Measurement of the Drell-Yan Four-fold Differential Cross Section with the ATLAS Experiment

Craig John Wells

University of Freiburg
Faculty of Mathematics and Physics
Institute of Physics

universität freiburg

Measurement of the Drell-Yan Four-fold Differential Cross Section with the ATLAS Experiment

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Fakultät für Mathematik und Physik
der
Albert-Ludwigs-Universität Freiburg

vorgelegt von
Craig John Wells

Juli 2025

Dekan

Prof. Dr. Michael Růžička

Betreuer

Prof. Dr. Beate Heinemann, Dr. Ludovica Aperio Bella, Dr. Christian Weiser

Gutachter der Disseration

Prof. Dr. Horst Fischer
Prof. Dr. Beate Heinemann

Datum der mündlichen Prüfung

16.05.2025

Prüfer der mündlichen Prüfung

Prof. Dr. Marco Gersabeck
Prof. Dr. Beate Heinemann
Prof. Dr. Heidi Rzehak

Abstract

With the discovery of the Higgs boson in 2012 at the Large Hadron Collider at CERN, the final elementary particle predicted by the Standard Model of particle physics was discovered. Since then the study of the Standard Model has focussed on precision measurements of various parameters and processes of the theory in the hope of measuring subtle differences between the theoretically predicted values and those observed in nature. This thesis presents one such measurement, that of the Drell-Yan cross section through its decomposition to a set of 8 angular coefficients, 9 harmonic polynomials and an unpolarised cross section. This decomposition allows for a measurement of the Drell-Yan cross section in the phase space defined by the Z boson mass, transverse momentum and rapidity which is unaffected by kinematic cuts on the decay leptons. This measurement uses the Run 2 dataset of the ATLAS experiment collected between 2015 and 2018 which encompasses 140 fb^{-1} of data recorded at $\sqrt{s} = 13$ TeV from proton-proton collisions. The angular coefficients and unpolarised cross section are measured in both the $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ decay channels as well as the combination of both channels which provides a more precise result than what can be found in the literature. The combination result is also compared to theoretical predictions where it is found to be in good agreement with the Standard Model.

A further Asimov sensitivity study on the effective leptonic weak mixing angle is presented through the reparameterisation of the A_4 angular coefficient previously measured. In this study the $Z \to e^+e^-$ channel is extended to include electrons with pseudorapidity values in the range $2.5 < |\eta| < 4.9$ where the sensitivity to the weak mixing angle is at its greatest. The calibration process to use these forward electrons for physics analysis purposes is described and results in a good agreement between the data and simulated Z boson mass spectra. The results of this sensitivity study show that the combination of the electron and muon channels may provide the most precise measurement of the weak mixing angle from a hadron collider and become competitive with the sensitivities achieved at lepton colliders.

Zusammenfassung

Mit der Entdeckung des Higgs-Bosons im Jahr 2012 am Large Hadron Collider am CERN wurde das letzte Elementarteilchen entdeckt, das von dem Standardmodell der Teilchenphysik vorausgesagt wurde. Seitdem hat sich die Erforschung des Standardmodells mit präzisen Messungen verschiedener Parameter und Prozesse der Theorie befasst, in der Hoffnung, subtile Unterschiede zwischen den theoretisch vorausgesagten Werten und denen, die in der Natur beobachtet werden, zu entdecken. Diese Dissertation stellt eine solche Messung vor, die des Drell-Yan-Querschnitts durch seine Zerlegung in 8 Winkelkoeffizienten, 9 harmonische Polynome und einen unpolarisierten Querschnitt. Diese Zerlegung sorgt für eine Messung des Drell-Yan-Querschnitts in dem Phasenraum, der von der Masse, dem Transversalimpuls und der Rapidität des Z-Bosons bestimmt wird, und von den kinematischen Auswahlkriterien der Zerfallsleptonen unbetroffen ist. Diese Messung verwendet den Run 2 Datensatz des ATLAS-Experiments, der zwischen den Jahren 2015 und 2018 in Proton-Proton-Kollisionen bei $\sqrt{s} = 13$ TeV gesammelt wurde und einer integrierten Luminosität von 140 fb^{-1} entspricht. Die Winkelkoeffizienten und der unpolarisierte Querschnitt werden in den $Z \to e^+e^-$ und $Z \to \mu^+\mu^-$ Zerfallskanälen und auch in der Kombination der beiden Kanäle gemessen, für ein Ergebnis, das präziser ist, als das, das in der momentanen Literatur zu finden ist. Das Kombinationsergebnis wird auch mit theoretischen Voraussagen verglichen, wobei es gut mit dem Standardmodell übereinstimmt.

Eine weitere Asimov-Sensitivitätsstudie zum effektiven leptonischen schwachen Mischungswinkel durch die Umparametrierung des vorher gemessenen A_4 Koeffizients wird auch vorgestellt. In dieser Studie wird der $Z \to e^+e^-$ Kanal erweitert, um Elektronen mit Pseudorapiditätwerten $2.5 < |\eta| < 4.9$ einzubeziehen, für die die Sensitivität zum schwachen Mischungswinkel am größten ist. Das Kalibrierungsverfahren, durch das die Vorwärtselektronen für eine Physikanalyse verwendet werden können, wird beschrieben, und die Ergebnisse zeigen eine gute Übereinstimmung zwischen den Daten und den simulierten Z-Boson Massespektren. Die Ergebnisse dieser Studie zeigen, dass die Kombination der Elektron- und Muonkanäle die präziseste Messung des schwachen Mischungswinkels von einem Hadronenbeschleuniger liefern könnte und mit den Sensitivitäten, die an Leptonenbeschleunigern erreicht worden sind, konkurrenzfähig werden kann.

Contents

Chapter 1

Introduction

The simple act of asking a question and answering it with rigorous experimental evidence has underpinned the scientific method for centuries. This has led to the greatest scientific discoveries in the history of humanity and has brought knowledge about how the universe works on every level; from the unfathomably large galactic scale, all the way down to the counterintuitive quantum mechanical subatomic world.

It is in the subatomic world in particular that especially rapid progress in the acquirement of knowledge has been made. Indeed, the past 130 years have gone from the discovery of the first elementary particle to the development of a Quantum Field Theory called the Standard Model that explains almost every aspect of particle physics. When the Higgs boson was discovered in 2012 by the ATLAS and CMS experiments, the final piece of experimental evidence for the Standard Model demanded by the scientific method had been found. It did not, however, mean the end of the field of particle physics due to the open questions that are not answered by the Standard Model. Whilst there have been many proposed extensions to the Standard Model which would answer many of these open questions, they all lack experimental evidence. This therefore begs the question, how well do we really know the Standard Model?

The answer is, of course, provided by following the scientific method. By revisiting previous measurements of the processes and parameters of the Standard Model and performing them at a higher level of precision, it may be possible to observe some subtle difference between what is predicted by the Standard Model and what is observed in nature. One such measurement would be that of the Drell-Yan cross section, which has been measured several times in the past by various experiments. Therefore, this thesis will present a measurement of the Drell-Yan cross section and compare it to the predictions of the Standard Model.

This thesis is structured as follows. Chapter 2 gives an overview of the theory and motivation behind this thesis, including the Standard Model and the Drell-Yan process whose cross section is to be measured. Chapter 3 describes the experimental setup necessary to measure the Drell-Yan cross section, namely the Large Hadron Collider at CERN and the ATLAS detector that sits along its ring. Chapter 4 describes how leptons are reconstructed, identified and calibrated with the ATLAS detector, with a particular focus on the calibration of forward electron objects. Chapter 5 describes the measurement of the Drell-Yan cross section beginning by defining all sources of signal and background, then describing the statistical methods used to measure the cross section and finally presenting the results. Chapter 6 presents the results of a study on the sensitivity to the weak mixing angle and Chapter 7 is the conclusion.

Author's Contribution As an individual, it would be nigh on impossible to design, construct, operate and analyse enough data from a particle detector in the modern age. Therefore, without the work of thousands of other physicists and technicians who work on the ATLAS experiment, this thesis would not have been possible. The author's contribution to the results of this thesis are summarised below.

In Chapter 4, the training, evaluation and implementation of the forward electron global scale energy calibration was carried out by the author. The uniformity correction and in situ calibration were also derived and evaluated by the author. In addition, the author also carried out parts of the software implementation of the in situ calibration.

For the measurement of the Drell-Yan cross section in Chapter 5, the author was responsible for the optimisation of the multijet backgrounds in both the central and central-forward electron channels. The author also carried out all other aspects for the fitting procedure and evaluation of the results in the single electron and muon channels. The compatibility study and combination fit result were performed wholly by the author.

When carrying out the sensitivity study to the weak mixing angle in Chapter 6, the author was responsible for the A_4 fits in the central-forward channel, the combination of the central electron and muon channels as well as the triple combination of all three analysis channels.

In addition, the author also contributed to the maintenance and development of the main software frameworks that were used to derive the results presented in Chapter 5 and the A_4 results in Chapter 6.

Chapter 2

Theory and Motivation

2.1 Particle Nature of the Universe

The idea that the macroscopic universe is composed of imperceptibly small particles is certainly not a new one. Millenia ago, the ancient Greeks devised a thought experiment involving the repeated halving of a plank of wood. They argued that eventually, one would eventually be left with a piece of wood so small that it could not be divided any further. They called this indivisible unit atomos, or atom in modern English. They hypothesised that atoms would come in four different elements; fire, earth, air and water. Out of these four elements, they posited that the entire universe could be formed.

In the intervening time period, the atoms of ancient Greece have since become Dmitri Mendeleev's periodic table of the elements and form the basis of the study of chemistry. It has also been observed experimentally that the atom is also not an elementary particle, meaning that it has an inner structure. Instead it is made up of an electrically positively charged nucleus, surrounded by a cloud of negatively charged electrons. The nucleus is itself made up of positively charged protons and neutral neutrons, which are themselves made up of elementary particles called up and down quarks. Thus, almost all of the observed matter in the universe is made up of combinations of these three particles.

This chapter will present an introduction to the modern day Standard Model of Particle Physics which describes the particles observed in nature and the interactions between them. The shortcomings of the Standard Model will also be discussed before an overview of the Drell-Yan process is given, the measurement of which is the overarching goal of this work.

2.2 Elementary Particles and Fundamental Forces

Experimental efforts spanning the late 19th to the early 21st centuries have discovered the existence of 17 distinct types of elementary particle in nature; each of which is depicted

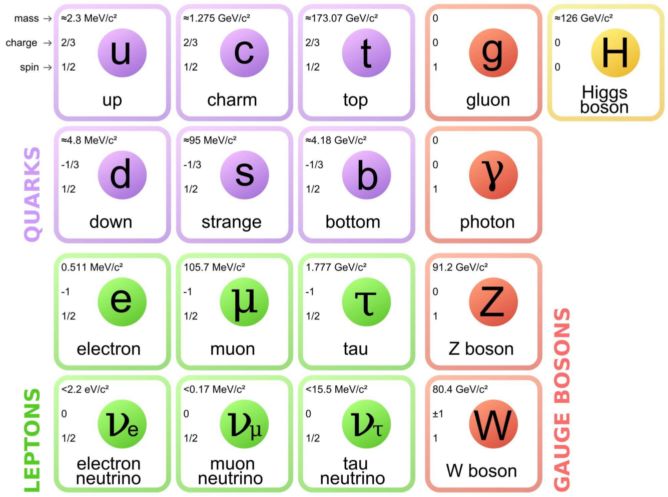

Figure 2.1: All known elementary particles of nature [20].

in Figure 2.1. These 17 particle types can be split into two main categories, depending on their intrinsic angular momentum (otherwise called spin): fermions and bosons. Fermions are particles that have half-integer spin, obey the Pauli Exclusion Principle and behave according to Fermi-Dirac statistics. Conversely bosons carry integer spin and behave according to Bose-Einstein statistics.

The fermions can be further categorised as either quarks or leptons. The six quarks carry a 'colour' charge and also have an electric charge of either +2/3 for up-type quarks and -1/3 for down-type quarks. The six leptons do not have colour charge and only the electron, muon and tau three carry an electric charge of -1. The remaining three leptons are collectively known as neutrinos, with each one carrying a lepton 'flavour' corresponding to one of electron, muon or tau. The quarks and leptons can be further organised into three 'generations', each with an increasing mass, which correspond to the columns of Figure 2.1.

The bosons can also be split into two categories, the spin-1 gauge bosons and the spin-0 scalar Higgs boson. The Higgs field (whose excitation is the Higgs boson) gives mass to all charged fermions as well as to the massive gauge bosons, whereas the gauge bosons allow the fermions to interact with each other. In turn this allows for the formation of fermion bound states which are the fundamental building blocks of the universe. Each gauge boson corresponds to one of the four fundamental forces of nature:

- the strong force
- electromagnetism
- the weak force
- gravity

Gravity is the weakest of all the fundamental forces, becoming noticeable only in interactions involving heavenly bodies. The weak force is then responsible for the β decay modes of atomic nuclei as well as flavour changing decays in the quark sector. The strong force binds the quarks and allows the formation of atomic nuclei to occur. Without the strong force, atomic nuclei would disintegrate due to the electromagnetic repulsion of the protons. Finally, electromagnetism is responsible for the formation of atoms, providing the necessary attractive force between negatively charged electrons and positively charged nuclei.

It is also known that 'antiparticles' exist as well - electric and colour charge conjugated versions of the particles already described. Most particles therefore have a corresponding partner antiparticle, which is denoted as the anti- version of its particle counterpart, e.g the antiparticle of the up quark is the anti-up quark. Some particles, such as the photon and Z boson, are their own antiparticle since they already carry no electric or colour charge. For the neturinos it is still unknown whether they have a distinct antiparticle, or if they are their own antiparticle.

2.3 The Standard Model of Particle Physics

So far, the treatment of elementary particles and the fundamental forces has been purely qualitative. In order to do calculations and predict the output of interactions between the particles, one must turn to a mathematical theoretical framework. In high energy physics, this framework is provided by a relativistic quantum field theory called the Standard Model of Particle Physics. In the Standard Model, each particle type is treated as a field, whose evolution is typically described using the Lagrangian formalism. The Standard Model Lagrangian, $\mathcal{L}_{SM}$, is typically written as the sum of the the Lagrangians of the individual sectors that make up the Standard Model as a whole:

$$\mathcal{L}_{SM} = \mathcal{L}_{EW} + \mathcal{L}_{QCD} + \mathcal{L}_{Higgs} + \mathcal{L}_{Yukawa} \tag{2.1}$$

In the language of group theory, the Standard Model can also be represented as:

$$SU(3) \otimes SU(2) \otimes U(1) \tag{2.2}$$

where $SU(3)(SU(2))$ is the special unitary group of 3×3 (2×2) matrices of unit determinant with matrix multiplication as the group operator. $U(1)$ is the group of 1×1 unitary matrices with matrix multiplication as the group operator i.e the complex numbers with magnitude 1 with ordinary multiplication.

Under this formalism, the Standard Model is therefore:

- minimal, as the fewest possible number of terms are used to describe its particle content and their interactions

- renormalisable, the predictions of the Standard Model are finite at any given energy scale
- invariant under transformations of the gauge group given by Equation (2.2)
- Lorentz invariant, meaning it holds in every inertial reference frame

2.3.1 Electroweak Sector

In the development of the Standard Model, the first step was to formulate the classical theory of electromagnetism as a quantum field theory. This was achieved in the 1950s with the $U(1)$ theory of quantum electrodynamics (QED) [57, 58]. In QED, the idea of a localised point particle is replaced with a quantum field and its excitations. The excitations of these quantum fields are what would be detected as particles in the lab. Interactions between the fields are mediated by a separate photon field, which manifests as a massless, electrically neutral, spin-1 gauge boson.

Whilst QED was being developed, the weak force was thought to be a separate force entirely. However, in the 1960s it was shown by Glashow, Salam and Weinberg [60, 82, 86] that at high energy scales the forces are unified into a single electroweak (EW) force. In the language of group theory, the EW force is represented by the direct product of $SU(2) \otimes U(1)$. The resultant non-Abelian group predicts the presence of four gauge boson fields, B_μ from the $U(1)$ symmetry and W^a_μ $(a = 1, 2, 3)$ from $SU(2)$. From this, the electroweak Lagrangian can be written:

$$\mathcal{L}_{EW} = \sum_f \overline{\psi}_f i\gamma^\mu D_\mu \psi_f - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} W^a_{\mu\nu} W^{\mu\nu}_a \tag{2.3}$$

where ψ_f is the field corresponding to fermion f and γ^μ is the appropriate Dirac gamma matrix. The operator D_μ is defined to be the covariant derivative, which maintains the local gauge invariance of the Lagrangian. In its explicit form, it is:

$$D_\mu = \partial_\mu + i\frac{g'}{2} Y B_\mu + i\frac{g}{2} T_j W^j_\mu \tag{2.4}$$

where g and g' are coupling constants, Y is the weak hypercharge, which is the conserved quantity in electroweak interactions according to Noether's theorem and analogous to the electric charge in QED. The T_j are 2×2 matrices and act as the generator of $SU(2)$. Specifically, matrix $T_j = \sigma_j/2$ where σ_j is the Pauli matrix j. In the Standard Model, the weak force interacts only with left handed fermions and right handed antifermions. This has the consequence that all the T_j corresponding to forbidden interactions are exactly zero. For the non-zero T_j, their eigenvalues correspond to the weak isospin, a quantum number which is conserved by weak interactions. The theory also predicts that the combination of the weak hypercharge and weak isospin result in the electric charge from electromagnetism:

$$Q = \frac{1}{2}Y + T_3 \tag{2.5}$$

with T_3 as the third component of the weak isospin.

The final components of $\mathcal{L}_{EW}$ are the $B_{\mu\nu}$ and $W^a_{\mu\nu}$ field strength tensors. The Abelian nature of $U(1)$ means that $B_{\mu\nu}$ can be simply defined as:

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \tag{2.6}$$

In comparison, $SU(2)$ is non-Abelian, which gives rise to self-interactions between the W_μ bosons. When the field strength tensor is explicitly written:

$$W^a_{\mu\nu} = \partial_\mu W^a_\nu - \partial_\nu W^a_\mu + g f^{abc} W^b_\mu W^c_\nu \tag{2.7}$$

The non-Abelian self-interactions manifest through the extra term containing the group structure constants f^{abc}, which are defined by the commutator $[T_a, T_b] = i f^{abc} T_c$.

In the electroweak sector, the parity (P) and charge conjugation (C) operations are of particular importance. Applying the parity transformation to a system has the effect of taking the mirror image of the process i.e applying the parity operator, $\hat{P}$, to a general function, f, of a position vector, $\underline{x}$ is equivalent to $\hat{P}f(\underline{x}) = f(-\underline{x})$. Instead of acting spatially on a particle, the charge conjugation operator turns a particle into its antiparticle equivalent and vice versa. It was originally thought that parity, charge conjugation and the combination of the two (CP) were fundamental symmetries of nature. This was, however, disproven in the mid-20th century. In 1956, Wu et al [90] showed that parity was not conserved in the weak decays of cobalt-60 nuclei, whereas in 1964, Cronin and Fitch [35] detected CP violating decays in the kaon system. Whilst CP violation is one of the three necessary Sakharov conditions [81] to generate an excess of matter over antimatter in the universe, the amount of CP violation observed within the electroweak sector is not large enough to generate the extremely large imbalance that is observed.

2.3.2 The Brout-Englert-Higgs Mechanism

In electroweak theory adding an explicit mass term to the Lagrangian for the gauge bosons is forbidden since this would violate gauge invariance. A separate mechanism is therefore required to give the $W^\pm$ and Z bosons the mass that they have been observed to have in nature. In the Standard Model, this is provided through the Brout-Englert-Higgs mechanism of spontaneous symmetry breaking, or simply the Higgs mechanism [56, 66].

In the Higgs mechanism a new complex scalar field, ϕ, and a corresponding potential $V(\phi)$ are introduced:

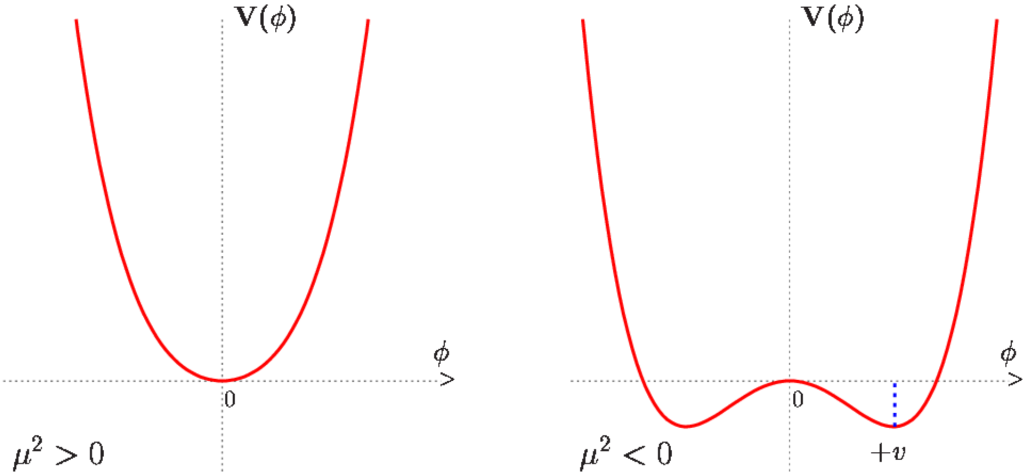

Figure 2.2: The shape of the Higgs potential in one dimension [52]. The left hand plot shows the shape of the Higgs potential when $\mu^2 > 0$ and no spontaneous symmetry breaking occurs. The right hand plot shows the case where $\mu^2 < 0$, meaning the symmetry is broken and the potential has a non-zero vacuum expectation value, v.

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \tag{2.8}$$

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda(\phi^\dagger \phi)^2 \tag{2.9}$$

where μ^2 and λ are two free parameters unconstrained by theory. In order for the vacuum to be stable i.e that it has is a minimum value as $V \rightarrow -\infty$, λ must be strictly greater than zero. This means that the shape of the potential depends on the behaviour of μ^2. As shown on the left hand side of Figure 2.2, if $\mu^2 > 0$, the minimum of the potential will occur when $\phi = 0$, representing a symmetric state. However, if $\mu^2 < 0$, the minimum of the potential no longer occurs at $\phi = 0$, instead at:

$$\phi(0) = \sqrt{\frac{-\mu^2}{2\lambda}} \equiv v \tag{2.10}$$

where v is the so-called vacuum expectation value of the field. At v, the symmetry of the ground state has been spontaneously broken and is now degenerate in ϕ, as is depicted in the right hand side of Figure 2.2. By performing a $SU(2)$ gauge transformation and considering a perturbation around the vacuum expectation, h, allows ϕ to be written as:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \tag{2.11}$$

Since the Higgs field is postulated to be a scalar field with an associated potential, the corresponding Lagrangian takes the form:

$$\mathcal{L}_{Higgs} = |D_\mu \phi|^2 - V(\phi) \tag{2.12}$$

where D_μ is the covariant derivative as previously defined for the electroweak sector in Equation (2.4). When the vacuum perturbed ϕ from Equation (2.11) is inserted into $\mathcal{L}_{Higgs}$ from Equation (2.12), the electroweak interaction from D_μ with ϕ will generate mass terms for the W_μ^a and B_μ bosons. The expansion of the potential will also reveal a mass term for the Higgs field, corresponding to the Higgs boson. In this way, the electroweak bosons and the Higgs have been given mass, whilst still retaining local gauge invariance.

