematic distributions for data and background in the $M_{\rm T2}$ sideband + inverted $\Delta\phi_{\rm min}$ control region for $H_{\rm T} > 1200$ GeV. The QCD background is from the rebalance and smear data-driven prediction. Non-QCD backgrounds are taken from Monte Carlo.

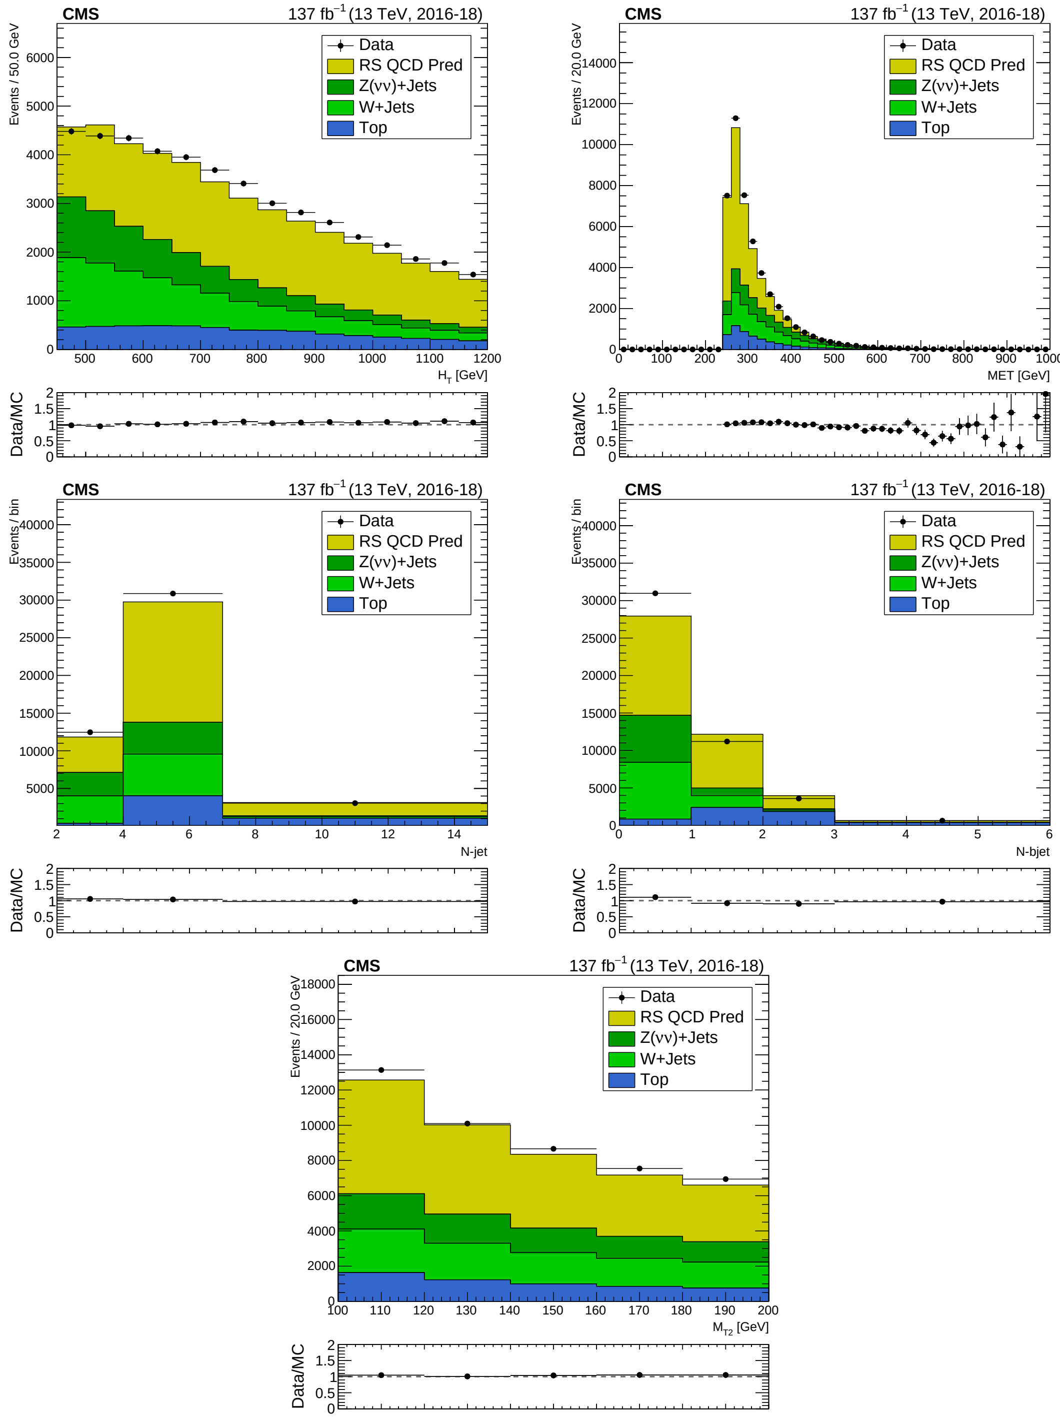

Figure 8.16: Comparison of kinematic distributions for data and background in the M_{T2} sideband + inverted $\Delta\phi_{\mathrm{min}}$ control region for $450 < H_{\mathrm{T}} < 1200$ GeV. The QCD background is from the rebalance and smear data-driven prediction. Non-QCD backgrounds are taken from Monte Carlo.

We can also look at total yields within each of the analysis topological regions in these 3 control regions. These are shown in Figures 8.17–8.19. In these plots, the electroweak background is data-driven using the same methods as in the main analysis. Where the available statistics do not permit such a data-driven estimate, the electroweak contribution is taken directly from MC. The gray bands represent the statistical error on the prediction combined with the systematics we assign, as discussed in Sec. 8.7. In all cases, data agrees with prediction within the assigned error.

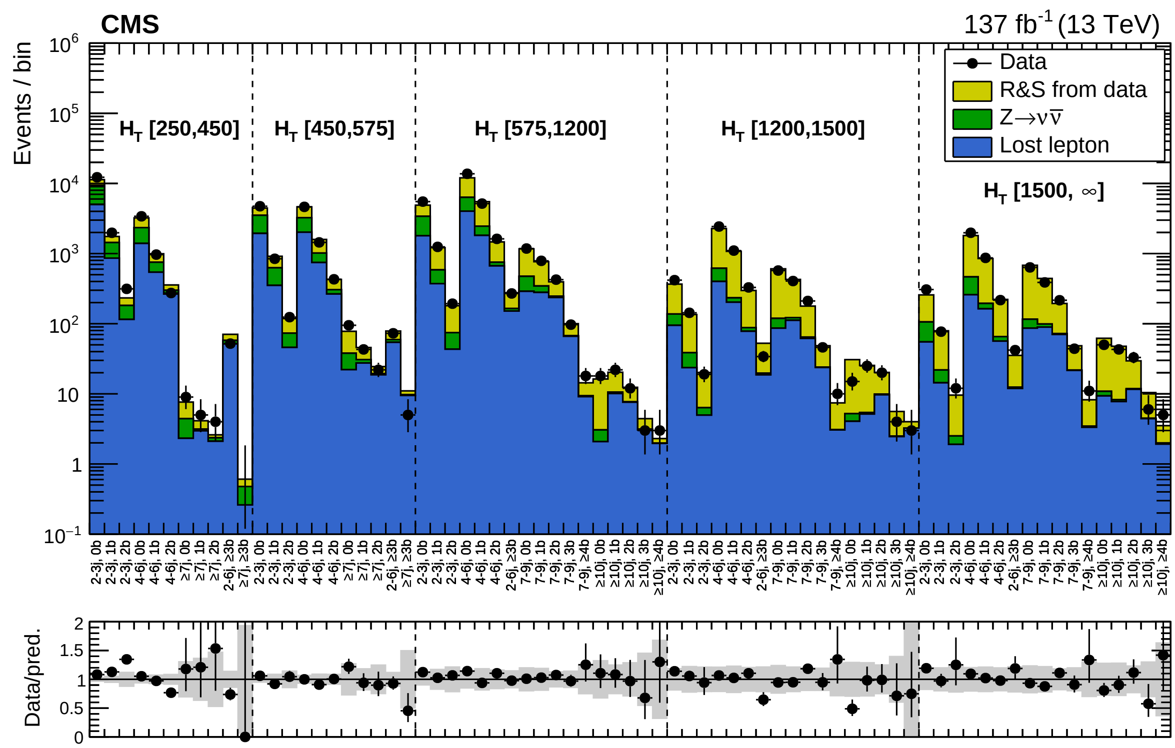

Figure 8.17: Data closure in the inverted-$\Delta\phi_{\mathrm{min}}$ control region. Electroweak backgrounds are data-driven, and the gray band in the ratio plot represents the statistical error on the prediction plus the systematic error we assign.

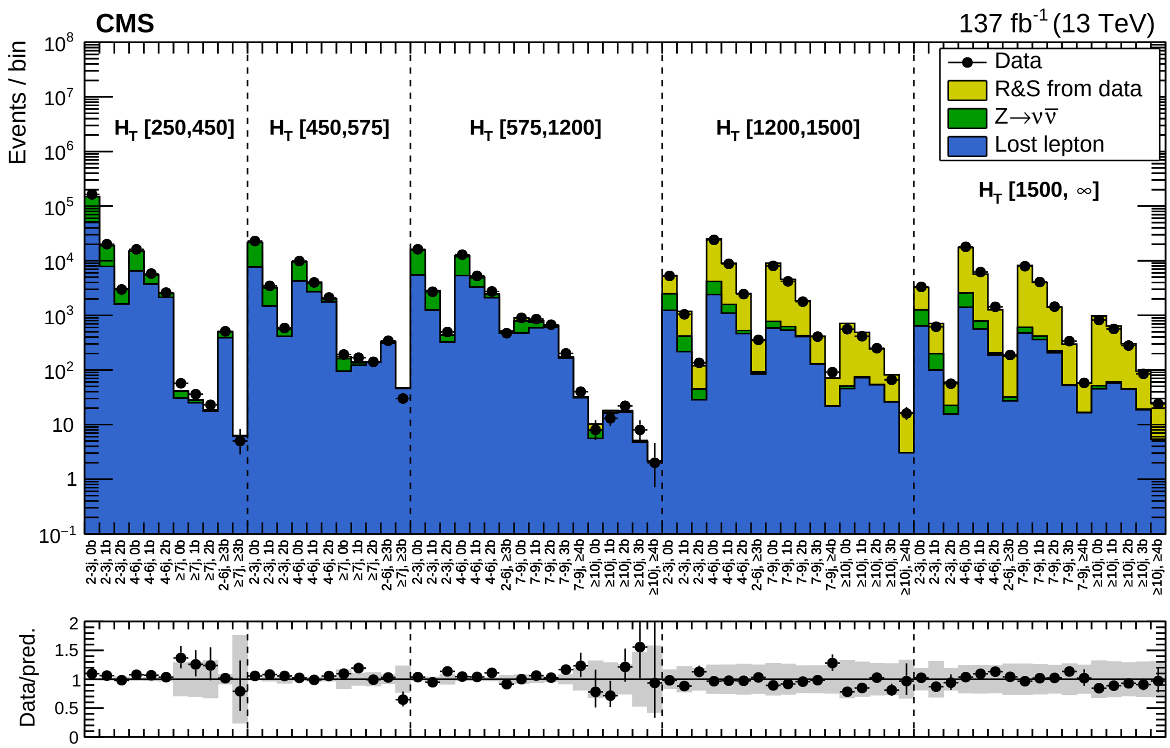

Figure 8.18: Data closure in the M_{T2} Sideband control region. Electroweak backgrounds are data-driven, and the gray band in the ratio plot represents the statistical error on the prediction plus the systematic error we assign.

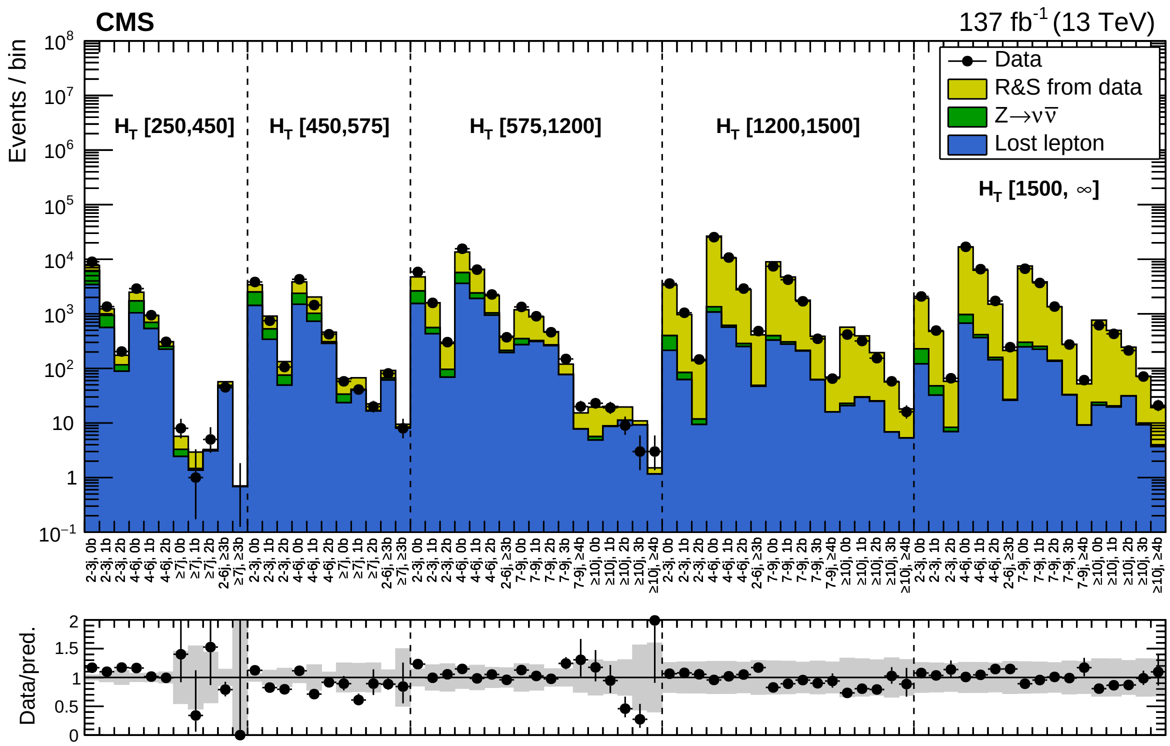

Figure 8.19: Data closure in the inverted-$\Delta\phi_{\mathrm{min}}$ plus M_{T2} Sideband control region. Electroweak backgrounds are data-driven, and the gray band in the ratio plot represents the statistical error on the prediction plus the systematic error we assign.

8.6 Extension to monojet regions

Rebalance and Smear is also used for estimating background from QCD events in the monojet signal regions. The methodology is exactly the same as for the multijet case. The procedure is validated in an inverted-$\Delta\phi_{\mathrm{min}}$ dijet control region, with exactly 2 jets, $H_{\mathrm{T}} > 250$ GeV and $p_{\mathrm{T}}^{\mathrm{miss}} > 250$ GeV. Fig. 8.20 shows QCD MC compared against the prediction from Rebalance and Smear applied to MC in this control region, both inclusively and in the $30 < p_{\mathrm{T}}(\mathrm{jet2}) < 60$ GeV sideband nearest the monojet signal region. Good agreement is seen in both cases.

Fig. 8.21 shows closure between data and the Rebalance and Smear method applied to data in the same dijet control region. Electroweak backgrounds are from MC. Again,

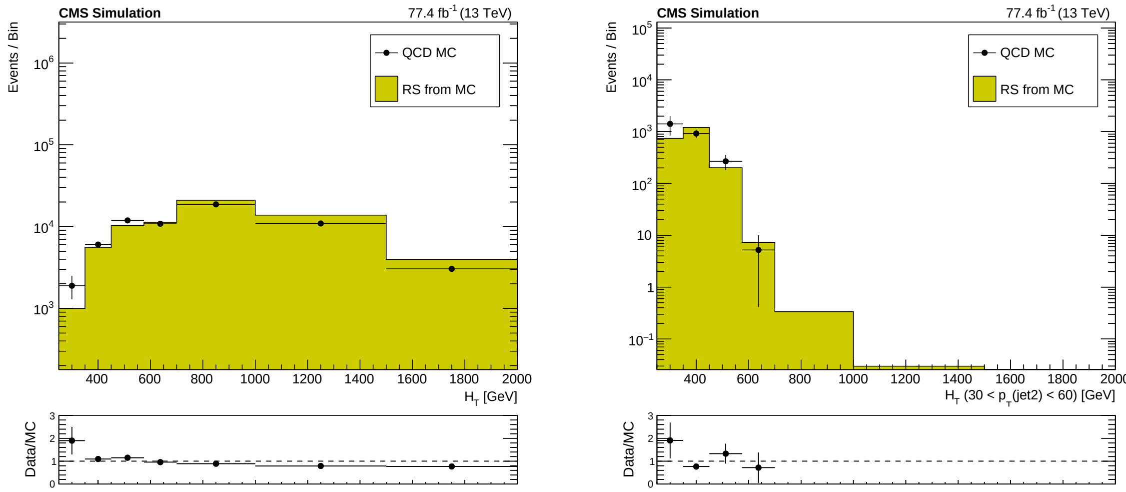

Figure 8.20: Closure between QCD MC and the prediction from Rebalance and Smear, in an inverted-$\Delta\phi_{\mathrm{min}}$ dijet control region. The left plot is inclusive in jet p_{T}, while the right is in a $30 < p_{\mathrm{T}}(\mathrm{jet2}) < 60$ GeV sideband that is adjacent to the monojet signal region.

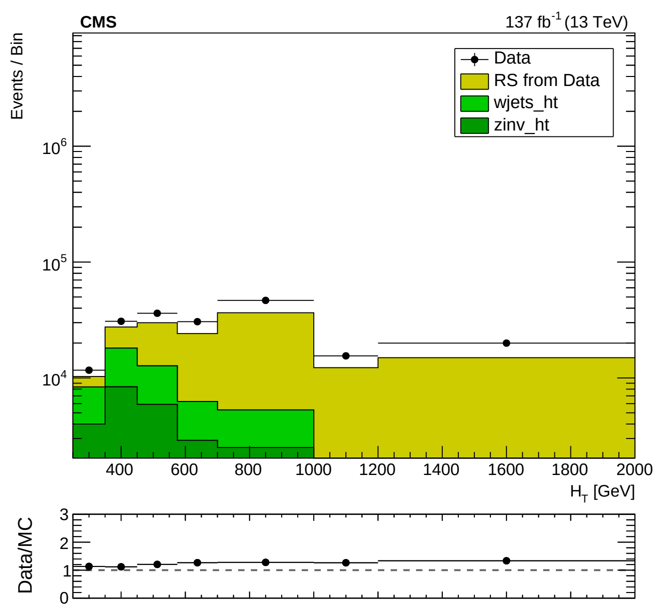

Figure 8.21: Closure between data and the prediction from Rebalance and Smear, in an inverted-$\Delta\phi_{\mathrm{min}}$ dijet control region. Electroweak backgrounds are from MC.

good agreement is seen.

8.7 Systematic uncertainties

Sections 8.7.1 through 8.7.4 study the effects of varying parameters in the jet response templates or the rebalancing procedure on the final estimate. In Sec. 8.7.5, we use these to assess systematic uncertainties on the final estimate.

8.7.1 Effect of modifying width of jet response core

Here we study in both MC and data how the rebalance and smear prediction is affected by increasing the width of the Gaussian core component of the jet response templates. The core of the response is identified by fitting a Gaussian to a narrow window around the peak response, as described in Sec. 3.4.2.

In MC, we test the effect by increasing the width of the Gaussian core by a uniform factor for all jet response templates. The procedure for performing this widening is again described in Sec. 3.4.2. Figure 8.22 shows the effect of increasing the width of the Gaussian core by 10% and 25% on the rebalance and smear predictions from MC in the analysis topological regions.

We also study this in data by varying the width of the core by the jet energy resolution (JER) smear factor uncertainties derived by the JetMET group, listed in Table 3.1. The result of varying the width up/down by this amount is shown in Figure 8.23. Variations in each H_{T}-region are shown in colored text near the top of the plot. These are used to derive a systematic in Sec. 8.7.5.

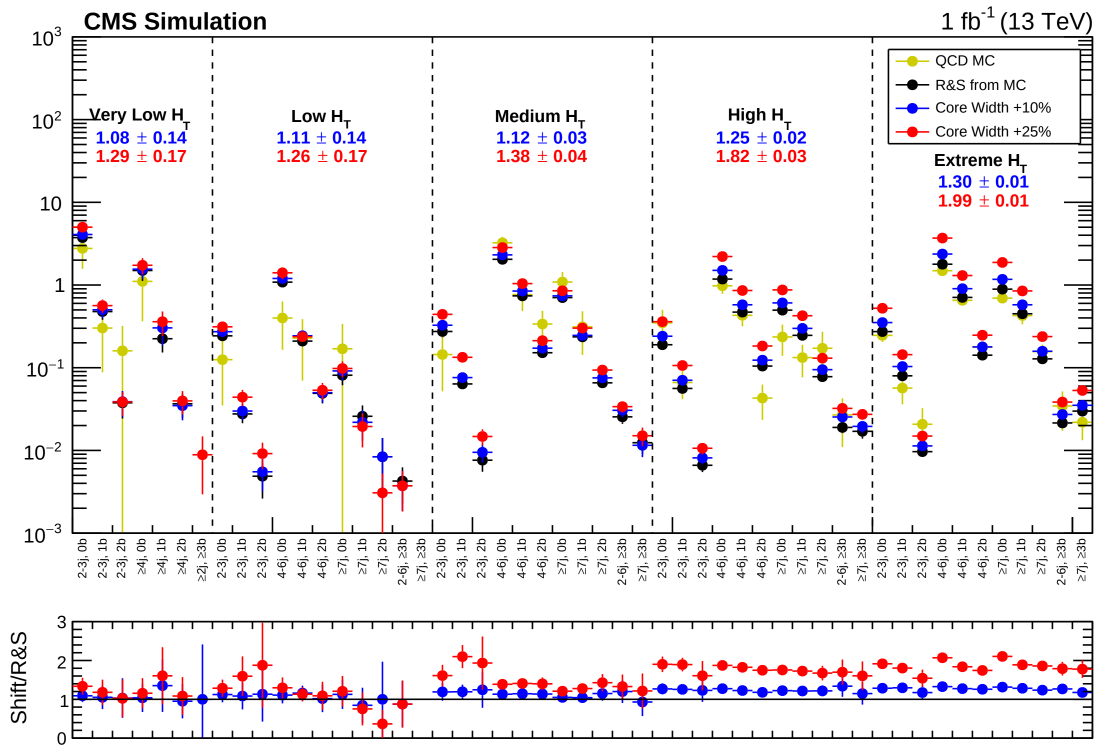

Figure 8.22: Comparison of yields in topological regions for QCD from MC (yellow points), standard R&S prediction from MC (black points), R&S with response core width + 10% (blue points), and R&S with response core width + 25% (red points). The bottom histogram shows the ratios of yields in topological regions for response core width + 10% and response core width + 25% with respect to the standard R&S prediction.

8.7.2 Effect of modifying size of jet response tail

Next, we study in Monte Carlo how the R&S method is affected by increasing the size of the non-Gaussian tail component of the jet response templates. As described in Sec. 3.4.2, the tail is defined as simply the subtraction of the fitted Gaussian core from the raw template. For this study, we increase the size of the tails by 25% and 50% for all templates, simply by scaling the tail components by a constant factor (while also scaling down the core normalization by an appropriate amount to preserve normalization). Figure 8.24 shows the result of this scaling in each of the analysis topological regions, and

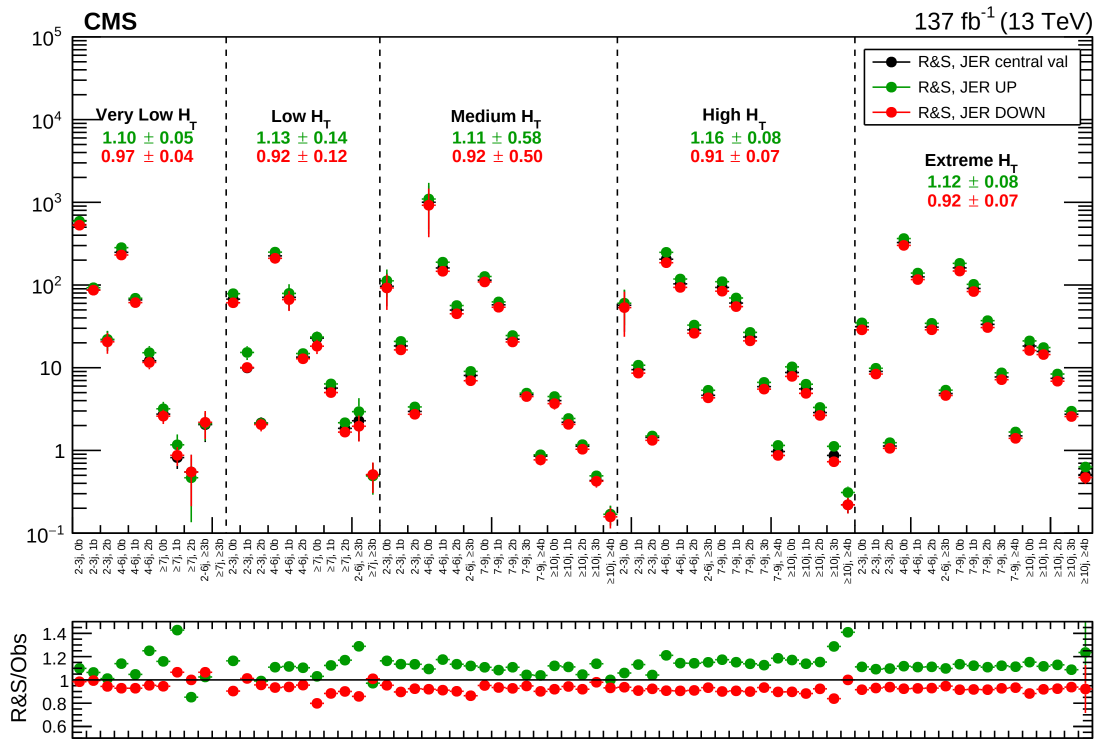

Figure 8.23: Comparison of yields in topological regions for R&S from data for nominal JER (black), JER varied UP (green), and JER varied DOWN (red). Variations with respect to the nominal yield in each H_T-region are shown in colored text.

Table 8.2 summarizes the change in yields for each H_T region.

Table 8.2: Effect of increasing the size of the jet response tails in each H_T region.

	tail+25% / standard R&S	tail+50% / standard R&S
very low H_T	1.17	1.34
low H_T	1.18	1.33
medium H_T	1.24	1.47
high H_T	1.25	1.50
extreme H_T	1.25	1.49

8.7.3 Effect of shifting mean of jet response

Next, we study in Monte Carlo how the R&S method is affected by shifting the mean of each jet response template by a constant amount, chose here to by +4%. This simulates

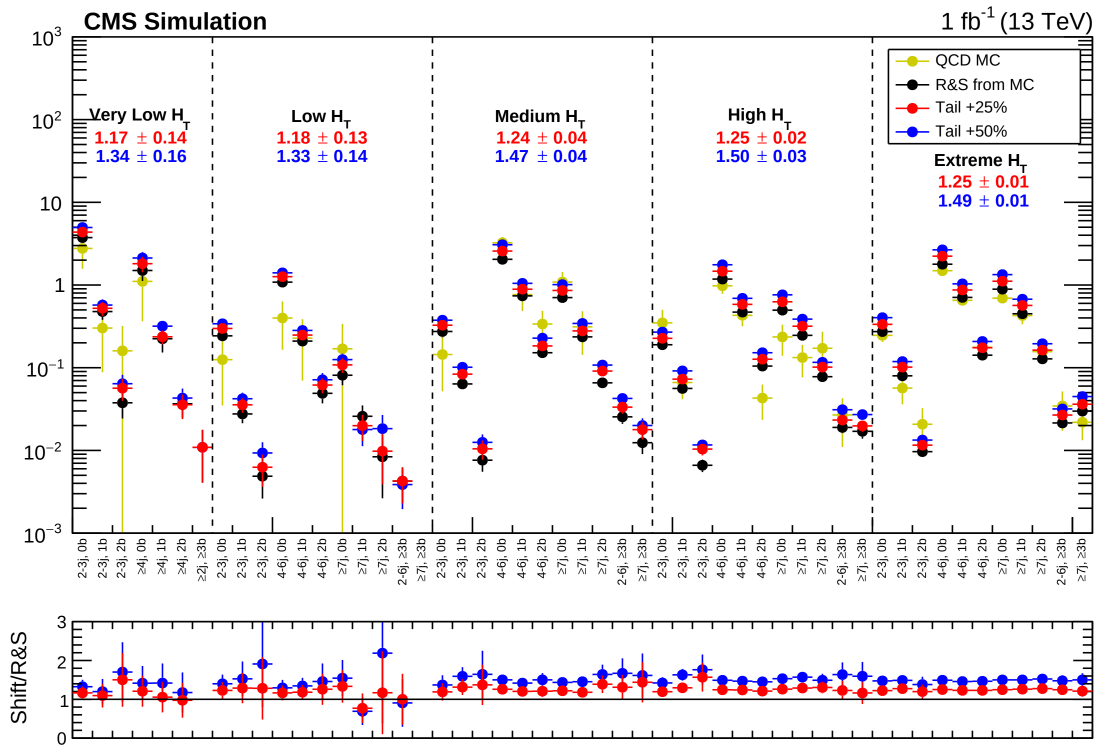

Figure 8.24: Comparison of yields in topological regions for QCD from MC (yellow points), standard R&S prediction from MC (black points), R&S with response tail size + 25% (red points), and R&S with response tail size + 50% (blue points). The bottom histogram shows the ratios of yields in topological regions for response tail size + 25% and response tail size + 50% with respect to the standard R&S prediction.

deriving the jet response templates on a sample that has a jet energy scale that is 4% higher than the sample on which the templates are used for a prediction. Figure 8.25 shows the effect of shifting the jet response mean by +4% on the R&S predictions from MC in the analysis topological regions. Table 8.3 summarizes the change in yields for each H_{T} region.

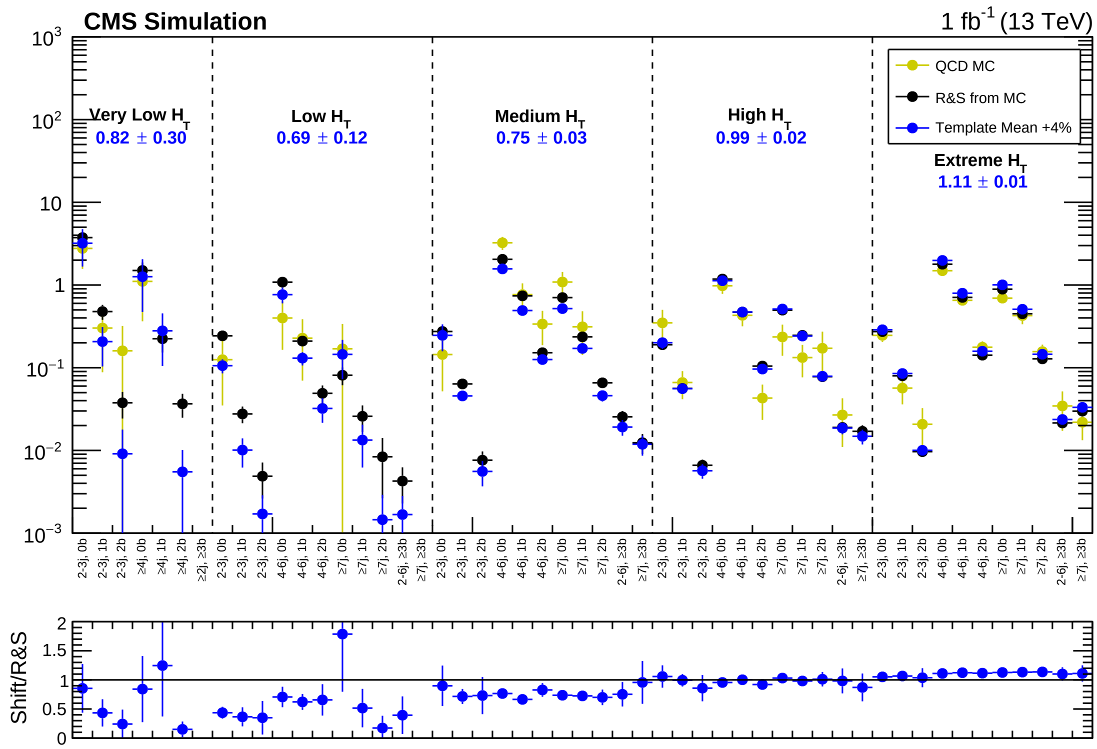

Figure 8.25: Comparison of yields in topological regions for QCD from MC (yellow points), standard R&S prediction from MC (black points), and R&S with response mean shifted higher by 4% (blue points). The bottom histogram shows the ratio of yields in topological regions for response mean shifted higher by 4% with respect to the standard R&S prediction.

8.7.4 Effect of modifying $\sigma_{\mathrm{T}}^{\mathrm{soft}}$

Finally, we check how the R&S prediction changes based on the value chosen for the $\sigma_{\mathrm{T}}^{\mathrm{soft}}$ parameter used in rebalancing. The parameter controls how tightly the $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ is constrained to be near 0 post-rebalancing. It is nominally chosen to be 20 GeV based on the width of the $p_{\mathrm{T}}^{\mathrm{miss}}$ distribution in minimum-bias data. Figure 8.26 shows the effect of changing $\sigma_{\mathrm{T}}^{\mathrm{soft}}$ by ±20% (i.e. to 15 or 25 GeV) on the R&S prediction from MC in the analysis topological regions. Table 8.4 summarizes the change in yields for each H_{T} region.

Table 8.3: Effect of shifting the mean of the jet response in each H_T region.

	mean+4% / standard R&S
very low H_T	0.82
low H_T	0.69
medium H_T	0.75
high H_T	0.99
extreme H_T	1.11

Table 8.4: Effect of changing $\sigma_\mathrm{T}^\mathrm{soft}$ in each H_T region.

	$\sigma_\mathrm{T}^\mathrm{soft} = 15$ GeV/ standard R&S	$\sigma_\mathrm{T}^\mathrm{soft} = 25$ GeV/ standard R&S
very low H_T	0.99	1.00
low H_T	1.08	0.97
medium H_T	1.12	1.05
high H_T	1.07	1.06
extreme H_T	1.05	1.04

8.7.5 Final systematic uncertainties assessed on estimate

We assign systematic uncertainties from three main sources: jet energy resolution uncertainty, template tail size uncertainty, and σ_T^{soft} uncertainty. As seen from Figures 8.23, 8.24, 8.26, the effect of varying these features has no strong $N_\mathrm{j}/N_\mathrm{b}$ dependence and varies mainly by H_T region. Therefore, we assign these systematics by H_T region and take them as correlated across all bins. Due to low statistics, we combine the Very Low and Low H_T regions when deriving systematics. For JER uncertainty, we take the maximum yield variation after varying JER up/down in data. For tail size uncertainty, we take the MC yield variation resulting from a 25% increase in tail size. And for $\sigma_\mathrm{T}^\mathrm{soft}$ variation, we take the maximum yield variation after varying $\sigma_\mathrm{T}^\mathrm{soft}$ by ± 5 GeV (20%). The results are given in Table 8.5.