Whilst the electroweak bosons have now been given mass, these postulated fields do not correspond to what is observed in nature of two massive, charged, spin-1 $W^\pm$ bosons; a massive, electrically neutral, spin-1 Z boson and the massless, electrically neutral, spin-1 photon, γ. The physical states can, however, be recovered with the correct interpretation of the W_μ^a and B_μ states. If linear superpositions of the W_μ^1 and W_μ^2 bosons are considered, then the behaviour of the W^+ and W^- bosons is seen:

$$W^\pm = \frac{1}{\sqrt{2}}\,(W_\mu^1 \mp i\,W_\mu^2) \tag{2.13}$$

By considering a rotation of θ_W, the physical photon and Z^0 bosons are interpreted as a mixing of the W_μ^3 and B_μ states:

$$\begin{pmatrix} Z^0 \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & -\sin\theta_W \\ \sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} \tag{2.14}$$

where θ_W is the weak mixing angle or Weinberg angle. By considering the effect of this rotation in the Lagrangian, at leading order in perturbation theory the masses of the bosons are predicted by the Standard Model to be:

$$M_H = \lambda\, v \tag{2.15}$$

$$M_W = \frac{g\,v}{2} = \frac{ev}{2\,\sin\theta_W} \tag{2.16}$$

$$M_Z = \frac{v}{2}\sqrt{g^2 + g'^2} = \frac{M_W}{\cos\theta_W} \tag{2.17}$$

$$M_\gamma = 0 \tag{2.18}$$

Since the interaction of the Higgs field with the EW bosons is contained within the covariant derivative, a separate term is necessary to generate the fermion masses through the Higgs mechanism. This interaction is described by the Yukawa Lagrangian, $\mathcal{L}_{Yukawa}$ [91]:

$$\mathcal{L}_{Yukawa} = -\sum_f Y_f(\overline{\psi}_{f,L}\,\phi\,\psi_{f,R} + \overline{\psi}_{f,R}\,\phi^\dagger\,\psi_L) \tag{2.19}$$

where f is any quark or charged lepton, ψ_L corresponds to the left handed projection of ψ and ψ_R the corresponding right handed component. Y_f is the so-called Yukawa coupling which determines the strength of the fermion field's interaction with the Higgs field. In general, the larger the size of the Yukawa coupling, the more massive the fermion will be. Once the Higgs field is expanded around v, the mass terms in the Lagrangian for the fermions will read:

$$m_f = Y_f \frac{v}{\sqrt{2}} \tag{2.20}$$

2.3.3 Quantum Chromodynamics

Quantum chromodynamics (QCD) is the mathematical theory underpinning the interactions of the strong force, which acts on particles that carry colour charge. In the Standard Model, QCD is represented by the SU(3) group, whose 8, 3×3 matrix generators correspond to the 8 Gell-Mann matrices, λ^a. The 8 Gell-Mann matrices also correspond to the 8 gauge bosons that are predicted to exist by QCD. The gauge boson of QCD is the gluon, a massless, electrically neutral particle, whose 8 different flavours correspond to different superpositions of the red, green and blue colour charges (along with their anticolour counterparts) that a QCD particle may carry. The QCD Lagrangian can be written as:

$$\mathcal{L}_{QCD} = \sum_q \overline{\psi}_q i\gamma^\mu D_\mu \psi_q - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a \tag{2.21}$$

where the sum q runs over the quark flavours and their corresponding fields ψ. Similar to $\mathcal{L}_{EW}$ the covariant derivative, D_μ is defined to be:

$$D_\mu = \partial_\mu - ig_s \frac{\lambda^a}{2} A^a_\mu \tag{2.22}$$

Finally, the gluon field strength tensor is denoted by $G^a_{\mu\nu}$ and defined to be:

$$G^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g_s f^{abc} A^b_\mu A^c_\nu \tag{2.23}$$

In both the covariant derivative and the field strength tensor, g_s is the strong gauge coupling parameter and A^a_μ the gluon field corresponding to colour state a. In the field strength tensor, f^{abc} are the structure constants of SU(3). The structure constants of $SU(3)$ are again determined by a commutator relation $[\lambda^a, \lambda^b] = 2i\, f^{abc}\, \lambda^c$.

When performing calculations with QCD, g_s is typically replaced by the related quantity α_S, the strong coupling constant, where α_S is simply a rescaling of g_s:

$$\alpha_S \equiv \frac{g_s}{4\pi} \tag{2.24}$$

The size of α_S is correspondingly larger than the coupling constants of QED (α_{QED}) and the weak force (α_W) to reflect that the strong force is the strongest of these fundamental forces. This relatively large magnitude of α_S leads to the phenomenon of confinement, where particles with colour charge may only exist in a colourless bound state. Since confinement is a property that is unique to QCD it is not possible for single quarks to be observed unlike isolated leptons. Although processes exist within the Standard Model that permit the production of individual quarks, they will quickly hadronise such that only the remnants of this process will be seen in a detector.

Although α_S is referred to as a constant its value actually varies depending on the energy scale of the process being investigated. At very high energy scales the decrease in magnitude of α_S allows for individual quarks to exist as free particles in a phase of matter known as the quark-gluon plasma. Inversely, at low energy scales the magnitude of α_S instead increases. This variance of α_S creates two different regimes of QCD; non-perturbative and perturbative QCD. At high energies the magnitude of α_S is small enough that perturbation theory can be used to make predictions about QCD as is also done for QED and EW interactions. However, since the magnitude of α_S is larger than α_{EW} and α_{QED} the next-to-leading-order (NLO) and next-to-next-to-leading-order (NNLO) corrections must often also be considered in order to return a theoretical prediction that resembles the measured data. At low energies α_S becomes so large that it is no longer possible to make predictions about QCD using perturbation theory. Instead, predictions must be made using the QCD Lagrangian directly as in lattice QCD [89].

Partons and Parton Distribution Functions

Unlike the particles of the Standard Model, protons are not elementary particles. Instead, the proton is, on average, a bound state of two up quarks and one down quark, which are known as valence quarks. The valence quarks interact strongly via the exchange of gluons to overcome the electrical repulsion that would otherwise cause the proton to be unstable. Through loop diagrams, it is also possible for exchanged gluons to become virtual quark-antiquark pairs which are known as sea quarks. The particles in the valence and sea quarks are individually known as partons. Each parton carries a

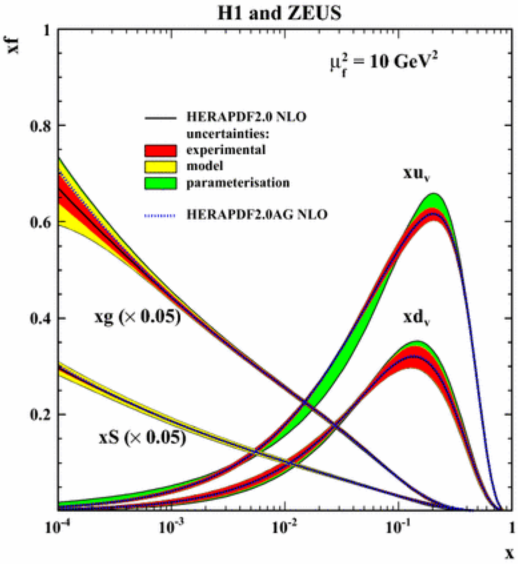

Figure 2.3: Combined H1 and ZEUS measurements of the proton PDF at an energy scale of $Q^2 = 10$ GeV^2 at the HERA collider [63]

fraction of the total proton momentum and the probability of finding an individual parton type a with momentum fraction x, at energy scale Q^2, is given through the parton distribution functions (PDF), $f_a(x, Q^2)$.

Due to the non-perturbative nature of QCD at the proton energy scale (1 GeV) the PDFs cannot be calculated from first principles. Instead, it is possible to predict the change of the PDFs at different energy scales through a process known as Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) evolution [53, 62, 19]. For DGLAP evolution to work, the PDFs themselves, must first be measured. This was the overarching goal of the HERA accelerator, through global fits to deep inelastic scattering events [63]. Through DGLAP evolution, the PDFs measured at HERA have been extrapolated to the energy scale of modern colliders, where they have been found to be in agreement with further measurements of the PDFs themselves.

Figure 2.3 shows that at the energy scale of $Q^2 = 10$ GeV^2, the largest momentum fractions are much more likely to be carried by the valence quarks than by sea quarks or gluons. The positions of these peaks change, however, with changing energy scales. At higher energy scales, it becomes more probable for heavier sea quarks to be resolved within the proton, and thus reducing the momentum fraction available to the valence quarks. Although not displayed in Figure 2.3., it is also possible to include the photon into the PDFs as even though quark interactions are dominated by the strong force, their electric charge still allows for QED interactions to occur.

2.3.4 Problems Unanswered by the Standard Model

The Standard Model is one of the most successful theories of physics that has ever been written down, having passed every experimental test that it has been posed with. Despite this success, it falls short in several areas and does not provide answers to many fundamental open questions within particle physics and cosmology.

The most glaring omission of the Standard Model is that of gravity. With experimental confirmation that all other forces of nature are communicated through the exchange of gauge bosons, it would appear logical that gravitational interactions would also behave in a similar manner. However, since the effects of gravity become easily visible only at cosmic scales, any gravitational interaction between two particles would be many orders of magnitude weaker than the weak force. Therefore, whilst theories exist that describe a graviton particle, lacking any experimental evidence means the particle nature of gravity remains purely hypothetical.

Within the Standard Model itself there are still unresolved problems; notably in the neutrino sector. According to the Standard Model, the neutrino is a massless particle, even though the observation of neutrino oscillations implies the existence of neutrino masses. If the neutrino is a Dirac particle (with a distinct antiparticle), it will gain its mass through the Higgs mechanism and also imply the existence of right-handed neutrinos that do not exist within the Standard Model. However, if the neutrino is its own antiparticle, i.e a Majorana particle, it will gain its mass through some other mechanism. One of the main experimental efforts in order to clarify this aspect of the nature of the neutrino is the search for the observation of neutrinoless double beta decay. In this scenario, the neutrino is a Majorana particle which acts as a propagator for the interaction $W^-W^- \to e^-e^-$. This process has, however, still not been experimentally observed, with the best estimate of the half-life for this process being more than 10^{26} years [54].

The particles of the Standard Model and their interactions also fail to answer several important cosmological queries. The rotational velocity of galaxies is larger than what would be expected given the amount of matter contained within them and assuming the validity of Einstein's Theory of General Relativity. This implies the presence of a type of matter which does not interact electromagnetically and is therefore not directly detectable on Earth. Since this matter cannot be detected through telescopes it is referred to as 'dark matter'. It is hypothesised that dark matter is either a single particle type or a collection of particles known as the 'dark sector' with only limited interactions with the Standard Model.

Galaxies not only rotate faster than expected but the expansion of the space between them is accelerating, rather than decelerating as would be expected from a purely gravitational interaction. The energy source causing this acceleration is unknown but is referred to as dark energy.

On a much more fundamental level, the observable universe appears to be almost entirely made of matter, with only very small amounts of naturally occurring antimatter. The Standard Model treats the two types of particle almost equally but does display a small preference for producing matter over antimatter. However, the rate of production of matter over antimatter within the Standard Model is not large enough to explain the extremely large imbalance that is observed.

Solutions to these problems have been proposed by theories that go beyond the Standard Model (BSM). However, all BSM theories currently share the single common factor of a lack of experimental evidence, without which these problems must remain unanswered. It is therefore of great importance that the parameters of the Standard Model and the cross sections of its processes are known as precisely as possible, in order to more easily identify the possible effects of new physics. One such example would be to measure the cross section of the Drell-Yan process.

2.4 The Drell-Yan Process

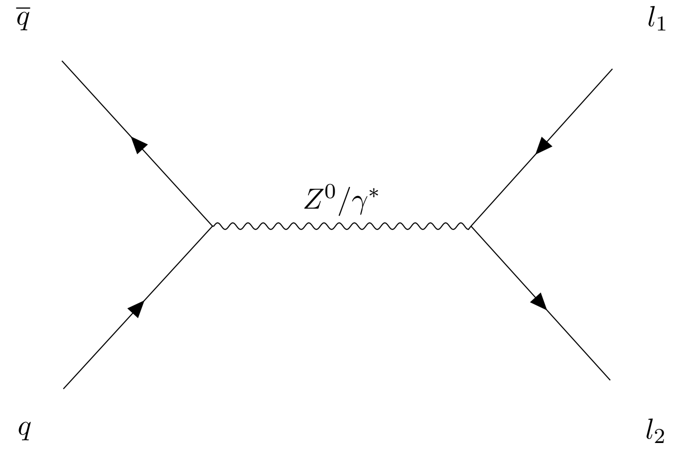

Figure 2.4: Feynman diagram of leading order Drell-Yan production at hadron colliders.

At hadron colliders the Drell-Yan process describes the production of a lepton pair from the annihilation of a quark/antiquark pair which is mediated by either a Z boson or virtual photon. Whilst the leading order Feynman diagram for this process that is shown in Figure 2.4 contains no QCD vertices, it is possible for the incoming quark and antiquark to radiate gluons which introduce QCD corrections to the underlying electroweak process. Similarly, the outgoing lepton pair can radiate photons which introduce further electroweak corrections although these corrections will be smaller than the QCD corrections due to the greater magnitude of α_S in comparison to α_{EW}. There can also be hadronic radiation from the proton remnants after the $q\overline{q}$ annihilation has occurred which must be considered if one wishes to calculate the cross section for this process.

Following the process of Ebert et al. [55], in the Collins-Soper reference frame (which will be defined in Section 2.4.1) the total differential cross section for the Drell-Yan process can be described through an angular decomposition of the form:

$$
\begin{aligned}
\frac{d\sigma}{dy^Z\, dp_T^Z\, dm_Z\, d\phi\, d\cos\theta} &= \frac{3}{16\pi}\frac{d\sigma^{U+L}}{dy^Z\, dp_T^Z\, dm_Z}\{(1+\cos^2\theta) + \frac{1}{2}A_0(1-3\cos^2\theta) + \sum_{i=1}^{7} A_i\, g_i(\theta,\phi)\} \\
&= \frac{3}{16\pi}\frac{d\sigma^{U+L}}{dy^Z\, dp_T^Z\, dm_Z}\{(1+\cos^2\theta) + \frac{1}{2}A_0(1-3\cos^2\theta) + A_1(\sin 2\theta\cos\phi) \\
&\quad + \frac{1}{2}A_2\sin^2\theta\cos 2\phi \ + A_3\sin\theta\cos\phi + A_4\cos\theta + A_5\sin^2\theta\sin 2\phi \\
&\quad + A_6\sin 2\theta\sin\phi + A_7\sin\theta\sin\phi\}
\end{aligned}
\tag{2.25}
$$

where the differential cross section, σ, is five-fold decomposed in terms of the Z rapidity, y^Z, its transverse momentum, p_T^Z, its mass, m_Z, and the Collins-Soper angles, ϕ and $\cos\theta$. The differential cross section is then parameterised as an unpolarised cross section, σ^{U+L}, angular coefficients that encode the boson kinematics, A_i, and spherical harmonics that depend on the decay angles of the produced leptons. Therefore, in this formalism the boson kinematics are completely separated from the angular components of the decay, as well as being valid to all orders of QCD.

From a theoretical standpoint, Equation (2.25) takes into account both the Z and virtual photon mediators as well as their interference. Experimentally, the contribution to the cross section from the photon and any interference can be minimised by selecting leptonic events where the invariant mass of the lepton pair is situated in a region around the Z mass. Since the Z is typically on shell in this region the contribution from pure Z production will dominate the measured cross section. The cross section decomposition would also break down if any QED interaction occurs between the incoming quarks and outgoing leptons. In addition, the decomposition would also break down for any loop diagrams involving the mediating boson.

2.4.1 Collins-Soper Frame

Before justifying Equation (2.25), it is first necessary to define the reference frame in which it holds, namely the Collins-Soper (CS) frame [47]. The CS frame is defined in the rest frame of the Z boson, meaning the decay leptons will be produced back to back. In the CS frame, the plane containing the two incoming partons is considered to be the x-z plane, with the z-axis pointing in the lab frame direction of the Z boson. The y-axis is the normal vector to the parton plane and points in the direction of the negatively charged lepton. Finally, the x-axis is chosen such that the x,y and z vectors form a right-handed coordinate system. The angle θ_{CS} is then defined to be the angle between the z-axis and the negative lepton, whilst ϕ_{CS} is the angle between x-axis and

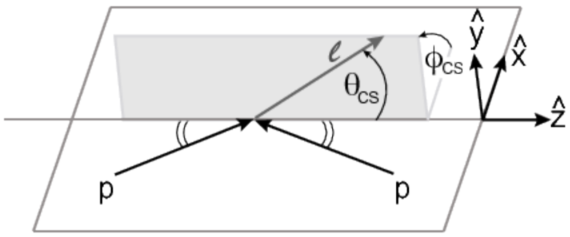

Figure 2.5: A sketch of the Collins-Soper frame, defining its Cartesian components as well as the angles θ_{CS} and ϕ_{CS} [6].

the negative lepton. A sketch outlining the directions of the unit vectors as well as θ_{CS} and ϕ_{CS} can be seen in Figure 2.5.

The cosine of θ_{CS} can be directly calculated through lab-frame quantities as

$$\cos\theta_{CS} = \frac{p_z^{ll}}{|p_z^{ll}|} \frac{(E_{l^-} + p_{z,l^-})(E_{l^+} - p_{z,l^+}) - (E_{l^-} - p_{z,l^-})(E_{l^+} + p_{z,l^+})}{m_Z\sqrt{m_Z^2 + p_{T,Z}^2}} \tag{2.26}$$

where p_z^{ll} is the dilepton longitudinal momentum, i.e the z momentum of the Z boson, $E_{l-(+)}$ is the total energy of the lepton (antilepton) and $p_{z,l}$ the corresponding longitudinal component of the momentum. Similarly ϕ_{CS} can be determined by:

$$\sin\phi_{CS} = \frac{2}{\sin\theta_{CS}} \frac{p_{y,l^-}p_{x,l^+} - p_{y,l^+}p_{x,l^-}}{p_{T,Z}m_Z} \tag{2.27}$$

where $p_{x,y}$ are the x and y components of the decay leptons three-momentum.

2.4.2 Angular Decomposition of the Cross Section

With the appropriate rest frame introduced, in order to derive Equation (2.25), it is first necessary to define the the matrix element corresponding to the Feynman diagram shown in Figure 2.4. When considering the Drell-Yan process and any further interactions of the proton remnants, this can be written as:

$$\mathcal{M}(pp \to V + X \to L + X) = \mathcal{M}^{\mu}_{V\to L}\,\langle X|J_{V_\mu}|pp\rangle \tag{2.28}$$

where pp is the two colliding protons, V is the vector boson mediating the interaction, X the hadronic radiation from the proton remnant, L the lepton pair that V decays to,

$\mathcal{M}^{\mu}_{V\to L}$ the matrix element corresponding to the Born level interaction shown in Figure 2.4, J_{V_μ} the electroweak current between the initial pp state and the hadronic X state. For the Z boson, the electroweak current is:

$$J^{\mu}_{Z} = -|e| \sum_{f} \overline{\psi}_f \gamma^{\mu} (v_f - a_f \gamma_5) \psi_f \tag{2.29}$$

where e is the elementary electric charge, the sum runs over the quark flavours and Q_f is the electric charge of quark f in units of $|e|$. The corresponding vector, v_f and axial, a_f couplings of the Z to each quark type is then given by:

$$v_f = \frac{T^f_3 - 2Q_f \sin^2\theta_W}{\sin 2\theta_W}, \quad a_f = \frac{T^f_3}{\sin 2\theta_W} \tag{2.30}$$

With all the components of Equation (2.28) defined, by performing the usual process of squaring the matrix element and integrating over phase space, one arrives at an expression for the differential cross section:

$$\frac{d\sigma}{d^3q\, d\mathcal{O}} = \frac{1}{E^2_{cm}} L_{\mu\nu}(q, \mathcal{O}) H^{\mu\nu}(q, P_a, P_b) \tag{2.31}$$

where $\mathcal{O}$ is any differential observable of the lepton pair, q is the total momentum of the vector boson (or equivalently the momentum of the measured lepton pair), E_{cm} the total centre of mass energy of the colliding pp pair, $P_{a,b}$ the momentum of incoming protons a and b. $L_{\mu\nu}$ is the so called leptonic tensor which encodes the information on the kinematics and decay products of the vector boson, whilst the hadronic tensor, $H^{\mu\nu}$ describes its production. The specific differentials in d^3q are the boson specific variables and in $d\mathcal{O}$ the CS angular variables, namely:

$$d^3q = \frac{1}{2} dm^Z dy^Z dp^Z_T \qquad d\mathcal{O} = d\cos\theta_{CS}\, d\phi_{CS} \tag{2.32}$$

where m is the mass of the boson, Y the rapidity and p^Z_T the momentum of the Z boson in the plane transverse to the z direction $(x - y)$ when measured using a right handed coordinate system.

By projecting the hadronic tensor onto a helicity matrix:

$$H_{\lambda\lambda'} = \varepsilon^*_{\mu}(\lambda)\, H^{\mu\nu} \varepsilon(\lambda') \tag{2.33}$$

where the $\lambda, \lambda' = \{+, -, 0\}$ encode the single longitudinal and two transverse polarisation modes open to the boson. The specific polarisation vectors, ε_μ, can then be written as:

$$\varepsilon_\mu(\pm) = \frac{1}{\sqrt{2}}(0, \pm 1, -i, 0) \tag{2.34}$$

$$\varepsilon_\mu(0) = (0, 0, 0, 1) \tag{2.35}$$

Using this parameterisation, upon contraction with the lepton tensor, the resulting differential cross section can be written as:

$$\frac{d\sigma}{d^3q\, d\mathcal{O}} = \sum_{\alpha \in \mathcal{P}} g_\alpha(\theta, \phi) \frac{3}{16\pi} \frac{d\sigma^\alpha}{d^3q} \tag{2.36}$$

where $\mathcal{P} = \{U + L, L, T, I, P, A, 7, 8, 9\}$ corresponds to the set of possible polarisations of the Z and g_α the spherical harmonics. The so-called helicity cross sections, $d\sigma^\alpha$, are related to components of the helicity projected hadronic tensor and describe one of the polarisation states [6]:

$$\begin{aligned}
\sigma^{U+L} &\propto H_{00} + H_{--} + H_{++} \\
\sigma^{L} &\propto H_{00} \\
\sigma^{T} &\propto \frac{1}{2}(H_{+-} + H_{-+}) \\
\sigma^{I} &\propto \frac{1}{4}(H_{+0} + H_{0+} - H_{-0} - H_{0-}) \\
\sigma^{P} &\propto H_{++} - H_{--} \\
\sigma^{A} &\propto \frac{1}{4}(H_{+0} + H_{0+} + H_{-0} + H_{0-}) \\
\sigma^{7} &\propto \frac{-i}{2}(H_{+-} - H_{-+}) \\
\sigma^{8} &\propto \frac{-i}{4}(H_{+0} - H_{0+} + H_{-0} - H_{0-}) \\
\sigma^{9} &\propto \frac{-i}{4}(H_{+0} - H_{0+} - H_{-0} + H_{0-})
\end{aligned} \tag{2.37}$$

By factorising out the unpolarised cross section, $d\sigma^{U+L}$, from each term in Equation (2.36), a set of 8 angular coefficients, A_i, can be defined as unitless, ratios of helicity cross sections:

$$
\begin{aligned}
A_0 &= 2\frac{d\sigma^L}{d\sigma^{U+L}} \\
A_1 &= 2\sqrt{2}\frac{d\sigma^I}{d\sigma^{U+L}} \\
A_2 &= 4\frac{d\sigma^T}{d\sigma^{U+L}} \\
A_3 &= 4\sqrt{2}\frac{d\sigma^A}{d\sigma^{U+L}} \\
A_4 &= 2\frac{d\sigma^P}{d\sigma^{U+L}} \\
A_5 &= 2\frac{d\sigma^7}{d\sigma^{U+L}} \\
A_6 &= 2\sqrt{2}\frac{d\sigma^8}{d\sigma^{U+L}} \\
A_7 &= 4\sqrt{2}\frac{d\sigma^9}{d\sigma^{U+L}}
\end{aligned}
\tag{2.38}
$$

By defining these angular coefficients, the Drell-Yan cross section can then be expanded fully as written in Equation (2.25).

Writing the cross section in this form has the advantage that the boson production is completely separated from its eventual decay. A precise measurement of the angular coefficients and the unpolarised cross section (which describe the boson kinematics), therefore provides a method of probing the boson production that is valid to all orders in QCD. On the other hand, the full phase space kinematics and angular properties of the lepton/antilepton pair are encoded through the spherical harmonics in the Collins-Soper frame.

2.4.3 Forwards-Backwards Asymmetry

As was shown by the work of Wu in 1956 [90], the weak force violates the conservation of parity. Whilst β decay is mediated by the W boson, parity violation is not limited to these interactions. Being an electroweak boson, the Z is also capable of violating parity conservation. When considering the production of Z bosons at particle colliders this parity violation manifests as a preferred direction for Z decay. This asymmetry is typically called the forwards-backwards asymmetry, A_{FB}, and is defined as:

$$
A_{FB} = \frac{N_F - N_B}{N_F + N_B} \tag{2.39}
$$

where N_F is the number of forward Z decays ($\cos\theta_{CS} > 0$) and N_B the number of Z decays in the backwards direction ($\cos\theta_{CS} < 0$). Equivalently, A_{FB} can also be defined through the contributions to the cross section from the forward and backwards regions, σ_F and σ_B respectively.

$$A_{FB} = \frac{\sigma_F - \sigma_B}{\sigma_F + \sigma_B} \tag{2.40}$$

Equation (2.40) therefore explicitly links the Drell-Yan cross section described by Equation (2.25) to the parity violation of the weak force. Since parity violation is a purely weak phenomenon, if the cross section is written at Born level, then A_{FB} can also be directly linked to the angular coefficient A_4. Without QCD interactions, Equation (2.25) becomes:

$$\frac{d\sigma}{d\cos\theta} = \frac{3}{16\pi}\frac{d\sigma^{U+L}}{dy^Z dm^Z}(1 + \cos^2\theta + A_4\cos\theta) \tag{2.41}$$

where the dependence on the rapidity, transverse momentum and mass of the Z is implicitly contained within A_4. By integrating Equation (2.41) with respect to $\cos\theta$, one arrives at an expression for $\sigma_{F,B}$ by using the appropriate limits of integration. When this is inserted into Equation (2.41), A_{FB} becomes:

$$A_{FB} = \frac{\int_0^1 \frac{d\sigma}{d\cos\theta} d\cos\theta - \int_{-1}^0 \frac{d\sigma}{d\cos\theta} d\cos\theta}{\int_{-1}^1 \frac{d\sigma}{d\cos\theta} d\cos\theta} = \frac{3}{8}A_4 \tag{2.42}$$

Although Equation (2.42) is derived without considering the higher order QCD corrections, if the full cross section of Equation (2.25) is considered then the further parity violating terms of A_3 and A_7 become relevant. However, A_3 becomes non-zero only at high values of p_T^Z and A_7 is a next to next to leading order (NNLO) QCD correction to the cross section so its contribution to the cross section and to A_{FB} is minimal. This means that even though Equation (2.42) is only a leading order approximation, it is still a valid way of calculating the forwards-backwards asymmetry.

Whilst Equation (2.42) describes the magnitude of the forwards-backwards asymmetry observed in the Drell-Yan process, it does not describe the physical origin of the asymmetry. Parity violation in the weak force arises due to the vector minus axial (V - A) couplings to particles when calculating matrix elements. Since A_{FB} is an example of parity violation, it is also caused by the (V - A) couplings of the weak force. In terms of the axial and vector couplings, A_{FB} can be explicitly written as [83]:

$$A_{FB} = \frac{6a_l a_q(8v_l v_q - Q_q K D_m)}{16(v_l^2 + a_l^2)(v_q^2 + a_q^2) - 8v_l v_q Q_q K D_m + Q_q^2 K^2 (D_m^2 + \Gamma_Z^2/m_{Z,ref}^2)} \tag{2.43}$$

where $a_{l,q}$ ($v_{l,q}$) are the axial (vector) couplings to the leptons and quarks as previously defined in Equation (2.30), Γ_Z is the decay width of the Z boson and $m_{Z,ref} = 91.1876$

GeV is the mass of the Z as given by the particle data group [72], Q_q is the charge of the quark in the $q\bar{q}$ pair producing the Z. K and D_m are defined to be:

$$K = \frac{8\sqrt{2}\pi\alpha}{G_F m_{Z,ref}^2} \qquad D_m = 1 - \frac{m_{Z,ref}^2}{m_{ll}^2} \tag{2.44}$$

where α is the fine structure constant, G_F the Fermi constant and m_{ll} the invariant mass of the decay lepton pair. The presence of the axial and vector couplings in Equation (2.43), mean that A_{FB} and A_4 are proportional to the weak mixing angle.

2.4.4 The Weak Mixing Angle

As introduced in Section 2.3.2, the weak mixing angle relates the mixing of the W^μ and B^μ fields of the electroweak theory with the physically observed $W^\pm$ and Z bosons. The weak mixing angle also relates the masses of the W and Z bosons after spontaneous symmetry breaking through Equation (2.17). Therefore, precision measurements of the weak mixing angle help to constrain the masses of the W and Z within the Standard Model and easier identify the effects of any new physics which may interfere with the masses of the Standard Model particles.

However, as currently defined, the weak mixing angle is valid only at leading order in electroweak theory. To include the effects of higher order corrections, it is convenient to define an effective weak mixing angle, θ_{eff}^l. This is related to the standard weak mixing angle by [40]:

$$\sin^2 \theta_{eff}^l = \kappa_l \sin^2 \theta_W \tag{2.45}$$

where the coefficient κ_l encodes the higher order corrections to the weak mixing angle. When expressed as a function of the vector and axial couplings, the effective weak mixing angle becomes [40]

$$\sin^2 \theta_{eff}^l = \frac{1}{4}\left(1 - \frac{v_l}{a_l}\right) \tag{2.46}$$

where the vector and axial couplings are in this case to the leptons.

This therefore means that by measuring the Drell-Yan cross section, not only can the QCD dynamics of the Z boson production be probed but also fundamental parameters of the Standard Model as well. This not only improves our knowledge of how the Standard Model works but also allows for greater constraints to be placed on the effect of new physics beyond the Standard Model when such searches are carried out.

Chapter 3

Experimental Setup

In order to measure the Drell-Yan process a dataset of $q\bar{q} \rightarrow Z \rightarrow l^{+}l^{-}$ must first be gathered. This chapter will describe how such a dataset was constructed between 2015 and 2018 using proton-proton collisions at the Large Hadron Collider and also explain the basic phenomenology of these collisions. In order to measure and record the $Z \rightarrow l^{+}l^{-}$ decays from the proton-proton collisions, the ATLAS detector and its various subsystems will also be described.

3.1 The Large Hadron Collider

Whilst the discovery of the first elementary particles of the Standard Model did not require the use of a collider, repeatable experiments at a specific energy scale necessitate the construction of particle accelerators. The study of particle physics has relied on many such machines colliding many different particle types in the past century with one such example being the Large Hadron Collider (LHC). Lying approximately 100m below the countryside surrounding Geneva, at 27 km in circumference and reaching a centre of mass energy $\sqrt{s} = 13$ TeV during Run 2 operations, the LHC is the largest and highest energy particle collider ever constructed. No other machine is capable of probing the regions of phase space at the energy frontier that can be accessed at the LHC.

3.1.1 The CERN Accelerator Complex

Whilst the LHC has collided both protons and heavy ions during its data taking periods, since the Drell-Yan process does not involve heavy ions only the proton-proton collisions at the LHC will be described. However, before protons can be collided in the LHC, they first have to be brought to the machine. This task is fulfilled by the CERN Accelerator Complex, whose 2015 - 2018 configuration can be seen in Figure. 3.1. At CERN, protons begin their life in a bottle of hydrogen gas, which is electrolysed to separate the protons from the electrons. The protons then enter the accelerator chain at LINAC 2 where they are accelerated from rest to 50 MeV. From here they enter the Proton Synchrotron Booster to reach 1.4 GeV. The Proton Synchrotron then boosts the energy to 26 GeV before injecting into the final pre-accelerator, the Super Proton Synchrotron. In the SPS the protons will reach 450 GeV before they are finally injected into the LHC.

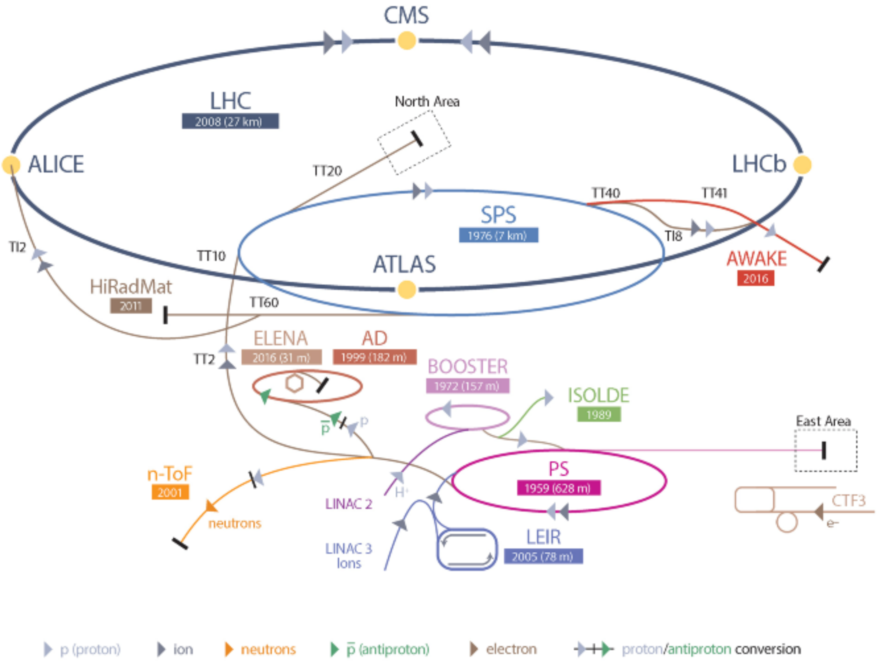

Figure 3.1: Schematic of the CERN accelerator complex during Run 2 of the LHC. [64]

The design of the LHC relies on tried and tested technology that has been developed in the past century, namely the usage of radio frequency (RF) cavities as well as superconducting magnets. The acceleration is provided by 16 superconducting RF cavities, located at multiple points along the LHC ring [29]. These RF cavities take protons from an initial energy of 450 GeV to the collision energy of 6.5 TeV. The particles must, however, be kept in a circular orbit within the LHC, a task that is performed by each of the 1,232 8T superconducting dipole magnets along the ring. As protons are electrically charged, the bunches must also be focussed whilst they travel through the ring, to prevent them from disintegrating. For this, a system of focussing and defocussing quadropole magnets is employed, which results in a net focussing effect and allows for collisions to occur at the interaction points.

At the four interaction points of the LHC a particle detector can be found. These four detectors correspond to the four main experiments at the LHC: ATLAS, CMS, ALICE and LHCb. ATLAS and CMS are general purpose particle detectors, whereas ALICE

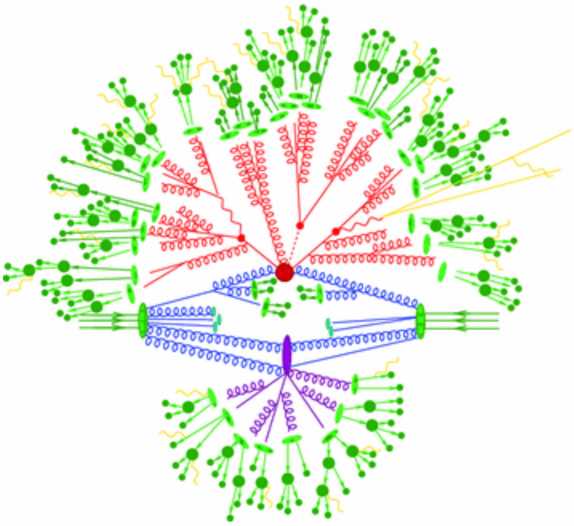

Figure 3.2: A cartoon schematic of a proton-proton collision [61]. The incoming partons are denoted by the green lines with arrows. Two of these partons interact in the hard process depicted as the dark red dot. From the hard process, a parton shower results in red. Further partons may interact in a multiple parton interaction denoted in purple, which itself produces a further parton shower. The end result of all showers is collection of stable particles which are shown in green. At any point, a particle with electric charge may also emit a photon, as show in yellow.

and LHCb are more specialised. The former focusses on the study of the quark-gluon plasma state of matter and the latter on the study of the decays of b-hadrons.

3.1.2 Phenomenology of Proton-Proton Collisions

As was outlined in Section 2.3.3 the proton is a composite particle made up of partons. With multiple partons within a single proton, the collision of two protons can quickly become a complicated process within a detector. In order to understand the physics produced by the collision it is therefore crucial to understand the collision process itself.

Collision Process

When colliding protons at the LHC, there are two possible types of interaction that can occur, within the constraints of the Standard Model:

1. parton-parton interactions that result in the exchange of gluons, photons, $W^{\pm}$, Z or Higgs bosons.
2. peripheral interactions from protons passing close by to each other but not colliding, resulting in an electroweak interaction.

Whilst both types of collision have been detected at the LHC , the dominant interaction type is the parton-parton interaction given the almost exactly head on nature of collisions at the LHC and much smaller photon PDF in comparison to the quark and gluon PDFs. Whilst the collision of two protons may sound trivial, the resultant process is itself complex, as shown in Figure 3.2.

As the incoming partons are almost exclusively QCD objects (quarks and gluons), there is a non-zero probability for them to radiate gluons before the actual interaction. Since quarks also carry electric charge, it is possible for them to radiate photons. This process is known as initial state radiation (ISR).

After any ISR has occurred, the actual interaction between the incoming partons occurs. This is known as the hard process and is what is most frequently studied at particle colliders. In the hard process the incoming partons may either annihilate or scatter from each other. During these processes at least one gauge boson may be produced. The decay products of a gauge boson can be used to reconstruct the physics of the hard process since the initial interaction will always be invisible to any detector. The decay products will themselves likely carry colour charge, electric charge or both. Similar to ISR, any radiation emitted from the decay products is known as final state radiation (FSR).

If the decay products are quarks, colour confinement will force them to hadronise into colour-neutral objects. From the initial quark, a shower of mesons and other hadronic particles is typically observed. These hadronic objects are collectively known as a jet. The size of the jet and the distance of its primary vertex from the collision point can be used to help identify the origin of the jet, either the flavour of quark or if it was originally a gluon.

The hard process, however, only accounts for two of the many partons inside the colliding protons. It is also possible for further partons to collide from the same two protons in what are known as multiple parton interactions. Such interactions and any further activity that is not attributable to the hard process is known as the underlying event. Activity which is recorded from further interactions that occur during the collision is known as pileup. Pileup and multiple parton interactions therefore make identifying the hard process in LHC collisions a considerable challenge.

Luminosity

Now that the process of how proton-proton collisions occur is understood, it is necessary to define a metric to evaluate the number of collisions happening at any given time. For this purpose, the instantaneous luminosity must first be introduced, $\mathcal{L}$. For a single pair of colliding bunches, this quantity is defined as:

$$\mathcal{L} = f\frac{n_1 n_2}{2\pi\sigma_x\sigma_y} \tag{3.1}$$

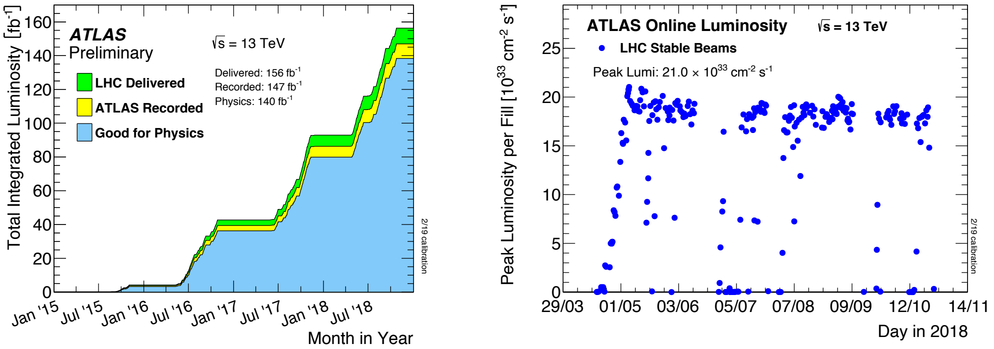

(a) ATLAS Run 2 Integrated Luminosity (b) ATLAS 2018 Instantaneous Luminosity

Figure 3.3: The integrated and maximum instantaneous luminosities for run 2 of the LHC and 2018 respectively, as measured by the ATLAS detector. Adapted from [46].

where f is the revolution frequency of the proton, n the number of particles in a given bunch and $\sigma_{x,y}$ the convolution of the beam width and height. The units of $\mathcal{L}$ are $\mathrm{m}^{-2}\ \mathrm{s}^{-1}$ but are more commonly quoted using the inverse barn, b^{-1}, where one barn corresponds to $10^{-28}\ \mathrm{m}^2$. The luminosity can also be simply integrated over time to give a measure of the number of collisions that happened during a data taking period.

$$L = \int \mathcal{L}\, dt \tag{3.2}$$

Similarly to the instantaneous luminosity, the integrated luminosity is typically quoted in inverse barns. For the ATLAS experiment in particular, approximately 60 seconds worth of integrated luminosity is referred to as a ‘lumiblock’.

Measuring the luminosity typically relies on some subsystem inside a particle detector that is sensitive to the total number of (visible) interactions happening at any given moment inside the detector. These subsystems are then calibrated using a method known as a van der Meer scan [85].

From 2015 to 2018, the LHC undertook its second long term data taking campaign (Run 2). During this period, the ATLAS detector managed to record a peak instantaneous luminosity of $21.0 \times 10^{33}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$. This was part of the 156 fb^{-1} of total integrated luminosity provided to ATLAS for physics analysis. However, since particle detectors are not 100% efficient and various subsystems must be initialised at the start of a data taking run, ATLAS was not able to use all of this luminosity. The amount of luminosity lost was, however, relatively low, with 147 fb^{-1} being recorded by ATLAS. Not all luminosity blocks are, however, of sufficiently high quality for physics analysis. During a run, detector components may not be performing optimally and human error in the operation of the detector mean that some data taking periods are ignored for physics analysis. Fortunately, these periods are relatively rare and 140 fb^{-1} of the data recorded by ATLAS is suitable for physics use [5]. These luminosity results can be seen in Figure 3.3a.

The luminosity is therefore a measure of how many total interactions occurred during a data taking period but this does not specify how many events were recorded for a given process, N_i. In order to calculate N, the luminosity must be weighted by the quantum mechanical probability of the interaction occurring for a given area. This weight is known as the cross section, σ. For a process i, N_i is thus given by:

$$N_i = \sigma_i L = \sigma_i \int \mathcal{L} \, dt \tag{3.3}$$

At the energy scales the LHC operates at, the cross section can be calculated using perturbation theory. For a general $pp \rightarrow X$ process, where two protons collide to form a new particle, X, the cross section can be factorised into a low energy component and a high energy component [48]. The low energy component describes the interactions of the partons within the proton itself, whereas the high energy component characterises the hard process. The cross section is then given by:

$$\sigma_{p_1 p_2 \rightarrow X} = \sum_{i,j \in \{q,g\}} \int dx_1 \; f_i(x_1, Q^2) \int dx_2 \; f_j(x_2, Q^2) \; \hat{\sigma}_{ij \rightarrow X}(\hat{s}, Q^2) \tag{3.4}$$

where i, j represent the quarks and gluons in the colliding protons, f is their corresponding PDF, x_k is the Björken x of parton k, the energy scale is Q^2, $\hat{\sigma}_{ij}$ is the cross section of the hard process forming particle X and $\hat{s} = x_1 x_2 s$ the squared centre of mass energy available to $\hat{\sigma}_{ij}$ from the total squared centre of mass energy s. $\hat{\sigma}_{ij}$ can itself be calculated by using the Feynman rules for QCD and inputting the appropriate value of α_S.

Equation 3.4 is applicable only to leading order QCD. The relatively large magnitude of α_S in comparison to the other coupling constants means that gluon radiation can have a large impact on the cross section for a given process. This has the consequence that higher order corrections to the cross section are often necessary in order for the theoretically calculated value to match the experimentally measured value at the LHC. These higher order corrections can take the form of extra α_S vertices in the Feynman diagram of the hard process or including higher order terms in the logarithmic expansions underpinning QFT. Although α_{EW} is smaller than α_S, the presence of photon ISR and FSR can also affect the cross section to such an extent that not only higher order corrections in QCD are necessary but also in EW corrections as well.

3.2 The ATLAS Detector

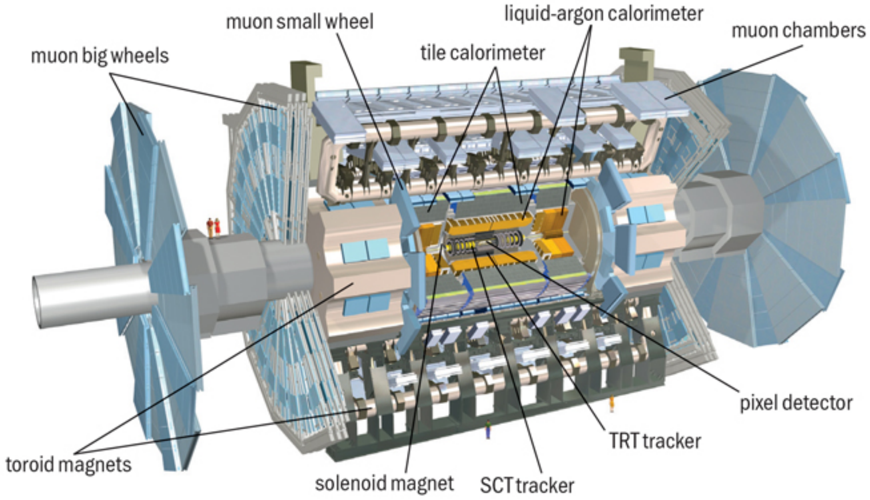

Figure 3.4: An overview image of the ATLAS detector, with many of the various subsystems labelled [3].