As the above systematics are integrated across all variables except H_T, we assign final systematics on the modeling of N_j and N_b shapes, derived from the N_j and N_b plots in the inverted-$\Delta\phi_\mathrm{min}$ control region (Figure 8.13). These systematics grow with N_j and N_b and are correlated within each H_T region. Table 8.6 lists the assigned systematics.

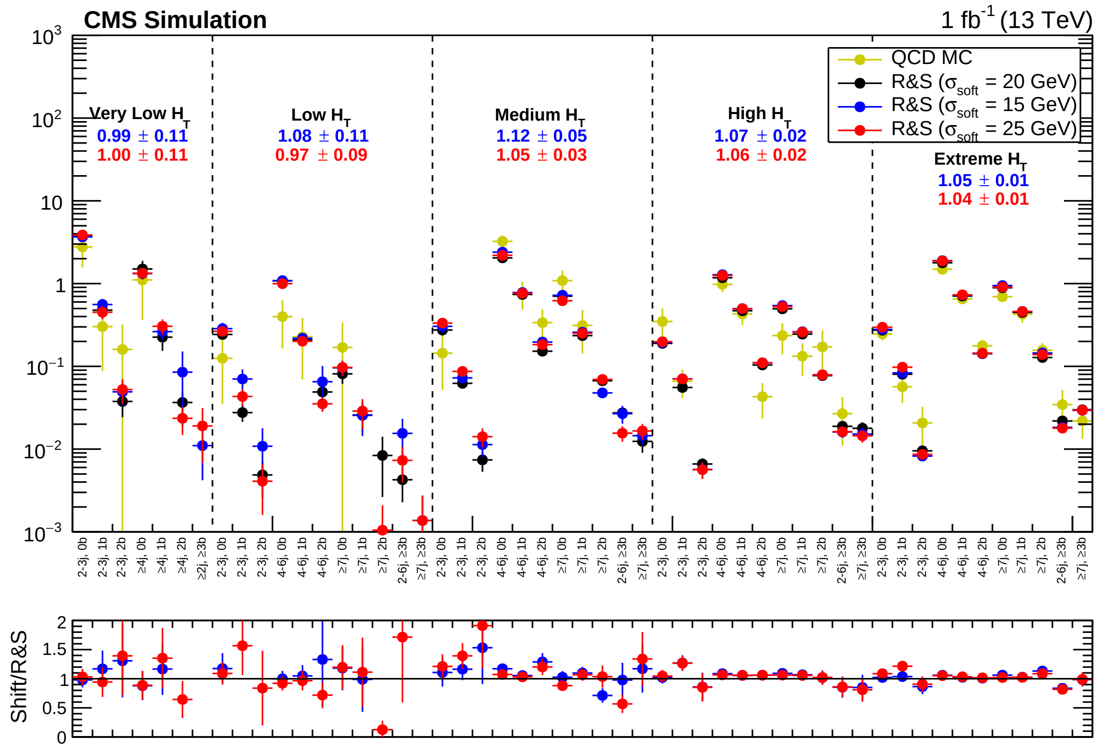

Figure 8.26: Comparison of yields in topological regions for QCD from MC (yellow points), standard R&S prediction from MC (black points), R&S with $\sigma_{\mathrm{T}}^{\mathrm{soft}} = 15$ GeV (blue points), and R&S with $\sigma_{\mathrm{T}}^{\mathrm{soft}} = 25$ GeV (red points). The bottom histogram shows the ratios of yields in topological regions for $\sigma_{\mathrm{T}}^{\mathrm{soft}} = 15$ GeV and $\sigma_{\mathrm{T}}^{\mathrm{soft}} = 25$ GeV with respect to the standard R&S prediction.

Table 8.5: Systematics assigned to each H_{T}-region based on uncertainties in JER, tail size modeling, and $\sigma_{\mathrm{T}}^{\mathrm{soft}}$. The "Total" column is the 3 systematics added in quadrature.

	JER	Tail size	σ_T^{soft}	Total
Very Low H_{T}	5%	17%	1%	18%
Low H_{T}	14%	17%	1%	22%
Medium H_{T}	18%	24%	12%	32%
High H_{T}	13%	25%	7%	29%
Extreme H_{T}	10%	25%	5%	27%
Monojet, 0b	10%	21%	17%	29%
Monojet, ≥ 1b	5%	8%	25%	27%

Table 8.6: Assigned systematics to each $N_{\mathrm{j}}/N_{\mathrm{b}}$ bin, based on observed discrepancies in $N_{\mathrm{j}}/N_{\mathrm{b}}$ shapes in the inverted-$\Delta\phi_{\mathrm{min}}$ control regions. The uncertainties are taken as completely correlated within each H_{T} region, and uncorrelated between H_{T} regions. The signs indicate the direction of correlation.

	$H_{\mathrm{T}} < 1200$	$H_{\mathrm{T}} \geq 1200$
2-3 jets	0%	2%
4-6 jets	1%	3%
2-6 jets	1%	3%
7-9 jets	−8%	−7%
$\geq$ 7 jets	−8%	−7%
$\geq$ 10 jets	−20%	−19%
0 b-jets	5%	4%
1 b-jets	−9%	−5%
2 b-jets	−9%	−5%
3 b-jets	−14%	−16%
$\geq$ 3 b-jets	−14%	−16%
$\geq$ 4 b-jets	−14%	−16%

Chapter 9

Results and Interpretation

The background estimation techniques described in the previous chapters are used to make predictions in each of the 282 signal regions bins. We then compare and fit to observed event counts, using a maximum-likelihood fit that takes into account all statistical and systematic uncertainties. Finally, the results of this fit are used to constrain a variety of BSM physics scenarios. All results are published in Ref. [72]. Sec. 9.1 presents the raw, pre-fit results, Sec. 9.2 briefly describes the fitting and limit-setting procedures, and Sec. 9.3 shows the resulting limits placed on all considered signal models.

9.1 Pre-fit results

Figures 9.1 to 9.4 show comparisons of the pre-fit background estimates and observed yields in all signal regions. The hatched regions in the upper panels (and the solid gray bands in the ratio panels) represent the total statistical and systematic uncertainty in the background prediction. For the monojet region, the x axis binning is $p_{\mathrm{T}}^{\mathrm{jet1}}$ (in GeV), and for the $N_{\mathrm{j}} \geq 2$ regions the x axis binning is M_{T2} (in GeV). Dashed lines separate the bins into categories of N_{j} and N_{b}. Observed yields are statistically consistent with the estimated SM background.

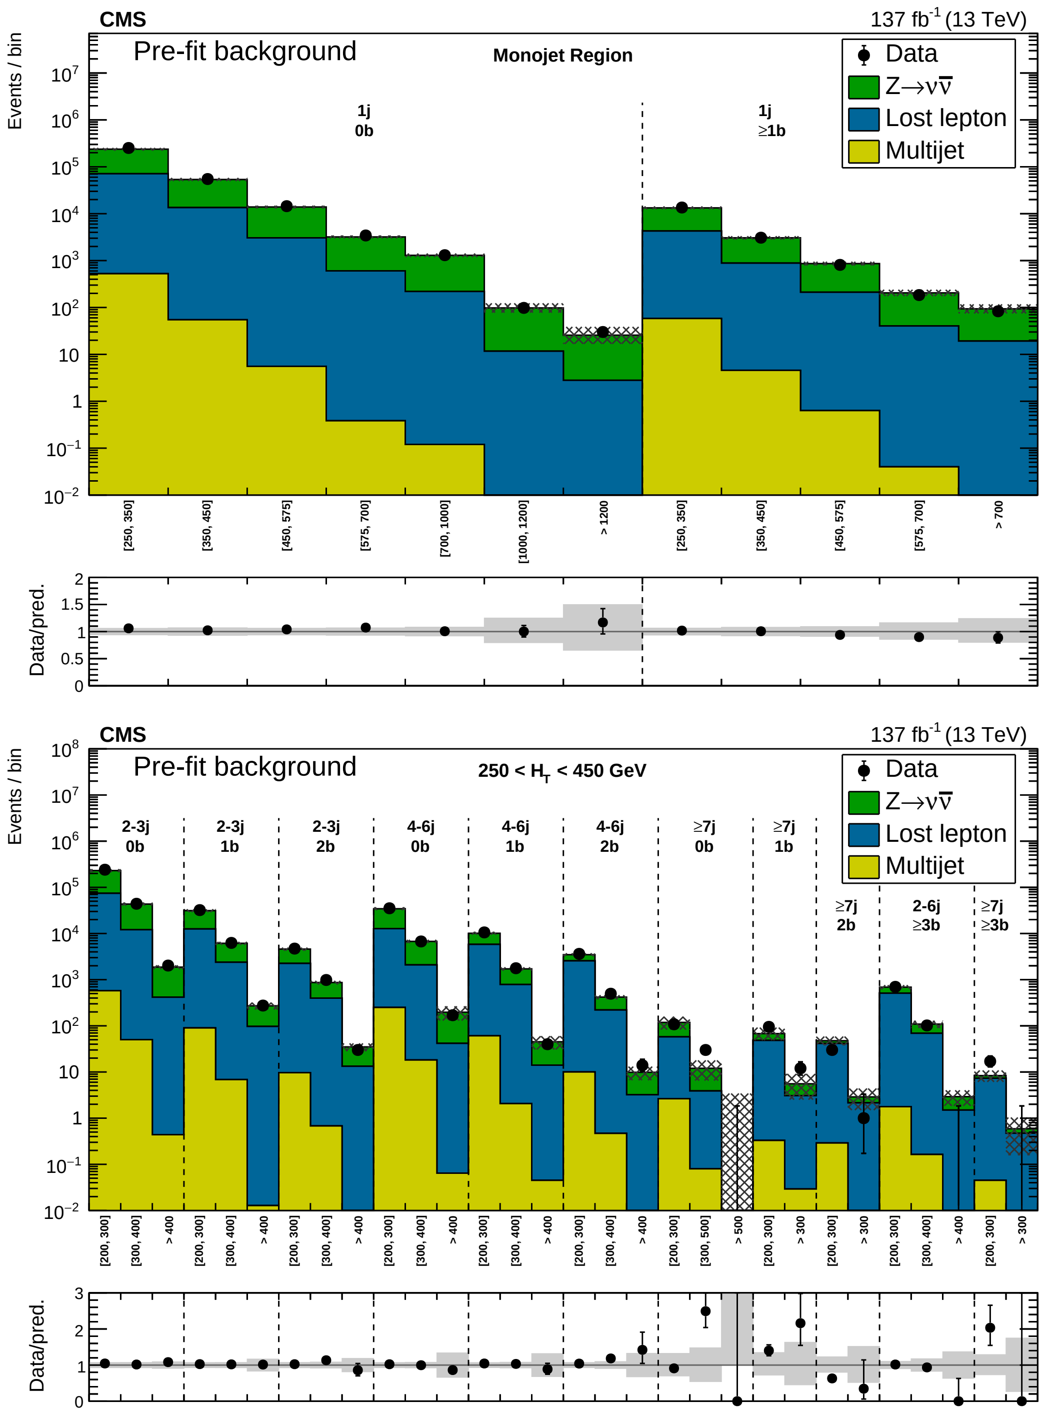

Figure 9.1: Expected (pre-fit) and observed yields for the full data set collected from 2016–18, corresponding to an integrated luminosity of 137 fb^{-1}. On top are the monojet signal regions, separated into N_b categories and with $p_\mathrm{T}^\mathrm{jet1}$ binning on the x axis. On the bottom are the $N_\mathrm{j} \geq 2$ signal regions with $250 \leq H_\mathrm{T} < 450$ GeV, separated into topological regions with the M_T2 binning on the x axis.

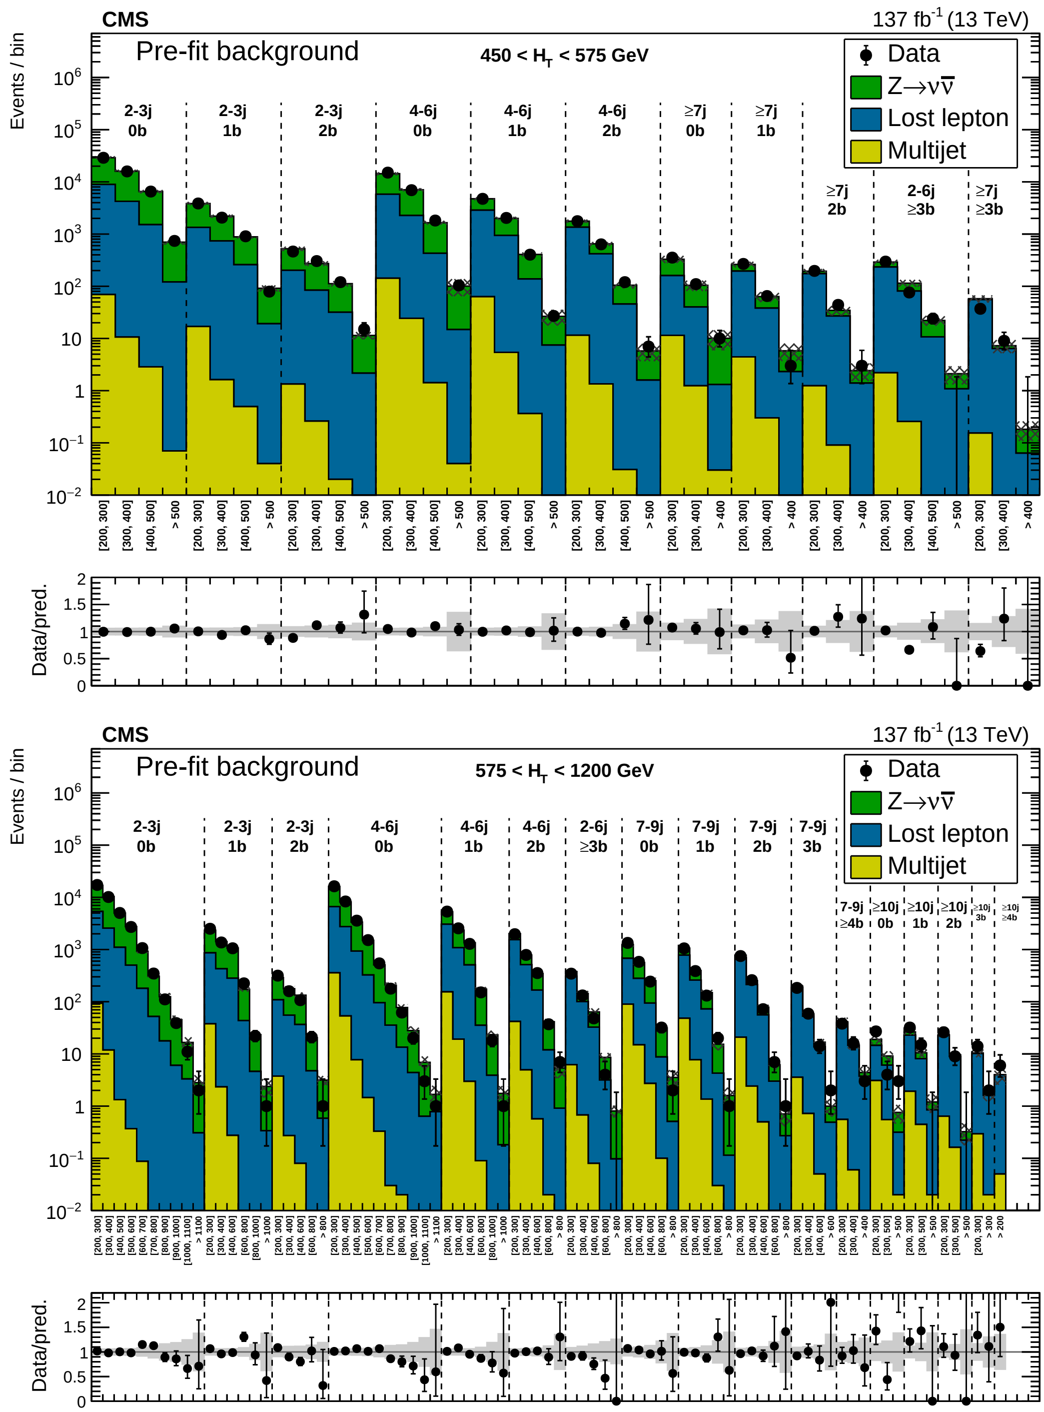

Figure 9.2: Expected (pre-fit) and observed yields for the full data set collected from 2016–18, corresponding to an integrated luminosity of 137 fb^{-1}. The top and bottom figures contain the $N_{\mathrm{j}} \geq 2$ signal regions for the $450 \leq H_{\mathrm{T}} < 575$ GeV and $575 \leq H_{\mathrm{T}} < 1200$ GeV regions, respectively. M_{T2} binning is shown on the x axis, with the dashed lines separating bins by topological region.

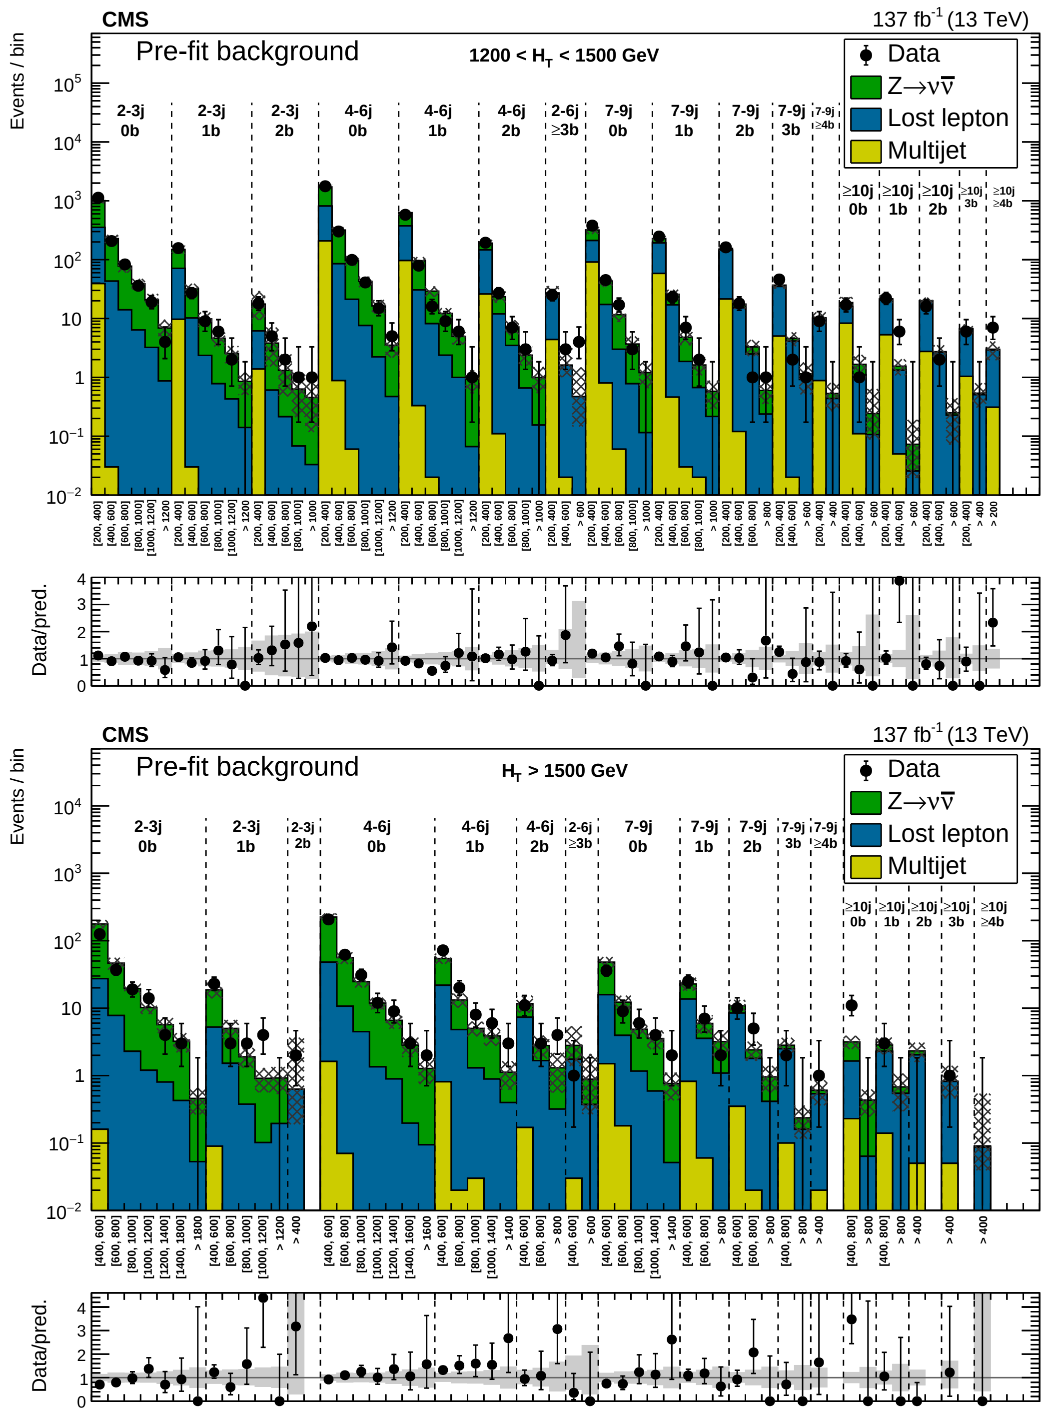

Figure 9.3: Expected (pre-fit) and observed yields for the full data set collected from 2016–18, corresponding to an integrated luminosity of 137 fb^{-1}. The top and bottom figures contain the $N_{\rm j} \geq 2$ signal regions for the $1200 \leq H_{\rm T} < 1500$ GeV and $H_{\rm T} \geq 1500$ GeV regions, respectively. $M_{\rm T2}$ binning is shown on the x axis, with the dashed lines separating bins by topological region.

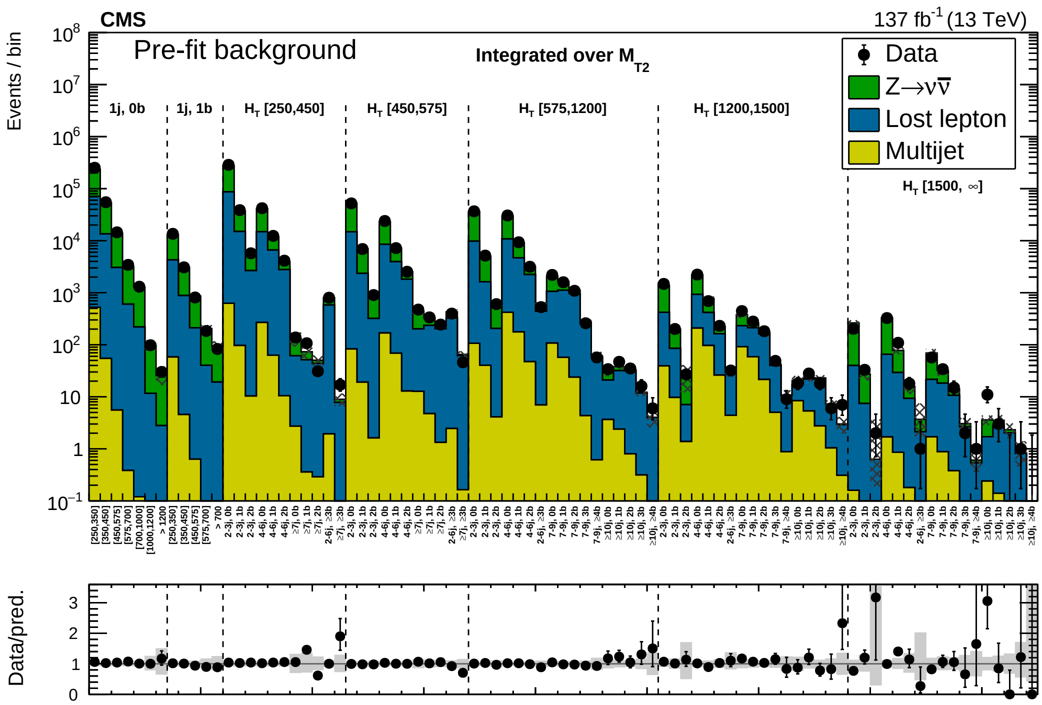

Figure 9.4: Expected (pre-fit) and observed yields for the full data set collected from 2016–18, corresponding to an integrated luminosity of 137 fb^{-1}. Here the results are shown integrated over M_{T2}, with each bin on the x axis corresponding to an $N_{\mathrm{j}}/N_{\mathrm{b}}$ topological region. Dashed lines separate the various H_{T} regions. The two leftmost regions contain the monojet signal regions, where the x axis binning is $p_{\mathrm{T}}^{\mathrm{jet1}}$.

9.2 Maximum-likelihood fits and the $\mathrm{CL_S}$ technique

The results presented in the previous section are *pre-fit*, meaning the shown backgrounds and uncertainties are straight from the estimation methods, without attempting to fit to the data in any way. For the purposes of evaluating the results and setting constraints on models of new physics, we apply a fitting procedure to test the compatibility of the observed results with SM predictions.

The expected background and signal yields are functions of *nuisance parameters*, which control the potential variations from all of the experimental and theoretical uncertainties. In practice these are modeled as log-normal constraints on the background and signal yields (except for statistical uncertainties on the observed control region event counts, which are modeled as gamma functions). Denoting the vector of all nuisance parameters as $\boldsymbol{\theta}$, and their joint probability distribution as $p(\boldsymbol{\theta})$, we can write the complete likelihood function as

$$\mathcal{L}(\text{data}|\mu, \boldsymbol{\theta}) = \prod_{j \in \text{bins}} \frac{[\mu s_j(\boldsymbol{\theta}) + b_j(\boldsymbol{\theta})]^{n_j}}{n_j!} e^{-[\mu s_j(\boldsymbol{\theta}) + b_j(\boldsymbol{\theta})]} p(\boldsymbol{\theta}), \tag{9.1}$$

where $s_j(\boldsymbol{\theta})$ and $b_j(\boldsymbol{\theta})$ are the predicted signal and background yields in bin j (and which are functions of the nuisance parameters), n_j is the observed event count in bin j, and μ is the signal strength parameter that we seek to set a constraint on. The likelihood is simply a product of Poisson probability terms with expectation $\mu s_j + b_j$ and observation n_j, and the PDF of the nuisance parameters, which is itself a product of log-normal and gamma functions.

To evaluate how well the estimated background fits the observed data, we could simply maximize this likelihood under the background-only hypothesis (i.e., set $\mu = 0$ in Eq. 9.1. However, we would like to interpret this in the context of particular signal models, so we need a method to evaluate the significance of the signal+background (S+B) assumption (the *alternative hypothesis*) with respect to the background-only (B-only) assumption (the *null hypothesis*).

To do this, we make use of the $\mathrm{CL_S}$ technique [73]. First, a test statistic is defined as

$$q_\mu = -2\log\frac{\mathcal{L}(\mu,\hat{\boldsymbol{\theta}}_\mu)}{\mathcal{L}(0,\hat{\boldsymbol{\theta}}_0)}, \tag{9.2}$$

where the likelihood $\mathcal{L}$ is as defined in Eq. 9.1, $\hat{\boldsymbol{\theta}}_0$ is the value of $\boldsymbol{\theta}$ that maximizes $\mathcal{L}$ for $\mu = 0$, and $\hat{\boldsymbol{\theta}}_\mu$ is the value of $\boldsymbol{\theta}$ that maximizes $\mathcal{L}$ for the given μ. This is known as a *profile likelihood*, as we have profiled away the dependence on the nuisance parameters $\boldsymbol{\theta}$ by maximizing over them. Note that the denominator is independent of μ, and is only there as a normalization.

We can construct distributions of q_μ for a given μ by generating random data, either under the B-only or μS+B hypotheses. The values of q_μ under the μS+B hypothesis tend to be lower on average than those under the B-only hypothesis, as the numerator will be larger than the denominator (see top right plot in Fig. 9.5 for an example). The distributions converge as $\mu \to 0$.

Using the actual observed data counts, we can then calculate the observed q_μ^{obs} for arbitrary μ. We then define

$$\mathrm{CL_S}(\mu) \equiv \frac{\mathrm{CL_{S+B}}}{\mathrm{CL_B}} \equiv \frac{p(q_\mu \geq q_\mu^{\mathrm{obs}} \mid \mu\mathrm{S{+}B})}{p(q_\mu \geq q_\mu^{\mathrm{obs}} \mid \mathrm{B})} \tag{9.3}$$

Note that $\mathrm{CL_S}(\mu)$ is exactly equal to 1 at $\mu = 0$, and is a strictly decreasing function of μ. One can think of $\mathrm{CL_S}$ as the relative likelihood that the observed counts would look as "background-like" as they do under the μS+B hypothesis versus the B-only hypothesis. When $\mathrm{CL_S}$ is very small, we can be pretty confident that the μS+B assumption is *not* supported by the data. Thus, we define an upper limit on μ at the $1-\alpha$ confidence level as the $\mu_{1-\alpha}$ that satisfies

$$\mathrm{CL_S}(\mu_{1-\alpha}) = \alpha \tag{9.4}$$

This process is illustrated in Fig. 9.5 for a toy experiment with 3 bins. We predict event counts of 10, 5, and 2, with an uncorrelated 20% uncertainty on each. Nominal expected signal with $\mu = 1$ is 2 events in each bin. We observe event counts of 7, 6, and 1. The top left panel shows the expected and observed yields in each bin. The top right panel shows distributions of $q_{\mu=1}$ under the B and S+B hypotheses, the observed value of $q_{\mu=1}^{\mathrm{obs}} = 3.75$, and shaded areas representing $\mathrm{CL_{S+B}}$ and $1 - \mathrm{CL_B}$ The bottom panel shows $\mathrm{CL_S}(\mu)$ as a function of μ, and the 95% confidence level upper limit on μ of 1.14, computed by finding the μ value at which $\mathrm{CL_S}(\mu)$ crosses below 0.05.

The $\mathrm{CL_S}$ method outlined above is used to set limits on all of the signal models in the following sections. Due to the very large number of bins in the M_{T2} analysis, it is computationally infeasible to compute distributions of the test statistic explicitly. Hence, we use an *asymptotic approximation* to derive the distributions necessary to calculate $\mathrm{CL_S}$, as described in [74]. This has been verified at a smaller number of signal points to give results consistent with the exact method. The calculations are performed with a CMS-produced wrapper around RooFit [75].

9.3 Interpretations

The statistical procedure described in the previous section is used to set 95% confidence level upper limits on the cross sections of a variety of signal models. For the interpretation of the results, simplified BSM physics models [24–26] are used. Simplified models are defined by sets of hypothetical particles and sequences of their production

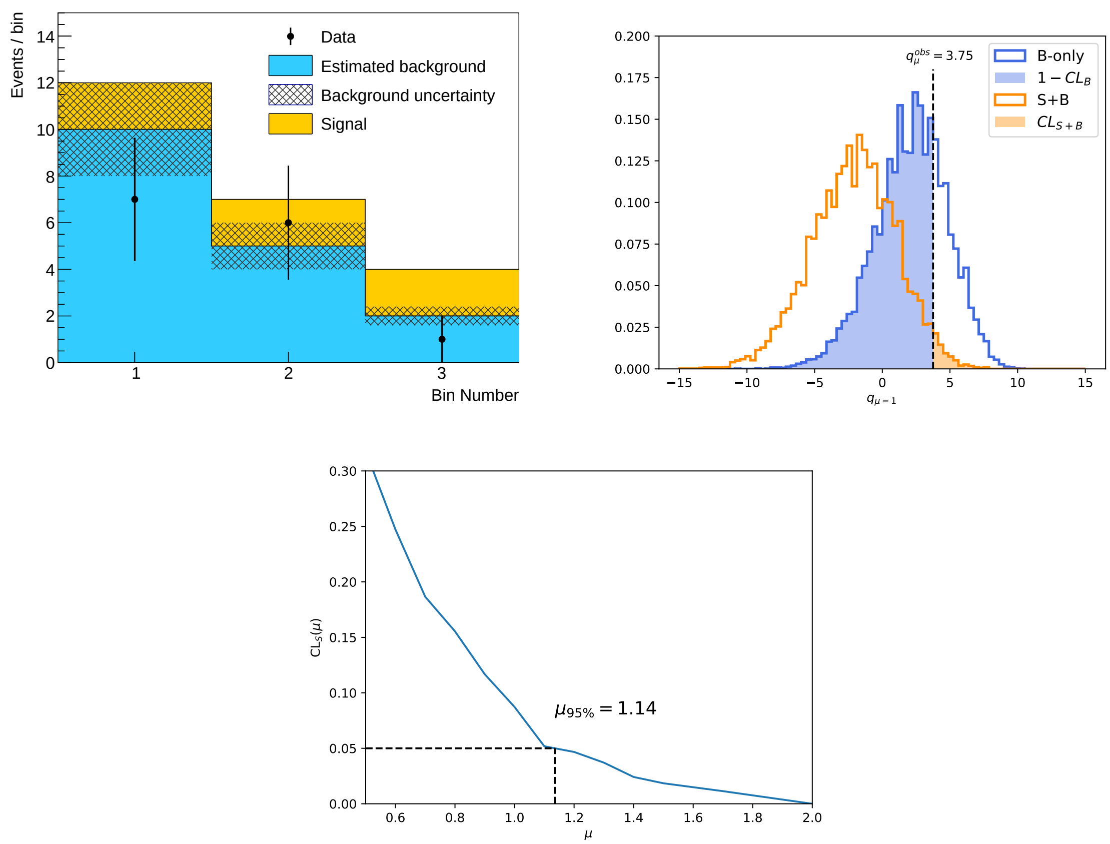

Figure 9.5: An example calculation of an upper limit on signal strength μ with a three-bin toy experiment. The estimated background, background uncertainty, and signal yields, as well as the observed counts, are shown in the top left. On the top right, distributions of $q_{\mu=1}$ are shown for both the B-only and S+B hypotheses. The dotted line shows the observed value of $q^{\mathrm{obs}}_{\mu=1} = 3.75$, and the blue and orange shaded areas represent $1 - \mathrm{CL_B}$ and $\mathrm{CL_{S+B}}$, respectively. The bottom plot shows $\mathrm{CL_S}(\mu) = \mathrm{CL_{S+B}}/\mathrm{CL_B}$ as a function of μ, and the observed 95% confidence level upper limit on μ of 1.14, found by locating the point where $\mathrm{CL_S}$ crosses below $1-0.95 = 0.05$.

and decay. The theoretical parameters are thus reduced to a small number of masses and cross sections, providing an effective tool to characterize potential signals of BSM physics.

The estimation of the three different background types and the corresponding uncertainties are described in Chapters 6 through 8. Estimated signal yields are taken from simulation, which has been produced using the CMS fast simulation framework [53]. Due to known mis-modeling of ISR jets in simulation, signal samples have been re-weighted based on the number of generator-level ISR jets, with weights derived from a $t\bar{t}$-enriched sample in data. In addition, due to known $p_{\mathrm{T}}^{\mathrm{miss}}$ tail mis-modeling in the fast simulation, estimated signal yields are computed using an average of the generator-level and reconstructed $p_{\mathrm{T}}^{\mathrm{miss}}$ values. Systematic uncertainties to account for both of these procedures are applied to the estimated signal yields in each bin.