A Toroidal Lhc ApparatuS (ATLAS), is one of the two general purpose detectors along the LHC ring. ATLAS was designed to give almost 4π solid angle coverage of the proton-proton collisions provided by the LHC. As such, ATLAS is well positioned to measure the properties of the Standard Model, as well as search for any new physics beyond the Standard Model in all regions of phase space that can be probed at the LHC. At 25m high, 44m long and weighing some 7000 tonnes, ATLAS is the largest particle detector ever constructed [3].

The main purpose of ATLAS is to reconstruct the kinematic properties of the particles produced at the interaction point and their decay products. To achieve this, the use of multiple subsystems is necessary, each of which are typically segmented radially and azimuthally in different layers of the detector. The cylindrical structure of ATLAS can be further split into three segments; a barrel whose components are fit concentrically around the beam pipe and two endcaps which sit perpendicular to the beam pipe.

Expanding radially outwards from the interaction point, the first of these subsystems is the Inner Detector (ID). The ID provides high granularity position and momentum measurements of charged particles. After extracting the tracking information, the next subsystem in ATLAS is the calorimeter. This provides a further measurement of position as well as the energy of the incident particles. After the calorimeter, ATLAS is further equipped with a muon spectrometer which specialises in measuring the position and momentum of muons. In order to deflect charged particles, and therefore measure their momentum, ATLAS also has several barrel and endcap magnets located between the ID and the calorimeter as well as between the calorimeter and the muon spectrometer.

3.2.1 Coordinate System

To describe the positions of the produced particles and subsystems within ATLAS, a right-handed spherical coordinate system is used. The origin of this coordinate system is taken to be the centre of the detector, inside the beam pipe, where interactions nominally occur. The z-axis is taken to be along the beam pipe. The x-axis is oriented towards the centre of the LHC ring and the y-axis points towards the surface. The azimuthal angle, ϕ is measured in the x-y transverse plane and the polar angle, θ, is taken with respect to the z-axis.

At the LHC, protons are collided head on with equal beam energies. This means that in the lab frame, the total momentum of the proton-proton system is zero. However, since the proton momentum is carried by the individual partons, it is not possible to infer the total z-momentum, p_z, of the parton-parton system that interact during the hard process. Nevertheless, due to the head on nature of collisions at the LHC, in the transverse plane the initial transverse momentum, p_T, is always zero. The conservation of momentum can therefore be applied in the transverse plane to every event reconstructed within ATLAS. This is of particular importance when reconstructing particles that do not interact with the detector, whose presence can be inferred due to missing transverse energy, missing E_T or MET.

Positions of particles in the detector are typically quoted using the pseudorapidity, η and the azimuthal angle ϕ. The pseudorapidity is defined as:

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right) \tag{3.5}$$

where θ is the polar angle with respect to the z-axis. Conveniently, the difference in pseudorapidity between two any two particles is a Lorentz invariant quantity, which is of particular use when dealing with the relativistic particles of the LHC. The pseudorapidity can be shown to be the ultrarelativistic limit of the standard rapidity, y, defined as:

$$y = \frac{1}{2}\ln\left(\frac{E+p_z}{E-p_z}\right) \tag{3.6}$$

where E is the total energy of the particle.

3.2.2 Inner Detector

When charged particles are produced at the interaction point, the first task of ATLAS is to measure their position and momentum. To do this, a collection of barrel and endcap trackers are used; collectively referred to as the Inner Detector (ID). Specifically, the ID is composed of the Pixel Detector, Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). The pixel and SCT systems provide coverage in the region

$|\eta| < 2.5$, whilst the TRT covers $|\eta| < 2.0$. A schematic overview of the various components of the ID can be found in Figure 3.5.

The pixel detector and SCT both work on the principle that an incoming charged particle will leave an ionisation trail along its path through a semiconducting material. This results in the formation of electron-hole pairs in the p-n junction, which will then drift towards a readout module. The amount of time that charge is detected for can be recorded and if the recorded time is above a noise threshold value, a hit in the detector is recorded. The time above this threshold value will be proportional to the charge of the original incident particle, which can then be inferred via calibration algorithms.

In contrast, the TRT works on the principle that a charged particle will induce the emission of transition radiation when passing through an inhomogeneous material. Similar to a semiconductor based tracking system, the emitted ionisation electrons can be directed to an anode and the corresponding current read out. Selection criteria can then be applied to discriminate between signal and noise.

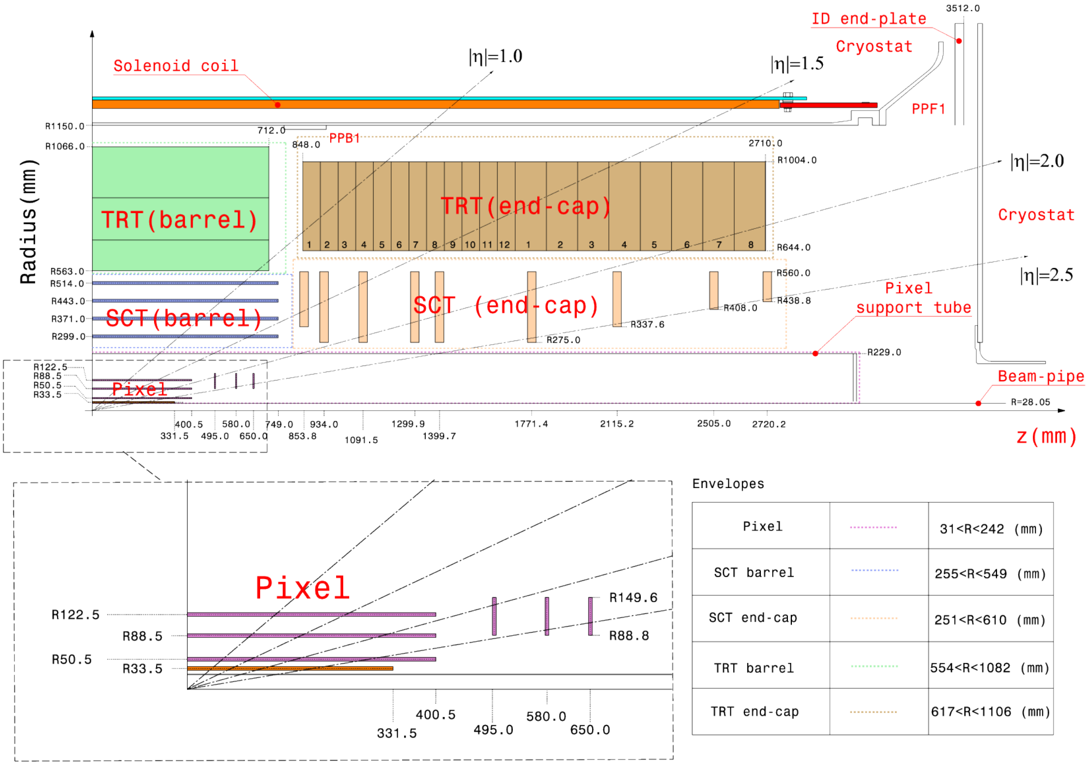

Envelopes

Pixel	┈┈┈	31<R<242 (mm)
SCT barrel	┈┈┈	255<R<549 (mm)
SCT end-cap	┈┈┈	251<R<610 (mm)
TRT barrel	┈┈┈	554<R<1082 (mm)
TRT end-cap	┈┈┈	617<R<1106 (mm)

Figure 3.5: A schematic overview of the ATLAS Inner Detector. The radial and z distances of each ID subsystem from the nominal interaction point are shown [38].

Detector	η Coverage	Unit Size	Intrinsic Resolution (μ^2)
IBL	$\lvert\eta\rvert < 3.0$	$50 \times 250\,\mu\mathrm{m}^2$	8×40
Pixel	$\lvert\eta\rvert < 2.5$	$50 \times 400\,\mu\mathrm{m}^2$	10×115
SCT	$\lvert\eta\rvert < 2.5$	80 μm	17
TRT	$\lvert\eta\rvert < 2.0$	4 mm	130

Figure 3.6: The tracking subsystems of the ATLAS inner detector.

Pixel Detector

At the very heart of ATLAS lies the pixel detector, whose first component is only 33.5 mm away from the interaction point. The pixel detector is formed of four cylindrical layers of silicon pixel modules in the barrel and a further three layers of circular disks in the endcap. These layers are segmented in ϕ and R (z) in the barrel (endcap), to provide localised position information about the incoming particle. Due to the proximity of the pixel detector to the interaction point, it has the corresponding highest granularity of all the ID modules. This granularity ranges from 50×250 $\mu\mathrm{m}^2$ in the first layer [34], to 50×400 $\mu\mathrm{m}^2$ [77] in the remaining three layers. In the barrel, the first layer of the pixel detector is known as the Insertable B-Layer (IBL) which was installed into ATLAS during Long Shutdown 1 for the start of Run 2 operations. The IBL was inserted to improve the reconstruction of secondary vertices associated with the production of b-quarks.

The high granularity and proximity of the pixel to the interaction point allows for a precise determination of the impact parameter of an incoming particle. This is crucial for the identification of displaced vertices within the tracking system which is itself important for the identification of jets originating from the hadronisation of b-quarks and tau leptons.

Semiconductor Tracker

Surrounding the pixel detector is the SCT, with four barrel layers and nine endcap disks at radii ranging from 299 mm to 560 mm from the interaction point. Unlike the pixel detector, the SCT is formed of strips of silicon rather than pixels, thus reducing the granularity of the measurement. Each SCT module is 6 × 8 cm large, with a detector strip every 80 μm. A hit in the SCT is defined as at least 0.1 fC of charge being deposited, with no timing information provided. The SCT provides the maximum tracking extent that can be achieved with ATLAS, up to $|\eta| = 2.5$ [10].

Transition Radiation Tracker

The final part of the ATLAS tracking system is the TRT, whose 73 layers span radii within the range $544 < R < 1082$ mm. Each layer of the TRT is made of straws 4 mm in diameter, that are filled with a gas mixture mainly consisting of argon but also containing carbon dioxide and oxygen. In the barrel, the straws are placed parallel to the beam whereas they are placed perpendicularly in the endcaps. The straws are embedded within fibres made of polyproplylene or polyethylene which induce the transition radiation for which the TRT is named.

The production of transition radiation is greatly enhanced when the incident particle is ultrarelativistic, meaning that the TRT can aid in particle identification. In particular, since many charged pions are produced in association with jets at the LHC, the TRT helps in distinguishing between the tracks left by these particles and those left by electrons. The larger amounts of transition radiation associated with electrons allow it to pass a more stringent set of criteria (High Level, HL, threshold) than the simple tracking hits requirement a pion would pass (low level, LL, threshold).

3.2.3 Calorimetry

After the measurement of the position and momentum of a charged particle, all that remains is to measure its energy. To do this, ATLAS employs a system of sampling calorimeters that alternate layers of absorbing and passive material to both measure the incident energy as well as gain information on the radial and longitudinal development of the resulting shower. In addition, the calorimeter systems allow for the detection of electrically neutral particles which are not visible in the ID.

To measure these energies both electromagnetic and hadronic calorimeters are used. The locations of each calorimeter can be seen in Figure 3.7a.

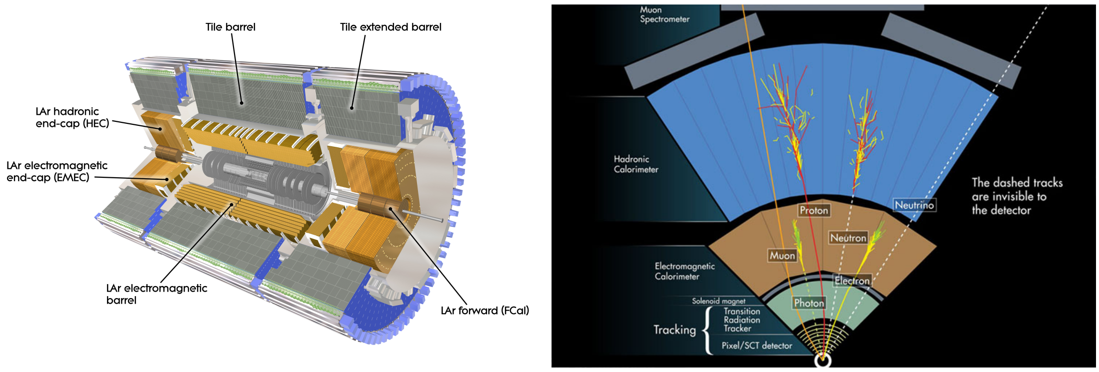

(a) Locations of each calorimeter system [3].

(b) Interactions with each layer by different types of particle [70].

Figure 3.7: An overview of the ATLAS calorimeter system and how it is used to differentiate between different particle types.

The calorimeter system can also help with particle identification due to the distances different particle types will travel through the detector before they begin showering. As shown in Figure 3.7b, electromagnetic objects (i.e electrons, positrons and photons) will be the first to decay within ATLAS. Any tracks in the ID that can be attributed to an EM deposit mean that the deposit most likely originated from a massive EM object, since the neutral photon leaves no trace in the ID. Correspondingly, the first set of calorimeters is the Electromagnetic Calorimeter (ECal).

Hadronic objects will also begin to shower within the ECal but their decay is primarily contained to the next calorimeter layer, the hadronic calorimeter (HCal). Since hadronic objects can interact with both the strong and the electromagnetic forces, their showers tend to be more complicated than those observed within the layers of the ECal.

The combination of the two overarching calorimeter systems allows ATLAS to measure the energies of almost all the particles of the Standard Model. The only exceptions are the neutrinos, whose interaction cross section is so low that measuring a statistically significant number of them across the lifespan of the LHC is highly unlikely, and muons, who travel unimpeded through this segment of the detector.

Electromagnetic Calorimeters

In ATLAS, the ECal is split into several sampling calorimeters which cover the full pseudorapidity range of $0 < |\eta| < 4.9$ with a high level of granularity. These subsystems are split into the central region, covering $|\eta| < 2.5$, and the forward region, where $2.5 < |\eta| < 4.9$. All of the subsystems of the ECal work using the same principle to measure the energy of incoming EM objects. Namely, the incoming object will interact with the passive layer of the calorimeter which will begin to create an electromagnetic shower, from which the total energy of the incoming particle can be inferred. For an incoming electron bremsstrahlung photons will be emitted, which will themselves undergo pair production and produce new electrons. This process will continue until the energy of each photon falls below the threshold necessary for pair production.

At this point any remaining electrons and photons will lose their energy through ionisation and excitation of the surrounding material. Whilst the shower is instigated by the passive layer, the energy is measured in the active layer. Ionisation will occur due to the shower in the active layer, where the resulting electrons create a drift current in the presence of an external high voltage electric field. The amplitude of this current is then proportional to the energy of the original incident particle and can hence be measured.

Whilst the sampling nature of the ECal means it is not possible to exactly measure total energy of the particle, it is possible to mitigate these effects through calibration from the calorimeter development process. Sampling calorimeters also have the advantage of being able to provide more granular information about the development of the EM shower which is of particular use in particle identification. For all subsystems of the ECal, the active material that is used is liquid argon but the passive material can vary depending on the subsystem.

The ECal is composed of four different subsystems to cover the pseudorapidity range $|\eta| < 4.83$: the electromagnetic barrel (EMB), the electromagnetic endcap outer and inner wheels (EMEC-OW/IW) and the forward calorimeter (FCal). The EMB and

Calorimeter	η Coverage	η/x Granularity	ϕ/y Granularity
EMB	$\lvert\eta\rvert < 1.475$	0.025 - 0.075 ($\Delta\eta$)	$\pi/128 - \pi/32$ ($\Delta\phi$)
EMEC - OW	$1.375 < \lvert\eta\rvert < 2.5$	0.025 - 0.05 ($\Delta\eta$)	$\pi/128 - \pi/32$ ($\Delta\phi$)
EMEC - IW	$2.5 < \lvert\eta\rvert < 3.2$	0.1 ($\Delta\eta$)	$\pi/128 - \pi/32$ ($\Delta\phi$)
FCal - 1	$3.1 < \lvert\eta\rvert < 4.83$	$1.5 - 3.0$ cm (Δx)	$1.3 - 2.6$ cm (Δy)

Table 3.1: Electromagnetic Calorimeters within ATLAS.

EMEC are sampling calorimeters using lead and liquid argon as the passive and active materials respectively. Each calorimeter also uses a $\Delta\eta\times\Delta\phi$ geometry for their individual cells. The cells of these calorimeters are also built with an accordion shape to prevent cracks between the individual cells. This also means that tracks within the detector will all pass through approximately the same amount of material. The EMB is formed of three concentric cylindrical layers around the ID in order to improve the quality of the reconstructed EM shower and covers the $|\eta| < 1.475$ region. The transition region between the EMB and EMEC ($1.37 < |\eta| < 1.52$) also contains large amounts of passive material which degrades the quality of the energy response. To help correct for this, a single presampler layer of liquid argon is used in the region $|\eta| < 1.8$.

The EMEC is then further split into two wheels, the outer wheel (OW) covering $1.375 < |\eta| < 2.5$ and the inner wheel (IW) covering $2.5 < |\eta| < 3.2$. Similar to the EMB, the OW has three layers except they are longitudinally segmented due to the perpendicular placement of the EMEC with respect to the beam pipe. The IW only has two layers and also lies outside of the ID tracking acceptance, meaning it is no longer possible to distinguish between hits caused by electrons and those caused by photons. This distinction separates the central region of ATLAS, $|\eta| < 2.5$, and the forward region, $|\eta| > 2.5$.

The final part of the ECal is the forward calorimeter, whose first module is used as an electromagnetic calorimeter in the region $3.1 < |\eta| < 4.83$. Unlike the EMB and EMEC, the FCal is a copper-liquid argon sampling calorimeter, that uses a $\Delta x \times \Delta y$ geometry for the cells. As can be seen in Table 3.1, the cells of the FCal are larger than the cells in the EMB and EMEC meaning that the FCal is particularly susceptible to the effects of pileup, as it becomes more likely that a single cell will measure the energy of multiple incident particles. Similarly to the transition region between the EMB and the EMEC, the transition region between the EMEC and FCal ($3.15 < |\eta| < 3.35$) also contains increased amounts of passive material. Therefore, the combination of the increased cell sizes and larger amounts of passive material result in a highly degraded energy response in this region.

Hadronic Calorimeters

In general, the interaction between a hadron and a calorimeter is more complicated than that of an electromagnetic object. Hadrons can interact with the nuclei of the absorber material in the calorimeter via nuclear excitation, spallation and neutron capture. In addition, hadrons that decay electromagnetically can also produce an EM shower similar to what would be seen in the ECal. These more complex showers result in the ATLAS hadronic calorimeter (HCal) being significantly larger than the ECal, in order to capture the much more diffuse energy deposits. The locations within ATLAS of the individual hadronic calorimeters can be seen in Figure 3.7a.

Calorimeter	η Coverage	η/x Granularity	ϕ/y Granularity
Tile Barrel	$\|\eta\| < 1.0$	0.1 - 0.2 ($\Delta\eta$)	$\pi/32$ ($\Delta\phi$)
Tile Extended Barrel	$0.8 < \|\eta\| < 1.7$	0.1 - 0.2 ($\Delta\eta$)	$\pi/32$ ($\Delta\phi$)
HEC	$1.5 < \|\eta\| < 3.2$	$0.1 - 0.2$ ($\Delta\eta$)	$\pi/32 - \pi/16$ ($\Delta\phi$)
FCal - 2/3	$3.1 < \|\eta\| < 4.9$	1.7 - 5.4 cm (Δx)	2.1 - 4.7 cm (Δy)

Table 3.2: The hadronic calorimeters within ATLAS

Centrally, the HCal is composed of the tile calorimeter and its two extended barrels. This is a three layer sampling calorimeter which uses steel as the passive material and scintillating tiles as the active material. The tile starts at a radius of 2.28m and extends to 4.25m. The tile barrel covers the range $|\eta| < 1.0$, with the extended barrel covering $0.8 < |\eta| < 1.7$. The cell size of the tile calorimeter varies from $\Delta\eta \times \Delta\phi = 0.1 \times 0.1 - 0.2 \times 0.2$.

The Hadronic Endcap Calorimeter (HEC) is the next hadronic calorimeter system within ATLAS, covering the η region $1.5 < |\eta| < 3.2$, resulting in slight overlaps with both the tile and Forward Calorimeter systems. This was a deliberate choice, made in order to reduce the drop in material density in the transition regions between calorimeters [3]. Similar to the ECal systems, the HEC also uses LAr as it's active material but uses copper as its absorber. The HEC has four layers split across two modules, whose cell sizes span the same range as those in the tile.

To cover the remaining pseudorapidity region of $3.1 < |\eta| < 4.9$, the second and third modules of the FCal are used as hadronic calorimeters. However, the last two FCal modules differ from the first in that they use tungsten as the passive material instead of copper. The cells retain the $\Delta x \times \Delta y$ geometry, with sizes $\Delta x = 1.7 - 5.4$ cm and $\Delta y = 2.1 - 4.7$ cm.

3.2.4 Muon Systems

Whilst the ATLAS subdetectors mentioned so far are excellent at measuring the position, energy and momentum of almost every particle type within the Standard Model, they are completely unsuited for studying muons. Although muons leave hits in the tracker and small energy deposits in the ECal and HCal, these are insufficient to fully reconstruct their properties. Instead, ATLAS has one final subsystem whose purpose is to reconstruct muon tracks and their momentum, the muon spectrometer (MS).

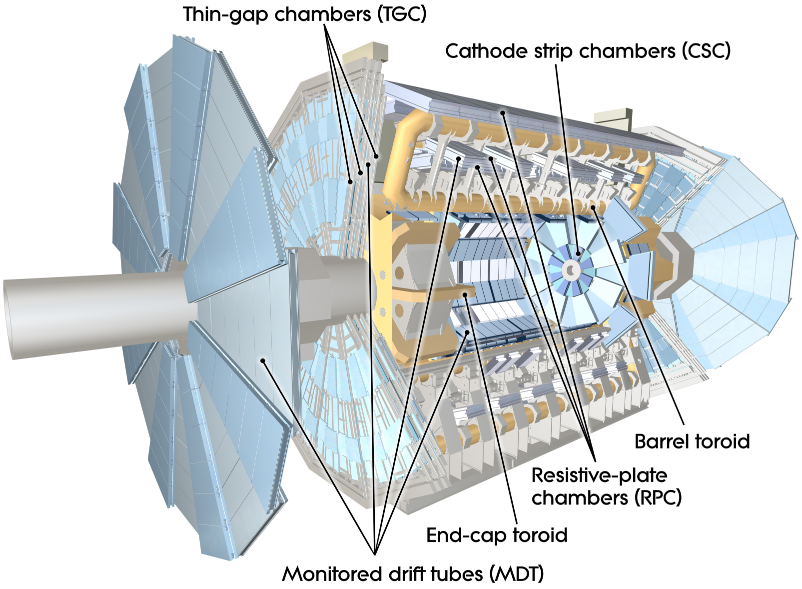

Figure 3.8: A schematic overview of the locations of the subsystems that made up the ATLAS muon spectrometer during Run 2 operations. [3]

Subdetector	Purpose	Location	$\|\eta\|$ Coverage	$\sigma_{z/R} \times \sigma_\phi$
MDT	Tracking	barrel & endcap	$\|\eta\| < 2.7$	35 μm × -
CSC	Tracking	endcap layer one	$2.0 < \|\eta\| < 2.7$	40μm × 5 mm
RPC	Trigger	barrel	$\|\eta\| < 1.05$	10 mm × 10 mm
TGC	Trigger	endcap	$1.0 < \|\eta\| < 2.4$	2-6 mm × 3-7 mm

Table 3.3: Muon Spectrometer subsystems

Monitored Drift Tubes

The monitored drift tubes (MDT) provide tracking information for muons within the range $|\eta| < 2.7$. These tubes are filled with a mixture of argon and carbon dioxide which will be ionised by muons travelling through the tube. The resultant ionisation electrons can then be collected and a hit in the system registered. The MDTs exist as three separate cylindrical shells in the barrel, at radii of approximately 5, 7.5 and 10 m from the interaction point. In the endcaps, the MDTs are instead positioned as four large wheels. These wheels are places at approximately $|z| = 7.4, 10.8, 14$ and 21.5 m from the interaction point. The MDTs themselves offer a spatial resolution of 80 μm, which is sufficient for this region due to the spatial dispersion of particles at large distances from the interaction point [3].

Cathode Strip Chambers

The pseudorapidity region $2.0 < |\eta| < 2.7$ receives a much higher background flux rate than other parts of the detector. To deal with this, the first layer of the MDT in the endcap is replaced with cathode strip chambers (CSC). In comparison to the MDTs, the CSCs have a finer resolution of 40 μm which is better suited for dealing with larger numbers of particles [3]. To provide the tracking information, the CSCs are made of multiwire proportional chambers with cathode planes divided into strips in perpendicular directions.

Resistive Plate Chambers

The muon spectrometer also requires separate detectors that can be used to provide a decision on whether to store the muons recorded; a trigger system. The MDTs and CSCs, however, do not provide information on a fast enough timescale in order to make such decisions. Instead, this task must be performed by tracking detectors that also provide a fast readout. In the barrel ($|\eta| < 1.05$) this is provided by Resistive Plate Chambers (RPC). The RPCs work on the same principal as the MDTs and provide a fast timing resolution of 1 ns, which comes at the expense of a much worsened spatial resolution of 10 mm.

Thin Gap Chambers

As with the precision muon trackers, an improved granularity is necessary in the endcaps to deal with increased background rates. Analogous to the RPCs, the Thin Gap chambers (TGC) perform this role within the muon wheels. The same working principle as the CSCs applies to the TGCs but the TGC is able to read out information with a temporal resolution of 4 ns at a spatial resolution of 3 mm.

3.2.5 Magnet Systems

Whilst the trackers of the ID provide the position of charged particles within ATLAS, alone they are not enough to be able to determine the momentum and charge of the particle. In order to measure these quantities ATLAS is equipped with a system of solenoid, toroid and endcap magnets.

Surrounding the inner detector is a solenoid magnet, that provides a uniform 2 T field. Whilst the solenoid sits in front of the calorimeter system and thus degrades the energy response of particles, its design contributes only 0.66 radiation lengths to the overall pre-calorimeter material budget.

The barrel and endcap toroidal magnets are used to further bend muons and thus measure their momentum. The barrel toroids sit outside of the calorimeter volume and provide a magnetic field of 0.5 T. The endcap toroids sit inside the end of the barrel toroids and produce a field of 1 T. The polarity of the magnetic field is also reversed in the toroids with respect to the solenoid magnet in order to improve the reconstruction of muon tracks within the muon spectrometer.

In order to produce these relatively large magnetic fields, all of the magnets within ATLAS are superconducting and operate at a temperature of approximately 4.5 K.

Whether a particle is positively or negatively electrically charged is inferred from the direction it bends within the magnetic field. The degree of bending caused by the magnet system can also be used to deduce the momentum of incident particles.

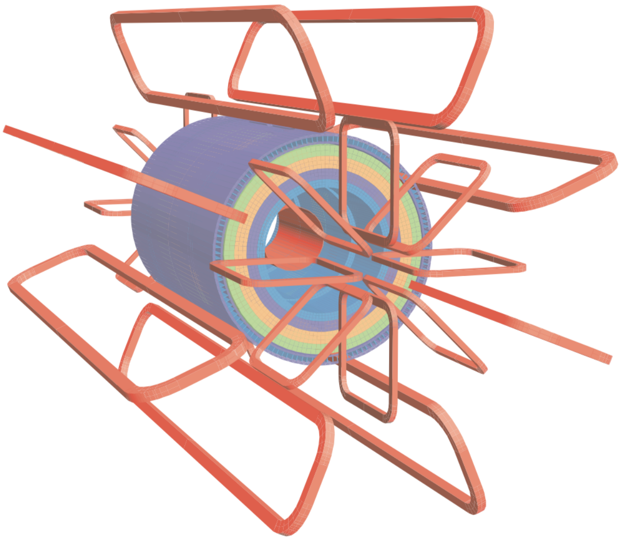

Figure 3.9: A schematic overview of the ATLAS magnet systems. The toroid magnets are shown outside of the calorimetry volume, with the solenoid in red within the barrel [3].

3.2.6 Trigger System

The complex system of trackers, calorimeters, muon chambers and magnets allows ATLAS to measure the properties of decay particles to very high levels of precision. However, since the LHC is a hadron collider this means that the ‘interesting’ parton collisions form only a small part of the total *pp* cross section. Figure 3.10 shows that within the Standard Model, σ_{pp} is five orders of magnitude larger than that of the production of the lightest vector boson which is itself orders of magnitude larger than some of the more exotic SM processes. Given that it is not possible to write out the result of every collision to long term storage due to the 40 MHz of LHC proton bunch crossing rate, it is critical that ATLAS records only the processes of interest. This task is fulfilled by the trigger system.

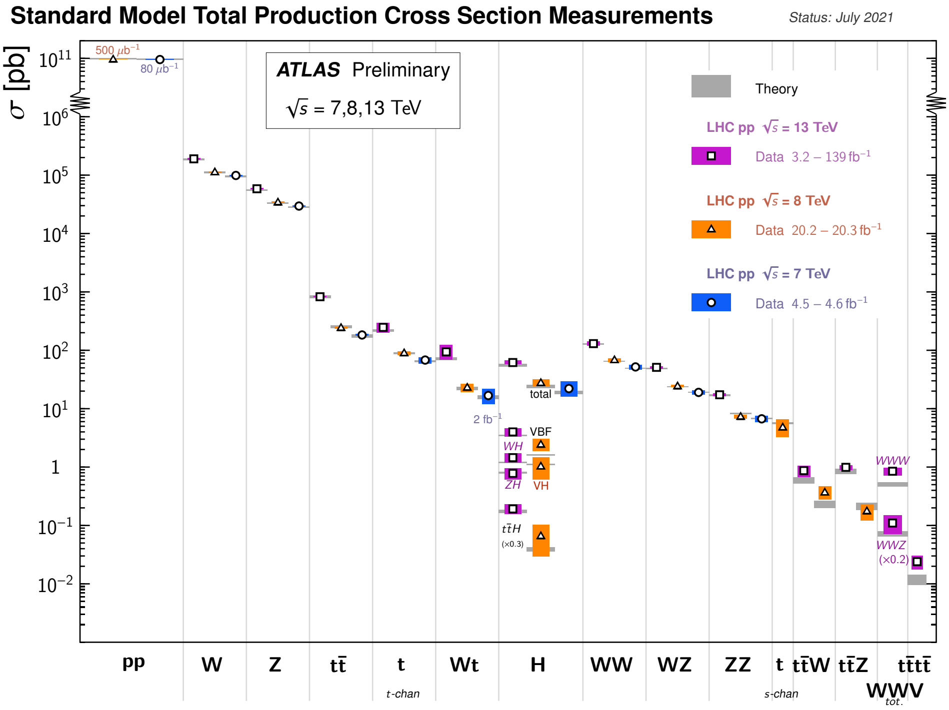

Figure 3.10: An overview of the cross sections of various Standard Model processes at the Run 1 and Run 2 centre of mass energies the LHC was operated at [44].

Level-1 Trigger

The Level-1 (L1) trigger is a hardware based trigger system. Since it has to deal with the full 40 MHz crossing rate provided by the LHC, the decisions made by the L1 trigger have to be fast with minimal amounts of deadtime decreasing its efficiency. Accordingly the L1 trigger provides a decision within approximately 2.5 μs. Due to this extremely short timescale, it is not possible for the L1 trigger to use information from the ID, due to the computational expense associated with the reconstruction of tracks from individual hits.

Instead, it uses low granularity information from both the calorimeter and muon systems. From the calorimeter, the L1 trigger is able to quickly identify electron, photon, tau lepton and jet candidates. The presence of these particles in the calorimeter indicate that a decay has likely taken place at the interaction point and so the event merits further study. The L1 trigger is also able to identify if any MET is present within an event.

From the muon system, the L1 trigger uses hits in the RPC and TGC as well as coincidence counts between the muon subdetectors as a selection criterion. This is done in order to reduce the number of events passing the trigger that did not originate from the interaction point. In the muon system, such particles are likely to be cosmic rays but they may also originate from out of time pileup as well.

The L1 trigger can also apply topological conditions to candidate particles using combined information from the calorimeter and muon detectors. This allows the trigger to consider quantities such as angular separation and invariant mass, which may indicate whether the event is from an ‘interesting’ physics process. Using all these inputs, the L1 trigger will reduce the incoming 40 MHz to 100 kHz of events to be processed by the High Level Trigger.

High Level Trigger

In comparison to the L1 trigger, the High Level Trigger (HLT) is a software based trigger which makes the final decision about whether an event is stored or thrown away. Since the HLT only has to deal with an incoming event rate of 100 kHz, the amount of time the HLT has to decide whether an event is to be kept is significantly longer than in the L1 trigger, with a decision being made in roughly 200 ms rather than 2.5 μs. This larger timespan allows tracking information to be integrated into the trigger decision, with fast reconstruction algorithms being used only with regions of interest highlighted by the L1 trigger. The HLT also incorporates higher granularity calorimeter and muon information in order to make a better informed decision than is possible with the L1 trigger.

After this extra information has been provided to the HLT, sets of predefined selection criteria, ‘triggers’, are used to determine whether an event is recorded or not. If an event passes all the criteria of a certain trigger, then all the raw information from the detector will be read out and stored on disk. From this raw information, the event can be reconstructed and calibrated offline in order to give the best estimate of exactly what happened at the interaction point. During Run 2, the HLT would read out at a rate of approximately 1.2 kHz [42].

3.2.7 Event Simulation

In order to determine whether the data measured with ATLAS contains new physics beyond the Standard Model, it is necessary to understand both the theory of how elementary particles interact and how this looks inside the detector. Whilst the former is handled by the Standard Model itself, the latter is handled by event simulation.

Event simulation starts by first defining the process of interest that is to be simulated, namely the hard process at the heart of a pp collision. The hard process is simulated using Monte Carlo methods to model the resultant outgoing decay particles from the process, given a certain partonic input. This simulation is done according to the PDFs, matrix elements within the Standard Model and the phase space available to the process. Since the information contained in the outgoing decay particles relies only on the understanding of the physics in the production process and does not depend on any interactions with a detector, this level of information can be referred to as ‘truth’ level. The generation of truth level Monte Carlo is typically handled by standard generators such as Powheg [33] for the calculation of leading order (LO) and next to leading order (NLO) cross sections.

After modelling the parton interaction of the hard process the parton shower must be modelled, which radiates energy away from the decay particles. Luckily, the factorisation theorem of QCD means that the soft scale QCD physics contained in these processes means that they can be added on top of the truth level results. It is common to use a generator such as PYTHIA [84] to generate the parton shower and match it with the processes generated at parton level. It is often the case that various combinations of Monte Carlo generators will be tested for both the calculation of matrix elements and matching parton showers to see which gives a result that is most accurate to data. There is often no single ‘best’ Monte Carlo generator, with some achieving better results in different regions of phase space than others.

The combination of the hadronised particles, their decay products as well as leptons and photons must also interact with any matter and fields they pass through if they are to be detected. The modelling of these interactions is typically handled by a program such as GEANT 4 [4]. This allows a full model of the detector to be simulated, whose input is the previously generated hard process and parton shower. From GEANT, the interaction of the incident particles with the detector can be modelled as if it were to actually take place. A simulation of the corresponding detector readout can be thus be generated. This is then passed through experiment specific reconstruction and calibration algorithms, providing a fully simulated Monte Carlo dataset for a given input physics process. Since calibration algorithms are not perfectly accurate, this level of simulation is known as ‘reconstructed level’ or ‘reco level’ to distinguish it from the previous truth level values.

Chapter 4

Lepton Performance

When studying the Drell-Yan process, only the decay products of the Z boson are experimentally accessible for analysis purposes due to the incredibly short lifetime of the Z boson. As described in Chapter 2, these decay products are more than sufficient to probe both the Z decay as well as the underlying $q\bar{q} \to Z$ production. Furthermore, the detector signature of the leptons considered in this analysis is very clean, with minimal hard process activity surrounding them and distinct detector deposits. These deposits can be used to deduce the path taken by the lepton, as well as its original energy. It is also important to properly identify the lepton so as to differentiate as much as possible between real leptons and objects that have a similar detector signature, such as low energy jets. Once a lepton has been reconstructed and identified, it is also necessary to ensure that its measured energy is as accurate as possible. This chapter will therefore outline the reconstruction, identification and calibration algorithms that underpin these performance goals.

4.1 Electron Reconstruction

The first step to producing an electron object for physics analysis is to reconstruct its path through the detector and to associate it with a deposit in the electromagnetic calorimeter [14]. Due to the pesudorapidity coverage of different detector subsystems, this process differs between central ($|\eta| < 2.5$) and forward ($2.5 < |\eta| < 4.9$) electrons.

4.1.1 Central Electrons

In the central region, electrons produced from the interaction point will mainly interact with two subsystems, the Inner Detector and the electromagnetic calorimeter, with a small amount of energy leaking into the hadronic calorimeter. This will be detected as a series of individual hits in each tracking layer, as well as a cluster of active cells from

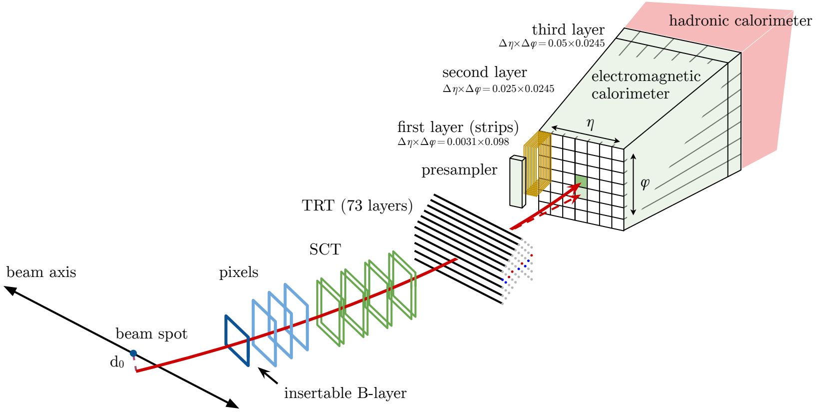

Figure 4.1: A schematic overview showing the path an electron (solid red) and bremsstrahlung photon (dashed red) will take through the tracking layers and electromagnetic calorimeter of ATLAS. In order to reconstruct the electron for physics analysis the path must be extracted from hits in the tracker and matched to a deposit in the electromagnetic calorimeter [2].

the development of the EM shower in the calorimeter. As the electron travels through the detector it is also likely to emit bremsstrahlung photons which will transport energy away from the electron. These effects must be properly accounted for if the reconstructed electron is to be an accurate representation of the one produced at the interaction point. Figure 4.1 gives an overview of this process, showing the interaction of electrons and bremsstrahlung photons with each detector layer.

However, what is not depicted in Figure 4.1 is the additional interactions between the electron and the support structure and cabling of the experiment. These interactions will cause the electron to lose energy that will not be measured in the calorimeter and hence must be corrected for through the electron calibration process which will be introduced in Section 4.4.

Topoclusters

The large number of charged particles from jets and pileup that are produced at every bunch crossing from the LHC mean that there are very high levels of activity within the Inner Detector. Therefore, in order to reduce the combinatoric complexity, electron reconstruction begins with deposits within the electromagnetic calorimeter rather than with track reconstruction. The developing EM shower will be spread across multiple layers and cells of the calorimeter, meaning that a clustering algorithm is necessary to decide which cells should be included in the description of the shower. As outlined by the ATLAS collaboration, clusters are constructed with a ‘4,2,0’ approach [7, 39]. This begins by identifying seed cells with a significance, $|\varsigma_{cell}^{EM}| > 4$, where the significance is defined to be:

$$\varsigma_{cell}^{EM} = \frac{E_{cell}^{EM}}{\sigma_{cell,noise}} \tag{4.1}$$

where E_{cell}^{EM} is the uncorrected energy measured in the cell and $\sigma_{cell,noise}$ is the expected amount of noise, both electronic and due to pileup, in the cell. To expand the seed into a topocluster, the significance criterion is relaxed to two and any such cells neighbouring the seed cell, including those in adjacent layers, are added to the cluster. To complete the process, any active cell, i.e with a significance greater than zero, neighbouring those in the cluster are added to the topocluster. In the process of building up the topocluster, multiple local maxima ($E_{cell} > 500$ MeV and no neighbouring cell has a larger energy) may be included within a single cluster. This typically happens when two incident particles are close together in the calorimeter. Ideally, a cluster should only contain the energy of a single incident particle so each cluster is split until it contains only a single maximum.

Track Reconstruction

The process of reconstructing the path taken by an electron through the Inner Detector begins by identifying hits in the silicon layers of the Inner Detector [2]. From these hits a pattern recognition algorithm is applied to identify the possible candidate tracks. If the candidate track has $p_T > 400$ MeV, its constituent hits are used as the input to the ATLAS Global χ^2 Track Fitter [49] in order to extract preliminary values for the track parameters. These preliminary values and their uncertainties are then improved and updated through a second electron specific fit in the form of a Gaussian Sum Filter (GSF) [36] fit. The primary advantage of the GSF over the χ^2 tracker is that it is better able to take into account the effects of bremsstrahlung of the track, which particularly helps in the extraction of the track curvature. The path of the GSF fitted track can then be extrapolated to the second layer of the calorimeter, where the corresponding pseudorapidity value, η_{track} must match the cluster pseudorapidity, $\eta_{cluster}$ within a tolerance of 0.05 in order for the track to be considered matched to the topocluster. Further constraints are also placed on the ϕ value, which must satisfy $-0.10 < \Delta\phi < 0.05$ where $\Delta\phi = -q \times (\phi_{track} - \phi_{cluster})$. Since these criteria may apply to multiple GSF fitted tracks, a further algorithm uses the differences between cluster and track η/ϕ as well as the number of hits across all silicon trackers and in the IBL in order to label a single track as the primary electron track.

Superclusters

As the electron is deflected by the magnetic field in the ID and interacts with the tracker, bremstrahlung photons will also be emitted. In order to properly reconstruct what occurred at the interaction point, the energy of the bremstrahlung photons must be added to the energy of the main electron topocluster. In order to do this nearby satellite clusters will be added to larger topoclusters to form superclusters [7]. The energy of

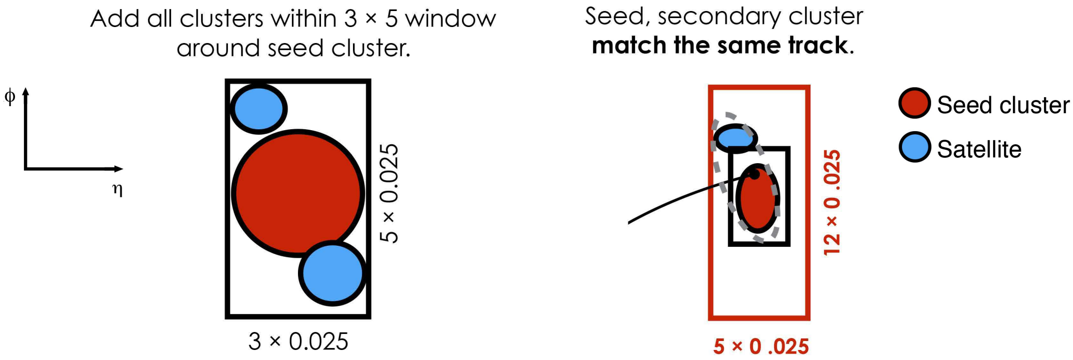

Figure 4.2: A 2D projection of the $\eta-\phi$ plane describing supercluster construction. First all satellite clusters around a seed cluster will be included and then finally any outlying satellite clusters that can be matched to the same track as the seed [7].

the supercluster is then taken to be the energy of the incident electron. To form a supercluster, all topoclusters are first sorted in descending order of their transverse energy, $E_T = \sqrt{p_T^2 + m^2}$, and are individually tested to see if they meet the criteria to become a seed cluster for the supercluster. Starting with the highest energy cell, if the candidate supercluster seed has an energy greater than 1.5 GeV and can also be matched to an ID track that has at least four hits in the silicon trackers then it will be considered a seed cluster. Any other topocluster within a $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ window of the seed barycentre will be classed as a satellite cluster and added to the supercluster. Additionally, any potential satellite cluster in a $\Delta\eta \times \Delta\phi = 0.125 \times 0.300$ window around the seed barycentre will be added to the supercluster if the best matching track to the potential satellite cluster matches that of the seed cluster. Figure 4.2 gives a cartoon representation of how the topoclusters are grouped together and matched together to form the superclusters. Once a supercluster has been identified, its constituent topoclusters are removed from the algorithm and the process is repeated until all possible superclusters have been reconstructed.

4.1.2 Forward Electrons

Unlike the other general purpose particle detector at the LHC, CMS, ATLAS is capable of reconstructing and calibrating electron objects outside of its tracking acceptance. Since no tracking criteria need to be taken into account the reconstruction of forward electrons is therefore simplified in comparison to their central counterparts. A forward electron is, therefore, simply a reconstructed topocluster following the same algorithm as outlined in Section 4.1.1. Whilst it would also be possible to apply the supercluster algorithm to forward topoclusters, the larger cell sizes of the EMEC and FCal shown in Table 3.1 would result in the inclusion of energy from separate electrons, rather than including mainly the bremsstrahlung of the incident electron.

In addition, since a forward electron is only a calorimeter deposit with no matching track, it is impossible to distinguish between forward photons and forward electrons. Therefore, when using forward electrons for physics purposes it is important to be able to match them to a specific resonance. For example, if a central and forward electron have a dielectron invariant mass close to that of the Z boson, it is much more likely that the forward electron is a true electron rather than a photon or some other object.

4.2 Electron Identification

Whilst the reconstruction algorithms are extremely efficient at flagging all possible electrons, they are not infallible. It is possible, and indeed highly likely, that jets or photons, amongst other sources of background, will fake the detector signature of an electron and be falsely classified as such. The ATLAS collaboration deals with this problem through the electron identification process [13]. This involves assigning a discriminant value to the candidate cluster that can be used to separate the signal from background.