A list of all sources of uncertainty on the signal yields is given in Table 9.1, along with the range of values for each uncertainty source over all analysis bins.

Table 9.1: Systematic uncertainties in the signal yields for the simplified models of BSM physics. The large statistical uncertainties in the simulated signal sample come from a small number of bins with low acceptance, which are typically not among the most sensitive bins contributing to a given model benchmark point.

Source	Range [%]
Integrated luminosity	2.3–2.5
Limited size of MC samples	1–100
Renormalization and factorization scales	5
ISR modeling	0–30
b tagging efficiency, heavy flavors	0–40
b tagging efficiency, light flavors	0–20
Lepton efficiency	0–20
Jet energy scale	5
Fast simulation $p_{\mathrm{T}}^{\mathrm{miss}}$ modeling	0–5

9.3.1 Supersymmetry models

A total of 11 R-parity conserving supersymmetry simplified models are considered, illustrated in Fig. 9.6. For each scenario of gluino (squark) pair production, the simplified models assume that all SUSY particles other than those shown in the corresponding diagram are too heavy to be produced directly, and that the gluino (squark) decays promptly. The models assume that each gluino (squark) decays with a 100% branching fraction into the decay products depicted in Fig. 9.6. For models where the decays of the two gluinos or squarks in the same diagram differ, a 1/3 (1/2) branching fraction for each of the three (two) decay modes is assumed. In particular, for the diagram of gluino pair production where the decays of the two gluinos differ, each gluino can decay via a $\widetilde{\chi}_2^0$, $\widetilde{\chi}_1^+$, or $\widetilde{\chi}_1^-$. For scenarios with top squarks decaying into top quarks, the polarization of the top quark can be model dependent and a function of the top squark and neutralino mixing matrices. To maintain independence of any particular model realization, events are generated with unpolarized top quarks. Signal cross sections are calculated at approximately NNLO+NNLL (next-to-next-to-leading-logarithm) order in α_s [40]. For direct light-flavor squark pair production we assume either one single squark, or eight degenerate squarks ($\widetilde{q}_L + \widetilde{q}_R$, with $\widetilde{q} = \widetilde{u}$, $\widetilde{d}$, $\widetilde{s}$, or $\widetilde{c}$). For direct bottom and top squark pair production, we assume one single squark.

Figure 9.7 shows the exclusion limits at 95% CL for direct gluino pair production where the gluinos decay to light-flavor quarks under three different decay scenarios. Exclusion limits for direct gluino pair production where the gluinos decay to bottom and top quarks are shown in Fig. 9.8, and those for the direct production of squark pairs are shown in Fig. 9.9. Three alternate decay scenarios are also considered for the direct pair production of top squarks, and their exclusion limits are shown in Fig. 9.10.

Table 9.2 summarizes the limits on the masses of SUSY particles excluded for the

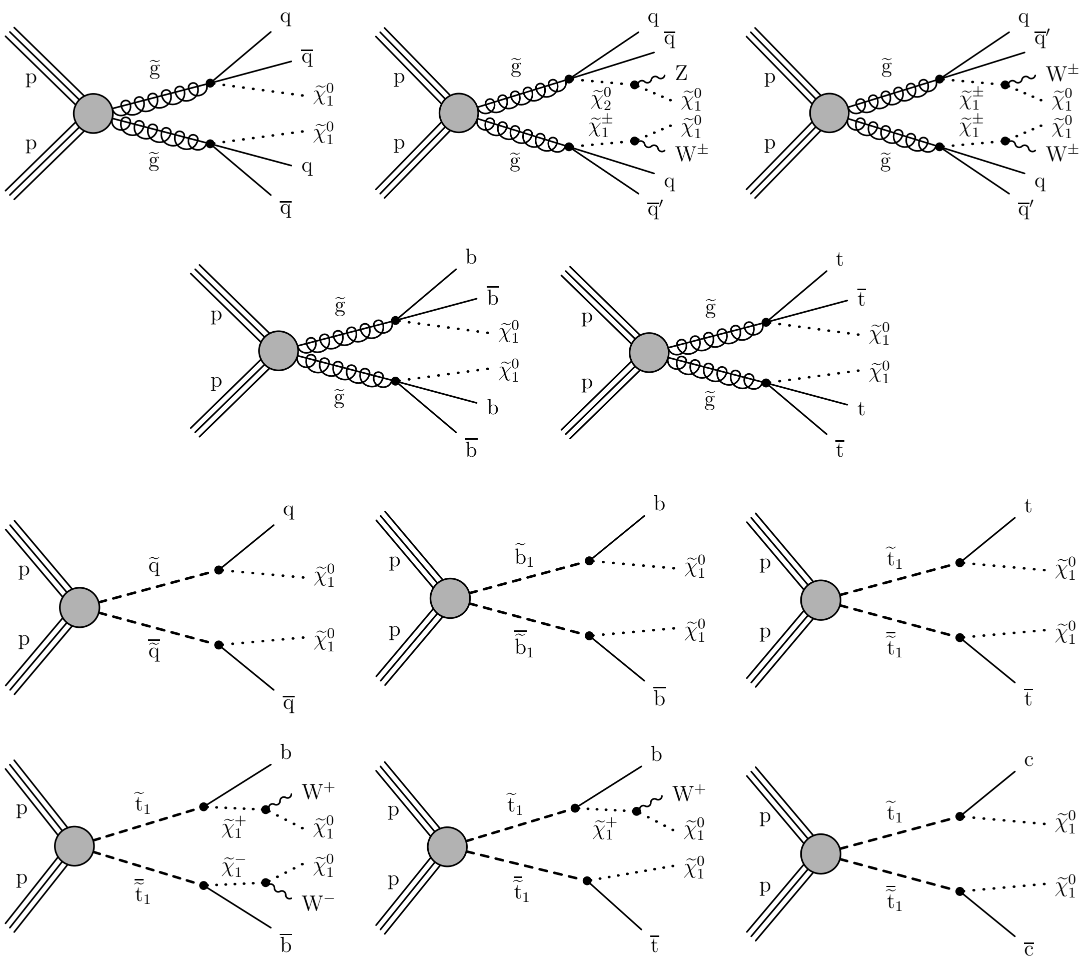

Figure 9.6: (Upper) Diagrams for three scenarios of gluino-mediated light-flavor squark pair production with different decay modes. For mixed decay scenarios, we assume equal branching fraction for each decay mode. (Upper middle) Diagrams for the gluino-mediated bottom squark and top squark pair production. (Lower middle) Diagrams for the direct pair production of light-flavor, bottom and top squark pairs. (Lower) Diagrams for three alternate scenarios of direct top squark pair production with different decay modes. For mixed decay scenarios, we assume equal branching fraction for each decay mode.

simplified model scenarios considered. These results extend the constraints on gluino and squark masses by about 100–350 GeV and on the $\widetilde{\chi}_1^0$ mass by 100–250 GeV with respect to the limits in the previous iteration of this analysis [68].

Table 9.2: Summary of the observed 95% CL exclusion limits on the masses of SUSY particles for different simplified model scenarios. The highest limits on the mass of the directly produced particles and on the mass of the $\widetilde{\chi}_1^0$ are quoted.

Simplified model	Highest limit on directly produced SUSY particle mass [GeV]	Highest limit on $\widetilde{\chi}_1^0$ mass [GeV]
Direct gluino pair production:		
$\widetilde{g} \to q\bar{q}\widetilde{\chi}_1^0$	1970	1200
$\widetilde{g} \to q\bar{q}Z\widetilde{\chi}_1^0$ or $\widetilde{g} \to q\bar{q}'W^{\pm}\widetilde{\chi}_1^0$	2020	1090
$\widetilde{g} \to b\bar{b}\widetilde{\chi}_1^0$	2250	1525
$\widetilde{g} \to t\bar{t}\widetilde{\chi}_1^0$	2250	1250
Direct squark pair production:		
Eight degenerate light squarks	1710	870
Single light squark	1250	525
Bottom squark	1240	700
Top squark	1200	580

9.3.2 Mono-ϕ model

The results of the inclusive M_{T2} search are also interpreted in the context of a BSM scenario where a colored scalar state ϕ is resonantly produced through coupling to quarks, and decays to an invisible massive Dirac fermion ψ and an SM quark. This is referred to as the mono-ϕ model. It has been proposed as an explanation of an excess in data in regions with low jet multiplicities, identified in the context of a reinterpretation [76,77] of the results of the previous inclusive M_{T2} search [68] as well as of other similar searches by the ATLAS [78,79] and CMS [80,81] Collaborations. The diagram for the hypothetical process is shown in Fig. 9.11.

Figure 9.12 shows the exclusion limits for the mono-ϕ model. Based on the LO cross section calculation, we obtain mass limits as large as 1660 and 925 GeV on m_ϕ and

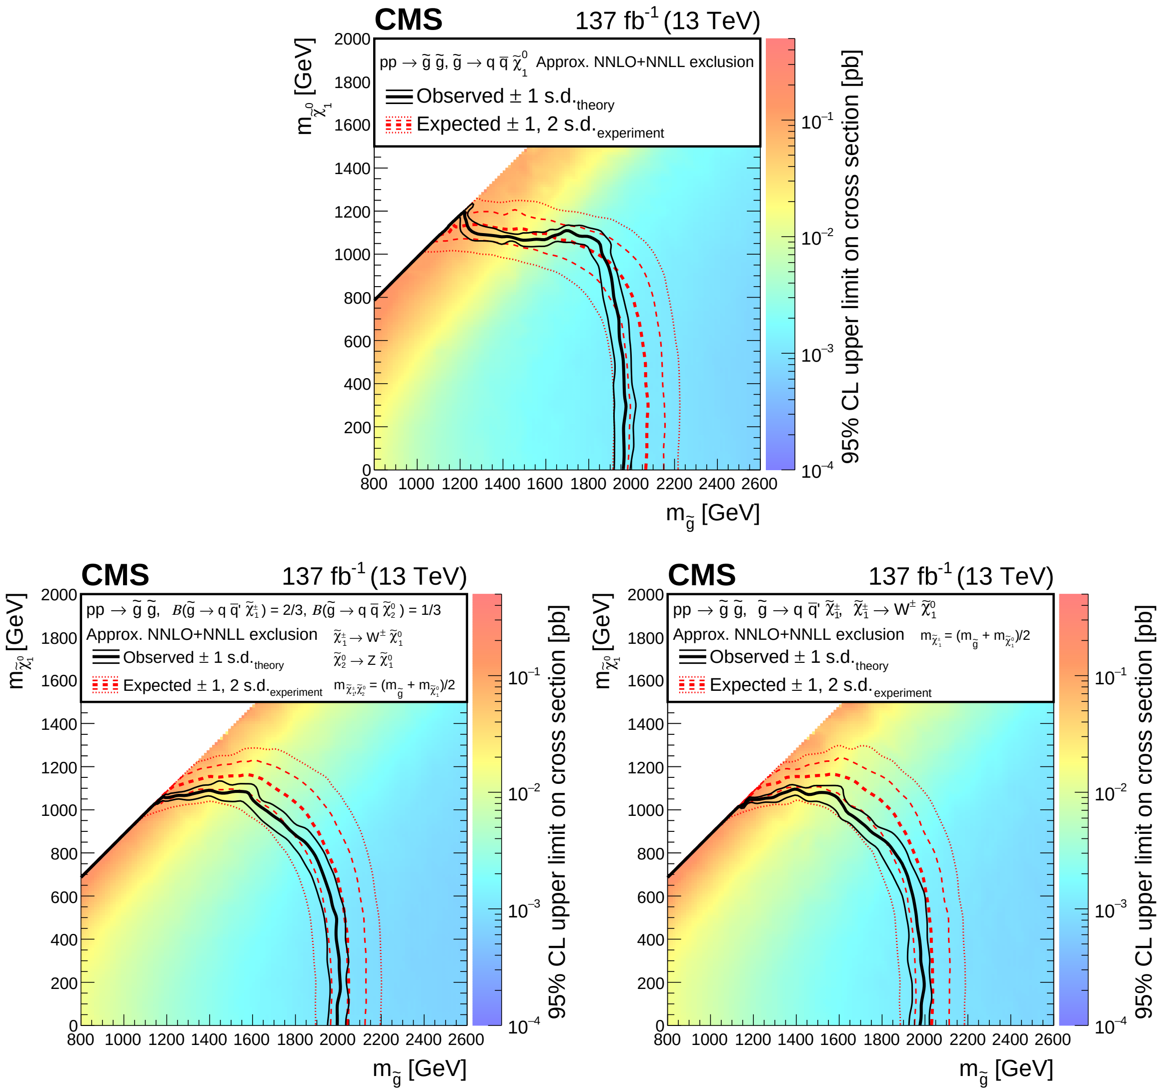

Figure 9.7: Exclusion limits at 95% CL for direct gluino pair production, where (upper) $\widetilde{g} \to q\bar{q}\widetilde{\chi}^0_1$, (lower left) $\widetilde{g} \to q\bar{q}\widetilde{\chi}^0_2$ and $\widetilde{\chi}^0_2 \to Z\widetilde{\chi}^0_1$, or $\widetilde{g} \to q\bar{q}'\widetilde{\chi}^{\pm}_1$ and $\widetilde{\chi}^{\pm}_1 \to W^{\pm}\widetilde{\chi}^0_1$, and (lower right) $\widetilde{g} \to q\bar{q}'\widetilde{\chi}^{\pm}_1$ and $\widetilde{\chi}^{\pm}_1 \to W^{\pm}\widetilde{\chi}^0_1$ (with $q = u,\ d,\ s,$ or c). For the scenarios where the gluinos decay via an intermediate $\widetilde{\chi}^0_2$ or $\widetilde{\chi}^{\pm}_1$, $\widetilde{\chi}^0_2$ and $\widetilde{\chi}^{\pm}_1$ are assumed to be mass-degenerate, with $m_{\widetilde{\chi}^{\pm}_1,\ \widetilde{\chi}^0_2} = 0.5(m_{\widetilde{g}} + m_{\widetilde{\chi}^0_1})$. The area enclosed by the thick black curve represents the observed exclusion region, while the dashed red lines indicate the expected limits and their ± 1 and ± 2 standard deviation ranges. The thin black lines show the effect of the theoretical uncertainties in the signal cross section. Signal cross sections are calculated at approximately NNLO+NNLL order in α_s [40], assuming 1/3 branching fraction ($\mathcal{B}$) for each decay mode in the mixed-decay scenarios, or unity branching fraction for the indicated decay.

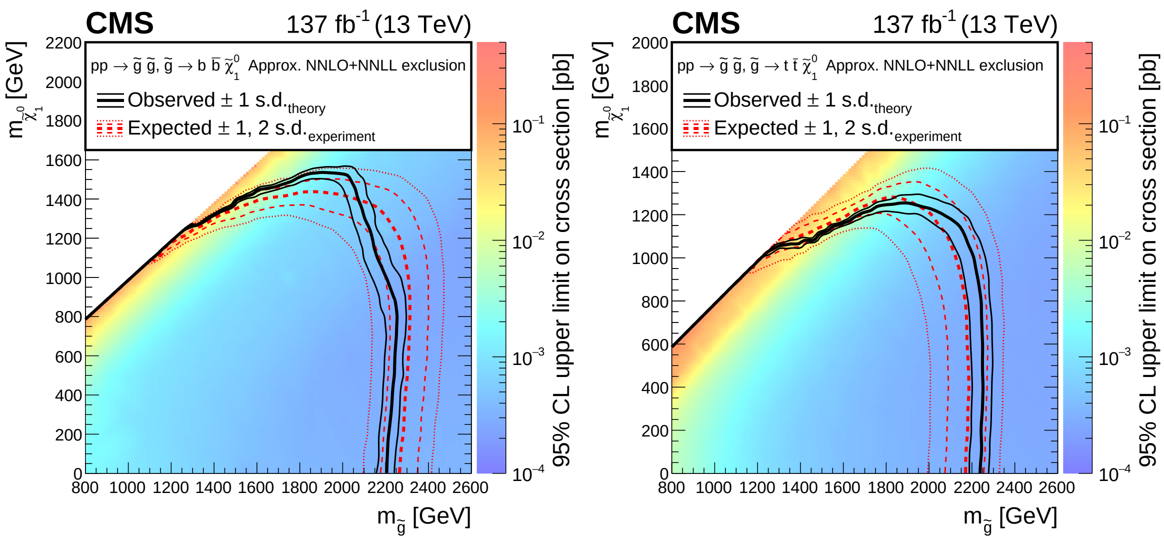

Figure 9.8: Exclusion limits at 95% CL for direct gluino pair production where the gluinos decay to (left) bottom quarks and (right) top quarks. The area enclosed by the thick black curve represents the observed exclusion region, while the dashed red lines indicate the expected limits and their ±1 and ±2 standard deviation ranges. The thin black lines show the effect of the theoretical uncertainties in the signal cross section. Signal cross sections are calculated at approximately NNLO+NNLL order in α_s [40], assuming unity branching fraction for the indicated decay.

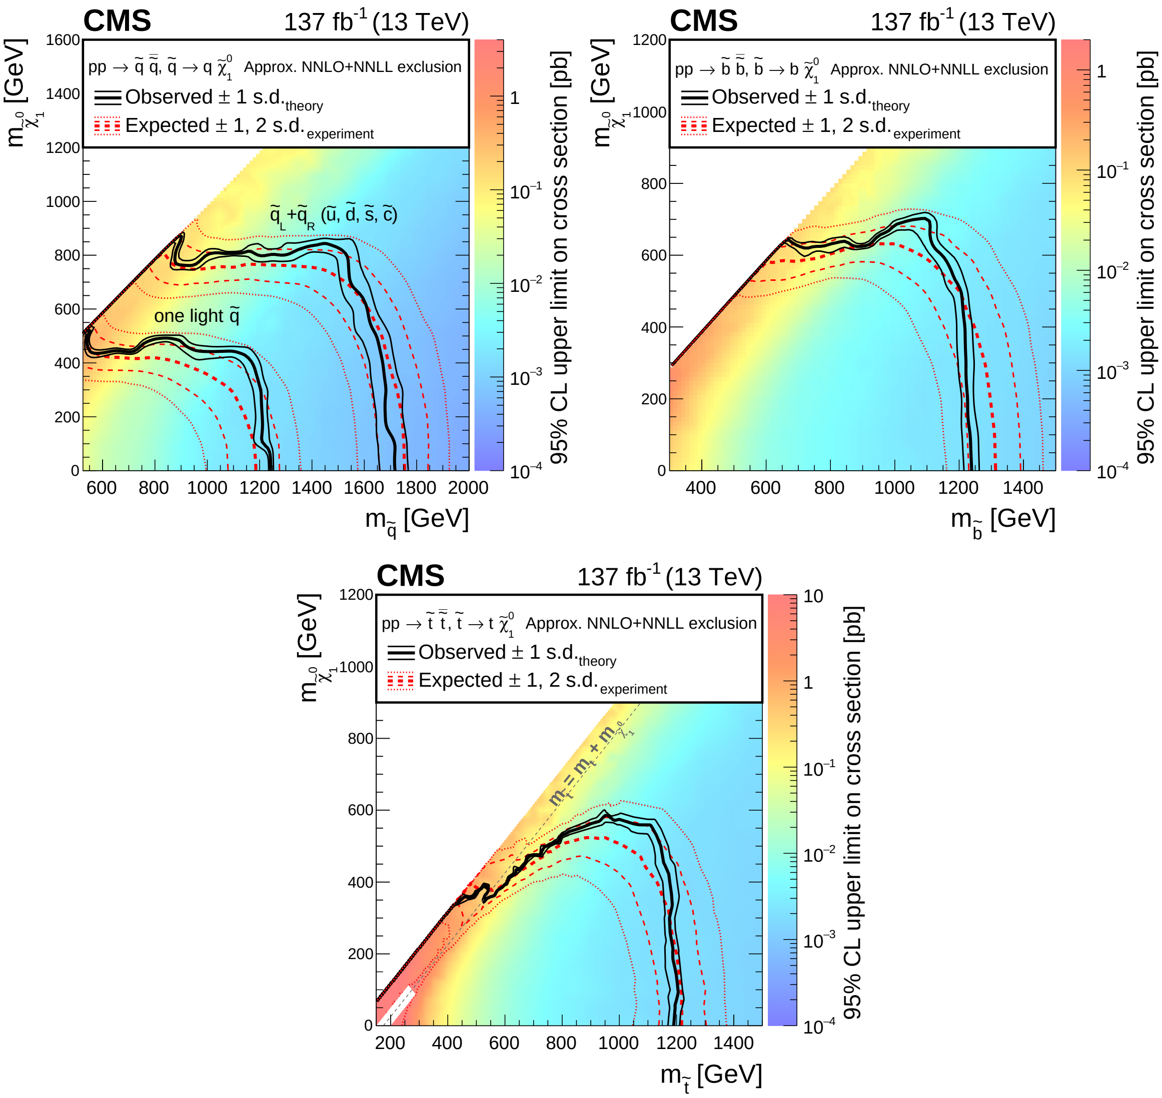

Figure 9.9: Exclusion limit at 95% CL for (upper left) light-flavor squark pair production, (upper right) bottom squark pair production, and (lower) top squark pair production. The area enclosed by the thick black curve represents the observed exclusion region, while the dashed red lines indicate the expected limits and their ±1 and ±2 standard deviation ranges. The thin black lines show the effect of the theoretical uncertainties in the signal cross section. The white diagonal band in the top squark pair production exclusion limit corresponds to the region $\left|m_{\tilde{t}} - m_t - m_{\tilde{\chi}_1^0}\right| < 25$ GeV and small $m_{\tilde{\chi}_1^0}$. Here the efficiency of the selection is a strong function of $m_{\tilde{t}} - m_{\tilde{\chi}_1^0}$, and as a result the precise determination of the cross section upper limit is uncertain because of the finite granularity of the available MC samples in this region of the $(m_{\tilde{t}}, m_{\tilde{\chi}_1^0})$ plane. In the same exclusion limit, the gray diagonal line corresponds to $m_{\tilde{t}} = m_t + m_{\tilde{\chi}_1^0}$. Signal cross sections are calculated at approximately NNLO+NNLL order in α_s [40], assuming unity branching fraction for the indicated decay.

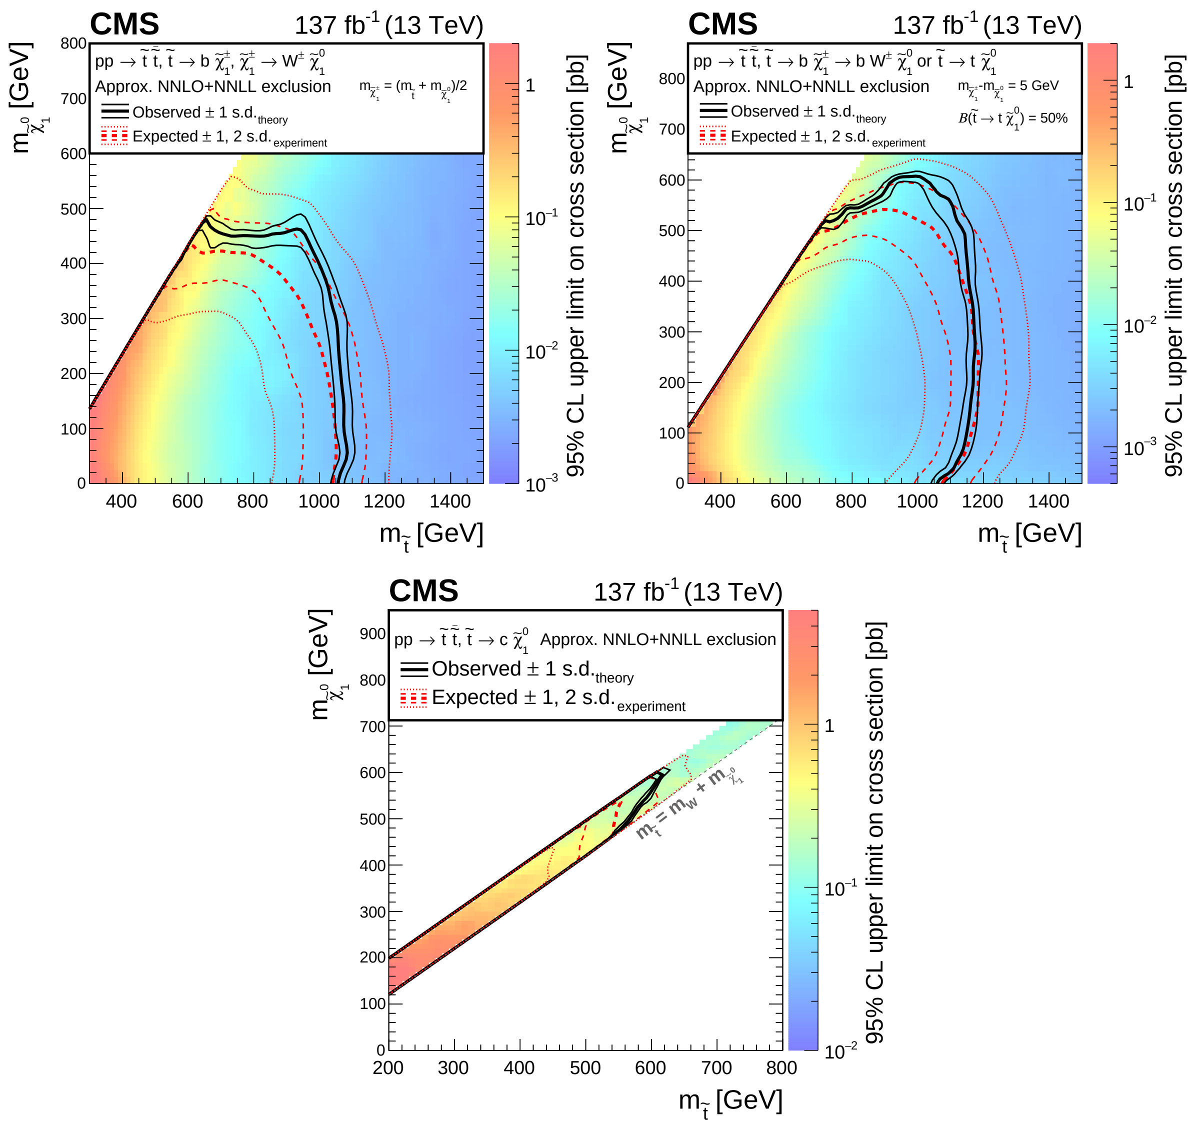

Figure 9.10: Exclusion limit at 95% CL for top squark pair production for different decay modes of the top squark. (Upper left) For the scenario where $pp \to \widetilde{t}\widetilde{t} \to b\bar{b}\widetilde{\chi}_1^{\pm}\widetilde{\chi}_1^{\pm}$, $\widetilde{\chi}_1^{\pm} \to W^{\pm}\widetilde{\chi}_1^0$, the mass of the chargino is chosen to be half way in between the masses of the top squark and the neutralino. (Upper right) A mixed-decay scenario, $pp \to \widetilde{t}\widetilde{t}$ with equal branching fractions for the top squark decays $\widetilde{t} \to t\widetilde{\chi}_1^0$ and $\widetilde{t} \to b\widetilde{\chi}_1^+$, $\widetilde{\chi}_1^+ \to W^+\widetilde{\chi}_1^0$, is also considered, with the chargino mass chosen such that $\Delta m\left(\widetilde{\chi}_1^{\pm}, \widetilde{\chi}_1^0\right) = 5$ GeV. (Lower) Finally, we also consider a compressed spectrum scenario where $pp \to \widetilde{t}\widetilde{t} \to c\bar{c}\widetilde{\chi}_1^0\widetilde{\chi}_1^0$. In this scenario, mass ranges are considered where $\widetilde{t} \to c\widetilde{\chi}_1^0$ branching fraction can be significant. The area enclosed by the thick black curve represents the observed exclusion region, while the dashed red lines indicate the expected limits and their ± 1 and ± 2 standard deviation ranges. The thin black lines show the effect of the theoretical uncertainties in the signal cross section.

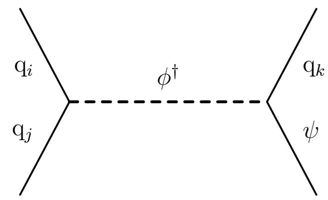

Figure 9.11: Diagram for the mono-ϕ model, where a colored scalar ϕ is resonantly produced, and then decays to an invisible massive Dirac fermion ψ and a SM quark.

m_ψ, respectively. In this model, the analysis of Refs. [76, 77] report best-fit parameters $(m_\phi, m_\psi) = (1250, 900)$ GeV and product of the cross section times branching ratio of about 0.3 pb. For this mass point, we find a modest (1.1 standard deviation) excess, and we set an upper limit on the product of cross section times branching ratio of about 0.6 (0.4 expected) pb.

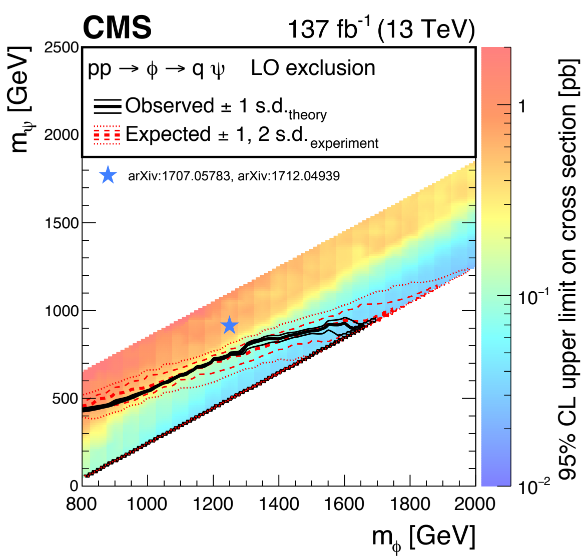

Figure 9.12: Exclusion limit at 95% CL for the mono-ϕ model. We consider the mass range where such a model could be interesting based on a reinterpretation of previous analyses [76, 77]. The area enclosed by the thick black curve represents the observed exclusion region, while the dashed red lines indicate the expected limits and their ±1 and ±2 standard deviation ranges. The thin black lines show the effect of the theoretical uncertainties in the signal cross section. The blue star at $(m_\phi,\ m_\psi) = (1250,\ 900)$ GeV indicates the best fit mass point reported in Refs. [76, 77]. Signal cross sections are calculated at LO order in α_s.

9.3.3 Leptoquark models

Finally, the search is interpreted in using models of leptoquark (LQ) pair production, similarly to a previous reinterpretation [82]. Leptoquarks, discussed in Sec. 1.3.2, are hypothetical particles that couple quarks and leptons. As they are necessarily colored, they also couple to gluons, giving the four main production modes shown in Fig. 9.13.

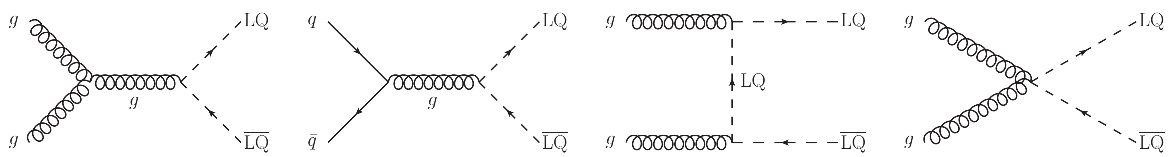

Figure 9.13: Diagrams for LQ pair production.

In the case of scalar leptoquarks ($\mathrm{LQ_S}$), the kinematics are identical to the pair production of squarks, where each squark decays to a quark and $\widetilde{\chi}^0_1$ with $m_{\widetilde{\chi}^0_1} = 0$. Differences in kinematic distributions introduced by vector leptoquarks ($\mathrm{LQ_V}$) are negligible, at least for the variables relevant to this analysis. Hence, the same signal simulation is used as for the pair production of squarks, with $m_{\widetilde{\chi}^0_1}$ set to 0, and rescaled to the theoretical LQ pair production cross section.

Figure 9.14 shows the 95% CL exclusion limits on the cross section times branching ratio of $\mathrm{LQ} \to q\nu$, as a function of LQ mass, where q is either a light-flavor quark, bottom quark, or top quark. Theory curves are the cross sections assuming 100% branching fraction to the relevant quark-neutrino pair, except for the pink curve in the bottom plot which assumes $\mathcal{B}(\mathrm{LQ_V} \to t\nu) = 50\%$ and $\mathcal{B}(\mathrm{LQ_V} \to b\tau) = 50\%$, corresponding to a proposed explanation [36] of various flavor physics anomalies.

Table 9.3 summarizes the limits on the masses of the leptoquarks excluded for the considered scenarios. These results extend the constraints on LQ masses by up to 200 GeV with respect to the limits of Ref. [82], providing the most stringent constraints to date on models of LQ pair production.

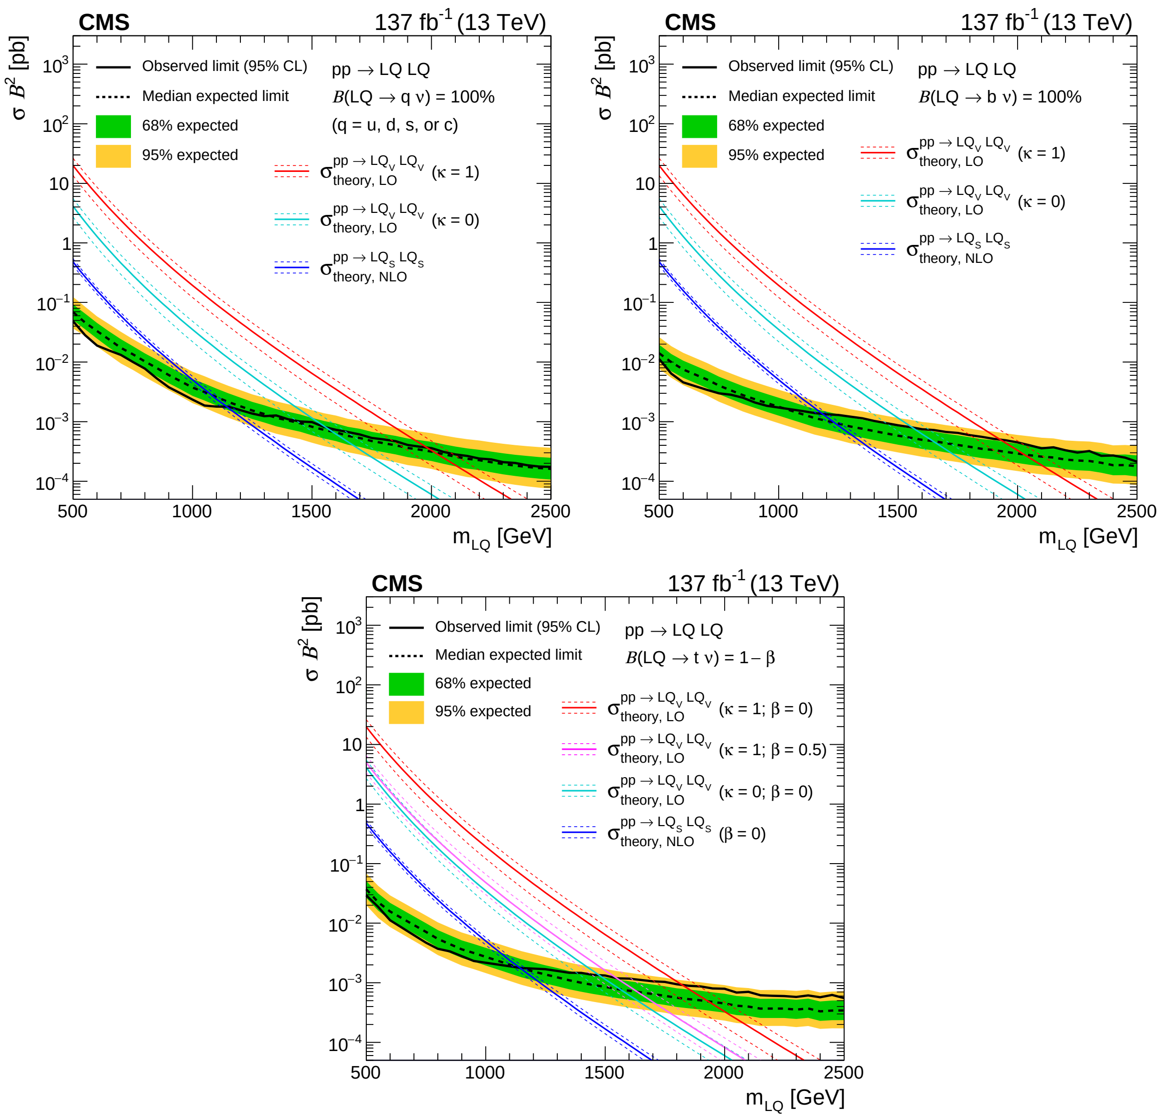

Figure 9.14: The 95% CL upper limits on the production cross sections as a function of LQ mass for LQ pair production decaying with 100% branching fraction ($\mathcal{B}$) to a neutrino and (upper left) a light quark (one of u, d, s, or c), (upper right) a bottom quark, or (lower) a top quark. The solid (dashed) black line represents the observed (median expected) exclusion. The inner green (outer yellow) band indicates the region containing 68 (95)% of the distribution of limits expected under the background-only hypothesis. The dark blue lines show the theoretical cross section for LQ_S pair production with its uncertainty. The red (light blue) lines show the same for LQ_V pair production assuming $\kappa = 1$ (0). (Lower) Also shown in magenta is the product of the theoretical cross section and the square of the branching fraction ($\sigma\mathcal{B}^2$), for vector LQ pair production assuming $\kappa = 1$ and a 50% branching fraction to $t\nu_\tau$, with the remaining 50% to $b\tau$. Signal cross sections are calculated at NLO (LO) in α_s for scalar (vector) LQ pair production.