The process of assigning a discriminant value to each electron uses multiple variables that describe the track, as well as the longitudinal and lateral shower development within the calorimeter. These variables are first transformed into probability density functions (pdf) separately for signal and background, $P_{S(B),i}$, using an adaptive kernel estimator to smooth histograms of each input variable, i. In order to extract the pdfs, $J/\psi \rightarrow e^+e^-$ and $Z \rightarrow e^+e^-$ events are used in the regions $4.5 < E_T < 15$ GeV and $E_T > 20$ GeV respectively. In the intermediate $15 < E_T < 20$ GeV bin, both J/ψ and Z decays are used to enhance statistics. From these pdfs, a likelihood that the cluster is either signal, S, or background, B, can be constructed:

$$L_{S(B)}(\mathbf{x}) = \prod_{i=1}^{n} P_{S(B),i}(x_i) \tag{4.2}$$

where $\mathbf{x}$ is the set of all input variables and x_i the value corresponding to variable i. The signal and background likelihoods can then be merged into a single discriminant value, d_L as:

$$d_L = \frac{L_S}{L_S + L_B} \tag{4.3}$$

In order to avoid large peaks at zero and one due to background and signal respectively, the shape of the discriminant is further transformed as:

$$d'_L = \frac{-\ln(d_L^{-1} - 1)}{15} \tag{4.4}$$

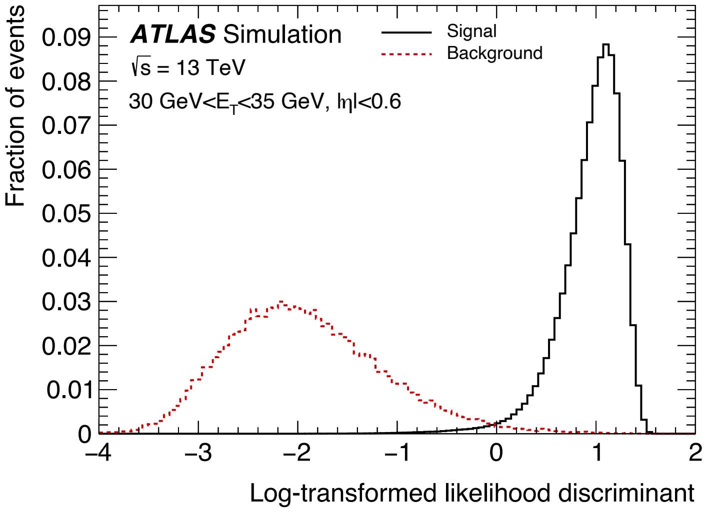

Figure 4.3: The transformed likelihood discriminant distribution for both signal and background [2]. The signal and background histograms are normalised to their respective total number of events.

As can be seen in Figure 4.3, when the discriminant is transformed, two distinct shapes are seen for signal and background. Therefore, by selecting a cut value and discarding all 'electrons' below this value, almost all fake electrons can be eliminated, leaving a relatively pure sample of true electrons. Multiple discriminant values can also be selected to define several so-called working points, to balance out choices between signal purity but also the corresponding reduction in total statistics. In ATLAS, the most commonly used are loose, medium and tight in increasing order of signal purity. For the full Run 2 period, in the $20 < E_T < 50$ GeV range these working points corresponded to signal identification efficiencies of 93%, 87% and 79% respectively. In the same energy range the background rejection was then 98.6%, 99.3% and 99.5% for loose, medium and tight respectively [13].

The likelihood method for calculating identification working points is used for both central and forward electrons but the variables used to construct the pdfs are different. For central electrons, variables that describe the track conditions, the matching of the track to the calorimeter deposits, the shower shape development in the calorimeter layers and the amount of hadronic leakage are all used [7]. The exact variables used are outlined in Table 4.1. However, for forward electrons, the variables used are those contained within Table 4.2; describing the position, energy deposited and shower shape development are used since there is no associated tracking information [39].

The shower shapes are of particular importance since they are used for both identification and calibration purposes. Concretely, they are geometric properties of the shower, describing its longitudinal and radial development and also angular properties, as shown

Type	Description	Name
Hadronic Leakage	Ratio of E_T in the first layer of the hadronic calorimeter to E_T of the EM cluster (used over the range $\|\eta\| < 0.8$ and $\|\eta\| > 1.37$)	R_{had_1}
	Ratio of E_T in the hadronic calorimeter to E_T of the EM cluster (used over the range $0.8 < \|\eta\| < 1.37$)	R_{had}
EM first layer	Total lateral shower width, $\sqrt{(\Sigma E_i(i - i_{max})^2)/(\Sigma E_i)}$, where i runs over all cells in a window of $\Delta\eta \approx 0.0625$ and i_{max} is the index of the highest-energy cell	w_{stot}
	Ratio of the energy difference between the maximum energy deposit the energy deposit in a secondary maximum in the cluster to the sum of these energies	E_{ratio}
	Ratio of the energy measured in the first layer of the electromagnetic calorimeter to the total energy of the EM cluster	f_1
EM second layer	Ratio of the sum of the energies of the cells contained in a 3×7 $\eta\times\phi$ rectangle (measured in cell units) to the sum of the cell energies in a 7×7 rectangle, both centred around the most energetic cell	R_η
	Lateral shower width, $\sqrt{(\Sigma E_i \eta_i^2)/(\Sigma E_i) - ((\Sigma E_i \eta_i)/(\Sigma E_i))^2}$, where E_i is the energy and η_i is the pseudorapidity of cell i and the sum is calculated within a window of 3×5 cells	$w_{\eta 2}$
	Ratio of the sum of the energies of the cells contained in a 3×3 $\eta\times\phi$ rectangle (measured in cell units) to the sum of the cell energies in a 3×7 rectangle, both centred around the most energetic cell	R_ϕ
EM third layer	Ratio of the energy in the third layer to the total energy in the EM calorimeter	f_3
Track conditions	Number of hits in innermost pixel layer	$n_{\text{innermost}}$
	Number of hits in the pixel detector	n_{Pixel}
	Total number of hits in the pixel and SCT detectors	n_{Si}
	Transverse impact parameter relative to the beam-line	d_0
	Significance of transverse impact parameter defined as the ratio of d0 to its uncertainty	$\|d_0/\sigma(d_0)\|$
	Momentum lost by the track between the perigee and the last measure- ment point divided by the momentum at perigee	$\Delta p/p$
	Likelihood probability based on transition radiation in the TRT	eProbabilityHT
Track-cluster matching	$\Delta\eta$ between the cluster position in the first layer of the EM calorimeter and the extrapolated track	$\Delta\eta_1$
	$\Delta\phi$ between the cluster position in the second layer of the EM calorimeter and the momentum-rescaled track, extrapolated from the perigee, times the charge q	$\Delta\phi_{res}$
	Ratio of the cluster energy to the measured track momentum	E/p

Table 4.1: Shower shape and tracking variables used for the identification of central electrons. Adapted from [7].

Type	Description	Name
Shower depth	Distance of the shower barycentre from the calorimeter front face measured along the shower axis	λ_{centre}
Longitudinal second moment	Second moment of the distance of each cell to the shower centre in the longitudinal direction (λ_i)	$\langle\lambda^2\rangle$
Transverse second moment	Second moment of the distance of each cell to the shower centre in the transverse direction (r_i)	$\langle r^2\rangle$
Maximum cell energy	Fraction of cluster energy in the most energetic cell	f_{max}
Normalised lateral moment	w_2 is the second moment of r_i setting $r_i = 0$ for the two most energetic cells, while w_{max} is the second moment of r_i setting $r_i = 4$ cm for the two most energetic cells and $r_i = 0$ for the others	$\frac{w_2}{w_2+w_{max}}$
Normalised longitudinal moment	l_2 is the second moment of λ_i setting $\lambda_i = 0$ for the two most energetic cells, while l_{max} is the second moment of λ_i setting $\lambda_i = 10$ cm for the two most energetic cells and $\lambda_i = 0$ for the others	$\frac{l_2}{l_2+l_{max}}$

Table 4.2: Shower shape variables used for the identification of forward electrons. Adapted from [39].

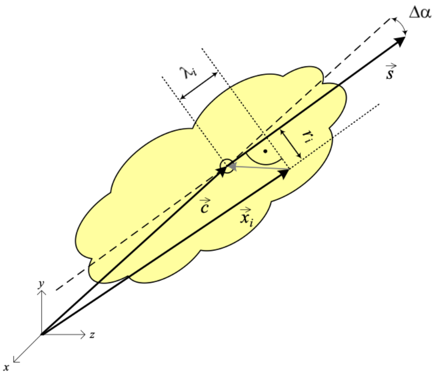

$\vec{c}$	*centre of gravity of cluster, measured from the nominal vertex ($x = 0, y = 0, z = 0$) in ATLAS*
$\vec{x}_i$	*geometrical centre of a calorimeter cell in the cluster, measured from the nominal detector centre of ATLAS*
$\vec{s}$	*particle direction of flight (shower axis)*
$\Delta\alpha$	*angular distance $\Delta\alpha = \angle(\vec{c}, \vec{s})$ between cluster centre of gravity and shower axis $\vec{s}$*
λ_i	*distance of cell at $\vec{x}_i$ from the cluster centre of gravity measured along shower axis $\vec{s}$ ($\lambda_i < 0$ is possible)*
r_i	*radial (shortest) distance of cell at $\vec{x}_i$ from shower axis $\vec{s}$ ($r_i \geq 0$)*

Figure 4.4: A cartoon depiction of an electromagnetic shower within ATLAS, with its various geometrical properties labelled and described [39].

in Figure 4.4. The shower shape variables are also complemented by their associated moments, which for a generic variable, x, is defined at order n as:

$$< x^n > = \frac{\sum_i E_i \, x_i^n}{\sum_i E_i} \tag{4.5}$$

where the sum over i takes every active cell in the EM cluster into account.

4.2.1 Shower Shape Corrections

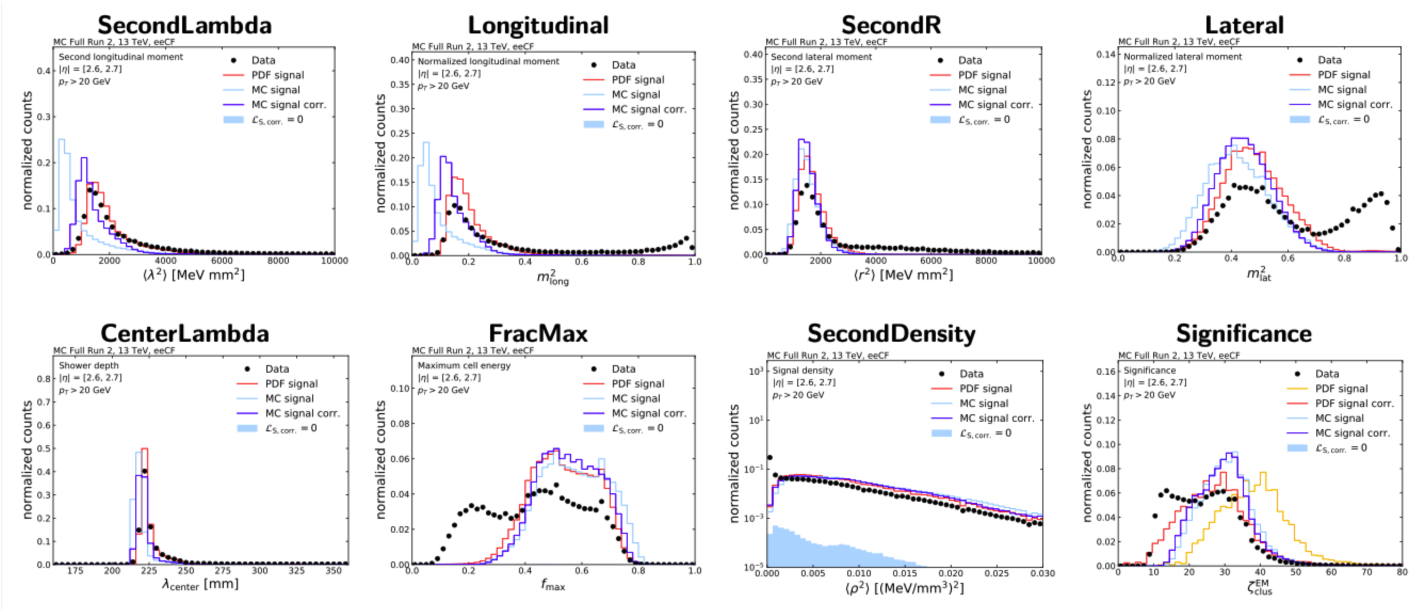

Figure 4.5: EMEC forward electron shower shapes pre- and post-correction compared with the corresponding data distribution [67].

In order to use the shower shapes for both forward electron identification and calibration purposes, it is imperative that the underlying Monte Carlo modelling is accurate. From previous work, it is, however, a known problem that the Monte Carlo modelling of the shower shapes in the EMEC is inaccurate [67] as is shown in Figure 4.5. Therefore, the shower shape variables must be corrected so that the identification and calibration algorithms can be accurately applied to data as well as to Monte Carlo. As was done in Run 1 [22], the correction procedure involves fitting a shift and smear parameter to reparameterised Monte Carlo distributions such that they match the data distributions better. In the FCal, the simulated shower shapes already show a good agreement with data and require no further correction [22].

4.3 Efficiency and Scale Factors

The imposition of all reconstruction, identification, isolation and trigger requirements naturally means that the total number of electrons available for physics analysis is much less than the number of electrons that are initially measured. The proportion of the remaining electrons to the initial amount is known as the efficiency ε [13]. The efficiency is itself made of various sources as shown in Equation (4.6):

$$\varepsilon_{total} = \varepsilon_{EMclus} \times \varepsilon_{reco} \times \varepsilon_{ID} \times \varepsilon_{iso} \times \varepsilon_{trigger} \tag{4.6}$$

where ε_{EMclus} is the ratio of the number of reconstructed EM clusters to the total number of true electrons. The reconstruction efficiency, ε_{reco}, is the number of electron candidates with an associated track (N_{reco}) divided by the total number of topoclusters ($N_{cluster}$). The identification efficiency, ε_{ID}, divides the number of electrons passing an ID working point (N_{ID}) by N_{reco}. The isolation efficiency is then defined as the number of candidates that pass all isolation, identification and reconstruction criteria (N_{iso}) divided by N_{ID}. The final requirement on an electron candidate would be to pass a HLT trigger, with efficiency ε_{trig}, which is defined to be the number of reconstructed, identified and isolated electron candidates, N_{trig} divided by N_{iso}.

In order to measure the efficiencies in both data and Monte Carlo, the tag-and-probe method is used [13]. This method uses a pair of particles produced by a resonance, a ‘tag’ and a ‘probe’, where the tag particle passes a tight identification working point and also satisfies the HLT trigger requirements so that there is a high degree of certainty that this particle is signal and not background. The probe satisfies much looser requirements than the tag, if any are applied at all. The efficiency is then the ratio of tag and probe pairs where the probe matches the criteria applied to the tag, to the total number of tag and probe pairs. Except for ε_{EMclus}, which can only be evaluated with Monte Carlo given its dependence on truth level information, all efficiencies in Equation (4.6) are measured via the tag-and-probe method. The tag-and-probe method is typically applied to $Z \to e^+e^-$ events since the electrons produced cover a large amount of the phase space probed by ATLAS.

Since the Monte Carlo modelling does not replicate the effects of data perfectly, the measured efficiencies are, however, not necessarily equal to those measured in data. To account for this, a scale factor (SF) is applied to all Monte Carlo which has the effect of equalising the efficiencies. The scale factor is defined as:

$$SF = \frac{\varepsilon_{data}}{\varepsilon_{MC}} \tag{4.7}$$

A different scale factor for each efficiency is derived such that when applied to Monte Carlo, the resulting scaling takes the form of Equation (4.6). The systematic uncertainties associated with each type of scale factor are also determined by changing the selection criteria for the tag-and-probe method as well as how the background subtraction is performed in data to minimise the contamination from non-prompt electrons. Different invariant mass ranges around the Z resonance are also used as a systematic when calculating the reconstruction efficiency. The full set of scale factors and their systematic uncertainties for central electrons are produced centrally by the ATLAS collaboration with more details in Reference [13]. Identification efficiency scale factors can also be calculated for forward electrons, incorporating the shower shape corrections of Section 4.2.1, however, these must be calculated at analysis level as they are not produced centrally by the ATLAS collaboration. The systematic variations on the forward identification scale factors are taken to be the statistical errors on each scale factor, as well as variations of the tag-and-probe method used to derive them.

4.4 Electron Calibration

Ideally, the ATLAS ECal would measure the total energy of an incoming electron and not be affected by pileup. However, the interactions of electrons with the Inner Detector as well as the passive material of the detector all result in a loss of energy. This means that the energy recorded in the ECal is not the true energy the electron had when it was produced at the interaction point. In addition, the passive material in the detector is not evenly distributed and issues with the high voltage systems in specific regions in the detector during data taking also occur. These introduce non-uniformities in the energy response of the electron which also must be corrected for, as an electron with a given energy should be reconstructed with exactly that energy, no matter where it is detected. Several calibration methods are applied to electrons which attempt to account for these issues and provide an electron energy value that is as close to the true energy as possible.

4.4.1 Central Electrons

The first step of the correction chain in Figure 4.6 is the layer intercalibration in data, where the data response for the presample layer scale and the ratio of the response of the first and second calorimeter layers, E_1/E_2, is corrected using muon data in low pileup runs as detailed in Reference [14].

The calibration then continues with the application of a global multivariate algorithm (MVA) calibration (step 2 in Figure 4.6), applied to both data and Monte Carlo, whose aim is to correct the reconstructed electron energy towards its true energy. This MVA takes the form of an ensemble of boosted decision trees (BDT), trained in bins segmented in both η and E_T in order to maximise the performance of E_{reco}/E_{true} across both the detector coverage and energy spectrum. The BDTs are trained on Monte Carlo, using variables that describe the energy, position and shower shape development of the electron, which has the effect of improving both the resolution (the width of the E_{reco}/E_{true} distribution) and energy scale response (the peak position of the E_{reco}/E_{true} distribution) that is measured. A full overview of the MVA calibration procedure is provided by the ATLAS collaboration [1].

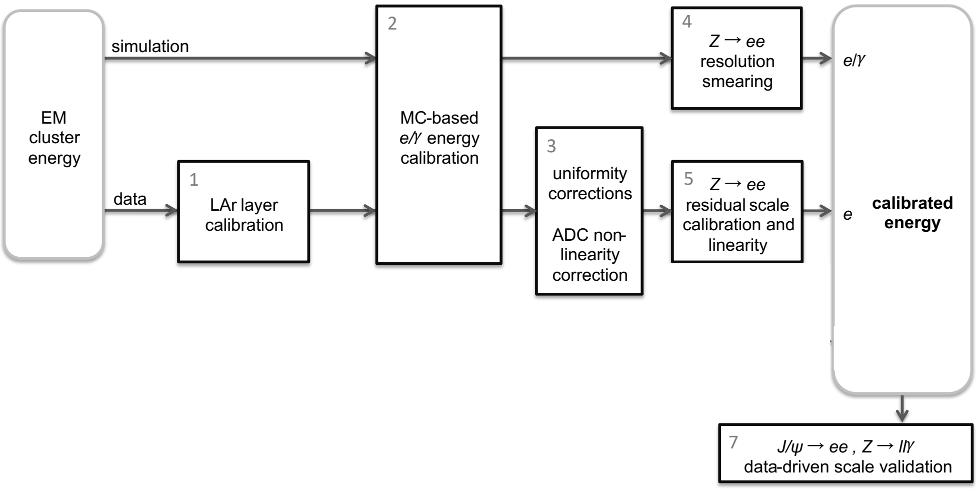

Figure 4.6: The calibration chain for central electrons in ATLAS. Adapted from [14]

Whilst the MVA performs excellently, its performance is not as good in data as it is in simulation. The distribution of passive material is known exactly in Monte Carlo and there are also no issues with the high voltage system that can affect the energy of the reconstructed electron. To counteract these effects, a separate uniformity correction is derived for data, which has the effect of flattening the distribution of $E_T/<E_T>$ versus the electron ϕ coordinate and hence harmonises the energy response across the calorimeter. A further correction is also applied to account for modelling differences between the calorimeter readout and the electrical current it is converted to, from which the cell energy is derived. After the application of these corrections, there will still be remaining differences in electron energy between data and Monte Carlo, which become apparent when looking at the Z resonance. The simulation tends to be too optimistic when modelling the electron energy resolution, hence a smearing factor is applied to the energy such that the resolution of the Z in Monte Carlo will match that in data. Due to detector effects, in data the peak of the Z is not in its true position. A uniform shift is therefore applied to the energy of the electrons in data which has the effect of matching the peak position of the Z with that of Monte Carlo. The combination of the Monte Carlo smearing and the data shifts is known as the residual in situ scale calibration. The effect of these combined calibrations and corrections is then validated using $J/\psi \rightarrow ee$ and $Z \rightarrow ll\gamma$ decays. With this procedure carried out, electrons can then be used for physics analysis. The systematic uncertainties from each of these calibration steps are applied to the measurement of the angular coefficients in Chapter 5 according to the treatment described in Reference [14].

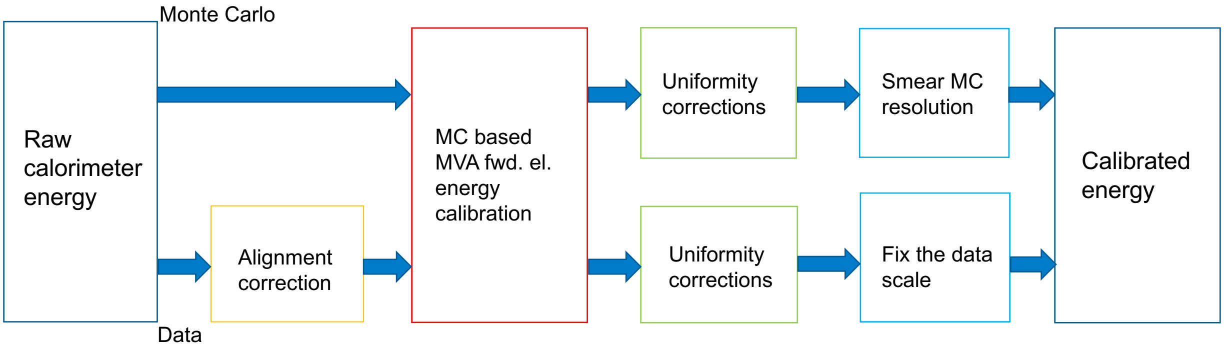

Figure 4.7: The new calibration chain for forward electrons developed for this analysis.

4.4.2 Forward Electrons

As with central electrons, forward electrons must also be calibrated so that the physics results obtained from their use are valid. However, since forward electrons are only used in a small fraction of physics analyses the ATLAS collaboration does not provide official corrections for them. This therefore meant that an entirely new forward electron calibration chain had to be developed at analysis level, which can be seen in Figure 4.7. Of particular importance is the MVA calibration that was developed by the author for the purpose of maximising sensitivity to the weak mixing angle [87]. This was the first time that an MVA calibration had been developed for the forward electrons and delivers large improvements to the electron energy scale and resolution as will be further outlined in Global Scale Energy Calibration. Alongside these improvements, the forward calibration chain also encompasses an alignment correction in data, a uniformity correction and finally an in situ calibration that matches the Z mass spectra in data and Monte Carlo.

Alignment Correction

When producing Monte Carlo samples for analysis purposes, the position of each sub-detector in ATLAS is known with perfect accuracy. However, data is produced from the real detector, whose components are not perfectly aligned. The small non-alignments in ATLAS affect the reconstruction of electrons and so must be accounted for. As outlined from Run 1 forward electron performance studies [22], in order to derive the necessary correction, a selection of low p_T^Z events are used, where the Z is produced approximately at rest. The electrons from such decays will be produced approximately back to back in the lab frame, which allows for the degree of misalignment to be quantified. If the central electron is considered to have been measured perfectly, then any bias in the phi coordinate difference between the central and forward electron will be caused by the forward electron alone. This bias, $d\phi$ can then be estimated through the first derivative of the ϕ coordinate of the forward electron plus an offset, $d\phi_0$ i.e

$$\phi = \tan^{-1}(y/x) \rightarrow d\phi = \frac{dx\cos\phi + dy\sin\phi}{r} + d\phi_0 \tag{4.8}$$

where x, y, and r are coordinates of the forward electron. The quantities, dx, dy and $d\phi_0$ are then fit to the 2D distribution of $< \Delta\phi_{ll} >$ in both data and Monte Carlo, separately for the EMEC and FCal in both the A and C sides of the detector. The fitting process is carried out iteratively in order to improve the closure of the fit, after which the data coordinates are corrected towards Monte Carlo.