Table 9.3: Summary of the observed 95% CL exclusion limits on the masses of LQs for the considered scenarios. The columns show scalar or vector LQ with the choice of κ, while the rows show the LQ decay channel. For mixed-decay scenarios, the assumed branching fractions ($\mathcal{B}$) are indicated.

	$\mathrm{LQ_S}$ mass [GeV]	$\mathrm{LQ_V},\ \kappa = 1$ mass [GeV]	$\mathrm{LQ_V},\ \kappa = 0$ mass [GeV]
$\mathrm{LQ} \rightarrow q\nu\ \ (q = u,\ d,\ s,\ \text{or}\ c)$	1140	1980	1560
$\mathrm{LQ} \rightarrow b\nu$	1185	1925	1560
$\mathrm{LQ} \rightarrow t\nu$	1140	1825	1475
$\mathrm{LQ} \rightarrow \begin{cases} t\nu\ \ (\mathcal{B} = 50\%) \\ b\tau\ \ (\mathcal{B} = 50\%) \end{cases}$	—	1550	1225

Chapter 10

milliQan: Searching for Millicharged Particles at the LHC

Here we report on a new experiment, milliQan, designed to search for elementary particles with charges much smaller than the electron charge. Since the initial proposal in 2015, a smaller-scale demonstrator was installed near the CMS interaction point and operated during the 2018 LHC running period to collect a data set corresponding to 37.5 fb^{-1} of proton-proton collisions. This data is used to exclude at 95% confidence level the existence of particles with masses between 20 and 4700 MeV and charges varying between $0.006e$ and $0.3e$, depending on mass. It is the first search at a hadron collider for particles with charges $\leq 0.1e$. Results have been submitted to the Physical Review D [83].

Section 10.1 briefly outlines some of the theoretical motivation for millicharged particles, Secs. 10.2–10.4 describe the design of the milliQan detector and the various calibrations performed to characterize its performance, Sec. 10.5 details the procedure for the simulation of signal and background processes, and Sec. 10.6 presents the design and results of the final analysis.

10.1 Motivation of a search for millicharged particles

One of the central mysteries of modern particle physics is the question of what makes up dark matter. It must consist of massive particles that interact at most very weakly with the SM, and no viable candidate exists among the currently known particles. Moreover, a decade's worth of data from the LHC has provided no evidence of any new particles that might provide an explanation.

We thus ask the question, what types of signatures might hypothetical dark matter particles produce that would escape detection at present experiments? One method of explaining dark matter is to add a new "dark sector" of particles beyond the SM that couples only weakly to the SM. As an example, we can add a "dark photon" A'_μ and a "dark fermion" ψ' charged under the new gauge field with charge e'. Allowing for kinetic mixing between A'_μ and the SM weak hypercharge field B_μ, the Lagrangian for this new dark sector can be written

$$\begin{aligned}\mathcal{L}_{\text{dark-sector}} =& -\frac{1}{4}A'_{\mu\nu}A'^{\mu\nu} \\ &+ i\bar{\psi}'(\gamma^\mu\partial_\mu + ie'\gamma^\mu A'_\mu + iM_{\text{mCP}})\psi' \\ &- \frac{\kappa}{2}A'_{\mu\nu}B^{\mu\nu},\end{aligned} \tag{10.1}$$

where the first line is the kinetic term for a massless dark photon, the second line contains the kinetic terms for a dark fermion with mass M_{mCP} as well as the interaction term with A'_μ, and the third line contains the mixing term between A'_μ and B_μ, with mixing strength parameter κ.

The mixing term can be eliminated by redefining $A'_{\mu\nu} \to A'_{\mu\nu} + \kappa B_{\mu\nu}$, resulting in an interaction term $\kappa e'\bar{\psi}'\gamma^\mu B_\mu\psi$ between ψ' and B_μ. Rewriting B_μ in terms of the physical photon and Z boson fields as $B_\mu = \cos\theta_w A_\mu - \sin\theta_w Z_\mu$, we find that the new dark fermion couples to the SM photon with electric charge $\kappa e'\cos\theta_w$, and couples to the SM

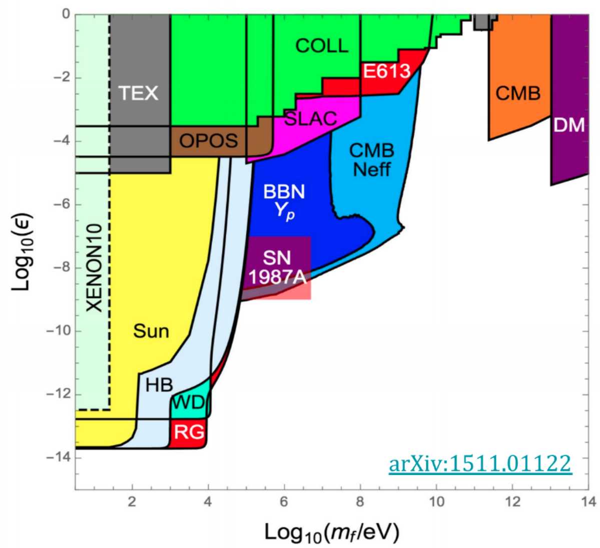

Figure 10.1: Existing exclusion limits for millicharged particles, coming from searches via colliders, solar effects, astronomical observations, and cosmological bounds. milliQan targets the unexcluded phase space with $\varepsilon \geq 10^{-3}$ and $10^{-1} < m_{\mathrm{mCP}} < 10^{2}$ GeV. (Image from [84])

Z with charge $\kappa e' \sin\theta_w$. The mixing strength κ must be small (otherwise the new dark sector would have been observed already), so we must have $\varepsilon \equiv \kappa e' \cos\theta_w / e \ll 1$, and we call ψ' a "millicharged particle" (mCP; note that the name is a bit of a misnomer because ε does not have to be exactly $\mathcal{O}(10^{-3})$)

Such millicharged particles have been searched for via a variety of methods, either directly through collider experiments or indirectly through solar effects, astronomical observations, or cosmological bounds. A summary of the present exclusion space in the mCP mass–charge plane is shown in Fig. 10.1, taken from [84].

There is a gap in the excluded phase space for $\varepsilon > 10^{-3}$ at the mass scales relevant at the LHC, roughly $10^{-1} < m_{\mathrm{mCP}} < 10^{2}$ GeV. mCPs at such masses and charges would be produced frequently at the LHC, but present experiments would not be able to detect them; direct sensitivity is lost for ε below a few times 10^{-1}, and a low cross section precludes missing energy searches. Therefore, a dedicated experiment is necessary to search for mCPs at the LHC.

10.2 Overview of the milliQan detector

The milliQan experiment, designed to search for mCPs using collisions at LHC P5, was proposed in 2015 [85,86]. It is located in a drainage gallery, elevated 43° above and 33 m from the CMS experiment, with 17 m of rock in between that naturally suppresses beam-based backgrounds. The proposed design consists of four stacked "layers" of arrays of plastic scintillator bars, with each scintillating bar coupled to a photomultiplier tube (PMT). The arrays are pointed at the interaction point (IP), such that a particle originating from a *pp* collision will generally pass through all four layers in a straight line. The bars are sensitive enough to detect individual photoelectrons produced by throughgoing mCPs, and requiring a simultaneous hit in all four layers drastically reduces background, which mostly consists of random overlap of pulses from PMT dark rate, environmental radiation, cosmic rays, and afterpulsing.

The full milliQan design is anticipated to consist of four layers of 20×20 scintillator arrays, each around 1 m^2 in total. In 2017–18, a smaller scale *demonstrator* was installed to study backgrounds and provide a proof-of-concept for the full-scale detector. This demonstrator consists of three 2×3 scintillator bar arrays, roughly 1% of the planned full detector.

The 18 scintillator bars each measure 5 cm $\times$ 5 cm $\times$ 80 cm, and are wrapped in layers of reflective and light-blocking materials to ensure optimal light-collection efficiency. 3D-printed plastic casings couple the bars to individual PMTs (two Hamamatsu R7725s, four Electron Tube 9814Bs, and twelve Hamamatsu R878s).

In addition to the scintillator bars, there are four 20 cm $\times$ 30 cm $\times$ 2.5 cm scintillator "slabs" placed at the front and rear of the detector as well as in between the layers, in

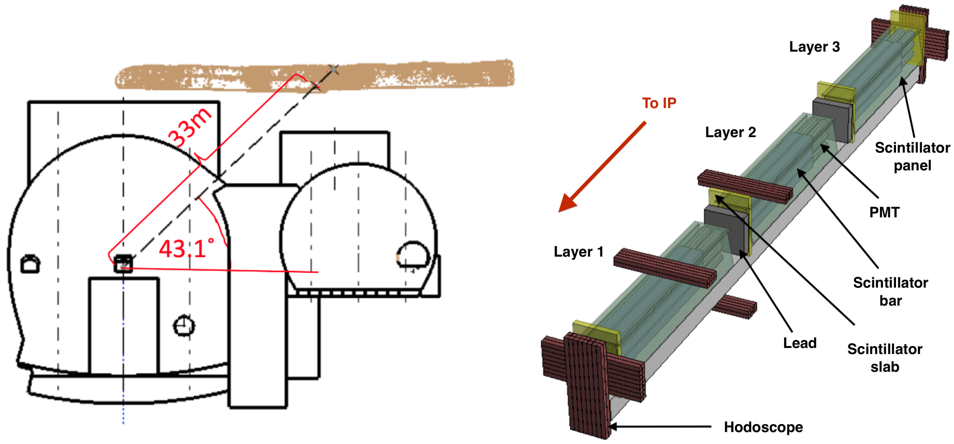

Figure 10.2: (left) Illustration of the location of the drainage gallery with respect to the CMS cavern. CMS is located in the large dome on the left; milliQan is elevated 43.1° above this, and 33 m away, with 17 m of rock in between. (right) 3D illustration of the demonstrator detector. The slabs are in yellow, the panels in translucent green, and the bars can be seen through the panels. The gray blocks in the layer gaps are lead bricks to block radiation.

order to tag/veto beam-based and cosmic particles. Additionally, each layer has three 18 cm × 102 cm × 0.7 cm scintillator "panels" covering the top and sides, in order to tag/veto cosmic particles and environmental radiation. Each of the slabs and panels is read out by a Hamamatsu R878 PMT.

There are 5 cm-thick lead bricks placed between each layer to reduce correlated pulses from radiation. All of the scintillators and lead bricks are mounted on a custom-designed aluminum support structure that can rotate the entire stack in multiple directions to facilitate alignment with the IP. A CERN engineering team performed such an alignment, so that the detector points to the IP to within a tolerance of just 1 cm over the 33 m distance. The location of the drainage gallery and an illustration of the demonstrator are shown in Fig. 10.2.

The 31 scintillator channels are read out by a pair of CAEN V1743 digitizers, which sample at 1.6 GHz and record a 640 ns waveform for each channel upon triggering.

The trigger can be configured to fire on single, double, or triple coincidence of peaks at an arbitrary threshold. For nominal data taking, the trigger is set to fire on a triple coincidence of pulses.

10.3 Bench tests for PMT calibration

The demonstrator detector makes use of three different “species” of PMTs, each of which has very different characteristics. Even PMTs within the same species differ somewhat in their behavior. It is therefore necessary to calibrate each PMT individually, so that data can be interpreted correctly and simulation can be adjusted to accurately model the behavior of the real detector.

Two separate calibrations are necessary. The first is the efficiency of the PMT to convert an optical photon into a photoelectron (PE). Once installed in the detector, this is intertwined with the efficiency of the scintillators themselves to generate and propagate the photons, so it is measured “in-situ”, and described in the following section. The other calibration involves the response of the PMT to a single photoelectron (SPE), which includes the pulse shape and pulse area distribution. The mean pulse area is also measured in-situ, as it can be affected by small magnetic fields and other effects, but we perform measurements of the pulse shapes and full area distributions in the lab using flashing LEDs, in order to (1) cross-check the in-situ measurements and (2) generate the necessary inputs for the pulse-injection step of the simulation, described in Sec. 10.5.

For the LED bench tests, we largely follow the method outlined in [87], which allows for the measurement of the non-Gaussian low-area tails of the SPE area distributions, arising from non-optimal trajectories within the PMT (e.g. the photon skipping the cathode and directly hitting the first dynode, or a photoelectron skipping a dynode stage). The laboratory setup is sketched in Fig. 10.3. A blue LED and a PMT are

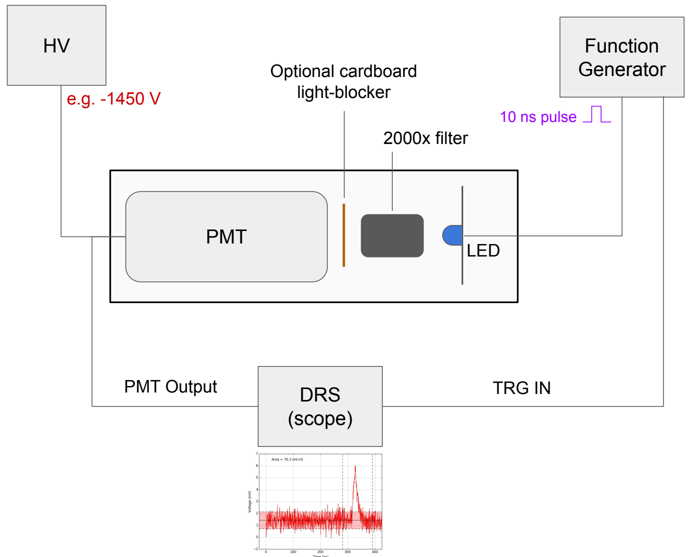

Figure 10.3: Diagram of the laboratory setup for LED bench tests of the PMTs. An LED and the PMT are mounted in a light-tight 3D-printed casing, with a 2000x optical filter in between. The LED is flashed in short ∼10 ns pulses, and the PMT signal is read out by a DRS evaluation board [88]. The board is triggered not by the PMT but by the LED pulse generator, so that even "blank" (0 PE) events are recorded. An optional cardboard light-blocker can be inserted between the LED and PMT, in order to collect a pure sample of 0 PE events.

placed in a light-tight enclosure, and the LED is flashed in short ∼10 ns bursts so that the PMT generates $\mathcal{O}(1)$ PE on average. The PMT is read out with a DRS board [88], which is triggered on the LED pulse so that even 0 PE events are recorded. An optional cardboard light-blocker can be inserted between the LED and PMT, to collect a pure sample of 0 PE events.

Upon each LED pulse, the PMT waveform is recorded. The pulse (when there is ≥ 1 PE) occurs at the same time in each event, so there is no peak-finding necessary and the area can be computed by integrating over a fixed time window. The offset voltage level and noise levels are estimated from the sideband before the pulse time window; this offset is subtracted off before integrating to get the area. Doing this for many events, one can build up a histogram of pulse areas, and if the LED intensity is set correctly one

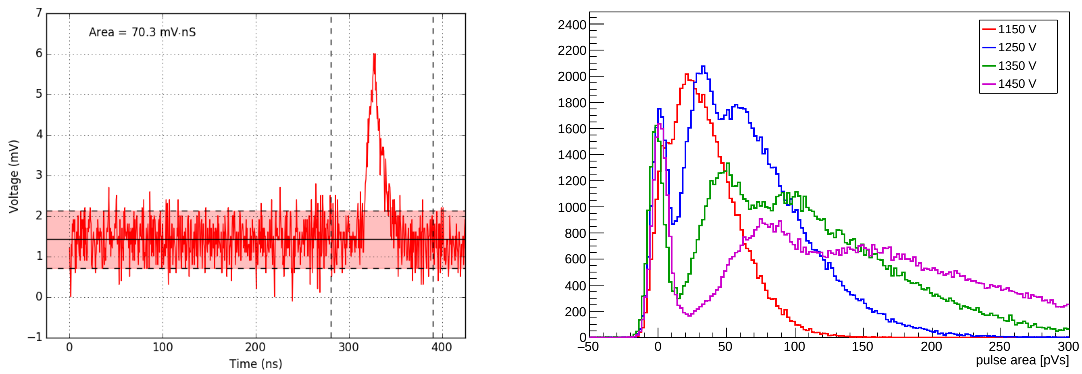

Figure 10.4: (left) An example SPE waveform from and R878 PMT. The light red band shows the measured offset (∼1.4 mV) and noise level (±0.6 mV). The vertical dashed lines indicate the fixed time window in which the waveform is integrated to compute the pulse area. (right) The pulse area distributions for a variety of HVs, for an R878 PMT. The LED intensity is set so that the average number of PE is near 1, so one can resolve distinct peaks for 0, 1, and 2 PE events in the area distributions. As the HV is increased, the peaks move to the right.

can observe distinct peaks in the distribution corresponding to 0, 1, 2, . . . PE events. Fig. 10.4 shows an example R878 waveform and the pulse area distributions for a few different high-voltages (HVs).

The actual calibration method makes use of the fact that the number of observed photoelectrons after the LED light is heavily filtered is a Poisson process. Two data sets are recorded: one with no cardboard light-blocker, and the LED intensity set so that there is $\mathcal{O}(1)$ PE on average per event, and one with the light-blocker inserted (but with the LED running at the same intensity, so that any electronic effects are still captured).

The LED-blocked data set is a pure sample of 0 PE events, and the resulting pulse area distribution (which is just a Gaussian near 0) can be scaled to fit the left edge of the area distribution of the non-blocked data set. It is important to fit only the left edge, and not the full 0-PE peak, as the right side can have contributions from sub-optimal 1 PE events. By taking the ratio of the areas of the scaled LED-blocked histogram and the non-blocked histogram, one can get the fraction f_0 of 0 PE events. Since N_{PE} is Poisson, we can then compute the mean and variance as

$$\mathrm{Var}(N_{\mathrm{PE}}) = \langle N_{\mathrm{PE}} \rangle = -\log(f_0). \tag{10.2}$$

Once we have this, it is straightforward to compute the mean and variance of the area distribution of an SPE pulse. Denoting A_{SPE} as the area distribution of an SPE pulse, A_{nb} as the area distribution of the non-blocked data set, and A_{b} as the area distribution of the LED-blocked data set, the mean and variance are given by

$$\langle A_{\mathrm{SPE}} \rangle = \frac{\langle A_{\mathrm{nb}} \rangle - \langle A_{\mathrm{b}} \rangle}{\langle N_{\mathrm{PE}} \rangle} \tag{10.3}$$

$$\mathrm{Var}(A_{\mathrm{SPE}}) = \frac{\mathrm{Var}(A_{\mathrm{nb}}) - \mathrm{Var}(A_{\mathrm{b}})}{\langle N_{\mathrm{PE}} \rangle} - \langle A_{\mathrm{SPE}} \rangle^2 \tag{10.4}$$

(the variance equation is not trivial, see [87] for details).

Fig. 10.5 shows an example of this process for an R878 and R7725. The dark blue histograms are the non-blocked area distributions, and the red histograms are the LED-blocked area distributions. The red histograms are scaled to fit the left edge of the blue histograms, and then used to find $\langle N_{\mathrm{PE}} \rangle$. The Eqs. 10.3 and 10.4 are used to compute the mean and variance of the SPE area distribution. Note in particular that the means are noticeably lower than the peaks from 1-PE events. This is due to the large non-Gaussian low-area tails in the SPE response, visible as the dotted histogram (just the subtraction of red from blue).

In addition to the mean SPE response that is used to calibrate each individual PMT, we are also interested in the full *distribution* of areas for an SPE, in order to use them for pulse injection, described in Sec. 10.5. To do this, we subtract the scaled LED-blocked area distribution from the nominal non-blocked distribution, to get the area distribution

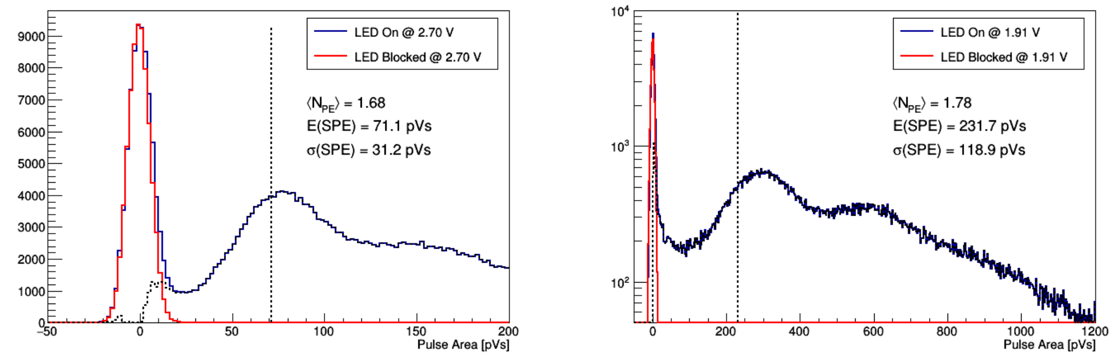

Figure 10.5: Example calibrations for an R878 (left) and R7725 (right). The dark blue histograms are pulse area distributions from the nominal LED setup, and the red histograms are area distributions with the LED blocked by cardboard. This red histogram is scaled to fit the left edge of the navy histogram in order to estimate the 0 PE component. Since the distribution of N_{PE} is Poisson, the fraction of 0 PE events can be used to compute $\langle N_{\text{PE}}\rangle$, from which one can get the mean and standard deviation of the SPE area distribution. The vertical dotted line indicates the computed mean, notably lower than the 1-PE peak due to the large non-Gaussian low-area tails.

of events with ≥1 PE. We then simultaneously fit the 1 PE and >1 PE distributions to this histogram. The SPE response is modeled as a Gaussian plus exponential left tail, with PDF given by

$$\psi_{\text{SPE}}(x) = G(\mu, \sigma)(1 - e^{-x/a})e^{-x/b}, \tag{10.5}$$

where $G(\mu, \sigma)$ is a standard Gaussian, and μ, σ, a, and b are free parameters to be fit. The multi-PE response is modeled as a sum of Gaussians, with means $n\mu$ and standard deviations $\sqrt{n}\,\sigma$, where n is the number of PE (we only consider up to $n = 3$). We then fit all four parameters simultaneously, and take ψ_{SPE} as the pulse area distribution for SPEs to be used for pulse injection. The fitted functions for all three types of PMTs are shown in Fig. 10.6.

One final piece of information that we need to extract from the LED tests are SPE pulse shape templates for all PMT types. Like the pulse area distributions described

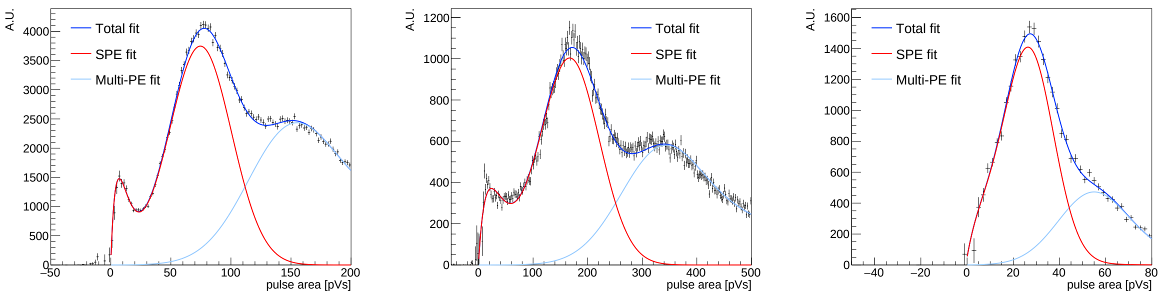

Figure 10.6: Measured area distributions for the R878, R7725, and ET (at 1450, 1400, and 1700 V), after subtracting off the 0-PE component. Black points are the raw 0 PE-subtracted area distributions. The red curve is the fitted SPE response (modeled as Guassian plus exponential left tail, given in Eq. 10.5), and the light blue curve is the fitted multi-PE response (one Gaussian each for 2- and 3-PE components). The red curve is used to draw random areas during pulse injection, described in Sec. 10.5.

above, these are used during pulse injection, described in Sec. 10.5. The procedure is as follows:

1. Collect a set of waveforms at low enough LED voltage so that $\langle N_{\mathrm{PE}} \rangle \sim 1$.
2. We want to average together many pulses to remove random noise, but there is fixed electronic noise that is the same in every waveform and does not get averaged away. To remove this, first average together a control sample of 0-PE events (chosen by selecting events whose pulse area is in the bulk of the 0-PE area distribution). These events have the same fixed noise but no actual signal.
3. Now, subtract this averaged control waveform from each raw SPE waveform (chosen by selecting events whose pulse area is near the peak of the SPE area distribution), and average the resulting waveforms together. This produces a smooth, noise-free template waveform. However, due to small $\mathcal{O}$(5–10 ns) time shifts of individual pulses, the template is "smeared", making it too short and wide.
4. To correct this, modify the above procedure by shifting each control-subtracted SPE waveform in time so that their peaks line up before averaging. The time of each

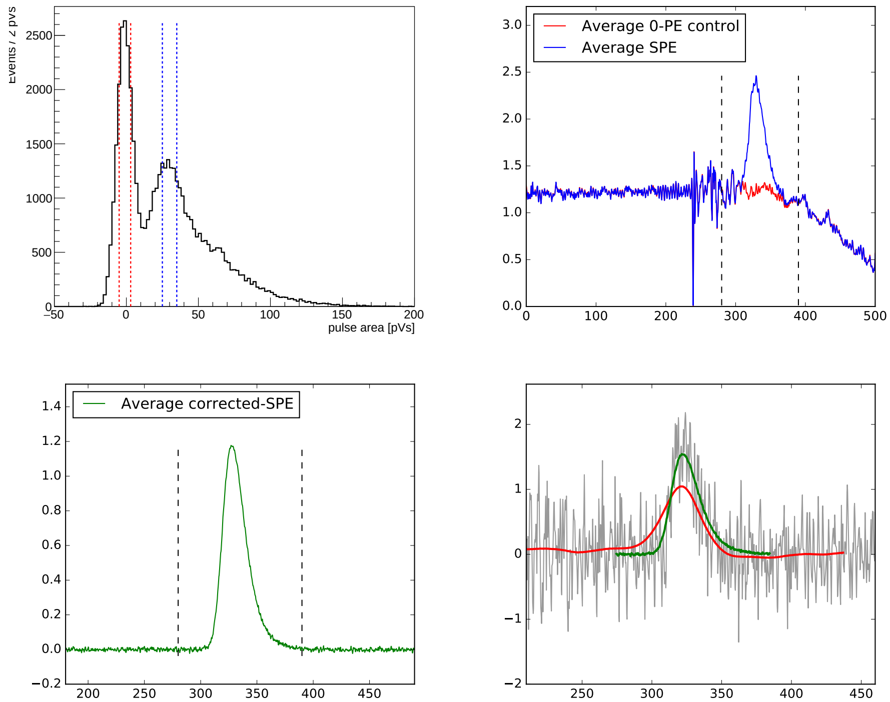

Figure 10.7: Illustration of the procedure for constructing pulse-shape templates, in this case for the R878. (top left) Area distribution of PMT pulses. Events between the red lines are selected for the 0-PE control sample, used to removed fixed electronic noise. Events between the blue lines are used as the SPE waveforms. (top right) The averaged SPE waveforms (blue) and 0-PE control sample (red). (bottom left) The control-subtracted waveform (blue minus red from the top right). (bottom right) Illustration of the timing procedure. The red curve is the convolution of the raw waveform (gray) with the template (green). The maximum is taken as the pulse time, and the SPE waveforms are shifted so that this is aligned before doing the final averaging.

pulse is computed by convolving the raw waveform with the non-time-corrected template from step #3, and finding the maximum point of this convolution.

This procedure is illustrated in Fig. 10.7 for the R878, and the resulting templates for all PMT types are shown in Fig. 10.8.

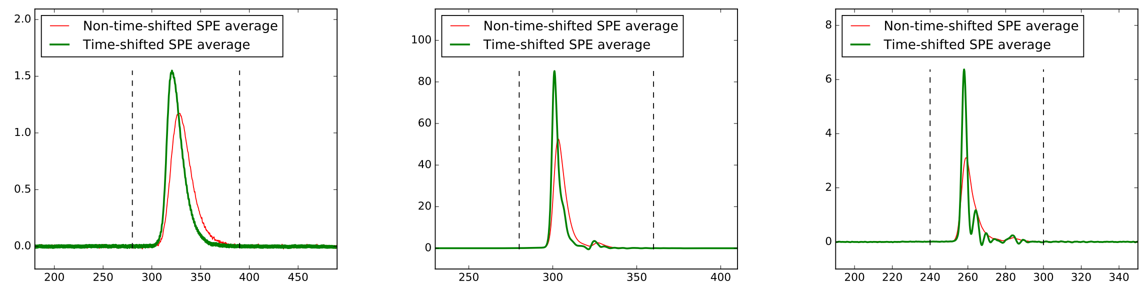

Figure 10.8: Pulse shape templates derived for the R878 (left), R7725 (center), and ET (right) PMTs. The green curves are the time-corrected templates, used for the pulse injection. The red curves are the templates before time correction, which are not correct and only shown here for illustration of the effect of time shifts.

10.4 In-situ calibrations

In addition to the LED-based bench calibrations described in the previous section, we also perform calibrations "in-situ" on the installed detector. These can be divided into four categories, the first two relating to "pulse size" and the second two relating to "pulse timing":

- SPE calibration: PMT-dependent characteristic size of a pulse from a single photoelectron
- $\langle N_{\mathrm{PE}} \rangle$ calibration: how many photoelectrons does a particle of a given charge produce in each channel?
- Channel-dependent timing calibration: per-channel time offsets due to cable lengths, scintillator shape differences, etc.
- Pulse size-dependent timing calibration: timing differences due to the pulse-finding algorithm responding differently to differently sized pulses

10.4.1 SPE and $\langle N_{\text{PE}} \rangle$ calibrations

First, we measure the mean pulse area of SPEs. These are expected to differ slightly from the areas measured with the LED tests, as small magnetic fields and other environmental effects can affect the response of the PMTs. Second, we perform a "charge calibration", or $\langle N_{\text{PE}} \rangle$ calibration, which measures the number of photoelectrons generated by a particle of a given charge traveling a given distance through the scintillator. This measurement is only possible in-situ as it involves properties of the scintillators and their couplings to the PMTs, in addition to the PMT-specific effects.

The SPE area is measured by selecting PMT afterpulses coming after a pulse from a vertical cosmic muon. Events are selected in which a top panel and three bars in a vertical line have hits consistent with a vertically traveling muon. The areas of all secondary pulses are plotted, and a Gaussian is fit to the peak. The mean of this Gaussian is taken as the peak SPE area for that channel.

The $\langle N_{\text{PE}} \rangle$ calibration is done differently for the slabs, panels, and bars. For the slabs, events with a throughgoing beam muon are identified, and the peak of the resulting pulse area distribution in each slab, scaled by the measured mean SPE area, is taken as the $\langle N_{\text{PE}} \rangle$ calibration. For the panels, a similar procedure is used, but with cosmic muons.

Performing the $\langle N_{\text{PE}} \rangle$ calibration for the bars is more challenging, as the PMTs saturate for both beam and cosmic muons so a direct measurement is not possible. Instead, an indirect method, in which pulse areas are extrapolated from lower voltages at which the PMTs do not saturate, is used. The procedure is as follows:

1. Measure the mean pulse area for cosmic pulses, at a variety of HVs low enough that the PMT does not saturate
2. Measure the mean SPE area at 2–3 points near the operating voltage of the PMTs
3. Empirically, the pulse areas scale with HV as a power law function. Jointly fit

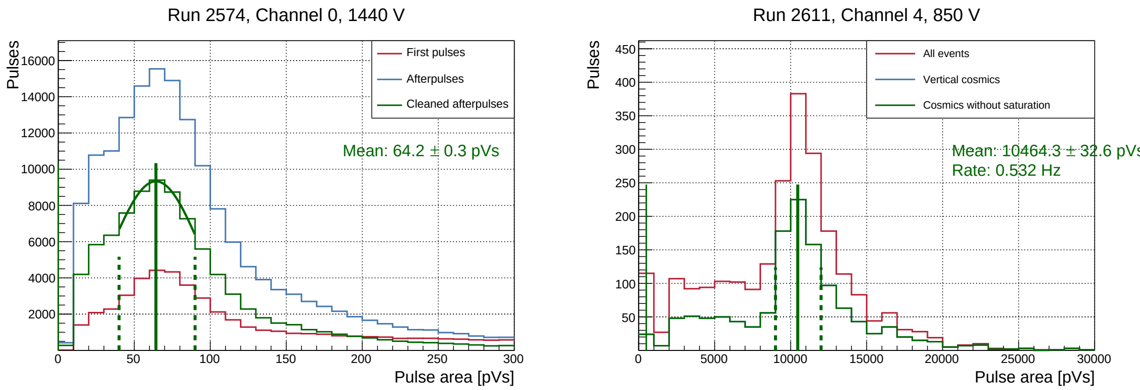

Figure 10.9: (left) Example measurement of the mean SPE area of Channel 0 (R878) at 1440 V. A selection of cleaned afterpulses is identified (green) and a Gaussian is fit to the peak. This measurement is used directly as the SPE calibration, and also as a component of the $\langle N_{\mathrm{PE}} \rangle$ calibration. (right) An example measurement of the mean cosmic pulse area for Channel 4, at a low HV of 850 V.

power law functions separately to the cosmic and SPE points, with the restriction that the exponent is the same. Then the ratio between these two functions is the mean N_{PE} of a cosmic pulse.

Example measurements of the mean SPE and cosmic pulse areas are shown in Fig. 10.9. The $\langle N_{\mathrm{PE}} \rangle$ calibration procedure is illustrated in Fig. 10.10. The top plot shows the measured cosmic (upper left) and SPE (middle right) areas, and the dashed lines show the fitted power law function through the SPE points. The bottom plots show an example validation of the procedure; on the left is a fit to the cosmic points, and on the right is an independent fit to a number of SPE points at higher voltages. The ratio of the fitted exponents is 1.01 ± 0.09, showing that the fit is consistent between the two sets of points as assumed. This is done independently for each channel; good agreement is seen in all cases.

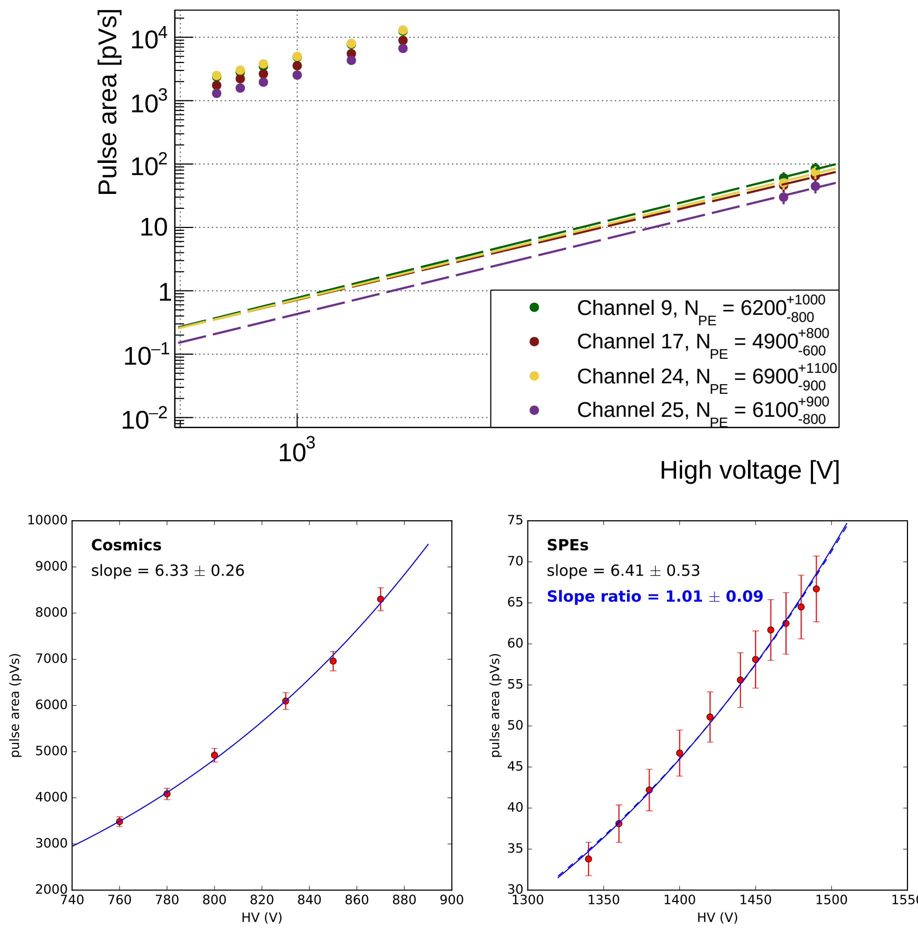

Figure 10.10: (top) Example $\langle N_{\mathrm{PE}} \rangle$ calibration of the four ET channels. The points in the upper left are the measured low-voltage cosmic areas, and the points on the right are the mean SPE areas near the operating voltage. Power law functions are jointly fit to the cosmic and SPE points, with the restriction that the exponent is the same; these fits are shown as dashed lines through the SPE points. The ratio of the two fits is taken as the $\langle N_{\mathrm{PE}} \rangle$ calibration. (bottom) A validation of the procedure for Channel 1. A key assumption of the method is that the power law exponents are the same for the cosmic and SPE points. Here, we fit each set separately and find that the exponent ratio is 1.01 ± 0.09.

10.4.2 Timing calibrations

A throughgoing particle will leave pulses at different times in different channels, due to time-of-flight, scintillator shape differences, varying cable lengths, and other effects. We calibrate these differences away such that a beam-based throughgoing particle traveling at the speed of light is expected to generate a pulse at the same time in each channel. This is done as follows:

1. Calibrate bars within each "slice" (left/right side of each layer), by tagging vertical cosmics that hit the top pannel and each bar in the slice. The mean cosmic pulse time is taken as the calibration.
2. Calibrate the two slices in each layer, by tagging cosmics that hit the top panel and at least one bar in each slice.
3. Calibrate the slabs by tagging beam-based muons passing through all four slabs.
4. Calibrate each layer by tagging beam-based muons and comparing to the time in the nearest slab.
5. Calibrate the panels by tagging cosmics that hit a panel and at least one bar in the same layer as the panel.

A validation of this procedure is shown in Fig. 10.11, in which the calibrated time difference between bar pulses in layer 3 and layer 1 for muon events is plotted. There is a single Gaussian distribution centered at 0 ns for beam muons, as expected if the calibration is done correctly.

In addition to the overall channel-dependent time calibrations, we also correct for the known effect of "time walk", in which small pulses are reconstructed at later times than large pulses. This is due to the algorithm for computing pulse time, which finds the time at which a signal rises above some fixed voltage threshold. Large pulses rise quicker than

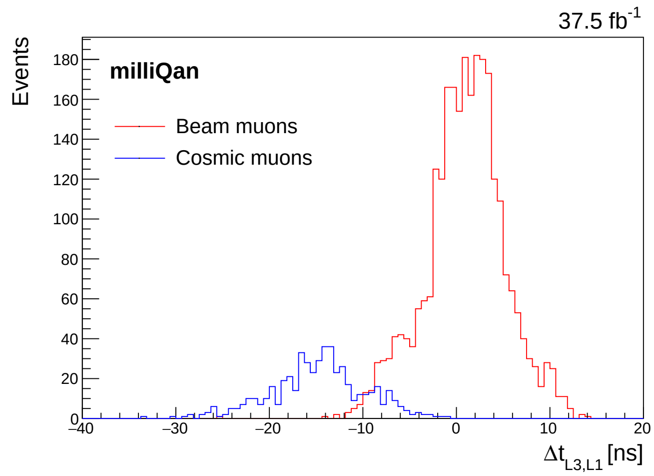

Figure 10.11: Calibrated time difference between the bar pulses in layer 3 and layer 1, for events in which a muon is tagged in all four slabs and bars in all three layers. Cosmic and beam muons are distinguished based on the time difference in the first and fourth slabs. The beam muon distribution is centered at 0 ns as expected, and has a resolution of ∼4 ns. The cosmic distribution is centered at ∼14 ns, as expected from time-of-flight considerations ($\Delta t = 2 \times 2$ m$/c = 13.3$ ns).

small pulses, and are hence assigned earlier times. This phenomenon is illustrated in Fig. 10.12, in which we plot the pulse size vs. time delay for pulses in bars neighboring a tagged throughgoing muon (likely caused by showering from the muon). One observes a distinct trend in time as the pulse size is varied.

The mean delay as a function of pulse size is extracted in bins of pulse area, and this is used to correct individual pulse times. This is done separately for each PMT species, as the effect highly depends on the specific pulse shape. Further, there are slight differences in the size of the effect between data and simulation, so separate calibrations are derived.

Finally, we also observe that the resolution (i.e. the spread in time for a given pulse size in Fig. 10.12) is slightly worse in data than in simulation. To account for this, we smear the reconstructed pulse times in simulation by an area-dependent factor such that the time resolutions match those measured in data.

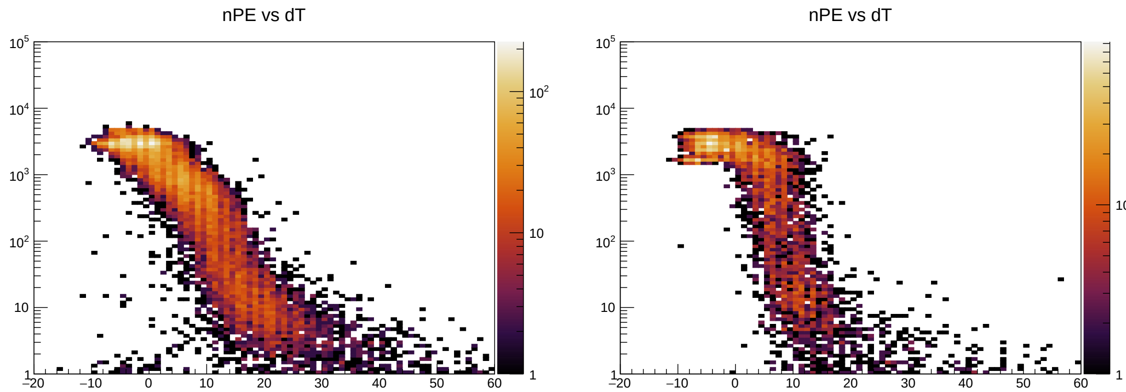

Figure 10.12: Pulse size in terms of N_{PE} (y axis) vs. time delay in ns (x axis), for R878s (left) and ETs (right). We tag events in which a beam muon passes through all four slabs and three bars in a straight line, and use pulses in neighboring bars caused by e.g. showering from the muon. Time delay is relative to the time in the preceding slab. One observes a delay of up to $\sim$20 ns between pulses of different sizes. The effect is more pronounced in the R878.

10.5 Simulation

Estimated signal yields are obtained via simulation, by generating events of mCP pair production, propagating to the milliQan detector, and simulating the detector and electronic responses. Additionally, muons from both beam-based processes and cosmic rays are simulated for calibration and background studies.

10.5.1 Event generation

Any process at the LHC that produces an e^+e^- pair via a virtual photon can also produce a $\chi^+\chi^-$ pair (we use χ as the symbol for an mCP). This includes the direct vector meson decays $V \to e^+e^-$, where $V = \rho,\ \omega,\ \phi,\ \psi$, or Υ, the Dalitz decays $A \to e^+e^-\gamma$, where $A = \pi^0,\ \eta$, or η', and the Dalitz decays $\omega \to e^+e^-\pi^0$ and $\eta' \to e^+e^-\omega$. Additionally, $\chi^+\chi^-$ pairs can be produced via the Drell-Yan process, as with couplings to the photon and Z of εe and $\varepsilon e \tan\theta_w$, respectively.

We generate Drell-Yan decays with the MadGraph5_aMC@NLO event generator [50], using the Lagrangian 10.1, with a cut on the invariant mass of the $\chi^+\chi^-$ pair of 2 GeV. The Drell-Yan production mode is subdominant when the mass of the mCP is below half the Υ mass, around 5 GeV.

For mCP pairs produced through meson decay, we perform the two-body/Dalitz decays manually and store the resulting mCP four-vectors. This requires two pieces of information for each process: (1) the branching ratio of the meson to a $\chi^+\chi^-$ pair (for Dalitz decays, we need the *differential* width as a function of the $\chi^+\chi^-$ invariant mass), and (2) the differential cross sections to produce the parent meson as a function of p_{T} and η..

Branching ratios for direct vector meson decays can be computed by simply scaling the e^+e^- BR by a phase space factor:

$$\frac{\Gamma(V \to \chi^+\chi^-)}{\Gamma(V \to e^+e^-)} = (Q/e)^2 \frac{(1-4x_\chi^2)^{1/2}(1+2x_\chi^2)}{(1-4x_e^2)^{1/2}(1+2x_e^2)}, \tag{10.6}$$

where $x_* = m_*/m_V$ (this comes from the Van Royen-Weisskopf formula [89]).

For Dalitz decays $A \to \chi^+\chi^- X$, we can write the differential width as a function of the $\chi^+\chi^-$ invariant mass as [90]

$$\begin{aligned}\frac{d\Gamma}{dq^2} =& \frac{C\alpha}{3\pi q^2}\left(1+\frac{2m_\chi^2}{q^2}\right)\sqrt{1-\frac{4m_\chi^2}{q^2}} \\ & \left[\left(1+\frac{q^2}{m_A^2-m_X^2}\right)-\frac{4m_A^2 q^2}{(m_A^2-m_X^2)^2}\right]^{3/2} \ |F(q^2)|^2 \ \Gamma(A \to X\gamma),\end{aligned} \tag{10.7}$$

where q^2 is the invariant mass of the $\chi^+\chi^-$ pair, C is 2 if X is a γ otherwise 1, and $F(q^2)$

is a form factor that can be approximated in the Vector Dominance Model as

$$|F(q^2)|^2 = \frac{m_\rho^4 + m_\rho^2\Gamma_\rho^2}{(m_\rho^2 - q^2)^2 + m_\rho^2\Gamma_\rho^2}, \tag{10.8}$$

where m_ρ and Γ_ρ are the mass and total width of the ρ meson.

Cross sections for the production of parent mesons are acquired in a variety of ways. For direct [91] and B meson-mediated [92] production of J/ψ and ψ' mesons, cross sections and p_{T} distributions (including uncertainties) are taken directly from theory calculations. Theoretical calculations of Υ production are not reliable at low p_{T}, so we use differential cross sections measured by experiment. For $p_{\mathrm{T}} > 20$ GeV, we use cross sections measured at $\sqrt{s} = 13$ TeV [93], and at lower p_{T} we use measurements from 7 TeV [94], rescaled using the measured ratio of 13 to 7 TeV cross sections at slightly higher rapidity.

Differential cross sections for all light-flavor mesons except ϕ mesons are computed by generating minimum bias events in PYTHIA8 [51], with the Monash 2013 tune [95]. This is the tune that gives best agreement with several measurements of light meson rates and p_{T} spectra at the LHC [96–99], albeit in most cases at a center of mass energies lower than 13 TeV. The MC spectra for η (ρ, ω) with $p_{\mathrm{T}} < 3$ (1) GeV are scaled down by factors as large as two, based on these experimental comparisons. ϕ production is modeled with the PYTHIA6 generator [100] using the DW tune [101], since this MC setup best reproduces ϕ meson data [102]. All PYTHIA minimum bias MC is normalized to a total inelastic pp cross section of 80 ± 10 mb based on a measurement by ATLAS [45] (with additional uncertainty to account for slight disagreement between experimental measurements). An additional 30% uncertainty, uncorrelated between production modes, is assessed on the total cross sections to account for the modeling of meson rates per minimum bias event.

A plot summarizing the cross section times branching ratio values for all $\chi^+\chi^-$ production modes is shown in Fig. 10.13. At low masses, production is dominated by Dalitz

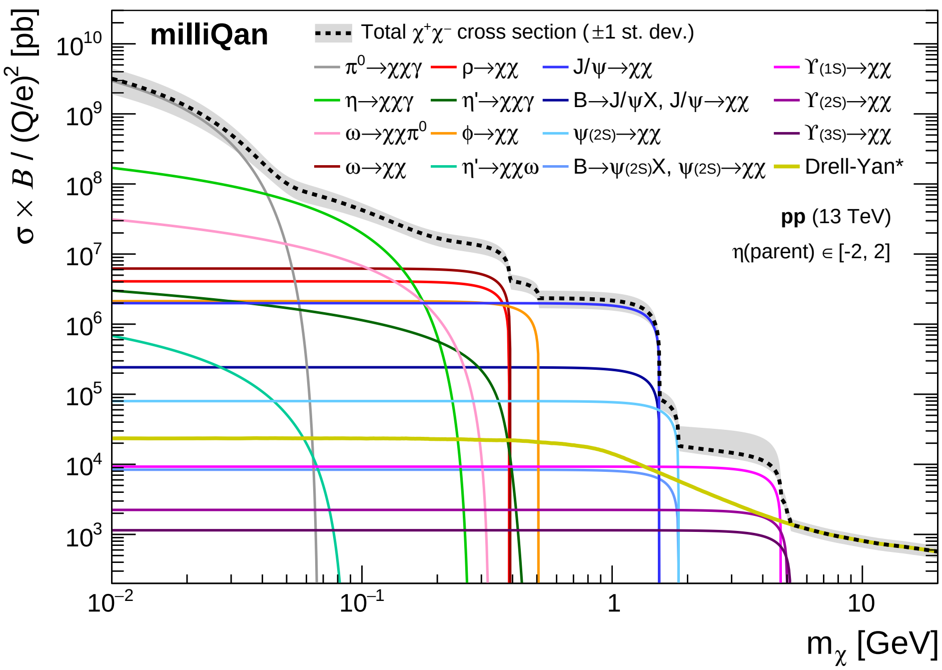

Figure 10.13: Cross section times branching ratio values for all considered production modes of $\chi^+\chi^-$ pairs, normalized to a charge of $Q/e = 1$. For non-DY modes, the mother particle must be within $|\eta| < 2$ (though the $\chi^\pm$ may have $|\eta| > 2$). For DY, the plotted cross section requires at least one mCP to have $|\eta| < 1$. The plateau in DY cross section below $m_\chi = 1$ GeV is due to a 2 GeV $\chi^+\chi^-$ invariant mass cut in MadGraph5. The gray band on the total cross section represents the total theoretical uncertainty in the cross section, adding in quadrature the uncertainties on each individual mode (with the exception of the 12.5% pp inelastic cross section uncertainty, which is correlated across all relevant modes).

decays of light mesons π^0, η. At around a few hundred MeV, direct decays of ω's and ρ's dominate. Near 1 GeV, J/ψ's provide almost all of the cross section, before the rate falls off substantially past half the J/ψ mass near 1.5 GeV. Υ's become the dominant production mode until they fall off at 5 GeV, after which only Drell-Yan contributes. Note that if this plot were continued, the rate would drop substantially again near half the Z mass at roughly 45 GeV.

In addition to mCPs, we are also interested in simulating both beam-based and cosmic muons in order to compare to data and calibrate/validate the simulation. Four beam-based muon production modes are considered: heavy-flavor (b or c) meson decays, light-flavor meson (usually $\pi^{\pm}$ or $K^{\pm}$) and tau decays, W boson decays, and Z boson decays. Differential cross sections for muons from heavy-flavor meson decay are taken from the FONLL tool [92] (the same that is used for $B \to \psi X$ decays for mCP generation). The overall $b\bar{b}$ cross section produced by this tool is known to be low based on a measurement from CMS [103], so the total rate is scaled up by a factor of 1.25 ± 0.25. Differential cross sections for light-flavor meson and tau decay are taken from PYTHIA8 MinBias generation with the CUETP8M1 tune. Rates from W and Z decay are taken from MadGraph5.

The composition and p_{T} distribution of muons that make it to the milliQan detector face (after the propagation described shortly) are shown in Fig. 10.14. Only muons with $p_{\mathrm{T}} \gtrsim 15$ GeV make it to the detector, as lower-energy muons are stopped short, either in CMS or the rock. Of the muons that do make it to milliQan, 53% are from b hadrons, 29% are from c hadrons, 14% are from light-flavor mesons or taus, 3% are from W bosons, and 1% are from Z bosons.

10.5.2 Propagation and detector response

Once mCP or muon four vectors are generated, they are propagated through a model of the CMS and its environment, including the magnetic field (see Fig. 2.5 for a colored map of the magnetic field that is used), detector material, and 17 m of rock between the IP and the drainage gallery.

The material geometry of CMS was extracted from a ROOT model, and simplified into a series of concentric iron cylinders, such that the total material budget is the same as that of the real detector. Iron is placed at radii

- $1.80 \leq r < 2.80$ m (tracker + HCAL)
- $3.15 \leq r < 3.50$ m (magnet)

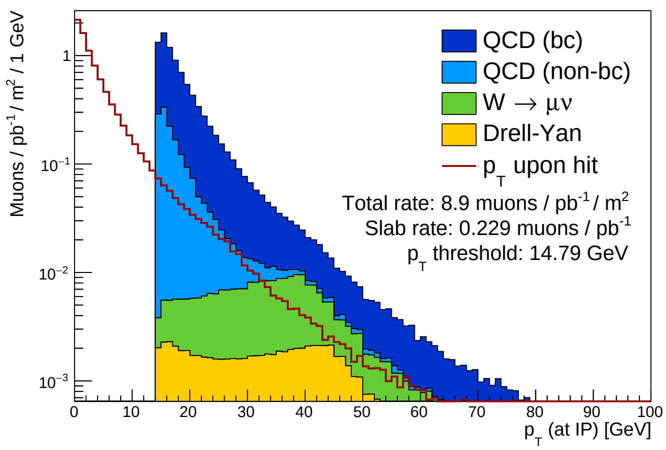

Figure 10.14: Rate of muons reaching the milliQan detector face as a function of production p_{T}, and broken down by production mode. There is a hard bound near 15 GeV due to energy loss in the material between the IP and milliQan. Above this threshold, the rate falls exponentially. The red curve shows the momentum upon intersection with milliQan, and is essentially the initial momentum distribution shifted left by ∼15 GeV.

- $3.85 \leq r < 4.05$ m (return yoke 1)
- $4.60 \leq r < 4.95$ m (return yoke 2)
- $5.35 \leq r < 5.95$ m (return yoke 3)
- $6.35 \leq r < 7.15$ m (return yoke 4)

In addition to the iron, solid rock is placed at $16 \leq r < 33$ m, until just before the face of the milliQan detector.

Particles are propagated with a fourth-order Runge-Kutta integrator, according to the relativistic kinematic equations

$$\frac{d\vec{x}}{dt} = \vec{v} = \frac{\vec{p}c^2}{E} = \frac{\vec{p}c^2}{\sqrt{(\vec{p}c)^2 + (mc^2)^2}} \tag{10.9}$$

$$\frac{d(\vec{p}c)}{dt} = 89.8755\, q\, \vec{v} \times \vec{B}, \tag{10.10}$$

where the form of the second equation is valid if B is in Tesla, t in ns, $\vec{v}$ in m/ns, q in units of e, and $\vec{p}c$ in MeV.

Multiple scattering is implemented using a small-angle Gaussian angle approximation, as described in the PDG [104]. At each time step, a θ_{RMS} is computed based on the particle momentum and charge, the radiation length of the material, and the distance traversed. Then two random deflection angles are drawn from a Gaussian of width θ_{RMS}, one for each transverse direction. The particle is also displaced a small amount in the transverse plane, in a way that is partially correlated with the angular deflection.

Finally, at each timestep the particle loses energy according to the Bethe equation

$$\left\langle -\frac{dE}{dx}\right\rangle = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}\log\frac{2m_ec^2\beta^2\gamma^2W_{\mathrm{max}}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2}\right], \qquad (10.11)$$

where all terms and parameters are defined in the PDG [104]. Note that energy loss is a stochastic process, and at each timestep the actual energy loss follows a highly skewed Landau distribution. However, the amount of material traversed in this application is quite large, so it should be sufficient to use the mean value at each timestep. At any rate, a systematic uncertainty is assessed on the interaction with matter that should cover for any inaccuracies in the modeling of energy loss.

Particles (either muons or mCPs) are propagated until 2 m before the face of the milliQan detector after which they are fed into a full Geant4 [52] simulation of the detector. This includes a full model of the drainage gallery, aluminum support structure, scintillator bars/slabs/panels with wrapping, lead bricks, and PMTs. Quantum efficiencies of the simulated PMTs are set individually on a per-channel basis so that the average number of PEs generated by a cosmic and/or beam muon match that measured data, from the calibrations described in Sec. 10.4. An illustration of a beam-based mCP traveling through the Geant4 model of the detector is shown in Fig. 10.15.

Figure 10.15: Rendering of a beam-based mCP traveling through a Geant4 model of the demonstrator, hitting all four slabs and three bars in a straight line. Light blue lines are generated optical photons.

The exact time of each detected photoelectron is saved, so that the Geant4 output events can be used for *pulse injection*, in which simulated PMT waveforms are overlaid on zero-bias (i.e. random trigger) data, in order to generate realistic detector signals that can be processed and analyzed with the same software as real data. For each simulated PE, a random SPE area is drawn from PMT-specific distributions derived via PMT bench tests, described in Sec. 10.3 (these distributions are corrected for slight differences between channels). Then a template SPE waveform, also derived from the PMT bench tests, is scaled to this area and overlaid on a random zero-bias event. These pulse-injected simulated events then look just like real data, with the exception that they do not model PMT saturation effects (however, this does not matter in practice).

10.5.3 Simulation validation and comparisons with data

Simulation is validated by comparing various predicted quantities and distributions to those measured in real data. The simplest comparison is the rate of muons coming from the beam and passing through all four slabs of the detector. This probes the accuracy of the muon cross sections and p_{T} distributions used for generation, as well as the propagation routine and material model between the IP and milliQan. The signal appears as four large pulses in the four slabs, with timing consistent with a particle coming from the beam (roughly, $t_4 - t_1 = 10 \pm 10$ ns); no requirement is placed on hits in the bars or panels.

In data, the rate is measured at 0.20 ± 0.01 muons/pb^{-1}, where the uncertainty comes primarily from the uncertainty in the total integrated luminosity. For the predicted rate from simulation, we must account for various systematic uncertainties. For the rate from heavy-flavor mesons, we scale the cross section by 1.25 ± 0.25 based on a comparison of FONLL with measured CMS data, so we take a 20% uncertainty. The rate from light-flavor mesons is sub-dominant but we take a 50% uncertainty. Finally, we assess a systematic uncertainty in the material modeling of CMS and the rock by varying the density of all materials up and down by 7%. This translates to a 25% uncertainty in muon rate due to the rapidly falling momentum spectrum. All together, the predicted rate from simulation is 0.25 ± 0.08 muons/pb^{-1}, agreeing with the rate measured in data to within one standard deviation.

We can also roughly measure the angle of muons as they pass through milliQan, which probes the accuracy of the multiple scattering and magnetic field modeling of the propagation. We do this by looking at various "patterns" of bars that record muon-like hits in each of the four-slab events. These patterns are shown in diagrams on the x axes of the plots in Fig. 10.16. The left plot shows milliQan from a top view, and the right plot from a side view. Black boxes mean we require a muon hit in the given position, white boxes mean we require *no* muon hit in the given position, and a hashed boxes means it can be either. Muon hits are OR'd over the orthogonal direction (so black actually means at least one hit over the relevant bars, and white means no muon hits in any of the relevant bars). We also make an adjustment to correct for imperfect modeling of muon "fake rates" (due to large muon-like hits from showering in a nearby bar) in simulation, which arises from the non-modeling of PMT saturation. The N_{PE} threshold in simulation to count a hit as muon-like is adjusted until the ratio of the middle bin in the right plot to the two neighboring bins is the same as in data (these all correspond to an angle of 0°, but the middle bin has a lower rate due to the higher probability of

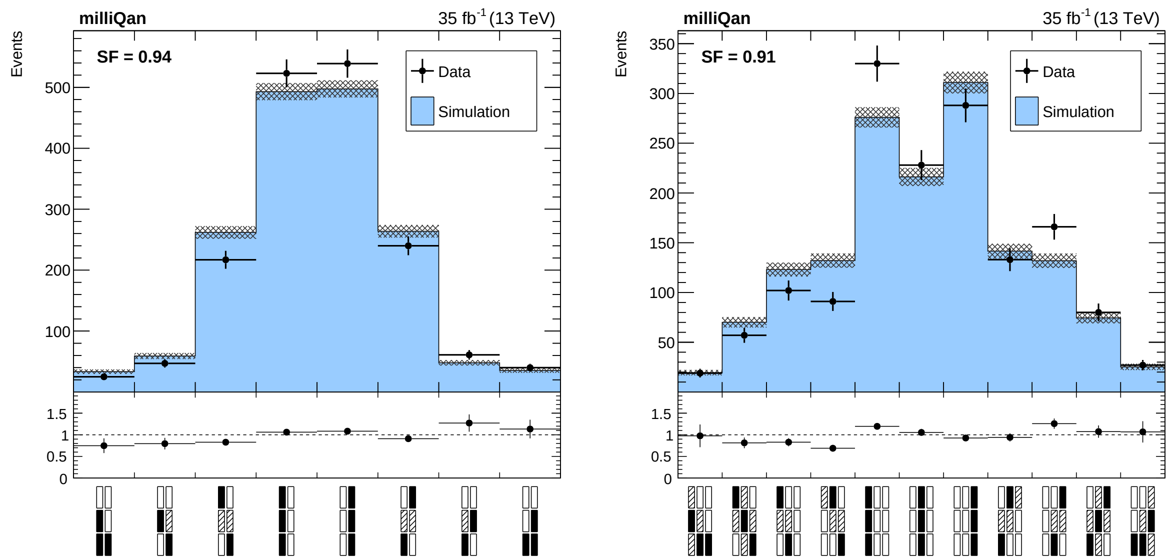

Figure 10.16: Rendering of a beam-based muon traveling through a GEANT4 model of the demonstrator, hitting all four slabs and three bars in a straight line. Light blue lines are generated optical photons.

a fake). After making this correction, we see that the observed angular distribution in data matches that from simulation quite well.

We validate the GEANT4 simulation by looking at N_{PE} distributions under various selections and comparing to data. Fig. 10.17 shows N_{PE} distributions in the bars for both data and simulation for four-slab beam muon events, either in the case that the muon passes through all slabs and no bars, or that the muon passes through a neighboring bar. The stacked histograms are colored by computing the fraction of PEs in each event that come from various production modes (generally, the muon knocks off an electron in some material, which then either directly leaves energy deposits or bremsstrahlungs a gamma ray that leaves deposits). Good agreement is seen everywhere, except at very low N_{PE} ($\lesssim 30$). Hits in this regime are predominantly come from a muon knocking off a very low-energy delta ray, which then bremsstrahlungs a low-energy gamma ray. The modeling of this in GEANT4 is likely not reliable, leading to the discrepancy. However,

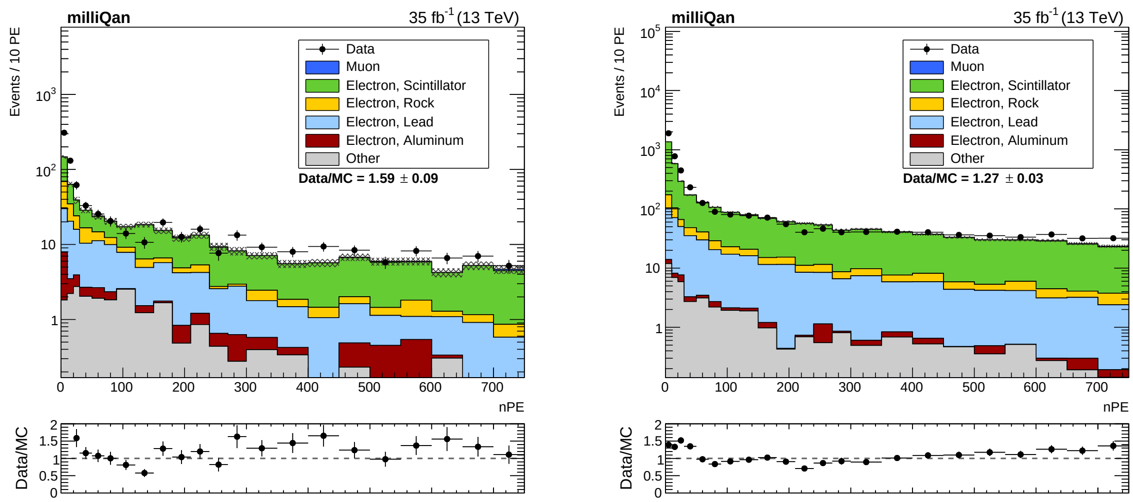

Figure 10.17: Comparison of N_{PE} distributions in the bars for events where (left) the muon hits all four slabs but no bars, and (right) the muon passes through a neighboring bar. Simulation is normalized by per-event weights computed from the cross sections of the various muon production modes described earlier. Good agreement is seen everywhere except at very low N_{PE} ($\lesssim 30$). Such small hits are generally produced by low-energy delta rays knocked off by the muon, which bremsstrahlung low-energy gamma rays which then leave small deposits. GEANT4 modeling of these low-energy delta rays is not perfect, and in any case this is not relevant for the simulation of mCPs. The colors indicate the original source of each PE.

this is of minimal concern, as such kinds of processes are irrelevant for the simulation of mCP events.

10.6 Background estimation and results

A number of sources may be expected to produce a pulse in each layer of the demonstrator and therefore provide a background:

- Dark rate overlap: each PMT has a dark current due to effects such as the thermal emission of electrons from the cathode. The simplest background source comes from random overlap of three such dark rate pulses. In addition, dark rate counts may overlap with a correlated double coincidence background from another source.

- Cosmic/beam muon showers: a large number of gammas, neutrons and electrons may be caused by an interaction of a cosmic ray muon with the rock in the demonstrator cavern. This may cause a pulse in each layer of the milliQan demonstrator. Such a background could also be expected from a beam muon which travels close to the demonstrator.
- Radiation: radiation in the cavern, scintillator bars or surrounding material can cause correlated deposits in several bars. The lead blocks placed between layers should reduce the probability of a three layer deposit arising from photons or electrons, however, neutrons will not be shielded.
- Afterpulses: afterpulses arising from correlated deposits may overlap and produce a triple coincidence signature in the demonstrator. The original correlated signature must not be triggered as in this case the afterpulses will fall in the readout deadtime and not be recorded.

10.6.1 Signal region selections

The backgrounds described above are largely reduced through the following baseline signal region selections:

- Activity in exactly 1 bar per layer, and these bars must be in a straight line path
- First pulse in each bar is the largest (i.e. no pre-pulses, but afterpulses are allowed)
- No activity in any panel
- No slab hit with $N_{\mathrm{PE}} > 250$
- Max/min bar $N_{\mathrm{PE}} < 10$
- max $\Delta t(\mathrm{bars}) < 15$ ns

Table 10.1: The five orthogonal signal regions used in the analysis. They are categorized based on number of slab hits (which must have $0.5 < N_{\mathrm{PE}} < 250$) and the minimum N_{PE} over the three bars. The final column lists the approximate mCP charge range probed by each signal region.

SR	# slab hits	min bar N_{PE}	Approx. targeted charge range
1	0	[2, 20]	≤0.014
2	0	> 20	0.014–0.03
3	1	[5, 30]	0.01–0.02
4	1	> 30	0.02–0.05
5	≥2	> 0.5	≥0.03

The first requirement reduces backgrounds that cause multiple hits per layer (mostly cosmic showers), as well as backgrounds that are uncorrelated in position between layers (almost all backgrounds). The panel veto reduces contributions from cosmic and beam-based showers, and the slab veto removes large pulses from throughgoing muons. The max/min bar N_{PE} and timing cuts reduce random backgrounds that are largely uncorrelated in both pulse size and time.

This baseline signal region is divided into five separate signal regions used for the final interpretation, based on the number of slab hits and minimum bar N_{PE}. These are listed in Table 10.1. The first two regions target small charges, roughly less than $\sim 0.02e$, as these generally produce no pulses in the slabs and small bar pulses. The second two regions target intermediate charges, roughly in the range 0.01 to $0.03e$, which produce $\mathcal{O}(1)$ slab hit and slightly larger bar pulses. The final region targets larger charges, which leave hits in two or more slabs.