Global Scale Energy Calibration

As with the central electron calibration, it is necessary to correct the reconstructed energy of a forward electron towards its true value. This was achieved through the training and validation of a multivariate algorithm (MVA), specifically an ensemble of boosted decision trees (BDT) [87]. Instead of training the BDTs using simulated $Z \rightarrow e^+e^-$ events, a single particle Monte Carlo sample of electrons, with pileup was produced using the Particle Gun generator. This Monte Carlo sample used was 'flat' in the sense that it has a uniform distribution of events in η, p_T and also in pileup. This meant that the training would not be biased towards any region of phase space and would also not be influenced by the physics resulting from a Z decay. To select the subset of events that would be used for training, the amount of cleaning applied to the Monte Carlo was minimal, namely:

- The reconstructed particles were matched to generator level particles.
- The raw reconstructed electron energy was not too far from the true value: $0.5 < E_{true}/E_{reco} < 3$
- The electron was reconstructed not too far from where it actually hit the calorimeter: $\Delta R^2 = (\eta_{cl} - \eta_{true})^2 + (\phi_{cl} - \phi_{true})^2 < 0.1$.

Target and Input Variables

Whilst the objective of the MVA is to predict the true energy of electrons, such an output would require the MVA to make predictions over several orders of magnitude which would lead to suboptimal performance of the MVA. Instead the MVA is given the target variable of E_{true}/E_{reco} such that the output of the MVA can be used as a correction factor to the raw reconstructed energy from the calorimeter i.e $E_{calib} = E_{reco} \times$ MVA Output. Since the MVA output acts as a proxy for E_{true}/E_{reco}, the overall calibrated energy should be the true energy.

The MVA training was split into three separate regions in order to optimise the training in each of them, with appropriate training variables used in each region. These three regions correspond to the EMEC, FCal and transition region between the two calorimeters, with the pseudorapidity coverage defined by the values in Table 3.1.

In each of these regions, the variables used for training were chosen in order to satisfy three main areas:

- Where was the cluster in the calorimeter?
- How did the electromagnetic shower develop in the calorimeter?
- How energetic was the cluster deposited in the calorimeter?

To cover these areas, the first implementation of the MVA utilised a set of common variables, as well as region specific variables. Specifically, the common variables were:

- The raw reconstructed energy of the cluster: E_{cl}
- Cell index: $\eta_{cl}/\Delta\eta$ where $\Delta\eta = 0.1\,(0.15)$ in the EMEC (FCal), the size of a single cell in the calorimeter. For the FCal where the geometry is not $\Delta\eta \times \Delta\phi$, the size of the cell is taken to be an approximate size when converted from $\Delta x \times \Delta y$.
- η with respect to the cell edge: $\eta_{cluster}$ mod $\Delta\eta = 0.1\,(0.15)$ in the EMEC (FCal)
- The second moment of the radial development of the EM shower: $< r^2 >$
- The distance from the shower centre to the calorimeter edge, measured along the shower axis: λ_{centre}
- The second moment of the energy density in the calorimeter: $< \rho^2 >$
- The actual number of interactions per beam crossing: μ
- The number of primary vertices that were reconstructed: npv

EMEC

In the EMEC, the position of the cluster can be parameterised through $\eta_{cluster}$ and $\phi_{cluster}$ due to the $\Delta\eta \times \Delta\phi$ geometry of the cells of the EMEC. In addition, further information on the shower shape development can be extracted by using the energy deposited in each layer of the EMEC The EMEC specific variables are therefore:

- Pseudorapidity of the cluster in the calorimeter frame: $\eta_{cluster}$.
- ϕ coordinate of the cluster in the calorimeter frame: $\phi_{cluster}$.
- ϕ coordinate of the cluster modulus $2\pi/16$
- The ratio of the maximum energy deposited in a single calorimeter cell in layer one to layer two of the EMEC: $E^{Max}_{S_1}/E^{Max}_{S_2}$

FCal

Due to the x-y geometry of the cells of the FCal it would not be appropriate to use $\eta_{cluster}$ directly as a training variable. Instead the centre X,Y and Z coordinates of the cluster can be used instead. From this it is still possible to use the cell index and $\eta \bmod \Delta\eta$ variables by using the appropriate spherical coordinate transformation to $\theta_{cluster}$, from which η can be calculated using the standard definition of the pseudorapidity in Equation (3.5).

Transition Region

Since the transition region is mainly within the FCal, the centre X,Y and Z coordinates are again used as for the FCal BDT. However, since there are no identification scale factors for this region, the shower shape variables were excluded from the training in this region.

Performance Metrics

The two most important metrics to be evaluated from the training are the scale and the resolution of the electron energy response.

The scale is evaluated through the peak position of a Gaussian fitted to the E_{reco}/E_{true} distribution. The fit is performed within the range [-1,+2] standard deviations of the E_{reco}/E_{true} distribution in order to negate the impact of any non-Gaussian tails in the distribution that are caused by bremsstrahlung upstream of the calorimeter.

The resolution is quantified via the interquartile efficiency (IQE) i.e the interquartile range of the E_{reco}/E_{true} distribution scaled by the appropriate value so that the result is a measure of the width of the distribution in comparison to a unit Gaussian. Mathematically it is defined as:

$$IQE = \frac{Q_3 - Q_1}{2\phi^{-1}(0.75)} \tag{4.9}$$

Where Q_1 and Q_3 are respectively the first and third quartiles of the distribution and ϕ^{-1} is the inverse cumulative distribution function for the normal distribution. By evaluating this function at 0.75, i.e at three-quarters into the Gaussian distribution, and multiplying it by two, the total width of a unit Gaussian is retrieved and thus an IQE value of one would correspond to a distribution that follows that of a unit Gaussian. In comparison to other methods that can be used to estimate the resolution, the IQE does not rely on any underlying fits to the distribution which reduces the sources of error in the value.

Initial Results

The initial approach was to keep everything as simple as possible; use only a single BDT to cover the entire η and E_T space. Figure 4.8 shows that even with a minimal approach, MVA calibration in the forward region can be highly successful. The E_{reco}/E_{true} distribution exhibits the expected behaviour of the shifted peak to one, whilst also becoming significantly thinner. This behaviour isn't just a global phenomenon but is also replicated when the peak position and IQE are plotted vs η. Whilst the degree to which the scale and improvement in IQE changes, depending on what part of the detector the

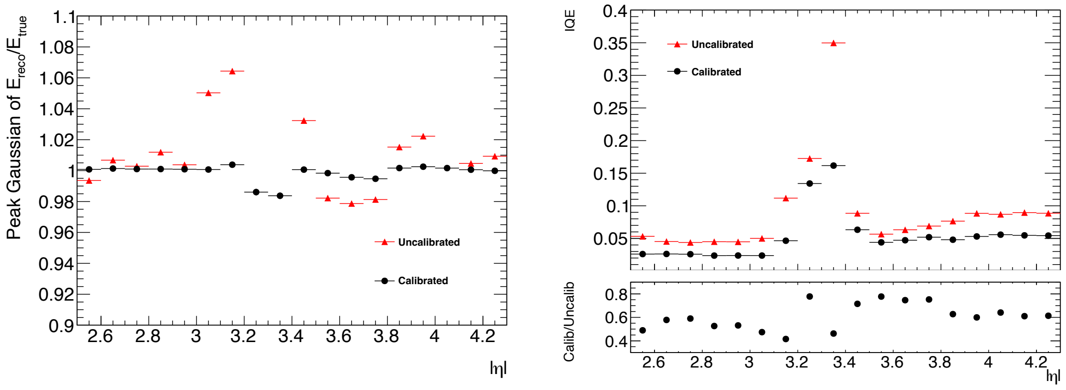

(a) Linearity vs η.

(b) Resolution vs η.

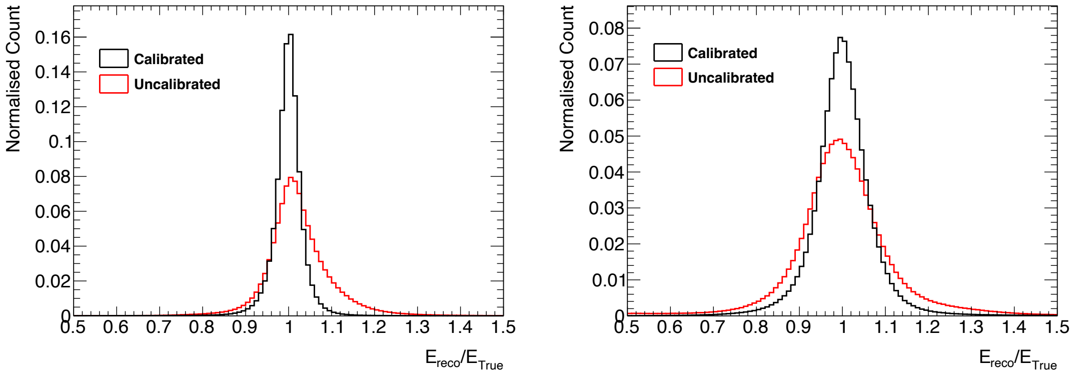

(c) Energy distribution in the EMEC.

(d) Energy distribution in the FCal.

Figure 4.8: Performance plots for the initial training showing the improvements in linearity and resolution, as well as the underlying E_{reco}/E_{true} distribution for both the uncalibrated and calibrated cases in the EMEC and FCal.

electron was reconstructed in, each η bin displays a large improvement in comparison to the uncalibrated case.

η Splits

The central electron calibration currently uses more than 100 BDTs, binned in η and p_T to optimise performance. Such optimisation, however, comes with the cost of highly increased complexity of the final software package. It was therefore decided to try and keep the BDT ensemble for the forward calibration as simple as possible, without sacrificing large gains in performance.

The simplest possible binning in η space is that where there is no binning, i.e a single BDT covers the entire η range, as was initially investigated. The next logical binning is to split the events based on what calorimeter they occur in, with one BDT for the EMEC and one for the FCal. In this simplified splitting regime, the transition region was included as part of the FCal training.

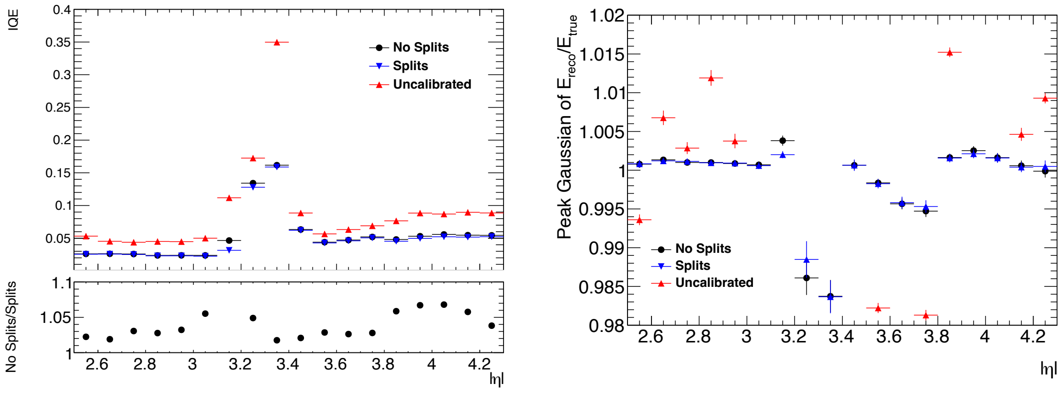

(a) Resolution vs η.

(b) Energy scale vs η.

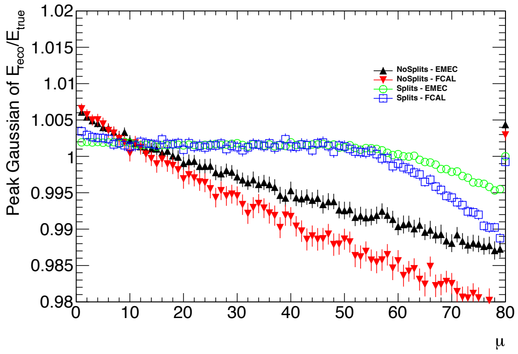

(c) Energy scale vs pileup

Figure 4.9: The performance plots comparing splitting the training into EMEC and FCal vs no splitting at all.

From Figure 4.9 it was found that whilst there was little change to the linearity when plotted against η, there was a 5% improvement in resolution in the FCal. Moreover, splitting the training also allowed for better performance in high pileup conditions by maintaining the quality of the energy scale for longer in comparison the no splits case.

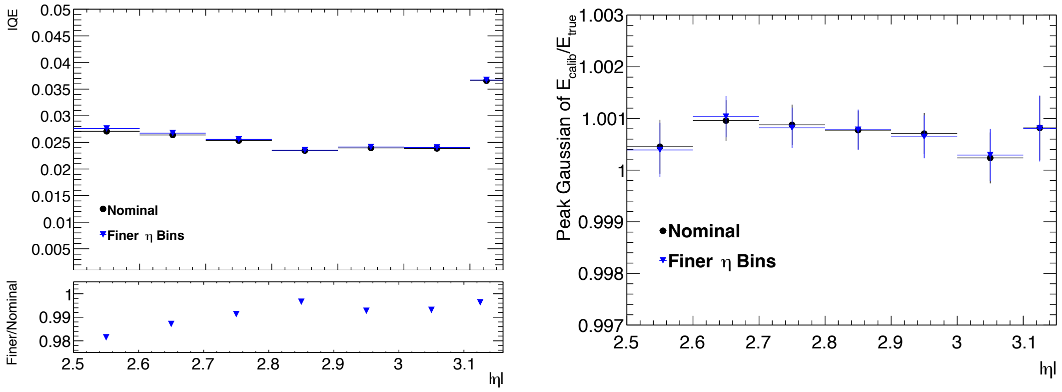

(a) EMEC resolution

(b) EMEC energy scale

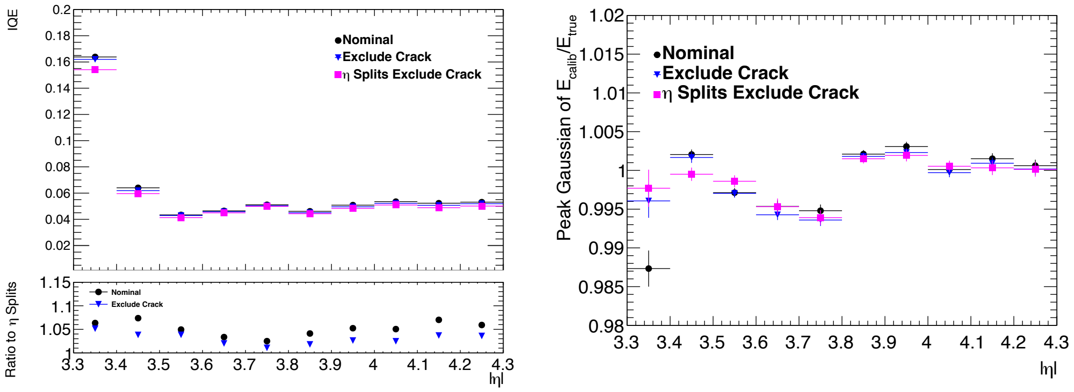

(c) FCal resolution

(d) FCal energy scale

Figure 4.10: Post calibration comparison of splitting the EMEC and FCal trainings into finer η bins. The results in the FCal are plotted excluding the transition region between the EMEC and FCal.

It was also decided to investigate a finer η splitting, where the bin edges were defined as:

- 2.5, 2.7, 2.9, 3.15 for the EMEC.
- 3.15, 3.35 for the transition region.
- 3.35, 3.5, 3.7, 3.9, 4.1, 4.3 for the FCal.

From Figure 4.10, it is clear that the finer η splitting is unnecessary in the EMEC since there is no gain in performance in either scale or resolution. Conversely, the FCal plots show that there are some improvements that can be made, with an approximate 5% improvement in resolution possible. The majority of these gains could however be obtained through a much simpler η splitting, where the FCal training was split into the transition region and the remaining FCal. The decision was therefore made to keep only the simpler η split of EMEC, transition region and FCal, rather than the finer binning.

E_T Splits

Splitting the events into $E_T^{Uncalib}$ bins was also investigated to see if an energy dependent training would improve the final results. It was decided to use splits of:

- 20, 40, 100 GeV for the EMEC
- 20, 40, 60, 80, 100 GeV for the FCal and transition region

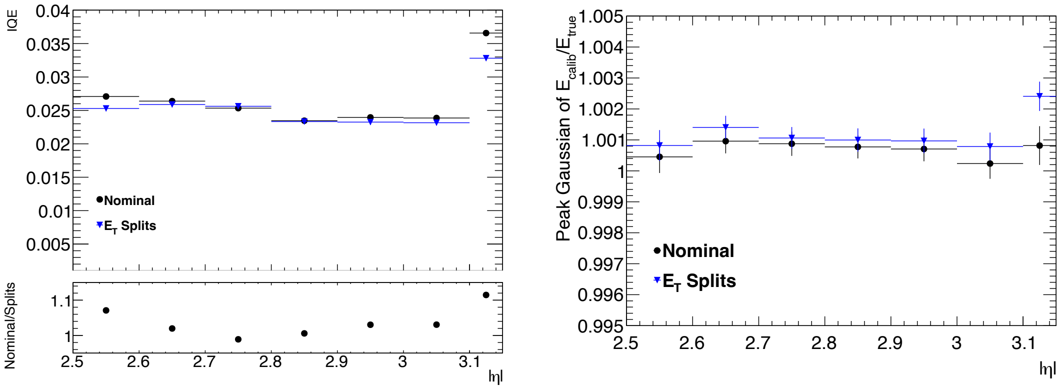

(a) EMEC Resolution (b) EMEC energy scale

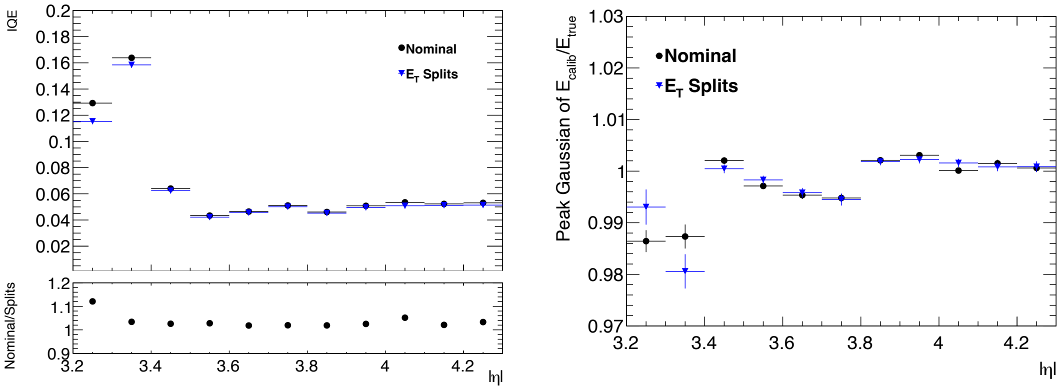

(c) FCal Resolution (d) FCal energy scale

Figure 4.11: Post calibration comparisons of splitting the separate EMEC and FCal trainings into finer E_T bins.

Like the case of finer η binning, splitting the events based on p_T has only a minimal effect. This is shown in Figure 4.11 where neither the linearity, nor the resolution is affected in the bulk of the EMEC or the FCal. The improvements are limited to the transition region only and since this region is not used for physics analysis, it would not be worthwhile to implement these improvements.

Shower Shape Dependence

The issue of the EMEC shower shape mismodelling is well known, however, the extent of the mismodelling depends on the shower shape variable. For example, Figure 4.5 shows that for $< r^2 >$ there is already a good agreement between ATLAS data and the raw MC signal for this variable. On the other hand, $< \lambda^2 >$ shows a highly inaccurate modelling for the raw MC signal. Even when applying the correction procedure, $< \lambda^2 >$ is still badly modelled.

However, as mentioned in Section 4.2.1, the FCal does not suffer from any shower shape mismodelling meaning they can be used without any issue. This EMEC modelling problem could be sidestepped entirely by not using any shower shape variables for the training. To test this, a training was carried out to see if removing the shower shape variables would have any effect on the scale and resolution that could be achieved.

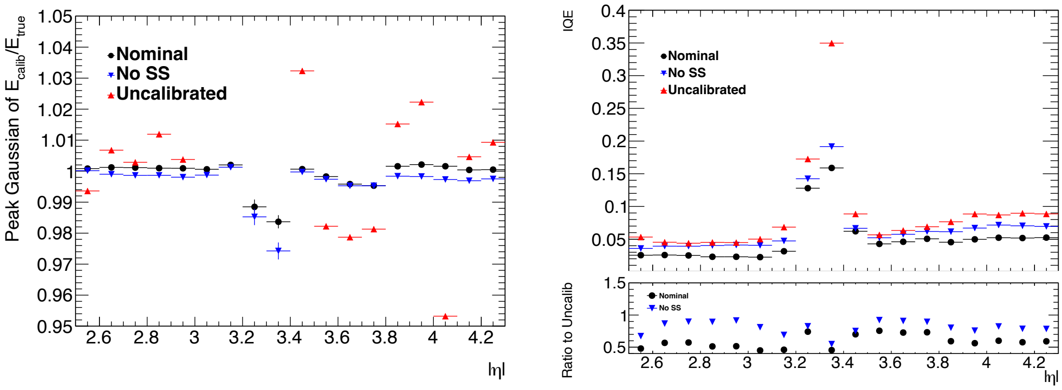

(a) Energy scale

(b) Resolution

Figure 4.12: Performance plots comparing training with and without shower shapes.

Figure 4.12 shows that whilst the scale is unaffected by the removal of the shower shapes, the resolution is greatly impacted. By removing the shower shapes, the resolution in the bulk of the EMEC is little better than the uncalibrated case. By including the shower shapes, there is an approximate factor of two improvement to the resolution. With such an improvement in resolution possible, it was decided to keep the shower shapes as training variables.