10.6.2 Background estimation method

The expected background counts in each region are estimated by an "ABCD" method, in which the observed counts in orthogonal non-pointing regions are scaled by transfer factors measured in inverted-Δt regions. Specifically, we define regions

- A: pointing, max Δt(bars) < 15 ns
- B: non-pointing, max Δt(bars) < 15 ns
- C: pointing, max Δt(bars) > 15 ns
- D: non-pointing, max Δt(bars) > 15 ns

where "pointing" refers to the requirement that the single bars in each layer form a straight line path, and max Δt(bars) is the maximum time difference between any two bars. Additional requirements on the number of slabs and minimum bar N_{PE} are added to each region for estimating the individual signal regions listed in Table 10.1.

Region A is the region whose background count we want to estimate, and B, C, and D are orthogonal control regions. Assuming that the timing distribution is not correlated with whether or not the pointing requirement is satisfied, we can estimate

$$N_{\text{A}} = \frac{N_{\text{C}}}{N_{\text{D}}} N_{\text{B}}. \tag{10.12}$$

We check that the method works in a beam-off control region; results are shown in Table 10.2. Good agreement between predicted and observed counts is seen in all regions. We must also verify that there are no significant beam-based backgrounds that would break the method. To do this, we compare predictions between beam-off and beam-on data, adjusted for total data collection time. Again we see good agreement, indicating that any beam-based backgrounds are negligible. In any case, we verify from simulation that adding beam-based backgrounds to the non-beam backgrounds preserves the validity of the ABCD method.

We additionally perform the validations as a function of minimum bar N_{PE}. Comparisons of minimum bar N_{PE} distributions, between predictions/observations in beam-off data as well as predictions between beam-off and beam-on data, are shown in Fig 10.18.

Table 10.2: Two validations of the ABCD background estimation method: comparisons of predicted vs. observed in a beam-off control region (shows that the method works with non-beam backgrounds), and comparisons of predictions from beam-off and beam-on data (shows that beam-based backgrounds are negligible). Uncertainties shown are statistical. Note that the beam-on data set corresponds to 6% more live-time than the beam-off data set.

SR	Beam off		Beam on
	Pred.	Obs.	Pred.
1	$121.2^{+6.0}_{-5.9}$	131	$124.2^{+5.6}_{-4.0}$
2	$47.4^{+5.2}_{-4.8}$	45	$49.9^{+5.5}_{-4.8}$
3	$7.8^{+2.5}_{-1.8}$	9	$10.7^{+3.2}_{-2.0}$
4	$2.7^{+2.1}_{-1.1}$	4	$2.4^{+1.8}_{-0.8}$
5	$0.8^{+1.4}_{-0.4}$	1	$0.0^{+0.9}_{-0.0}$

This is done for the 0- and 1-slab regions, but not the ≥2 slab region as there are not enough events.

10.6.3 Systematic uncertainties

We assess a number of systematic uncertainties both on the background estimate and the signal yield prediction. On the background estimate, the primary uncertainty arises from the statistical uncertainty in the ABCD prediction, especially for the ≥1 slab regions where statistics are lower. These uncertainties are given in the final column of Table 10.2. Additionally, we assess a systematic based on the small prediction/observation disagreements in the beam-off data.

For the signal yield prediction, we assess a systematic uncertainty in the differential cross section of each production mode individually, as well as an overall proton-proton inelastic cross section uncertainty used for the light meson production modes. Details are given in Sec. 10.5.

A material modeling systematic is assigned by varying the density of material between

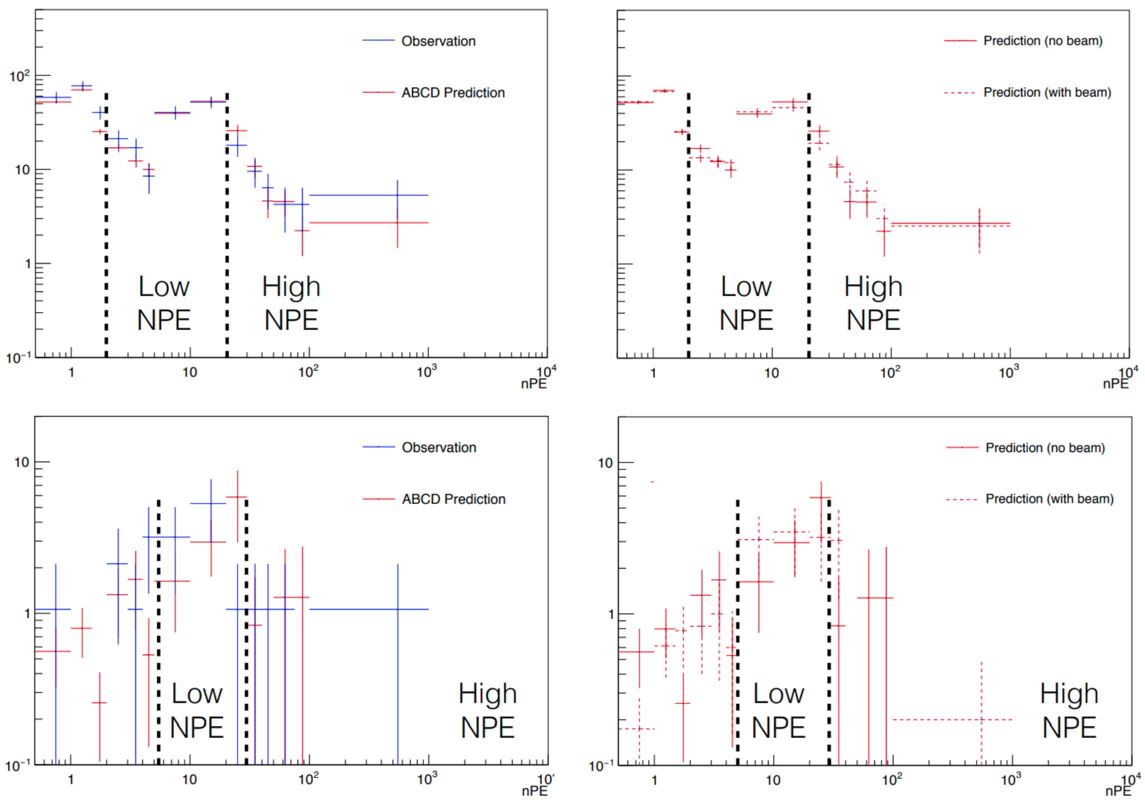

Figure 10.18: ABCD method validations as a function of minimum bar N_{PE}. The left plots show beam-off predictions/observations, and the right plots compare predictions from beam-on and beam-off data. On the top are the 0 slab regions, and the bottom the 1 slab regions. The $\geq$2 slab region does not have enough statistics to allow a full N_{PE} distribution. Vertical dashed lines indicate the boundaries of the signal regions (regions 1 and 2 for the top plots, and 3 and 4 for the bottom plots). Integrated predictions and observations for these regions are listed in Table 10.2.

the CMS IP and milliQan up and down by 7%, re-running the signal simulation and observing the change in signal region yields (7% is chosen based on uncertainty in both the distance and density of the intervening rock). This is negligible for all but the highest charges considered (>0.1e), and even for these it remains a sub-leading source of uncertainty.

We also account for uncertainties in the per-channel $\langle N_{\mathrm{PE}} \rangle$ calibrations, but this is non-trivial to propagate to final signal region yields. The uncertainty in each channel's calibration comes from a number of sources, some correlated and some not:

- Statistical uncertainty from fit: 10–20%, uncorrelated everywhere
- Correction of SPE mean: 1–5%; arises from the fact that the mode of the SPE area distribution (measured in-situ) differs from the mean (measured in the lab) by up to 10%, due to a large left tail in the area distributions. We correct for this, and take half the difference as an uncertainty. This is correlated between bars with the same PMT species.
- B/no-B differences: typically <10%, but up to 20%. Arises from differences in measured SPE values between magnetic field on/off runs. Uncorrelated everywhere.
- Low-pass filter: 5–20%, only on R878s (and fully correlated between them). A low-pass filter is applied to the waveforms for calibration, but not in the main analysis. We take the difference in measured SPE values with/without the filter as a systematic. R7725s and ETs are unaffected as their SPE pulses are much larger and cleaner.
- Global fully-correlated 5% uncertainty, from observed differences in the power law fits between cosmic and SPE points.

To translate these per-channel uncertainties into signal region yield uncertainties, we generate 1000 "calibration toys" in which the PMT efficiencies are varied according to the model above, taking into account all correlations. Signal region yields are recorded for each toy, and the covariance matrix of the yields is diagonalized to generate N_{SR} independent uncertainties, fully accounting for correlations between the various signal region yields. Scatterplots of signal region yields from such toys for an example (mass, charge) point are shown in Fig. 10.19.

Finally, we assess a systematic uncertainty in the time smearing correction applied to signal simulation. As the time resolution discrepancies between simulation and data have an unknown origin, the full size of the correction is used as a systematic uncertainty.

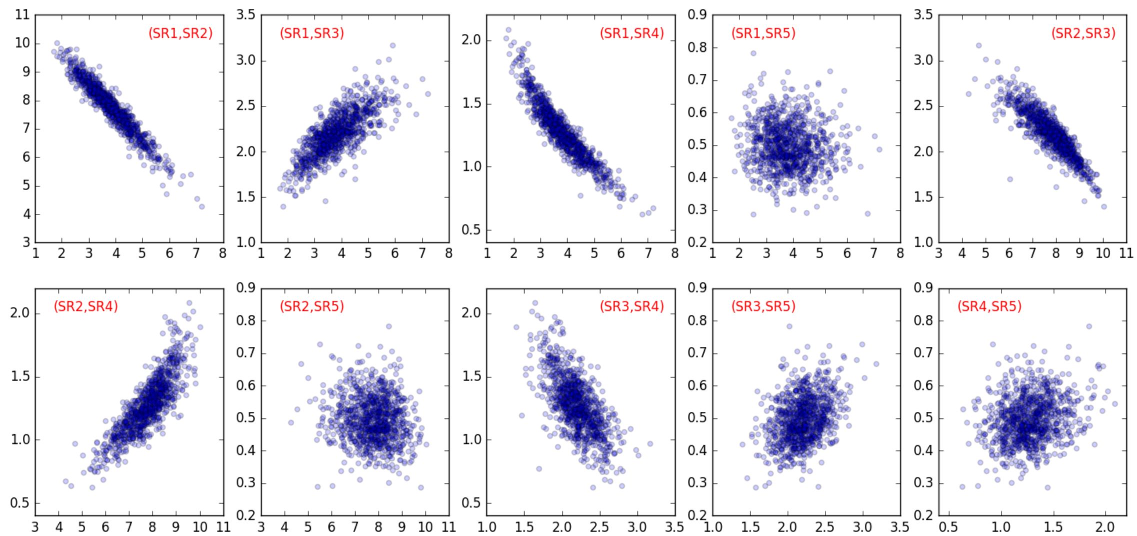

Figure 10.19: Scatterplots of signal region yields in 1000 $\langle N_{\mathrm{PE}} \rangle$ calibration toys, for an example signal point $m = 1.0$ GeV, $q = 0.014e$. Diagonalizing the covariance matrix gives five independent uncertainties, accounting for all correlations between the various regions (e.g. the dominant uncertainty in this example shifts events from regions 1 and 3 to 2 and 4, i.e. between the low bar N_{PE} and high bar N_{PE} regions, arising from correlated calibration uncertainties between the bars).

10.6.4 Results and limits

Predictions are made from the beam-on data set, representing 37.5 fb^{-1} of proton-proton collision data taken in 2018, using the ABCD method described in Sec. 10.6.2. In addition to the statistical uncertainty from the limited size of the control regions, an additional systematic uncertainty is assigned based on small disagreements between prediction and observation in the beam-off validation (Table 10.2). Predictions and their total uncertainties are listed in Table 10.3, along with the observed event counts. The observed counts agree with predictions within the uncertainty in all cases.

The results are interpreted using the signal production model discussed in Sec. 10.5, with the systematic uncertainties on signal yields listed in Sec. 10.6.3. Under the signal plus background hypothesis, a modified frequentist approach is used to determine observed upper limits at 95% confidence level on the cross section to produce a pair of

Table 10.3: Predictions and observations in the five orthogonal signal regions. Uncertainties shown in the predictions include both statistical uncertainty from the limited size of the control regions, and systematic uncertainty from small disagreements in the beam-off validation. The observed event counts are consistent with the predicted counts in all cases. The final three columns give estimated signal yields for three benchmark (m_{mCP}, Q) points near the exclusion boundary.

					Signal yields (m_{mCP} [GeV], Q/e)		
SR	# slab hits	min bar N_{PE}	Pred.	Obs.	(0.05, 0.007)	(1.0, 0.02)	(3.0, 0.1)
1	0	[2, 20]	124 ± 11	129	47.4	0.4	0
2	0	> 20	$49.9^{+6.0}_{-5.4}$	52	0	16.5	0
3	1	[5, 30]	$10.7^{+3.6}_{-2.6}$	8	1.1	0.5	0
4	1	> 30	$2.4^{+2.1}_{-1.1}$	4	0	8.7	0
5	≥ 2	> 0.5	$0.0^{+0.9}_{-0.0}$	1	0	2.0	4.2

mCPs, as a function of mass and charge. The approach uses the the $\mathrm{CL_S}$ technique, described in Sec. 9.2. The observed upper limits are evaluated through the use of asymptotic formulae.

Figure 10.20 shows the exclusion at 95% confidence level in mass and charge of the mCP. The existence of new particles with masses between 20 and 4700 MeV is excluded for charges varying between $0.006e$ and $0.3e$, depending on mass. Compared to existing constraints, the milliQan demonstrator provides new sensitivity for mCP masses above 700 MeV.

The successful operation of the milliQan demonstrator and the search carried out have shown the feasibility of a dedicated detector for millicharged particles at the LHC, and provided important lessons for the design of the full detector planned to operate during LHC Run 3.

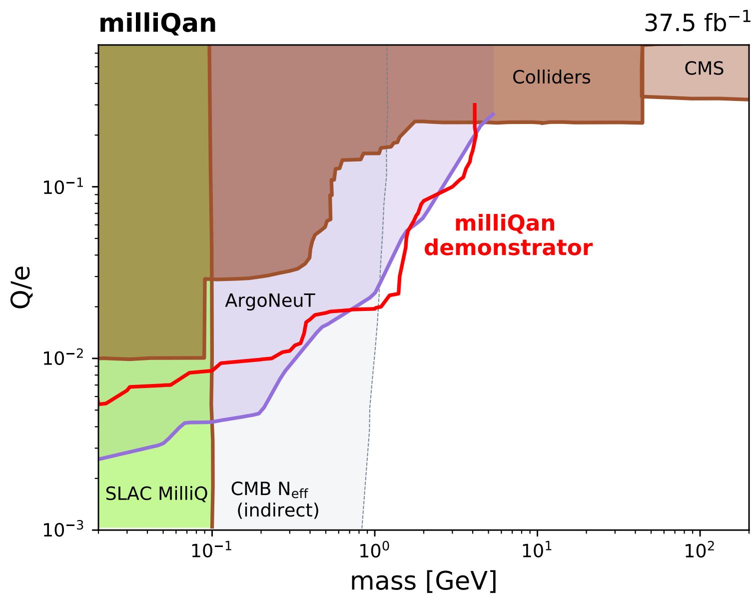

Figure 10.20: Exclusion in the mCP (m, Q) plane at 95% confidence level compared to existing constraints from colliders, CMS [105, 106], ArgoNeuT [107], and SLAC MilliQ [108], as well as the indirect constraint from the CMB relativistic degrees of freedom [109].

Appendix A

Detailed Signal Regions and Results for the M_{T2} Analysis

Table A.1: Predictions and observations for the 12 search regions with $N_j = 1$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$N_j = 1$						
N_j, N_b	p_T^{jet1} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
1j, 0b	250–350	$70700 \pm 400 \pm 4100$	$167000 \pm 1000 \pm 11000$	$530 \pm 20 \pm 160$	$\mathbf{238000} \pm 1000 \pm 14000$	**251941**
	350–450	$13440 \pm 130 \pm 790$	$40100 \pm 500 \pm 3100$	$55 \pm 5 \pm 16$	$\mathbf{53600} \pm 500 \pm 3700$	**54870**
	450–575	$3050 \pm 50 \pm 180$	$10850^{+230}_{-220} \pm 690$	$5.6 \pm 1.1 \pm 1.6$	$\mathbf{13910} \pm 230 \pm 840$	**14473**
	575–700	$603^{+20}_{-19} \pm 38$	$2590^{+110}_{-100} \pm 160$	$0.38 \pm 0.06 \pm 0.11$	$\mathbf{3200} \pm 110 \pm 190$	**3432**
	700–1000	$220 \pm 13 \pm 16$	$1076^{+70}_{-66} \pm 66$	$0.12 \pm 0.03 \pm 0.03$	$\mathbf{1295}^{+71}_{-67} \pm 79$	**1304**
	1000–1200	$11.7^{+4.1}_{-3.2} \pm 0.9$	$86^{+23}_{-19} \pm 6$	< 0.01	$\mathbf{98}^{+24}_{-19} \pm 7$	**98**
	>1200	$2.8^{+2.7}_{-1.5} \pm 0.6$	$23^{+12}_{-8} \pm 2$	< 0.01	$\mathbf{26}^{+13}_{-9} \pm 2$	**30**
1j, ≥1b	250–350	$4210 \pm 110 \pm 260$	$9030 \pm 230 \pm 630$	$58 \pm 10 \pm 17$	$\mathbf{13310}^{+260}_{-250} \pm 820$	**13549**
	350–450	$878 \pm 38 \pm 56$	$2180^{+110}_{-100} \pm 170$	$4.6 \pm 0.4 \pm 1.3$	$\mathbf{3060} \pm 110 \pm 220$	**3078**
	450–575	$211^{+16}_{-15} \pm 13$	$651^{+57}_{-53} \pm 44$	$0.63 \pm 0.18 \pm 0.18$	$\mathbf{863}^{+59}_{-55} \pm 53$	**810**
	575–700	$40.3^{+6.0}_{-5.5} \pm 2.5$	$164^{+30}_{-26} \pm 11$	$0.04 \pm 0.02 \pm 0.02$	$\mathbf{205}^{+31}_{-26} \pm 13$	**184**
	>700	$19.2^{+5.7}_{-4.6} \pm 1.3$	$74^{+21}_{-16} \pm 7$	< 0.01	$\mathbf{94}^{+21}_{-17} \pm 7$	**83**

Table A.2: Predictions and observations for the 30 search regions with $250 < H_{\mathrm{T}} < 450$ GeV. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$250 < H_{\mathrm{T}} < 450$ GeV						
N_{j}, N_{b}	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 0b	200–300	$73700 \pm 500 \pm 5000$	$156000 \pm 1000 \pm 12000$	$580 \pm 20 \pm 140$	$\mathbf{231000} \pm 1000 \pm 16000$	**240867**
	300–400	$12030 \pm 200 \pm 820$	$31300 \pm 200 \pm 2500$	$50 \pm 5 \pm 10$	$\mathbf{43400} \pm 300 \pm 3200$	**44074**
	>400	$417^{+51}_{-47} \pm 28$	$1450 \pm 10 \pm 140$	$0.44 \pm 0.09 \pm 0.09$	$\mathbf{1870} \pm 50 \pm 160$	**2022**
2-3j, 1b	200–300	$12450 \pm 170 \pm 820$	$18700 \pm 300 \pm 1500$	$90 \pm 8 \pm 21$	$\mathbf{31300} \pm 300 \pm 2200$	**32120**
	300–400	$2380 \pm 80 \pm 160$	$3750 \pm 60 \pm 310$	$6.9 \pm 1.0 \pm 1.5$	$\mathbf{6130} \pm 100 \pm 430$	**6258**
	>400	$97 \pm 8 \pm 39$	$174 \pm 3 \pm 17$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{271}^{+9}_{-8} \pm 45$	**275**
2-3j, 2b	200–300	$2240 \pm 70 \pm 150$	$2340^{+110}_{-100} \pm 200$	$9.7 \pm 1.1 \pm 2.3$	$\mathbf{4600}^{+130}_{-120} \pm 320$	**4709**
	300–400	$398^{+34}_{-32} \pm 27$	$469^{+21}_{-20} \pm 39$	$0.68 \pm 0.17 \pm 0.15$	$\mathbf{868}^{+40}_{-38} \pm 61$	**984**
	>400	$13.3 \pm 2.3 \pm 5.4$	$21.7^{+1.0}_{-0.9} \pm 2.2$	< 0.01	$\mathbf{35.0} \pm 2.5 \pm 6.0$	**30**
2-6j, ≥3b	200–300	$507^{+32}_{-31} \pm 38$	$179^{+35}_{-30} \pm 27$	$1.77 \pm 0.46 \pm 0.46$	$\mathbf{688}^{+47}_{-43} \pm 54$	**699**
	300–400	$69 \pm 6 \pm 15$	$40.0^{+7.8}_{-6.6} \pm 6.0$	$0.16 \pm 0.12 \pm 0.04$	$\mathbf{109}^{+10}_{-9} \pm 16$	**102**
	>400	$1.50 \pm 0.80 \pm 0.61$	$1.43^{+0.28}_{-0.24} \pm 0.25$	< 0.01	$\mathbf{2.92}^{+0.85}_{-0.83} \pm 0.67$	**0**
4-6j, 0b	200–300	$12500 \pm 180 \pm 800$	$21600 \pm 300 \pm 1800$	$250 \pm 17 \pm 58$	$\mathbf{34400} \pm 400 \pm 2400$	**35187**
	300–400	$2070 \pm 80 \pm 130$	$4660 \pm 70 \pm 410$	$18.2 \pm 3.6 \pm 3.8$	$\mathbf{6750} \pm 110 \pm 510$	**6725**
	>400	$42 \pm 5 \pm 17$	$155 \pm 2 \pm 64$	$0.06 \pm 0.03 \pm 0.01$	$\mathbf{197} \pm 5 \pm 67$	**170**
4-6j, 1b	200–300	$5750 \pm 100 \pm 380$	$4300 \pm 150 \pm 360$	$61 \pm 7 \pm 15$	$\mathbf{10120} \pm 180 \pm 680$	**10564**
	300–400	$784^{+43}_{-42} \pm 52$	$928^{+32}_{-31} \pm 84$	$2.07 \pm 0.29 \pm 0.45$	$\mathbf{1710} \pm 50 \pm 120$	**1769**
	>400	$14.0 \pm 2.5 \pm 5.7$	$31 \pm 1 \pm 13$	$0.04 \pm 0.02 \pm 0.01$	$\mathbf{45} \pm 3 \pm 14$	**40**
4-6j, 2b	200–300	$2550^{+70}_{-60} \pm 170$	$921^{+68}_{-63} \pm 87$	$10.0 \pm 1.5 \pm 2.2$	$\mathbf{3480} \pm 90 \pm 230$	**3621**
	300–400	$220^{+23}_{-21} \pm 15$	$198^{+15}_{-14} \pm 20$	$0.47 \pm 0.15 \pm 0.11$	$\mathbf{419}^{+27}_{-25} \pm 31$	**496**
	>400	$3.2 \pm 0.8 \pm 1.3$	$6.6 \pm 0.5 \pm 2.7$	< 0.01	$\mathbf{9.8} \pm 0.9 \pm 3.1$	**14**
≥7j, 0b	200–300	$55^{+15}_{-13} \pm 4$	$61^{+23}_{-17} \pm 26$	$2.64 \pm 0.39 \pm 0.57$	$\mathbf{119}^{+28}_{-22} \pm 27$	**108**
	300–500	$3.8^{+2.1}_{-2.0} \pm 0.8$	$8.1^{+3.1}_{-2.3} \pm 4.3$	$0.08 \pm 0.04 \pm 0.02$	$\mathbf{12.0}^{+3.7}_{-3.1} \pm 4.4$	**30**
	>500	$0.0^{+3.2}_{-0.0} \pm 0.0$	$0.0^{+1.2}_{-0.0} \pm 0.0$	< 0.01	$\mathbf{0.0}^{+3.4}_{-0.0} \pm 0.0$	**0**
≥7j, 1b	200–300	$48.0^{+9.1}_{-8.2} \pm 3.5$	$19^{+19}_{-11} \pm 10$	$0.33 \pm 0.14 \pm 0.09$	$\mathbf{68}^{+21}_{-13} \pm 11$	**95**
	>300	$3.0 \pm 1.4 \pm 1.2$	$2.5^{+2.4}_{-1.3} \pm 1.7$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{5.6}^{+2.8}_{-1.9} \pm 2.1$	**12**
≥7j, 2b	200–300	$41.3^{+7.7}_{-7.0} \pm 3.1$	$6.0^{+5.8}_{-3.2} \pm 3.7$	$0.29 \pm 0.14 \pm 0.06$	$\mathbf{47.6}^{+9.7}_{-7.7} \pm 5.0$	**30**
	>300	$2.15^{+0.78}_{-0.76} \pm 0.87$	$0.74^{+0.72}_{-0.40} \pm 0.57$	< 0.01	$\mathbf{2.9}^{+1.1}_{-0.9} \pm 1.1$	**1**
≥7j, ≥3b	200–300	$7.3^{+1.7}_{-1.5} \pm 0.9$	$1.0^{+1.0}_{-0.6} \pm 1.1$	$0.04 \pm 0.04 \pm 0.01$	$\mathbf{8.4}^{+1.9}_{-1.6} \pm 1.5$	**17**
	>300	$0.47 \pm 0.35 \pm 0.20$	$0.12^{+0.11}_{-0.06} \pm 0.14$	< 0.01	$\mathbf{0.59}^{+0.37}_{-0.35} \pm 0.24$	**0**

Table A.3: Predictions and observations for the 28 search regions with $450 < H_{\mathrm{T}} < 575$ GeV, $N_{\mathrm{j}} < 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$450 < H_{\mathrm{T}} < 575$ GeV, $N_{\mathrm{j}} < 7$						
N_{j}, N_{b}	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 0b	200–300	$8860 \pm 110 \pm 640$	$20100 \pm 200 \pm 1300$	$69 \pm 13 \pm 16$	$\mathbf{29100} \pm 300 \pm 1900$	**28956**
	300–400	$4230 \pm 80 \pm 300$	$11770 \pm 140 \pm 790$	$10.6 \pm 0.8 \pm 2.4$	$\mathbf{16000} \pm 200 \pm 1000$	**15876**
	400–500	$1510 \pm 60 \pm 110$	$5020 \pm 60 \pm 360$	$2.86 \pm 0.62 \pm 0.60$	$\mathbf{6540} \pm 80 \pm 440$	**6527**
	>500	$121^{+24}_{-21} \pm 9$	$580 \pm 7 \pm 63$	$0.07 \pm 0.03 \pm 0.02$	$\mathbf{701}^{+25}_{-22} \pm 68$	**740**
2-3j, 1b	200–300	$1326 \pm 43 \pm 88$	$2500 \pm 80 \pm 170$	$17.0 \pm 8.4 \pm 3.8$	$\mathbf{3840}^{+100}_{-90} \pm 240$	**3859**
	300–400	$737 \pm 35 \pm 49$	$1464^{+49}_{-48} \pm 99$	$1.62 \pm 0.20 \pm 0.43$	$\mathbf{2200} \pm 60 \pm 140$	**2065**
	400–500	$259^{+25}_{-23} \pm 19$	$626^{+21}_{-20} \pm 45$	$0.49 \pm 0.10 \pm 0.12$	$\mathbf{885}^{+32}_{-31} \pm 58$	**907**
	>500	$19.1^{+2.8}_{-2.7} \pm 7.8$	$72.4 \pm 2.4 \pm 7.9$	$0.04 \pm 0.02 \pm 0.02$	$\mathbf{92} \pm 4 \pm 11$	**79**
2-3j, 2b	200–300	$201 \pm 15 \pm 13$	$322^{+31}_{-28} \pm 25$	$1.34 \pm 0.62 \pm 0.47$	$\mathbf{524}^{+35}_{-32} \pm 35$	**463**
	300–400	$83.8^{+9.6}_{-9.1} \pm 9.1$	$188^{+18}_{-17} \pm 15$	$0.26 \pm 0.07 \pm 0.07$	$\mathbf{272}^{+21}_{-19} \pm 20$	**304**
	400–500	$31.8^{+4.1}_{-4.0} \pm 6.7$	$80.4^{+7.7}_{-7.1} \pm 6.6$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{112}^{+9}_{-8} \pm 10$	**120**
	>500	$2.16^{+0.67}_{-0.66} \pm 0.88$	$9.3^{+0.9}_{-0.8} \pm 1.1$	< 0.01	$\mathbf{11.4} \pm 1.1 \pm 1.4$	**15**
2-6j, ≥3b	200–300	$232^{+17}_{-16} \pm 15$	$57^{+17}_{-13} \pm 7$	$2.20 \pm 0.70 \pm 0.80$	$\mathbf{291}^{+24}_{-21} \pm 19$	**297**
	300–400	$81^{+12}_{-11} \pm 6$	$33.6^{+9.9}_{-7.8} \pm 4.3$	$0.26 \pm 0.08 \pm 0.08$	$\mathbf{115}^{+16}_{-14} \pm 8$	**76**
	400–500	$10.7^{+2.1}_{-2.0} \pm 2.3$	$11.4^{+3.4}_{-2.7} \pm 1.5$	< 0.01	$\mathbf{22.1}^{+4.0}_{-3.4} \pm 2.8$	**24**
	>500	$1.08 \pm 0.58 \pm 0.44$	$1.03^{+0.30}_{-0.24} \pm 0.17$	< 0.01	$\mathbf{2.11}^{+0.65}_{-0.62} \pm 0.48$	**0**
4-6j, 0b	200–300	$5660 \pm 90 \pm 370$	$8560 \pm 170 \pm 600$	$143 \pm 7 \pm 35$	$\mathbf{14360} \pm 190 \pm 890$	**15047**
	300–400	$2250 \pm 60 \pm 150$	$4790^{+100}_{-90} \pm 350$	$24.3 \pm 2.6 \pm 6.2$	$\mathbf{7060} \pm 110 \pm 460$	**6939**
	400–500	$428^{+32}_{-30} \pm 28$	$1220 \pm 20 \pm 110$	$1.42 \pm 0.21 \pm 0.52$	$\mathbf{1650} \pm 40 \pm 130$	**1817**
	>500	$14.8 \pm 2.2 \pm 6.0$	$86 \pm 2 \pm 35$	$0.04 \pm 0.02 \pm 0.01$	$\mathbf{101} \pm 3 \pm 36$	**104**
4-6j, 1b	200–300	$2810 \pm 60 \pm 190$	$1880 \pm 80 \pm 130$	$63 \pm 15 \pm 19$	$\mathbf{4750} \pm 100 \pm 300$	**4736**
	300–400	$937 \pm 36 \pm 63$	$1054^{+45}_{-43} \pm 78$	$5.4 \pm 0.4 \pm 1.4$	$\mathbf{2000} \pm 60 \pm 130$	**2039**
	400–500	$138^{+17}_{-16} \pm 10$	$269 \pm 11 \pm 25$	$0.36 \pm 0.10 \pm 0.10$	$\mathbf{407}^{+20}_{-19} \pm 31$	**403**
	>500	$7.5 \pm 2.2 \pm 3.0$	$19.1 \pm 0.8 \pm 7.9$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{26.5} \pm 2.3 \pm 8.5$	**27**
4-6j, 2b	200–300	$1343^{+38}_{-37} \pm 89$	$414^{+39}_{-35} \pm 33$	$11.5 \pm 1.0 \pm 3.3$	$\mathbf{1770} \pm 50 \pm 110$	**1767**
	300–400	$418^{+24}_{-23} \pm 29$	$232^{+22}_{-20} \pm 19$	$1.35 \pm 0.35 \pm 0.39$	$\mathbf{651}^{+32}_{-31} \pm 43$	**636**
	400–500	$45.6^{+3.9}_{-3.8} \pm 9.6$	$59.1^{+5.5}_{-5.1} \pm 5.9$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{105}^{+7}_{-6} \pm 12$	**120**
	>500	$1.59 \pm 0.89 \pm 0.65$	$4.2 \pm 0.4 \pm 1.7$	< 0.01	$\mathbf{5.8} \pm 1.0 \pm 1.9$	**7**

Table A.4: Predictions and observations for the 12 search regions with $450 < H_T < 575$ GeV, $N_j \geq 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$450 < H_T < 575$ GeV, $N_j \geq 7$							
N_j, N_b	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
≥7j, 0b	200–300	$149^{+17}_{-16} \pm 13$	$169^{+31}_{-27} \pm 34$	$11.5 \pm 0.8 \pm 3.0$	$\mathbf{329}^{+36}_{-31} \pm 38$	**354**
	300–400	$38.9^{+5.8}_{-5.6} \pm 8.2$	$64^{+12}_{-10} \pm 17$	$1.24 \pm 0.42 \pm 0.32$	$\mathbf{104}^{+13}_{-12} \pm 20$	**110**
	>400	$1.28 \pm 0.82 \pm 0.52$	$8.8^{+1.6}_{-1.4} \pm 3.8$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{10.1}^{+1.8}_{-1.6} \pm 3.8$	**10**
≥7j, 1b	200–300	$191^{+13}_{-12} \pm 15$	$67^{+19}_{-15} \pm 15$	$4.4 \pm 0.5 \pm 1.2$	$\mathbf{262}^{+23}_{-19} \pm 23$	**268**
	300–400	$37.8^{+3.4}_{-3.3} \pm 8.0$	$25.3^{+7.2}_{-5.7} \pm 7.3$	$0.30 \pm 0.07 \pm 0.08$	$\mathbf{63}^{+8}_{-7} \pm 11$	**65**
	>400	$2.31 \pm 0.69 \pm 0.94$	$3.5^{+1.0}_{-0.8} \pm 1.5$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{5.8}^{+1.2}_{-1.0} \pm 1.8$	**3**
≥7j, 2b	200–300	$173^{+12}_{-11} \pm 13$	$19.9^{+5.7}_{-4.5} \pm 5.2$	$1.24 \pm 0.18 \pm 0.33$	$\mathbf{194}^{+13}_{-12} \pm 15$	**197**
	300–400	$26.8 \pm 2.6 \pm 5.7$	$7.6^{+2.2}_{-1.7} \pm 2.4$	$0.09 \pm 0.04 \pm 0.03$	$\mathbf{34.6}^{+3.4}_{-3.1} \pm 6.3$	**44**
	>400	$1.40 \pm 0.44 \pm 0.57$	$1.02^{+0.29}_{-0.23} \pm 0.46$	< 0.01	$\mathbf{2.42}^{+0.53}_{-0.49} \pm 0.73$	**3**
≥7j, ≥3b	200–300	$55.4^{+4.8}_{-4.7} \pm 7.3$	$2.3^{+0.7}_{-0.5} \pm 1.1$	$0.15 \pm 0.06 \pm 0.06$	$\mathbf{57.8}^{+4.8}_{-4.7} \pm 7.4$	**37**
	300–400	$6.4 \pm 1.2 \pm 1.5$	$0.86^{+0.25}_{-0.20} \pm 0.46$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{7.3} \pm 1.2 \pm 1.6$	**9**
	>400	$0.06 \pm 0.01 \pm 0.03$	$0.12 \pm 0.03 \pm 0.06$	< 0.01	$\mathbf{0.18}^{+0.04}_{-0.03} \pm 0.07$	**0**

Table A.5: Predictions and observations for the 20 search regions with $575 < H_\mathrm{T} < 1200$ GeV, $N_\mathrm{j} < 7$, $N_\mathrm{b} = 0$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$575 < H_\mathrm{T} < 1200$ GeV, $N_\mathrm{j} < 7$, $N_\mathrm{b} = 0$							
N_j, N_b	M_T2 [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 0b	200–300	$5270 \pm 60 \pm 370$	$11550 \pm 160 \pm 790$	$93 \pm 20 \pm 30$	$\mathbf{16900} \pm 200 \pm 1100$	**17256**
	300–400	$2560 \pm 50 \pm 180$	$7770^{+110}_{-100} \pm 540$	$11.9 \pm 1.3 \pm 4.4$	$\mathbf{10340}^{+120}_{-110} \pm 680$	**10145**
	400–500	$1101^{+32}_{-31} \pm 77$	$3900 \pm 50 \pm 280$	$1.33 \pm 0.24 \pm 0.41$	$\mathbf{5000} \pm 60 \pm 340$	**5021**
	500–600	$502^{+24}_{-23} \pm 35$	$2250 \pm 30 \pm 170$	$0.37 \pm 0.07 \pm 0.12$	$\mathbf{2760} \pm 40 \pm 200$	**2706**
	600–700	$180^{+16}_{-15} \pm 13$	$746 \pm 10 \pm 73$	$0.09 \pm 0.03 \pm 0.03$	$\mathbf{926}^{+19}_{-18} \pm 80$	**1066**
	700–800	$52.1^{+7.3}_{-6.5} \pm 5.5$	$256 \pm 3 \pm 36$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{308}^{+8}_{-7} \pm 38$	**347**
	800–900	$17.7^{+2.6}_{-2.3} \pm 2.2$	$107 \pm 1 \pm 20$	< 0.01	$\mathbf{125} \pm 3 \pm 21$	**111**
	900–1000	$6.0 \pm 0.9 \pm 1.3$	$39.4 \pm 0.5 \pm 8.5$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{45.4}^{+1.1}_{-1.0} \pm 8.7$	**39**
	1000–1100	$3.3^{+1.1}_{-1.0} \pm 1.0$	$13.3 \pm 0.2 \pm 3.9$	< 0.01	$\mathbf{16.6} \pm 1.1 \pm 4.1$	**11**
	>1100	$0.31^{+0.09}_{-0.08} \pm 0.12$	$2.5 \pm 0.0 \pm 1.1$	< 0.01	$\mathbf{2.8} \pm 0.1 \pm 1.1$	**2**
4-6j, 0b	200–300	$6280 \pm 70 \pm 420$	$9470 \pm 160 \pm 650$	$360 \pm 20 \pm 110$	$\mathbf{16100} \pm 180 \pm 1000$	**16292**
	300–400	$2700 \pm 50 \pm 180$	$5410 \pm 90 \pm 380$	$53 \pm 1 \pm 17$	$\mathbf{8160} \pm 100 \pm 520$	**8330**
	400–500	$927^{+28}_{-27} \pm 62$	$2420 \pm 40 \pm 180$	$7.7 \pm 0.4 \pm 2.4$	$\mathbf{3350} \pm 50 \pm 230$	**3576**
	500–600	$324^{+17}_{-16} \pm 22$	$1171^{+20}_{-19} \pm 100$	$1.46 \pm 0.12 \pm 0.46$	$\mathbf{1500} \pm 30 \pm 110$	**1516**
	600–700	$95.4^{+9.4}_{-8.7} \pm 6.4$	$413 \pm 7 \pm 47$	$0.33 \pm 0.06 \pm 0.10$	$\mathbf{509}^{+12}_{-11} \pm 50$	**543**
	700–800	$35.6^{+5.0}_{-4.5} \pm 3.6$	$171 \pm 3 \pm 27$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{206}^{+6}_{-5} \pm 27$	**178**
	800–900	$13.4^{+2.0}_{-1.8} \pm 1.6$	$64 \pm 1 \pm 11$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{77} \pm 2 \pm 11$	**62**
	900–1000	$4.39^{+0.78}_{-0.73} \pm 0.93$	$23.6 \pm 0.4 \pm 5.3$	< 0.01	$\mathbf{28.0}^{+0.9}_{-0.8} \pm 5.4$	**20**
	1000–1100	$0.64 \pm 0.16 \pm 0.20$	$6.3 \pm 0.1 \pm 2.0$	< 0.01	$\mathbf{6.9} \pm 0.2 \pm 2.0$	**3**
	>1100	$0.78 \pm 0.58 \pm 0.32$	$0.89^{+0.02}_{-0.01} \pm 0.40$	< 0.01	$\mathbf{1.68} \pm 0.58 \pm 0.52$	**1**

Table A.6: Predictions and observations for the 27 search regions with $575 < H_{\mathrm{T}} < 1200$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} \geq 1$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$575 < H_{\mathrm{T}} < 1200$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} \geq 1$						
N_{j}, N_{b}	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 1b	200–300	$826^{+27}_{-26} \pm 54$	$1480^{+60}_{-50} \pm 100$	$38 \pm 15 \pm 12$	$\mathbf{2340} \pm 60 \pm 140$	**2499**
	300–400	$426^{+21}_{-20} \pm 28$	$994^{+38}_{-37} \pm 69$	$2.33 \pm 0.26 \pm 0.84$	$\mathbf{1422}^{+43}_{-42} \pm 90$	**1366**
	400–600	$282^{+18}_{-17} \pm 20$	$788^{+30}_{-29} \pm 55$	$0.27 \pm 0.06 \pm 0.10$	$\mathbf{1071}^{+35}_{-34} \pm 69$	**1057**
	600–800	$43.5^{+3.2}_{-3.1} \pm 6.5$	$129 \pm 5 \pm 12$	< 0.01	$\mathbf{172} \pm 6 \pm 15$	**225**
	800–1000	$4.6 \pm 0.7 \pm 1.3$	$18.8 \pm 0.7 \pm 3.3$	< 0.01	$\mathbf{23.4} \pm 1.0 \pm 3.6$	**22**
	>1000	$0.34 \pm 0.08 \pm 0.14$	$2.05 \pm 0.08 \pm 0.90$	< 0.01	$\mathbf{2.38} \pm 0.11 \pm 0.91$	**1**
2-3j, 2b	200–300	$105.1^{+9.2}_{-8.7} \pm 7.6$	$181^{+20}_{-18} \pm 15$	$3.8 \pm 0.5 \pm 1.3$	$\mathbf{290}^{+22}_{-20} \pm 20$	**316**
	300–400	$55.0^{+6.7}_{-6.3} \pm 7.5$	$122^{+14}_{-12} \pm 10$	$0.27 \pm 0.06 \pm 0.10$	$\mathbf{177}^{+15}_{-14} \pm 14$	**159**
	400–600	$36.5^{+4.6}_{-4.3} \pm 5.5$	$97^{+11}_{-10} \pm 8$	$0.08 \pm 0.03 \pm 0.03$	$\mathbf{133}^{+12}_{-11} \pm 11$	**107**
	600–800	$4.7 \pm 0.8 \pm 1.3$	$15.8^{+1.8}_{-1.6} \pm 1.6$	< 0.01	$\mathbf{20.6}^{+1.9}_{-1.8} \pm 2.2$	**21**
	>800	$0.59 \pm 0.19 \pm 0.24$	$2.56^{+0.29}_{-0.26} \pm 0.45$	< 0.01	$\mathbf{3.14}^{+0.35}_{-0.32} \pm 0.52$	**1**
4-6j, 1b	200–300	$2900 \pm 50 \pm 200$	$2220^{+80}_{-70} \pm 150$	$154 \pm 16 \pm 50$	$\mathbf{5270} \pm 90 \pm 330$	**5335**
	300–400	$1066 \pm 29 \pm 74$	$1267^{+44}_{-42} \pm 89$	$19.2 \pm 0.9 \pm 6.2$	$\mathbf{2350} \pm 50 \pm 150$	**2547**
	400–600	$504^{+22}_{-21} \pm 35$	$840^{+29}_{-28} \pm 61$	$2.98 \pm 0.21 \pm 0.93$	$\mathbf{1347}^{+36}_{-35} \pm 88$	**1284**
	600–800	$35.3^{+5.9}_{-5.2} \pm 2.6$	$138 \pm 5 \pm 14$	$0.09 \pm 0.03 \pm 0.03$	$\mathbf{174}^{+8}_{-7} \pm 16$	**151**
	800–1000	$3.89^{+0.83}_{-0.77} \pm 0.82$	$19.3^{+0.7}_{-0.6} \pm 4.3$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{23.2}^{+1.1}_{-1.0} \pm 4.5$	**18**
	>1000	$0.18 \pm 0.07 \pm 0.07$	$1.57 \pm 0.05 \pm 0.65$	< 0.01	$\mathbf{1.75} \pm 0.09 \pm 0.65$	**1**
4-6j, 2b	200–300	$1500 \pm 30 \pm 100$	$473^{+36}_{-33} \pm 36$	$42 \pm 2 \pm 13$	$\mathbf{2020} \pm 50 \pm 130$	**1968**
	300–400	$508 \pm 20 \pm 35$	$270^{+20}_{-19} \pm 21$	$4.9 \pm 0.3 \pm 1.6$	$\mathbf{783}^{+29}_{-28} \pm 50$	**788**
	400–600	$167 \pm 12 \pm 12$	$179^{+14}_{-13} \pm 14$	$0.57 \pm 0.08 \pm 0.18$	$\mathbf{346}^{+18}_{-17} \pm 23$	**354**
	600–800	$11.9^{+1.3}_{-1.2} \pm 2.5$	$29.5^{+2.2}_{-2.1} \pm 3.5$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{41.4}^{+2.6}_{-2.4} \pm 4.6$	**37**
	>800	$0.91 \pm 0.23 \pm 0.37$	$4.4 \pm 0.3 \pm 1.8$	< 0.01	$\mathbf{5.4} \pm 0.4 \pm 1.9$	**7**
2-6j, ≥3b	200–300	$299^{+17}_{-16} \pm 22$	$73^{+15}_{-13} \pm 10$	$6.2 \pm 0.4 \pm 2.1$	$\mathbf{379}^{+22}_{-21} \pm 28$	**345**
	300–400	$100 \pm 10 \pm 7$	$43.5^{+8.8}_{-7.4} \pm 6.2$	$0.68 \pm 0.09 \pm 0.24$	$\mathbf{144}^{+14}_{-12} \pm 11$	**132**
	400–600	$32.5^{+6.3}_{-5.6} \pm 2.5$	$31.2^{+6.3}_{-5.3} \pm 4.4$	$0.08 \pm 0.03 \pm 0.03$	$\mathbf{63.8}^{+8.9}_{-7.7} \pm 5.8$	**48**
	600–800	$3.16^{+0.95}_{-0.90} \pm 0.68$	$5.4^{+1.1}_{-0.9} \pm 0.8$	< 0.01	$\mathbf{8.6}^{+1.4}_{-1.3} \pm 1.1$	**4**
	>800	$0.10 \pm 0.03 \pm 0.04$	$0.71^{+0.14}_{-0.12} \pm 0.15$	< 0.01	$\mathbf{0.81}^{+0.15}_{-0.12} \pm 0.16$	**0**

Table A.7: Predictions and observations for the 34 search regions with $575 < H_\mathrm{T} < 1200$ GeV, $N_\mathrm{j} \geq 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$575 < H_\mathrm{T} < 1200$ GeV, $N_\mathrm{j} \geq 7$						
N_j, N_b	M_T2 [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
7-9j, 0b	200–300	$589^{+27}_{-26} \pm 39$	$573^{+47}_{-43} \pm 64$	$90 \pm 10 \pm 28$	$\mathbf{1252}^{+55}_{-52} \pm 93$	**1340**
	300–400	$265^{+19}_{-18} \pm 18$	$279^{+23}_{-21} \pm 42$	$14.9 \pm 0.5 \pm 4.7$	$\mathbf{559}^{+29}_{-28} \pm 51$	**581**
	400–600	$92^{+10}_{-9} \pm 6$	$159^{+13}_{-12} \pm 28$	$2.72 \pm 0.18 \pm 0.85$	$\mathbf{253}^{+16}_{-15} \pm 30$	**243**
	600–800	$8.6 \pm 1.2 \pm 1.8$	$22.8^{+1.9}_{-1.7} \pm 6.4$	$0.10 \pm 0.03 \pm 0.03$	$\mathbf{31.6}^{+2.2}_{-2.1} \pm 6.8$	**32**
	>800	$0.51 \pm 0.16 \pm 0.21$	$3.0 \pm 0.2 \pm 1.3$	< 0.01	$\mathbf{3.5} \pm 0.3 \pm 1.3$	**2**
7-9j, 1b	200–300	$733 \pm 21 \pm 52$	$278^{+28}_{-25} \pm 33$	$48 \pm 3 \pm 16$	$\mathbf{1059}^{+35}_{-33} \pm 73$	**1052**
	300–400	$252^{+13}_{-12} \pm 18$	$135^{+14}_{-12} \pm 21$	$7.7 \pm 0.4 \pm 2.5$	$\mathbf{395}^{+19}_{-17} \pm 32$	**387**
	400–600	$71.3^{+6.9}_{-6.5} \pm 5.2$	$77^{+8}_{-7} \pm 14$	$1.36 \pm 0.13 \pm 0.45$	$\mathbf{150} \pm 10 \pm 16$	**131**
	600–800	$4.26^{+0.73}_{-0.71} \pm 0.90$	$11.0^{+1.1}_{-1.0} \pm 3.1$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{15.3}^{+1.3}_{-1.2} \pm 3.3$	**20**
	>800	$0.11 \pm 0.04 \pm 0.05$	$1.48^{+0.15}_{-0.13} \pm 0.63$	< 0.01	$\mathbf{1.60}^{+0.15}_{-0.14} \pm 0.63$	**1**
7-9j, 2b	200–300	$675 \pm 20 \pm 51$	$82^{+8}_{-7} \pm 10$	$20.9 \pm 3.0 \pm 6.7$	$\mathbf{777}^{+22}_{-21} \pm 56$	**750**
	300–400	$211 \pm 11 \pm 16$	$39.8^{+4.0}_{-3.6} \pm 6.4$	$2.42 \pm 0.19 \pm 0.79$	$\mathbf{253}^{+12}_{-11} \pm 19$	**259**
	400–600	$55.4^{+5.5}_{-5.2} \pm 4.2$	$22.7^{+2.3}_{-2.1} \pm 4.2$	$0.50 \pm 0.07 \pm 0.16$	$\mathbf{78.6}^{+5.9}_{-5.6} \pm 6.6$	**72**
	600–800	$3.00^{+0.63}_{-0.62} \pm 0.64$	$3.25^{+0.32}_{-0.30} \pm 0.93$	$0.01 \pm 0.01 \pm 0.01$	$\mathbf{6.3} \pm 0.7 \pm 1.2$	**7**
	>800	$0.27 \pm 0.20 \pm 0.11$	$0.44 \pm 0.04 \pm 0.19$	< 0.01	$\mathbf{0.71} \pm 0.20 \pm 0.22$	**1**
7-9j, 3b	200–300	$185 \pm 8 \pm 18$	$11.3^{+1.1}_{-1.0} \pm 1.9$	$3.6 \pm 0.2 \pm 1.2$	$\mathbf{200} \pm 8 \pm 18$	**184**
	300–400	$52.0 \pm 3.8 \pm 5.0$	$5.5 \pm 0.5 \pm 1.2$	$0.72 \pm 0.12 \pm 0.26$	$\mathbf{58.3}^{+3.9}_{-3.8} \pm 5.3$	**59**
	400–600	$13.6 \pm 1.8 \pm 1.3$	$3.13^{+0.31}_{-0.29} \pm 0.82$	$0.05 \pm 0.02 \pm 0.02$	$\mathbf{16.8} \pm 1.8 \pm 1.6$	**14**
	>600	$0.49 \pm 0.21 \pm 0.20$	$0.51 \pm 0.05 \pm 0.21$	< 0.01	$\mathbf{1.00} \pm 0.21 \pm 0.29$	**2**
7-9j, ≥4b	200–300	$38.8 \pm 3.1 \pm 7.4$	$2.01^{+0.20}_{-0.18} \pm 0.71$	$0.55 \pm 0.08 \pm 0.19$	$\mathbf{41.3}^{+3.2}_{-3.1} \pm 7.4$	**38**
	300–400	$14.5^{+2.0}_{-1.9} \pm 2.8$	$0.98^{+0.10}_{-0.09} \pm 0.43$	$0.06 \pm 0.02 \pm 0.02$	$\mathbf{15.6}^{+2.0}_{-1.9} \pm 2.8$	**16**
	>400	$3.75^{+0.98}_{-0.97} \pm 0.70$	$0.65 \pm 0.06 \pm 0.35$	< 0.01	$\mathbf{4.40}^{+0.98}_{-0.97} \pm 0.79$	**3**
≥10j, 0b	200–300	$11.5 \pm 1.6 \pm 1.0$	$4.4^{+0.4}_{-0.3} \pm 2.3$	$3.1 \pm 0.8 \pm 1.1$	$\mathbf{19.0} \pm 1.8 \pm 2.8$	**27**
	300-500	$5.6 \pm 1.0 \pm 0.5$	$3.0 \pm 0.2 \pm 1.7$	$0.55 \pm 0.08 \pm 0.20$	$\mathbf{9.1} \pm 1.0 \pm 1.8$	**4**
	>500	$0.30 \pm 0.11 \pm 0.12$	$0.44^{+0.04}_{-0.03} \pm 0.24$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{0.76} \pm 0.11 \pm 0.27$	**3**
≥10j, 1b	200–300	$21.0 \pm 1.8 \pm 1.6$	$3.5 \pm 0.3 \pm 1.9$	$1.92 \pm 0.18 \pm 0.72$	$\mathbf{26.4} \pm 1.8 \pm 2.7$	**32**
	300-500	$7.7 \pm 1.0 \pm 0.6$	$2.4 \pm 0.2 \pm 1.4$	$0.45 \pm 0.07 \pm 0.17$	$\mathbf{10.5} \pm 1.1 \pm 1.6$	**15**
	>500	$0.83^{+0.42}_{-0.41} \pm 0.07$	$0.36^{+0.04}_{-0.03} \pm 0.20$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{1.20}^{+0.42}_{-0.41} \pm 0.22$	**0**
≥10j, 2b	200–300	$21.8 \pm 1.8 \pm 1.6$	$1.05 \pm 0.10 \pm 0.66$	$0.64 \pm 0.08 \pm 0.24$	$\mathbf{23.5} \pm 1.8 \pm 1.8$	**26**
	300-500	$8.8 \pm 1.2 \pm 0.6$	$0.69^{+0.07}_{-0.06} \pm 0.45$	$0.16 \pm 0.04 \pm 0.06$	$\mathbf{9.6}^{+1.3}_{-1.2} \pm 0.8$	**9**
	>500	$0.22 \pm 0.13 \pm 0.02$	$0.10 \pm 0.01 \pm 0.06$	< 0.01	$\mathbf{0.32} \pm 0.13 \pm 0.07$	**0**
≥10j, 3b	200–300	$9.9 \pm 1.3 \pm 1.2$	$0.25 \pm 0.02 \pm 0.20$	$0.29 \pm 0.05 \pm 0.12$	$\mathbf{10.4} \pm 1.3 \pm 1.2$	**14**
	>300	$1.59 \pm 0.50 \pm 0.18$	$0.19 \pm 0.02 \pm 0.16$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{1.80} \pm 0.50 \pm 0.25$	**2**
≥10j, ≥4b	>200	$3.9 \pm 1.2 \pm 0.8$	$0.00^{+0.17}_{-0.00} \pm 0.00$	$0.05 \pm 0.02 \pm 0.02$	$\mathbf{4.0} \pm 1.2 \pm 0.8$	**6**

Table A.8: Predictions and observations for the 12 search regions with $1200 < H_{\mathrm{T}} < 1500$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} = 0$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$1200 < H_{\mathrm{T}} < 1500$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} = 0$						
N_{j}, N_{b}	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 0b	200–400	$315 \pm 15 \pm 21$	$656^{+51}_{-47} \pm 73$	$39 \pm 16 \pm 12$	$\mathbf{1009}^{+55}_{-52} \pm 85$	**1128**
	400–600	$43.0^{+5.2}_{-4.7} \pm 4.9$	$185^{+14}_{-13} \pm 30$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{228}^{+15}_{-14} \pm 31$	**207**
	600–800	$14.1^{+2.1}_{-2.0} \pm 1.7$	$64 \pm 5 \pm 17$	< 0.01	$\mathbf{78} \pm 5 \pm 17$	**83**
	800–1000	$6.4^{+1.1}_{-1.0} \pm 1.3$	$32.5^{+2.5}_{-2.3} \pm 7.6$	< 0.01	$\mathbf{38.9}^{+2.7}_{-2.5} \pm 7.8$	**36**
	1000–1200	$3.23^{+0.61}_{-0.59} \pm 0.99$	$17.5 \pm 1.3 \pm 5.2$	< 0.01	$\mathbf{20.7}^{+1.5}_{-1.4} \pm 5.3$	**19**
	>1200	$0.87^{+0.14}_{-0.13} \pm 0.35$	$6.0^{+0.5}_{-0.4} \pm 2.6$	< 0.01	$\mathbf{6.9} \pm 0.5 \pm 2.6$	**4**
4-6j, 0b	200–400	$606^{+21}_{-20} \pm 41$	$909^{+63}_{-59} \pm 90$	$208 \pm 12 \pm 64$	$\mathbf{1720}^{+70}_{-60} \pm 130$	**1768**
	400–600	$84.3^{+7.4}_{-6.9} \pm 5.8$	$234^{+16}_{-15} \pm 34$	$0.88 \pm 0.09 \pm 0.27$	$\mathbf{319}^{+18}_{-17} \pm 36$	**301**
	600–800	$21.1^{+3.2}_{-2.9} \pm 2.3$	$75 \pm 5 \pm 17$	$0.06 \pm 0.02 \pm 0.02$	$\mathbf{96} \pm 6 \pm 17$	**99**
	800–1000	$7.6^{+1.2}_{-1.1} \pm 1.1$	$35.2^{+2.4}_{-2.3} \pm 8.0$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{42.7}^{+2.7}_{-2.5} \pm 8.2$	**41**
	1000–1200	$2.23^{+0.36}_{-0.33} \pm 0.61$	$14.1^{+1.0}_{-0.9} \pm 4.2$	< 0.01	$\mathbf{16.3} \pm 1.0 \pm 4.2$	**15**
	>1200	$0.47^{+0.10}_{-0.09} \pm 0.19$	$3.0 \pm 0.2 \pm 1.3$	< 0.01	$\mathbf{3.5} \pm 0.2 \pm 1.3$	**5**

Table A.9: Predictions and observations for the 25 search regions with $1200 < H_{\mathrm{T}} < 1500$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} \geq 1$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$1200 < H_{\mathrm{T}} < 1500$ GeV, $N_{\mathrm{j}} < 7$, $N_{\mathrm{b}} \geq 1$						
N_{j}, N_{b}	M_{T2} [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 1b	200–400	$61.5^{+7.2}_{-6.5} \pm 4.2$	$78^{+19}_{-16} \pm 10$	$9.7 \pm 0.7 \pm 3.0$	$\mathbf{149}^{+21}_{-17} \pm 12$	**157**
	400–600	$10.1 \pm 1.4 \pm 1.0$	$21.9^{+5.4}_{-4.4} \pm 3.8$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{32.0}^{+5.6}_{-4.6} \pm 4.1$	**27**
	600–800	$2.36^{+0.36}_{-0.35} \pm 0.41$	$7.5^{+1.9}_{-1.5} \pm 2.0$	< 0.01	$\mathbf{9.8}^{+1.9}_{-1.6} \pm 2.1$	**9**
	800–1000	$0.78^{+0.16}_{-0.15} \pm 0.19$	$3.84^{+0.95}_{-0.78} \pm 0.93$	< 0.01	$\mathbf{4.62}^{+0.97}_{-0.79} \pm 0.96$	**6**
	1000–1200	$0.43 \pm 0.08 \pm 0.14$	$2.13^{+0.53}_{-0.43} \pm 0.64$	< 0.01	$\mathbf{2.56}^{+0.54}_{-0.44} \pm 0.66$	**2**
	>1200	$0.14^{+0.05}_{-0.04} \pm 0.06$	$0.71^{+0.18}_{-0.14} \pm 0.31$	< 0.01	$\mathbf{0.86}^{+0.18}_{-0.15} \pm 0.31$	**0**
2-3j, 2b	200–400	$4.8^{+2.0}_{-1.6} \pm 0.3$	$11^{+11}_{-6} \pm 2$	$1.38 \pm 0.13 \pm 0.43$	$\mathbf{18}^{+11}_{-6} \pm 2$	**18**
	400–600	$0.61^{+0.30}_{-0.25} \pm 0.07$	$3.2^{+3.1}_{-1.7} \pm 0.7$	< 0.01	$\mathbf{3.8}^{+3.1}_{-1.8} \pm 0.7$	**5**
	600–800	$0.21^{+0.11}_{-0.09} \pm 0.04$	$1.1^{+1.1}_{-0.6} \pm 0.4$	< 0.01	$\mathbf{1.3}^{+1.1}_{-0.6} \pm 0.4$	**2**
	800–1000	$0.07^{+0.04}_{-0.03} \pm 0.02$	$0.56^{+0.55}_{-0.31} \pm 0.18$	< 0.01	$\mathbf{0.63}^{+0.55}_{-0.31} \pm 0.18$	**1**
	>1000	$0.03 \pm 0.02 \pm 0.01$	$0.42^{+0.41}_{-0.23} \pm 0.18$	< 0.01	$\mathbf{0.46}^{+0.41}_{-0.23} \pm 0.18$	**1**
2-6j, ≥3b	200–400	$22.6^{+4.7}_{-4.2} \pm 1.8$	$0.0^{+6.6}_{-0.0} \pm 0.0$	$4.4 \pm 0.2 \pm 1.5$	$\mathbf{27.0}^{+8.1}_{-4.2} \pm 2.4$	**25**
	400–600	$1.58^{+0.51}_{-0.48} \pm 0.34$	$0.0^{+1.6}_{-0.0} \pm 0.0$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{1.6}^{+1.7}_{-0.5} \pm 0.3$	**3**
	>600	$0.47^{+0.27}_{-0.26} \pm 0.19$	$0.00^{+0.94}_{-0.00} \pm 0.00$	< 0.01	$\mathbf{0.47}^{+0.98}_{-0.26} \pm 0.19$	**4**
4-6j, 1b	200–400	$278^{+15}_{-14} \pm 20$	$254^{+33}_{-30} \pm 28$	$97 \pm 2 \pm 30$	$\mathbf{629}^{+36}_{-33} \pm 50$	**579**
	400–600	$30.3^{+4.0}_{-3.7} \pm 2.7$	$65^{+9}_{-8} \pm 10$	$0.33 \pm 0.06 \pm 0.10$	$\mathbf{96}^{+9}_{-8} \pm 11$	**79**
	600–800	$8.2^{+1.4}_{-1.3} \pm 1.0$	$21.0^{+2.8}_{-2.5} \pm 4.8$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{29.2}^{+3.1}_{-2.8} \pm 5.0$	**16**
	800–1000	$2.36^{+0.56}_{-0.54} \pm 0.50$	$9.8^{+1.3}_{-1.1} \pm 2.3$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{12.2}^{+1.4}_{-1.3} \pm 2.4$	**9**
	1000–1200	$1.00 \pm 0.24 \pm 0.31$	$4.0 \pm 0.5 \pm 1.2$	< 0.01	$\mathbf{5.0}^{+0.6}_{-0.5} \pm 1.2$	**6**
	>1200	$0.07 \pm 0.02 \pm 0.03$	$0.86^{+0.11}_{-0.10} \pm 0.37$	< 0.01	$\mathbf{0.92}^{+0.11}_{-0.10} \pm 0.37$	**1**
4-6j, 2b	200–400	$120.4^{+9.1}_{-8.7} \pm 9.8$	$45^{+18}_{-13} \pm 5$	$26.0 \pm 0.6 \pm 8.1$	$\mathbf{191}^{+20}_{-16} \pm 15$	**194**
	400–600	$11.9 \pm 1.4 \pm 1.5$	$11.5^{+4.6}_{-3.4} \pm 1.8$	$0.11 \pm 0.03 \pm 0.04$	$\mathbf{23.4}^{+4.8}_{-3.7} \pm 2.6$	**27**
	600–800	$3.49 \pm 0.83 \pm 0.75$	$3.7^{+1.5}_{-1.1} \pm 1.0$	< 0.01	$\mathbf{7.2}^{+1.7}_{-1.4} \pm 1.3$	**7**
	800–1000	$0.66 \pm 0.16 \pm 0.20$	$1.73^{+0.69}_{-0.51} \pm 0.48$	< 0.01	$\mathbf{2.38}^{+0.71}_{-0.54} \pm 0.53$	**3**
	>1000	$0.15 \pm 0.04 \pm 0.06$	$0.84^{+0.34}_{-0.25} \pm 0.36$	< 0.01	$\mathbf{1.00}^{+0.34}_{-0.25} \pm 0.36$	**0**

Table A.10: Predictions and observations for the 31 search regions with $1200 < H_\mathrm{T} < 1500$ GeV, $N_\mathrm{j} \geq 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$1200 < H_\mathrm{T} < 1500$ GeV, $N_\mathrm{j} \geq 7$						
N_j, N_b	M_T2 [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
7-9j, 0b	200–400	$120.4^{+9.8}_{-9.2} \pm 9.0$	$108^{+26}_{-21} \pm 21$	$91 \pm 3 \pm 29$	$\mathbf{319}^{+28}_{-24} \pm 38$	**379**
	400–600	$16.5^{+1.9}_{-1.8} \pm 2.0$	$25.8^{+6.3}_{-5.1} \pm 5.7$	$0.80 \pm 0.09 \pm 0.25$	$\mathbf{43.1}^{+6.5}_{-5.4} \pm 6.3$	**45**
	600–800	$2.94 \pm 0.42 \pm 0.63$	$8.6^{+2.1}_{-1.7} \pm 2.1$	$0.06 \pm 0.02 \pm 0.02$	$\mathbf{11.6}^{+2.1}_{-1.8} \pm 2.2$	**17**
	800–1000	$0.77^{+0.14}_{-0.13} \pm 0.24$	$2.90^{+0.70}_{-0.58} \pm 1.00$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{3.7}^{+0.7}_{-0.6} \pm 1.0$	**3**
	>1000	$0.11 \pm 0.03 \pm 0.05$	$1.09^{+0.26}_{-0.22} \pm 0.50$	< 0.01	$\mathbf{1.21}^{+0.27}_{-0.22} \pm 0.50$	**0**
7-9j, 1b	200–400	$133.8^{+8.0}_{-7.7} \pm 9.8$	$36^{+13}_{-10} \pm 8$	$58 \pm 2 \pm 18$	$\mathbf{228}^{+15}_{-13} \pm 23$	**247**
	400–600	$16.6^{+2.9}_{-2.7} \pm 1.3$	$8.7^{+3.2}_{-2.4} \pm 2.1$	$0.46 \pm 0.07 \pm 0.14$	$\mathbf{25.8}^{+4.3}_{-3.6} \pm 2.7$	**23**
	600–800	$1.83^{+0.43}_{-0.41} \pm 0.28$	$2.9^{+1.1}_{-0.8} \pm 0.8$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{4.8}^{+1.1}_{-0.9} \pm 0.8$	**7**
	800–1000	$0.65^{+0.24}_{-0.23} \pm 0.18$	$0.95^{+0.34}_{-0.26} \pm 0.34$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{1.62}^{+0.42}_{-0.35} \pm 0.39$	**2**
	>1000	$0.22 \pm 0.19 \pm 0.09$	$0.36^{+0.13}_{-0.10} \pm 0.17$	< 0.01	$\mathbf{0.58}^{+0.23}_{-0.21} \pm 0.19$	**0**
7-9j, 2b	200–400	$124.0^{+7.6}_{-7.4} \pm 9.1$	$9.9^{+3.6}_{-2.7} \pm 2.5$	$21.4 \pm 0.5 \pm 6.9$	$\mathbf{155} \pm 8 \pm 12$	**162**
	400–600	$15.0^{+2.8}_{-2.6} \pm 1.3$	$2.41^{+0.87}_{-0.66} \pm 0.67$	$0.12 \pm 0.03 \pm 0.04$	$\mathbf{17.5}^{+3.0}_{-2.7} \pm 1.5$	**18**
	600–800	$2.47^{+0.78}_{-0.76} \pm 0.53$	$0.81^{+0.29}_{-0.22} \pm 0.26$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{3.29}^{+0.83}_{-0.79} \pm 0.60$	**1**
	>800	$0.24 \pm 0.11 \pm 0.10$	$0.36^{+0.13}_{-0.10} \pm 0.16$	< 0.01	$\mathbf{0.60}^{+0.17}_{-0.15} \pm 0.19$	**1**
7-9j, 3b	200–400	$30.0 \pm 2.6 \pm 3.2$	$1.89^{+0.68}_{-0.52} \pm 0.64$	$5.0 \pm 0.3 \pm 1.8$	$\mathbf{36.9}^{+2.7}_{-2.6} \pm 3.8$	**46**
	400–600	$4.1^{+1.1}_{-1.0} \pm 0.6$	$0.45^{+0.16}_{-0.12} \pm 0.18$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{4.6}^{+1.1}_{-1.0} \pm 0.6$	**2**
	>600	$0.92^{+0.50}_{-0.49} \pm 0.38$	$0.23^{+0.08}_{-0.06} \pm 0.11$	< 0.01	$\mathbf{1.15} \pm 0.50 \pm 0.40$	**1**
7-9j, ≥4b	200–400	$9.1 \pm 1.6 \pm 1.8$	$0.26^{+0.10}_{-0.07} \pm 0.23$	$0.88 \pm 0.10 \pm 0.32$	$\mathbf{10.3} \pm 1.6 \pm 1.9$	**9**
	>400	$0.44^{+0.24}_{-0.23} \pm 0.08$	$0.10^{+0.04}_{-0.03} \pm 0.09$	< 0.01	$\mathbf{0.53} \pm 0.24 \pm 0.12$	**0**
≥10j, 0b	200–400	$7.7^{+1.2}_{-1.1} \pm 0.8$	$2.7^{+0.6}_{-0.5} \pm 2.8$	$8.3 \pm 0.9 \pm 3.0$	$\mathbf{18.7}^{+1.6}_{-1.5} \pm 4.1$	**17**
	400–600	$1.00 \pm 0.32 \pm 0.22$	$0.56^{+0.13}_{-0.11} \pm 0.62$	$0.11 \pm 0.03 \pm 0.04$	$\mathbf{1.66}^{+0.35}_{-0.34} \pm 0.66$	**1**
	>600	$0.10^{+0.35}_{-0.04} \pm 0.04$	$0.14^{+0.08}_{-0.03} \pm 0.14$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{0.24}^{+0.36}_{-0.05} \pm 0.15$	**0**
≥10j, 1b	200–400	$15.2 \pm 1.8 \pm 1.4$	$1.1^{+0.4}_{-0.3} \pm 1.2$	$5.3 \pm 0.2 \pm 1.9$	$\mathbf{21.6}^{+1.9}_{-1.8} \pm 2.7$	**22**
	400–600	$1.27^{+0.38}_{-0.36} \pm 0.11$	$0.22^{+0.08}_{-0.06} \pm 0.26$	$0.05 \pm 0.02 \pm 0.02$	$\mathbf{1.55}^{+0.39}_{-0.37} \pm 0.29$	**6**
	>600	$0.03 \pm 0.02 \pm 0.01$	$0.05^{+0.10}_{-0.01} \pm 0.05$	< 0.01	$\mathbf{0.07}^{+0.11}_{-0.02} \pm 0.05$	**0**
≥10j, 2b	200–400	$16.9 \pm 1.8 \pm 1.5$	$0.44^{+0.16}_{-0.12} \pm 0.50$	$2.7 \pm 0.2 \pm 1.0$	$\mathbf{20.1} \pm 1.8 \pm 1.9$	**16**
	400–600	$2.62^{+0.71}_{-0.68} \pm 0.30$	$0.09 \pm 0.03 \pm 0.11$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{2.73}^{+0.71}_{-0.68} \pm 0.32$	**2**
	>600	$0.23 \pm 0.15 \pm 0.10$	$0.02^{+0.08}_{-0.01} \pm 0.02$	< 0.01	$\mathbf{0.25}^{+0.17}_{-0.15} \pm 0.10$	**0**
≥10j, 3b	200–400	$5.58^{+0.86}_{-0.85} \pm 0.61$	$0.12^{+0.11}_{-0.03} \pm 0.16$	$1.04 \pm 0.10 \pm 0.42$	$\mathbf{6.74}^{+0.87}_{-0.86} \pm 0.76$	**6**
	>400	$0.51 \pm 0.22 \pm 0.06$	$0.03^{+0.11}_{-0.01} \pm 0.04$	< 0.01	$\mathbf{0.54}^{+0.25}_{-0.22} \pm 0.08$	**0**
≥10j, ≥4b	>200	$2.59 \pm 0.82 \pm 0.62$	$0.10^{+0.13}_{-0.03} \pm 0.13$	$0.31 \pm 0.06 \pm 0.13$	$\mathbf{3.00}^{+0.83}_{-0.82} \pm 0.65$	**7**

Table A.11: Predictions and observations for the 30 search regions with $H_\mathrm{T} > 1500$ GeV, $N_\mathrm{j} < 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$H_\mathrm{T} > 1500$ GeV, $N_\mathrm{j} < 7$						
N_j, N_b	M_T2 [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
2-3j, 0b	400–600	$27.2^{+4.4}_{-3.9} \pm 2.5$	$150^{+14}_{-13} \pm 19$	$0.16 \pm 0.04 \pm 0.05$	$\mathbf{177}^{+15}_{-13} \pm 20$	**125**
	600–800	$7.8^{+1.4}_{-1.2} \pm 0.8$	$38.7^{+3.6}_{-3.3} \pm 8.4$	< 0.01	$\mathbf{46.5}^{+3.9}_{-3.6} \pm 8.6$	**37**
	800–1000	$2.29^{+0.39}_{-0.34} \pm 0.35$	$17.2^{+1.6}_{-1.5} \pm 3.4$	< 0.01	$\mathbf{19.5}^{+1.7}_{-1.5} \pm 3.4$	**19**
	1000–1200	$1.20^{+0.21}_{-0.19} \pm 0.26$	$9.0 \pm 0.8 \pm 1.8$	< 0.01	$\mathbf{10.2}^{+0.9}_{-0.8} \pm 1.9$	**14**
	1200–1400	$0.80^{+0.16}_{-0.14} \pm 0.22$	$4.9^{+0.5}_{-0.4} \pm 1.3$	< 0.01	$\mathbf{5.7}^{+0.5}_{-0.4} \pm 1.4$	**4**
	1400–1800	$0.43^{+0.09}_{-0.08} \pm 0.15$	$2.80^{+0.26}_{-0.24} \pm 0.98$	< 0.01	$\mathbf{3.23}^{+0.28}_{-0.26} \pm 0.99$	**3**
	>1800	$0.05 \pm 0.02 \pm 0.02$	$0.41^{+0.04}_{-0.03} \pm 0.19$	< 0.01	$\mathbf{0.46} \pm 0.04 \pm 0.19$	**0**
2-3j, 1b	400–600	$5.2^{+1.1}_{-1.0} \pm 0.6$	$13.4^{+4.9}_{-3.7} \pm 1.9$	$0.09 \pm 0.03 \pm 0.03$	$\mathbf{18.7}^{+5.0}_{-3.8} \pm 2.1$	**23**
	600–800	$1.52^{+0.43}_{-0.41} \pm 0.27$	$3.5^{+1.3}_{-1.0} \pm 1.0$	< 0.01	$\mathbf{5.0}^{+1.3}_{-1.0} \pm 1.0$	**3**
	800–1000	$0.38 \pm 0.09 \pm 0.10$	$1.53^{+0.55}_{-0.42} \pm 0.35$	< 0.01	$\mathbf{1.90}^{+0.56}_{-0.43} \pm 0.37$	**3**
	1000–1200	$0.10 \pm 0.03 \pm 0.03$	$0.81^{+0.29}_{-0.22} \pm 0.24$	< 0.01	$\mathbf{0.91}^{+0.29}_{-0.22} \pm 0.24$	**4**
	>1200	$0.19 \pm 0.06 \pm 0.08$	$0.73^{+0.26}_{-0.20} \pm 0.31$	< 0.01	$\mathbf{0.92}^{+0.27}_{-0.21} \pm 0.32$	**0**
2-3j, 2b	>400	$0.63^{+0.49}_{-0.36} \pm 0.26$	$0.0^{+3.0}_{-0.0} \pm 0.0$	< 0.01	$\mathbf{0.6}^{+3.0}_{-0.4} \pm 0.3$	**2**
2-6j, ≥3b	400–600	$1.72^{+0.73}_{-0.68} \pm 0.42$	$1.1^{+2.4}_{-0.9} \pm 0.3$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{2.8}^{+2.5}_{-1.1} \pm 0.6$	**1**
	>600	$0.37^{+0.19}_{-0.18} \pm 0.16$	$0.5^{+1.2}_{-0.4} \pm 0.2$	< 0.01	$\mathbf{0.9}^{+1.2}_{-0.5} \pm 0.2$	**0**
4-6j, 0b	400–600	$46.4^{+5.6}_{-5.1} \pm 3.6$	$176^{+15}_{-14} \pm 23$	$1.62 \pm 0.13 \pm 0.46$	$\mathbf{224}^{+16}_{-15} \pm 24$	**207**
	600–800	$10.6^{+2.3}_{-1.9} \pm 1.2$	$45.5^{+4.0}_{-3.7} \pm 9.9$	$0.07 \pm 0.03 \pm 0.02$	$\mathbf{56}^{+5}_{-4} \pm 10$	**62**
	800–1000	$4.5^{+1.1}_{-1.0} \pm 0.5$	$20.3^{+1.8}_{-1.6} \pm 3.9$	< 0.01	$\mathbf{24.8}^{+2.1}_{-1.9} \pm 4.1$	**31**
	1000–1200	$1.35^{+0.30}_{-0.26} \pm 0.24$	$10.6 \pm 0.9 \pm 2.1$	< 0.01	$\mathbf{11.9}^{+1.0}_{-0.9} \pm 2.2$	**12**
	1200–1400	$0.89^{+0.27}_{-0.25} \pm 0.23$	$5.7 \pm 0.5 \pm 1.5$	< 0.01	$\mathbf{6.6}^{+0.6}_{-0.5} \pm 1.6$	**9**
	1400–1600	$0.20 \pm 0.05 \pm 0.07$	$2.64^{+0.23}_{-0.21} \pm 0.92$	< 0.01	$\mathbf{2.84}^{+0.24}_{-0.22} \pm 0.92$	**3**
	>1600	$0.09 \pm 0.03 \pm 0.04$	$1.18 \pm 0.10 \pm 0.51$	< 0.01	$\mathbf{1.27}^{+0.11}_{-0.10} \pm 0.51$	**2**
4-6j, 1b	400–600	$21.0^{+3.7}_{-3.3} \pm 2.0$	$32.6^{+7.0}_{-5.8} \pm 5.5$	$0.81 \pm 0.09 \pm 0.23$	$\mathbf{54.5}^{+7.9}_{-6.7} \pm 6.3$	**72**
	600–800	$4.79^{+0.91}_{-0.83} \pm 0.62$	$8.4^{+1.8}_{-1.5} \pm 2.3$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{13.2}^{+2.0}_{-1.7} \pm 2.5$	**20**
	800–1000	$1.27^{+0.26}_{-0.24} \pm 0.27$	$3.71^{+0.79}_{-0.66} \pm 0.92$	$0.03 \pm 0.02 \pm 0.01$	$\mathbf{5.01}^{+0.84}_{-0.71} \pm 0.97$	**8**
	1000–1400	$0.89^{+0.21}_{-0.20} \pm 0.28$	$3.00^{+0.64}_{-0.54} \pm 0.93$	< 0.01	$\mathbf{3.89}^{+0.68}_{-0.57} \pm 0.98$	**6**
	>1400	$0.40^{+0.34}_{-0.33} \pm 0.16$	$0.72^{+0.15}_{-0.13} \pm 0.31$	< 0.01	$\mathbf{1.12}^{+0.37}_{-0.36} \pm 0.36$	**3**
4-6j, 2b	400–600	$7.2^{+1.2}_{-1.1} \pm 1.1$	$4.3^{+2.9}_{-1.9} \pm 1.4$	$0.17 \pm 0.04 \pm 0.05$	$\mathbf{11.7}^{+3.2}_{-2.2} \pm 1.9$	**11**
	600–800	$1.66^{+0.41}_{-0.40} \pm 0.46$	$1.12^{+0.76}_{-0.48} \pm 0.55$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{2.79}^{+0.86}_{-0.63} \pm 0.73$	**3**
	>800	$0.32 \pm 0.13 \pm 0.13$	$0.99^{+0.67}_{-0.43} \pm 0.52$	< 0.01	$\mathbf{1.31}^{+0.68}_{-0.45} \pm 0.54$	**4**

Table A.12: Predictions and observations for the 21 search regions with $H_\mathrm{T} > 1500$ GeV, $N_\mathrm{j} \geq 7$. For each of the background predictions, the first uncertainty listed is statistical (from the limited size of data control samples and Monte Carlo samples), and the second is systematic.

$H_\mathrm{T} > 1500$ GeV, $N_\mathrm{j} \geq 7$						
N_j, N_b	M_T2 [GeV]	Lost lepton	$Z \to \nu\bar{\nu}$	Multijet	Total background	Data
7-9j, 0b	400–600	$14.3^{+1.8}_{-1.7} \pm 1.7$	$32.3^{+7.5}_{-6.2} \pm 4.3$	$1.50 \pm 0.13 \pm 0.44$	$\mathbf{48.1}^{+7.7}_{-6.4} \pm 5.0$	**36**
	600–800	$3.77^{+0.56}_{-0.55} \pm 0.69$	$8.3^{+1.9}_{-1.6} \pm 2.2$	$0.18 \pm 0.04 \pm 0.05$	$\mathbf{12.3}^{+2.0}_{-1.7} \pm 2.3$	**9**
	800–1000	$1.16^{+0.18}_{-0.17} \pm 0.30$	$3.70^{+0.86}_{-0.71} \pm 0.83$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{4.86}^{+0.88}_{-0.73} \pm 0.90$	**6**
	1000–1400	$0.58 \pm 0.11 \pm 0.19$	$2.96^{+0.69}_{-0.57} \pm 0.86$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{3.55}^{+0.69}_{-0.58} \pm 0.89$	**4**
	>1400	$0.05 \pm 0.01 \pm 0.02$	$0.71^{+0.17}_{-0.14} \pm 0.30$	< 0.01	$\mathbf{0.76}^{+0.17}_{-0.14} \pm 0.30$	**2**
7-9j, 1b	400–600	$12.8^{+2.5}_{-2.3} \pm 1.6$	$9.2^{+4.2}_{-3.0} \pm 1.4$	$0.82 \pm 0.09 \pm 0.24$	$\mathbf{22.9}^{+4.9}_{-3.8} \pm 2.3$	**25**
	600–800	$3.49^{+0.94}_{-0.89} \pm 0.76$	$2.4^{+1.1}_{-0.8} \pm 1.0$	$0.06 \pm 0.02 \pm 0.02$	$\mathbf{5.9}^{+1.4}_{-1.2} \pm 1.2$	**7**
	>800	$1.09^{+0.34}_{-0.32} \pm 0.45$	$2.10^{+0.96}_{-0.69} \pm 0.93$	< 0.01	$\mathbf{3.2}^{+1.0}_{-0.8} \pm 1.0$	**2**
7-9j, 2b	400–600	$8.1^{+1.8}_{-1.6} \pm 1.0$	$2.4^{+1.1}_{-0.8} \pm 0.4$	$0.35 \pm 0.06 \pm 0.10$	$\mathbf{10.9}^{+2.1}_{-1.8} \pm 1.2$	**10**
	600–800	$1.78^{+0.54}_{-0.52} \pm 0.40$	$0.62^{+0.28}_{-0.20} \pm 0.25$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{2.41}^{+0.61}_{-0.56} \pm 0.49$	**5**
	>800	$0.40^{+0.19}_{-0.18} \pm 0.17$	$0.55^{+0.25}_{-0.18} \pm 0.25$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{0.96}^{+0.31}_{-0.26} \pm 0.30$	**0**
7-9j, 3b	400–800	$2.40^{+0.74}_{-0.72} \pm 0.29$	$0.32^{+0.15}_{-0.10} \pm 0.12$	$0.10 \pm 0.03 \pm 0.03$	$\mathbf{2.82}^{+0.76}_{-0.72} \pm 0.32$	**2**
	>800	$0.16 \pm 0.09 \pm 0.07$	$0.08^{+0.04}_{-0.03} \pm 0.04$	< 0.01	$\mathbf{0.24} \pm 0.09 \pm 0.08$	**0**
7-9j, ≥4b	>400	$0.52^{+0.23}_{-0.22} \pm 0.08$	$0.07^{+0.03}_{-0.02} \pm 0.06$	$0.02 \pm 0.01 \pm 0.01$	$\mathbf{0.61}^{+0.23}_{-0.22} \pm 0.10$	**1**
≥10j, 0b	400–800	$1.41 \pm 0.38 \pm 0.33$	$1.52^{+0.35}_{-0.29} \pm 0.34$	$0.23 \pm 0.05 \pm 0.08$	$\mathbf{3.17}^{+0.52}_{-0.48} \pm 0.49$	**11**
	>800	$0.05 \pm 0.02 \pm 0.02$	$0.37^{+0.09}_{-0.07} \pm 0.17$	$0.01 \pm 0.01 \pm 0.00$	$\mathbf{0.43}^{+0.09}_{-0.08} \pm 0.17$	**0**
≥10j, 1b	400–800	$2.16^{+0.71}_{-0.69} \pm 0.25$	$0.56^{+0.25}_{-0.18} \pm 0.16$	$0.14 \pm 0.04 \pm 0.05$	$\mathbf{2.85}^{+0.76}_{-0.71} \pm 0.31$	**3**
	>800	$0.55 \pm 0.30 \pm 0.22$	$0.13^{+0.06}_{-0.04} \pm 0.07$	< 0.01	$\mathbf{0.68}^{+0.31}_{-0.30} \pm 0.23$	**0**
≥10j, 2b	>400	$1.98^{+0.69}_{-0.67} \pm 0.24$	$0.30^{+0.14}_{-0.10} \pm 0.12$	$0.05 \pm 0.02 \pm 0.02$	$\mathbf{2.33}^{+0.70}_{-0.68} \pm 0.28$	**0**
≥10j, 3b	>400	$0.77 \pm 0.35 \pm 0.09$	$0.00^{+0.45}_{-0.00} \pm 0.00$	$0.05 \pm 0.03 \pm 0.02$	$\mathbf{0.82}^{+0.57}_{-0.35} \pm 0.09$	**1**
≥10j, ≥4b	>400	$0.09 \pm 0.05 \pm 0.01$	$0.00^{+0.45}_{-0.00} \pm 0.00$	< 0.01	$\mathbf{0.09}^{+0.45}_{-0.05} \pm 0.01$	**0**

References

[1] Wikipedia (user: MissMJ). `https://commons.wikimedia.org/w/index.php?curid=4286964`. Public Domain. Accessed 2020-01-05.

[2] L. Maiani, F. Piccinini, A. Polosa, *et al.*, "Diquark-Antidiquarks with Hidden or Open Charm and the Nature of $X(3872)$." *Phys. Rev. D* **71** (2005) 014028, [hep-ph/0412098].

[3] LHCb collaboration, "Observation of a narrow pentaquark state, $P_c(4312)^+$, and of two-peak structure of the $P_c(4450)^+$." *Phys. Rev. Lett.* **122** (2019) 222001, [arXiv:1904.03947].

[4] Y. Hsiao and C. Geng, "Identifying Glueball at 3.02 GeV in Baryonic B Decays." *Phys. Lett. B* **727** (2013) 168, [arXiv:1302.3331].

[5] Super-Kamiokande Collaboration, "Search for Proton Decay via $p \to e^+\pi^0$ and $p \to \mu^+\pi^0$ in a Large Water Cherenkov Detector." *Phys. Rev. Lett.* **102** (2009) 141801, [arXiv:0903.0676].

[6] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons." *Phys. Rev. Lett.* **13** (1964), no. 9, 321.

[7] P.W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons." *Phys. Rev. Lett.* **13** (1964), no. 16, 508.

[8] G.S. Guralnik, C.R. Hagen, and T.W.B. Kibble, "Global Conservation Laws and Massless Particles." *Phys. Rev. Lett.* **13** (1964), no. 20, 585.

[9] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC." *Phys. Lett. B* **716** (2012) 1, [arXiv:1207.7214].

[10] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC." *Phys. Lett. B* **716** (2012) 30, [arXiv:1207.7235].

[11] M.E. Peskin and D.V. Schroeder, *An Introduction To Quantum Field Theory.* Westview Press, 1st ed., 1995.

[12] S.L Glashow, “Partial-symmetries of weak interactions.” *Nucl. Phys.* **22** (1961), no. 4, 579.

[13] S. Weinberg, “A Model of Leptons.” *Phys. Rev. Lett.* **19** (1967) 1264.

[14] A. Salam, “Weak and Electromagnetic Interactions.” *Conf. Proc. C680519* (1968) 367.

[15] M. Tanabashi *et al.* (Particle Data Group), “Review of Particle Physics, chap. 14: Neutrino Masses, Mixing, and Oscillations.” *Phys. Rev. D* **98** (2018) 030001.

[16] L. Canetti, M. Drewes, M. Shaposhnikov, “Matter and Antimatter in the Universe.” *New J. Phys.* **14** (2012) 095012, [arXiv:1204.4186].

[17] A.V Zasov, A.S. Saburova, A.V. Khoperskov, and S.A. Khoperskov, “Dark matter in galaxies.” *Physics-Uspekhi* **80** (2017), no. 1, 3, [arXiv:1710.10630].

[18] D. Hanneke, S. Fogwell Hoogerheide, and G. Gabrielse, “Cavity Control of a Single-Electron Quantum Cyclotron: Measuring the Electron Magnetic Moment.” *Phys. Rev. A* **83** (2010) 052122, [arXiv:1009.4831].

[19] T. Blum, A. Denig, I. Logashenko, *et al.*, “The Muon $(g-2)$ Theory Value: Present and Future.” (2013) [arXiv:1311.2198].

[20] E. Graverini for the ATLAS, CMS, and LHCb Collaborations, “Flavour anomalies: a review.” *J. Phys.: Conf. Ser.* **1137** (2019) 012025, [arXiv:1311.2198].

[21] Heavy Flavor Averaging Group (HFLAV), “Average of $R(D)$ and $R(D^*)$ for Summer 2018.” `https://hflav-eos.web.cern.ch/hflav-eos/semi/summer18/RDRDs.html`. Accessed 2020-01-21.

[22] S.P. Martin, “A Supersymmetry Primer.” (1997) [hep-ph/9709356].

[23] A. Djouadi, S. Rosier-Lees, M. Bezouh, *et al.*, “The Minimal Supersymmetric Standard Model: Group Summary Report.” (1999) [hep-ph/9901246].

[24] J. Alwall, P.C. Schuster, and N. Toro, “Simplified models for a first characterization of new physics at the LHC.” *Phys. Rev. D* **79** (2009) 075020.

[25] J. Alwall, M. Le, M. Lisanti, J.G. Wacker, “Model-independent jets plus missing energy searches.” *Phys. Rev. D* **79** (2009) 015005.

[26] LHC New Physics Working Group Collaboration, “Simplified models for LHC new physics searches.” *J. Phys. G* **39** (2012) 105005.

[27] H. Georgi and S.L. Glashow, “Unity of All Elementary-Particle Forces.” *Phys. Rev. Lett.* **32** (1974) 438.

[28] H. Fritzsch and P. Minkowski, "Unified interactions of leptons and hadrons." *Annals of Physics* **93** (1975) 193.

[29] J.C. Pati and A. Salam, "Unified lepton-hadron symmetry and a gauge theory of the basic interactions." *Phys. Rev. D* **8** (1973) 1240.

[30] S. Dimopoulos, "Technicolored signatures." *Nucl. Phys. B* **168** (1980) 69.

[31] E. Farhi and L. Susskind, "Technicolor." *Phys. Rept.* **74** (1981) 277.

[32] K.D. Lane and M.V. Ramana, "Walking technicolor signatures at hadron colliders." *Phys. Rev. D* **44** (1991) 2678.

[33] B. Schrempp and F. Schrempp, "Light leptoquarks." *Phys. Lett. B* **153** (1985) 101.

[34] B. Gripaios, "Composite leptoquarks at the LHC." *J. High Energy Phys.* **02** (2010) 045, [arXiv:0910.1789].

[35] R. Barbier, C. Brat, M. Besanon, *et al.*, "*R*-parity violating supersymmetry." *Phys. Rept.* **420** (2005) 1, [hep-ph/0406039].

[36] D. Buttazzo, A. Greljo, G. Isidori, and D. Marzocca, "B-physics anomalies: a guide to combined explanations." *J. High Energy Phys.* **11** (2017) 044, [arXiv:1706.07808].

[37] Phys.org, "HERA experiments H1 and ZEUS publish most precise picture of the proton." `phys.org/news/2015-07-hera-h1-zeus-publish-precise.html`. Accessed 2020-01-23.

[38] L.A. Harland-Lang, A.D. Martin, P. Motylinski, and R.S. Thorne, "Parton distributions in the LHC era: MMHT 2014 PDFs." *Eur. Phys. J. C* **75** (2015) 204, [arXiv:1412.3989].

[39] CMS Collaboration, "Summaries of CMS cross section measurements." `http://cern.ch/go/pNj7`. Accessed 2020-01-22.

[40] W. Beenakker, C. Borschensky, M. Krämer, *et al.*, "NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation." *J. High Energy Phys.* **12** (2016) 133, [arXiv:1607.07741].

[41] OpenStreetMap. `https://wiki.openstreetmap.org/wiki/CERN`. Accessed 2020-01-27.

[42] S. Charley, "The LHC does a dry run." Symmetry Magazine (Mar. 20, 2015). Image originally from Maximillien Brice, CERN.

[43] CERN. `http://public-archive.web.cern.ch/public-archive/en/research/AccelComplex-en.html`. Accessed 2020-01-27.

[44] CERN, "LHC: the guide." https://home.cern/resources/brochure/accelerators/lhc-facts-and-figures. Accessed 2020-01-27.

[45] ATLAS Collaboration, "Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC." *Phys. Rev. Lett.* **117** (2016) 182002, [arXiv:1606.02625].

[46] M. Bajko, F. Bertinelli, N. Catalan Lasheras, *et al.*, "Report of the Task Force on the Incident of 19th September 2008 at the LHC." Tech. Rep. LHC-PROJECT-Report-1168, CERN, March 2009.

[47] T. Sakuma and T. McCauley, "Detector and Event Visualization with SketchUp." *J. Phys.: Conf. Ser.* **513** (2014) 022032.

[48] CMS Collaboration, "CMS Technical Design Report, Volume I: Detector Performace and Software." Tech. Rep. CMS-TDR-8-1, CERN, 2006.

[49] CMS Collaboration, "CMS Technical Design Report, Volume II: Physics Performance." *J. Phys. G* **34** (2007) 995–1579.

[50] J. Alwall, R. Frederix, S. Firxione, *et al.*, "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations." *J. High Energy Phys.* **07** (2014) 079, [arXiv:1405.0301].

[51] T. Sjöstrand, S. Mrenna, and P. Skands, "An Introduction to PYTHIA 8.2." *Comput. Phys. Commun.* **191** (2015) 159, [arXiv:1410.3012].

[52] GEANT4 Collaboration, "GEANT4–a simulation toolkit." *Nucl. Instrum. Methods A* **506** (2003) 250.

[53] A. Giammanco, "The Fast Simulation of the CMS Experiment." *J. Phys.: Conf. Ser.* **513** (2014) 022012.

[54] CMS Collaboration, "Technical Proposal for the Phase-II Upgrade of the CMS Detector." Tech. Rep. CMS-TDR-15-02, CERN, 2015.

[55] CMS Collaboration, "A MIP Timing Detector for the CMS Phase-2 Upgrade." Tech. Rep. CMS-TDR-020, CERN, 2019.

[56] N. Moffat, R. Bates, M. Bullough, *et al.*, "Low Gain Avalanche Detectors (LGAD) for particle physics and synchrotron applications." *JINST* **13** (2018) C03014.

[57] CMS Collaboration, "Particle-flow reconstruction and global event description with the CMS detector." *JINST* **12** (2017) P10003, [arXiv:1706.04965].

[58] M. Cacciari and G.P. Salam, "The anti-k_t jet clustering algorithm." *J. High Energy Phys.* **04** (2008) 063, [arXiv:0802.1189].

[59] CMS Collaboration, "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV." *JINST* **13** (2018) P05011, [arXiv:1712.07158].

[60] D. Guest, J. Collado, P. Baldi, *et al.*, "Jet flavor classification in high-energy physics with deep neural networks." *Phys. Rev. D* **94** (2016) 112002, [arXiv:1607.08633].

[61] CMS Collaboration, "Pileup Jet Identification." Tech. Rep. CMS-PAS-JME-13-005, CERN, 2013.

[62] CMS Collaboration, "Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV." *JINST* **12** (2017) P02014, [arXiv:1607.03663].

[63] CMS Collaboration, "Measurements of inclusive W and Z cross sections in pp collisions at $\sqrt{s} = 7$ TeV." *J. High Energy Phys.* (2011), no. 01, 80, [arXiv:1012.2466].

[64] C. Lester and D. Summers, "Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders." *Phys. Lett. B* **463** (1999) 99, [hep-ph/9906349].

[65] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker." *JINST* **9** (2014) P10009, [arXiv:1405.6569].

[66] CMS Collaboration, "Pileup Removal Algorithms." Tech. Rep. CMS-PAS-JME-14-001, CERN, 2014.

[67] CMS Collaboration, "Search for new physics with the M_{T2} variable in all-jets final states produced in pp collisions at $\sqrt{s} = 13$ TeV." *J. High Energy Phys.* **10** (2016) 6, [arXiv:1603.04053].

[68] CMS Collaboration, "Search for new phenomena with the M_{T2} variable in the all-hadronic final state produced in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Eur. Phys. J. C* **77** (2017) 710, [arXiv:1705.04650].

[69] CMS Collaboration, "Measurements of $t\bar{t}$ cross sections in association with b jets and inclusive jets and their ratio using dilepton final states in pp collisions at $\sqrt{s} = 13$ TeV." *Phys. Lett. B* **776** (2018) 355, [arXiv:1705.10141].

[70] CMS Collaboration, "Search for top squark pair production in pp collisions at $\sqrt{s} = 13$ TeV using single lepton events." *J. High Energy Phys.* (2017), no. 10, 019, [arXiv:1706.04402].

[71] F. James and M. Roos (CERN), "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations." *Comput. Phys. Commun.* **10** (1975) 343.

[72] CMS Collaboration, "Searches for physics beyond the standard model with the M_{T2} variable in hadronic final states with and without disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Eur. Phys. J. C* **80** (2020) 3, [arXiv:1909.03460].

[73] A. L. Read, "Presentation of search results: the $\mathrm{CL_S}$ technique." *J. Phys. G* **28** (2002) 2693.

[74] G. Cowan, K. Cramer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics." *Eur. Phys. J. C* **71** (2011) 1554.

[75] "RooFit." `https://root.cern.ch/roofit`. Accessed 2019-12-20.

[76] P. Asadi, M.R. Buckley, A. DiFranzo, A. Monteux, and D. Shih, "Digging Deeper for New Physics in the LHC Data." *J. High Energy Phys.* **11** (2017) 194, [arXiv:1707.05783].

[77] P. Asadi, M.R. Buckley, A. DiFranzo, A. Monteux, and D. Shih, "An Update on the LHC Monojet Excess." *J. High Energy Phys.* **03** (2018) 130, [arXiv:1712.04939].

[78] ATLAS Collaboration, "Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector." *J. High Energy Phys.* **01** (2018) 126, [arXiv:1711.03301].

[79] ATLAS Collaboration, "Search for squarks and gluinos in final states with jets and missing transverse momentum using 36 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector." *Phys. Rev. D* **97** (2018) 112001, [arXiv:1712.02332].

[80] CMS Collaboration, "Search for supersymmetry in multijet events with missing transverse momentum in proton-proton collisions at 13 TeV." *Phys. Rev. D* **96** (2017) 032003, [arXiv:1704.07781].

[81] CMS Collaboration, "Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV." *Phys. Rev. D* **97** (2018) 092005, [arXiv:1712.02345].

[82] CMS Collaboration, "Constraints on models of scalar and vector leptoquarks decaying to a quark and a neutrino at $\sqrt{s} = 13$ TeV." *Phys. Rev. D* **98** (2018) 032005, [arXiv:1805.10228].

[83] A. Ball, G. Beauregard, J. Brooke, *et al.*, "Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Submitted to Phys. Rev. D* (2020) [arXiv:2005.06518].

[84] N. Vinyoles and H. Vogel, "Minicharged Particles from the Sun: A Cutting-Edge Bound." *J. Cosmology and Astroparticle Phys.* **03** (2016) 002, [arXiv:1511.01122].

[85] A. Haas, C.S. Hill, E. Izaguirre, and I. Yavin, "Looking for milli-charged particles with a new experiment at the LHC." *Phys. Lett. B* **746** (2015) 117, [arXiv:1410.6816].

[86] A. Ball, J. Brooke, C. Campagnari, *et al.*, "A Letter of Intent to Install a Milli-charged Particle Detector at LHC P5." (2016) [arXiv:1607.04669].

[87] R. Saldanha, L. Grandi, Y. Guardincerri, *et al.*, "Model Independent Approach to the Single Photoelectron Calibration of Photomultiplier Tubes." *Nucl. Instrum. Methods A* **863** (2017) 35, [arXiv:1602.03150].

[88] DRS board developed by the Paul Scherrer Institut. `https://www.psi.ch/en/drs/evaluation-board`. Accessed 2020-02-13.

[89] R. Van Royen and V.F. Weisskopf, "Hadron Decay Processes and the Quark Model." *Nuovo Cim. A* **50** (1967) 617.

[90] L.G. Landsberg, "Electromagnetic Decays of Light Mesons." *Phys. Rept.* **128** (1985) 301.

[91] Y.-Q. Ma, K. Wang, and K.-T. Chao, "J/$\psi(\psi')$ production at the Tevatron and LHC at $\mathcal{O}(\alpha_s^4 v^4)$ in nonrelativistic QCD." *Phys. Rev. Lett.* **106** (2011) 042002.

[92] M Cacciari, S. Firxione, N. Houdeau, M.L. Mangano, P. Nason, and G. Ridolfi, "Theoretical predictions for charm and bottom production at the LHC." *J. High Energy Phys.* **10** (2012) 137, [arXiv:1205.6344].

[93] CMS Collaboration, "Measurement of quarkonium production cross sections in pp collisions at $\sqrt{s} = 13$ TeV." *Phys. Lett. B* **780** (2018) 251–272, [arXiv:1710.11002].

[94] ATLAS Collaboration, "Measurement of Upsilon production in 7 TeV pp collisions at ATLAS." *Phys. Rev. D* **87** (2013), no. 5, 052004, [arXiv:1211.7255].

[95] P. Skands, S. Carrazza, and J. Rojo, "Tuning PYTHIA 8.1: the Monash 2013 Tune." *Eur. Phys. J. C* **74** (2014), no. 8, 3024, [arXiv:1404.5630].

[96] ALICE Collaboration, "Production of $\omega(782)$ in pp collisions at $\sqrt{s} = 7$ TeV." Tech. Rep. ALICE-PUBLIC-2018-004, CERN, 2018.

[97] ALICE Collaboration, "Production of the $\rho(770)^0$ meson in pp and Pb-Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV." *Phys. Rev. C* **99** (2019), no. 6, 064901, [arXiv:1805.04365].

[98] ALICE Collaboration, "π^0 and η meson production in proton-proton collisions at $\sqrt{s} = 8$ TeV." *Eur. Phys. J. C* **78** (2018), no. 3, 263, [arXiv:1708.08745].

[99] CMS Collaboration, "Measurement of charged pion, kaon, and proton production in proton-proton collisions at $\sqrt{s} = 13$ TeV." *Phys. Rev. D* **96** (2017), no. 11, 112003, [arXiv:1706.10194].

[100] T. Sjostrand, S. Mrenna, and P. Z. Skands, "PYTHIA 6.4 Physics and Manual." *J. High Energy Phys.* **05** (2006) 026, [hep-ph/0603175].

[101] M. Albrow, M. Begel, D. Bourilkov, *et al.*, "Tevatron-for-LHC Report of the QCD Working Group." (2006) [hep-ph/0610012].

[102] ATLAS Collaboration, "The differential production cross section of the ϕ (1020) meson in $\sqrt{s} = 7$ TeV pp collisions measured with the ATLAS detector." *Eur. Phys. J. C* **74** (2014), no. 7, 2895, [arXiv:1402.6162].

[103] CMS Collaboration, "Measurement of the total and differential inclusive B^+ hadron cross sections in pp collisions at $\sqrt{s} = 13$ TeV." *Phys. Lett. B* **771** (2017) 435, [arXiv:1609.00873].

[104] M. Tanabashi *et al.* (Particle Data Group), "Review of Particle Physics, chap. 33: Passage of Particles Through Matter." *Phys. Rev. D* **98** (2018) 030001.

[105] CMS Collaboration, "Search for fractionally charged particles in pp collisions at $\sqrt{s} = 7$ TeV." *Phys. Rev. D* **87** (2013) [arXiv:1210.2311].

[106] CMS Collaboration, "Searches for long-lived charged particles in pp collisions at $\sqrt{s}$ =7 and 8 TeV." *J. High Energy Phys.* **07** (2013) [arXiv:1305.0491].

[107] R. Acciarri, C. Adams, J. Asaadi, *et al.* (ArgoNeuT), "Improved Limits on Millicharged Particles Using the ArgoNeuT Experiment at Fermilab." *Phys. Rev. Lett.* **124** (2020) [arXiv:1911.07996].

[108] A. A. Prinz, R. Baggs, J. Ballam, *et al.*, "Search for Millicharged Particles at SLAC." *Phys. Rev. Lett.* **81** (1998) 1175, [hep-ex/9804008].

[109] C. Brust, D.E. Kaplan, and M.T. Walters, "New Light Species and the CMB." *J. High Energy Phys.* **12** (2013) 058, [arXiv:1303.5379].