However, it is important to note that at this point, the EMEC BDT was trained and evaluated using uncorrected shower shapes that did not match the data. Although the simulated $< r^2 >$, $< \rho^2 >$ and λ_{centre} distributions have a good modelling of the data,

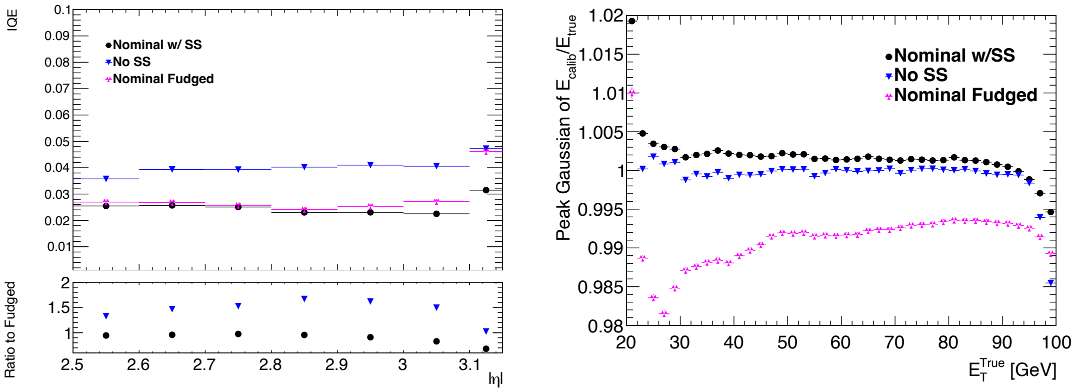

(a) Resolution (b) Energy scale

Figure 4.13: Performance plots showing the resolution and linearity (vs E_T^{True}) for the trainings with shower shapes (nominal), without shower shapes and after evaluating the shower shapes training with corrected shower shapes (nominal fudged on the plots). The resolution can be seen to vary from the nominal case at high η whereas the linearity exhibits a large discrepancy between the two cases.

any mismodelling may result in a suboptimal performance of the BDT when it is applied to data. Therefore, it is important to test the BDT using Monte Carlo where the shower shapes have been corrected as outlined in Section 4.2.1 i.e Monte Carlo that acts as a proxy for ATLAS data. The resultant scale and resolution could then be evaluated, the results of which can be seen in Figure 4.13.

Ideally, there would be no difference between the BDT outputs when evaluating on either corrected or uncorrected Monte Carlo. Whilst this was approximately true for the resolution, a discrepancy was found in the energy scale. When plotted against the true transverse energy, E_T^{True}, the bulk of the distribution exhibited a 1% deviation from unity in linearity (the energy scale as a function of the transverse energy) as shown in Figure 4.13b. This non-linearity increased to a maximal value of 2% at low E_T, which represents a significant deviation from the true energy scale.

Since the non-linearity is caused by the shower shape variables, then removing some of them from the training will correspondingly decrease the size of the non-linearity. Figure 4.14 shows that by removing λ_{centre} and $< \rho^2 >$, retaining only $< r^2 >$, the performance was impacted only slightly, with no change in the linearity and only a 10% reduction in resolution in comparison to the calibrated, three shower shapes case.

Figure 4.15 shows that when this BDT is evaluated using corrected Monte Carlo, the size of the non-linearity is reduced as would be expected. At low E_T, the non-linearity was reduced to 1.5% and within the bulk to 0.5%. The decision was therefore made to keep $< r^2 >$ as the only shower shape variable used in the EMEC. However, the size of these non-linearities mean that applying the EMEC BDT to data would still be suboptimal,

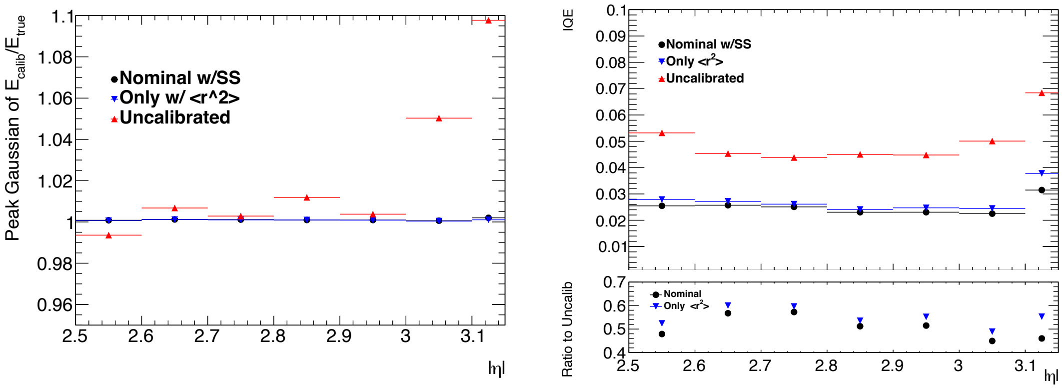

(a) Energy scale

(b) Resolution

Figure 4.14: Performance plots showing the effect of removing $< r^2 >$ from the training.

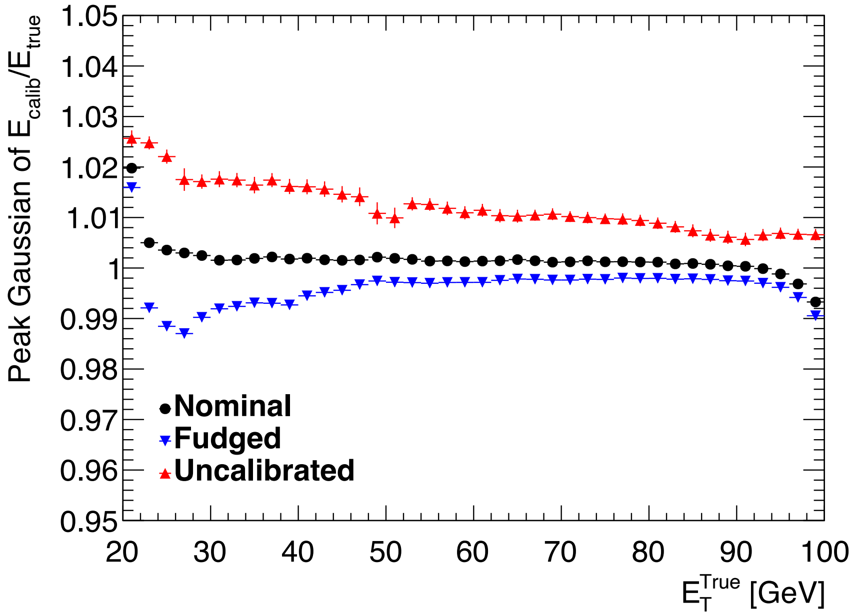

Figure 4.15: Energy scale in the EMEC as a function of E_T^{True} when retaining only $< r^2 >$ as a shower shape variable for the uncalibrated, evaluated on uncorrected Monte Carlo (nominal) and evaluated on corrected Monte Carlo (fudged) cases.

particularly in the 25 GeV E_T region.

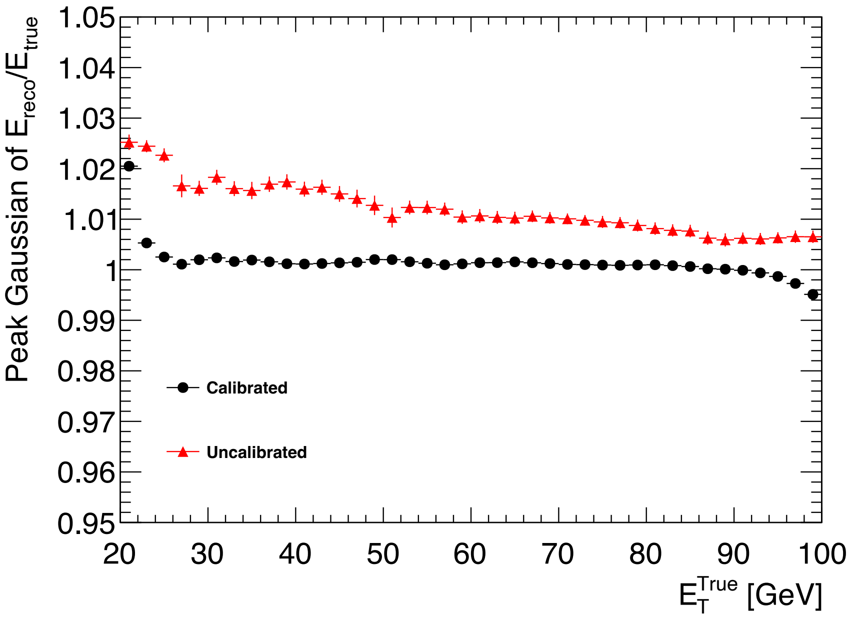

Figure 4.16: Energy scale in the EMEC as a function of E_T^{True} when training and evaluating the BDT performance with corrected Monte Carlo (calibrated). The uncalibrated performance is left unchanged by correcting the shower shapes.

Therefore, a further EMEC BDT was both trained on and evaluated using corrected Monte Carlo. As can be seen in Figure 4.16 this eliminates the non-linearity in the $22 < E_T^{True} < 50$ GeV region and returns to a similar level of performance that can be seen in Figure 4.13b for the case when training and evaluating the BDT performance using uncorrected Monte Carlo.

Overall Single Particle Performance

The final implementation of the MVA is therefore three BDTs covering the EMEC, FCal and transition region respectively with the input variables as defined in Table 4.3. In addition, the EMEC BDT is trained using corrected Monte Carlo in order to negate non-linearities that would have a significant effect when evaluating on data.

As shown in Figure 4.17, the MVA shows significant improvements to both the linearity and resolution across the forward phase space. In the EMEC, an almost perfect energy scale is observed, with variances from unity on the permille level, and resolution improvements between 40% and 50%. Whilst the FCal does not produce a final result as good as the EMEC, the scale is still within 1% of unity and the resolution improves by 30% to 40%. The performance in the transition region is also greatly improved. When considering the combined performance of the entire BDT ensemble, the energy scale

EMEC ($\lvert\eta\rvert < 3.15$)	FCal ($\lvert\eta\rvert > 3.35$)	Transition Region ($3.15 < \lvert\eta\rvert < 3.35$)
E_{raw}	E_{raw}	E_{raw}
η_{cl} ϕ_{cl}	$CentreX/Y/Z$	$CentreX/Y/Z$
$\eta\ Modulus\ \Delta\eta$	$\eta\ Modulus\ \Delta\eta$	$\eta\ Modulus\ \Delta\eta$
$\lfloor\eta/\Delta\eta\rfloor$	$\lfloor\eta/\Delta\eta\rfloor$	$\lfloor\eta/\Delta\eta\rfloor$
μ	μ	μ
npv	npv	npv
$< r^2 >$	$< r^2 >$	–
$E^{Max}_{S_1}/E^{Max}_{S_2}$	$< \rho^2 >$	–
$\phi\ Modulus\ \frac{2\pi}{16}$	λ_{centre}	–

Table 4.3: The input variables used for training the BDTs. N.B μ is defined here as the actual number of interactions per bunch crossing.

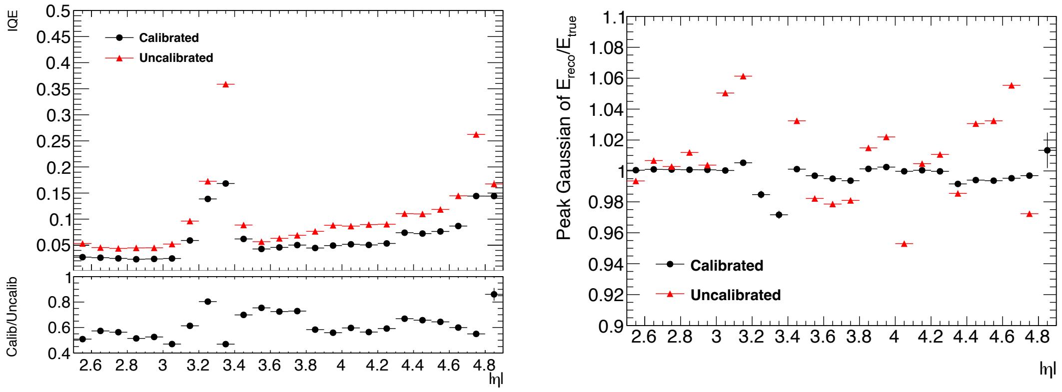

(a) Resolution

(b) Energy scale

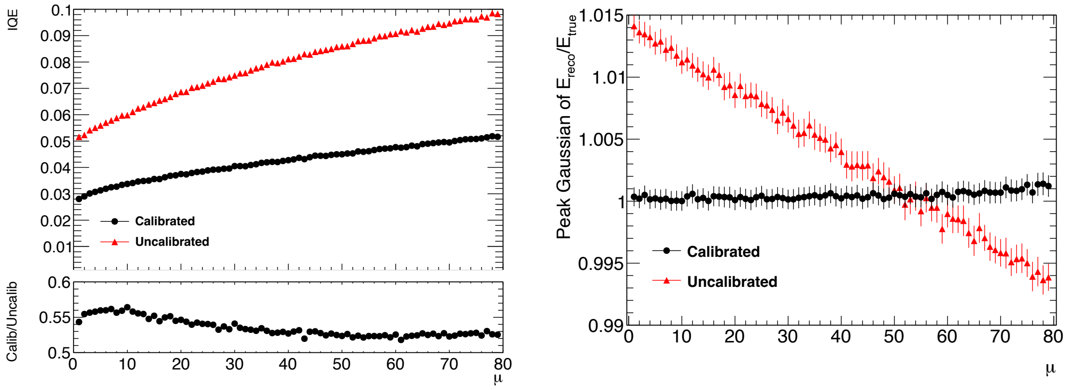

(c) Resolution as a function of pileup.

(d) Energy scale as a function of pileup.

Figure 4.17: The energy scale and resolution both pre- and post calibration for the final implementation.

remains constant as a function of both pileup and transverse energy, thus demonstrating the robustness of the MVA.

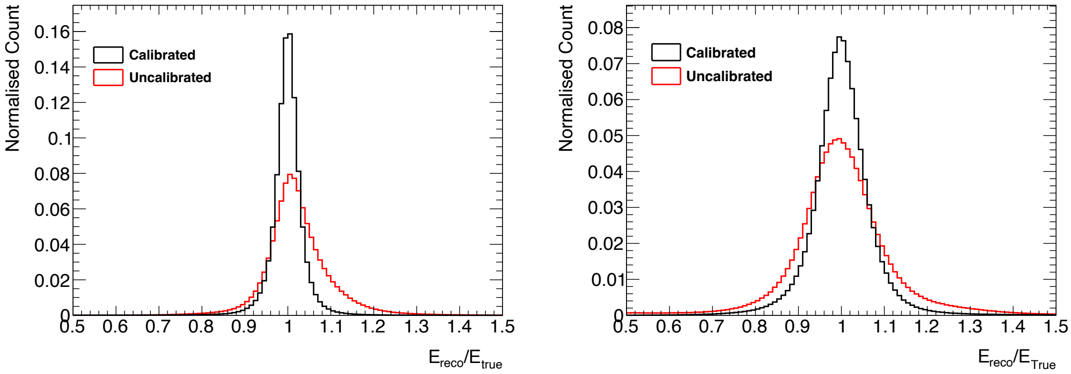

(a) EMEC

(b) FCal

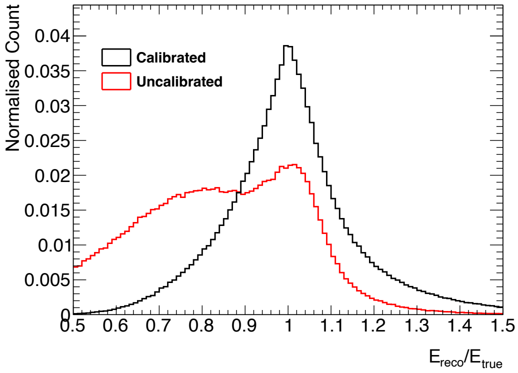

(c) Transition Region

Figure 4.18: The distributions of E_{reco}/E_{true} in each of the three regions of the forward phase space.

In Figure 4.18, the pre- and post-calibration E_{reco}/E_{true} dsitributions from each BDT can be seen. Although the performance differs in each BDT, each BDT delivers clear improvements to the scale and resolution as well as vastly reducing the non-Gaussian nature of each distribution.

Application to $Z \rightarrow e^+e^-$ Decays

To further evaluate the performance of the BDTs, the MVA was applied to Monte Carlo simulated $Z \rightarrow e^+e^-$ decays where one electron is reconstructed centrally and the other in the forward region. Since the MVA was trained using a Monte Carlo sample that populates each region of the forward phase space equally, it is important to test that the MVA will also be able to provide similar improvements when it is applied to forward electrons in a phase space that corresponds to that of an actual physics process.

Before applying the MVA to the Z Monte Carlo, the following selection criteria were applied to reduce the contribution of misidentified particles:

- the invariant mass of the electron pair was between $\pm$ 30 GeV of the reference Z boson mass.
- the reconstructed transverse momentum of the forward electron was greater than 20 GeV.
- the forward electron must pass the loose identification working point.
- $\Delta R < 0.1$ between the true and reconstructed forward electrons.

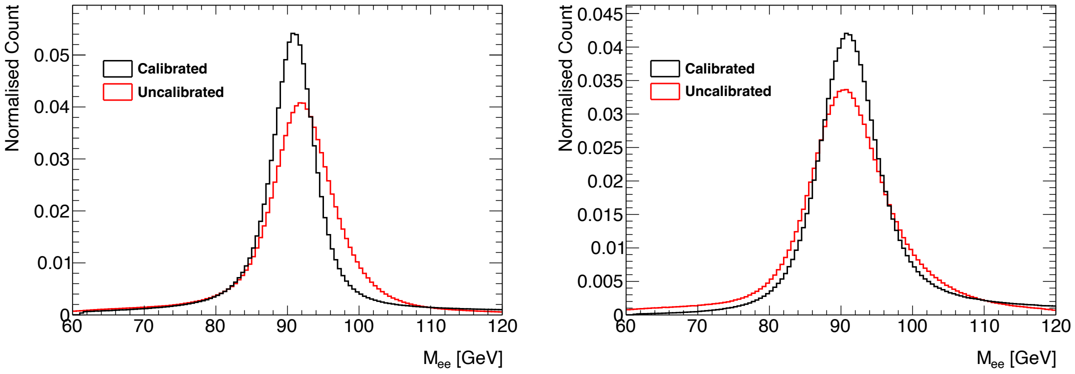

(a) EMEC

(b) FCal

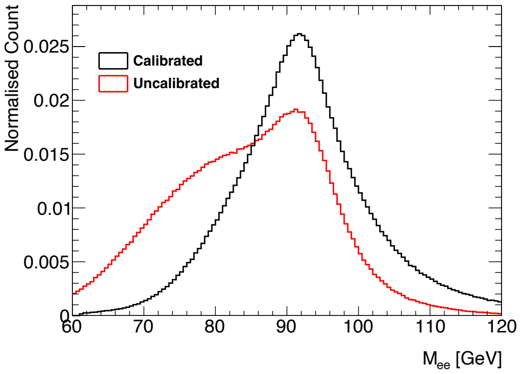

(c) Transition Region

Figure 4.19: The pre- and post-calibration Z boson mass spectrum in the CF channel, shown for three separate regions where the forward electron is reconstructed.

Figure 4.19 shows the shape of the Z boson mass spectra in the EMEC, FCal and transition region both pre- and post-calibration. In each calorimeter and the transition region, the Z peak is shifted towards its true position and the resolution is also improved. The starkest improvement is that seen in the transition region, where a clear Z resonance

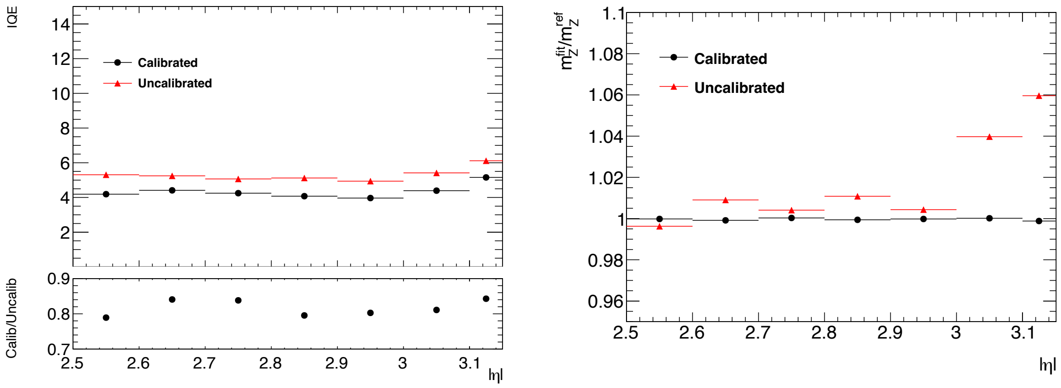

(a) EMEC Resolution (b) EMEC energy scale

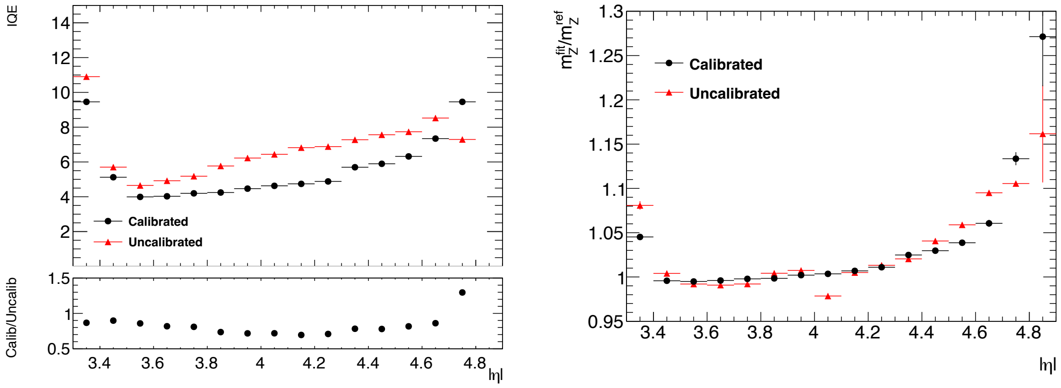

(c) FCal Resolution (d) FCal energy scale

Figure 4.20: The pre- and post-calibration improvements to the Z boson mass resolution and scale as a function of the pseudorapidity. The energy scale plots show the fitted mass of the Z boson divided by the reference value given by the Particle Data Group [72].

emerges from a highly degraded mass spectrum when the MVA calibration is turned on. However, since it is highly unlikely that events in the transition region will be used for physics analysis the performance of the transition region will not be investigated further.

Using the same performance metrics as previously, the resolution is improved by 15 - 20% in the EMEC and 10 - 30% in the FCal. In addition, the energy scale in the EMEC almost perfectly matches the true Z peak and also in the FCal below a pseudorapidity of 4.2. Above this point the energy scale begins to deviate from the true Z mass for both the uncalibrated and calibrated cases. Above a pseudorapidity of 4.7, the MVA calibration breaks down and returns a worse energy scale than the uncalibrated case and so the decision was made to not apply the MVA in this region.

Uniformity Correction

Since the MVA is provided with positional information of forward electrons, it should be able to account for any non-uniformities in material distribution within ATLAS that affect the energy reconstruction. When the average Z mass is plotted against the phi coordinate of the forward electron, the performance of the MVA in correcting for these effects can be evaluated. Figures 4.21, 4.22 and 4.23 show that for each data taking period, the MVA is able to reduce the modulations in the energy response very well in Monte Carlo but a clear modulation remains in the data. Since the modelling of the passive material used in Monte Carlo simulations is only an approximation of that of the real detector, the training of the MVA cannot take into account these modelling inaccuracies. In addition, data taking can also be affected by issues with the high voltage system that is used to translate the calorimeter output to an energy value. These issues are non-existent in Monte Carlo and hence the MVA cannot properly account for them. The combination of the mismodelling of passive material and high voltage issues mean that the performance of the MVA is, correspondingly, suboptimal when applied to data in comparison to Monte Carlo. If the calibration process were to stop here, Figures 4.21, 4.22 and 4.23 could be interpreted as showing that the mass of the Z boson depends on which region of ATLAS is being investigated; something which is entirely unphysical. These modulations must therefore be corrected before forward electrons can be used for physics analysis.

In general, the correction process works by correcting the cell energies in a given η slice towards the median energy response, to account for local variations of the passive material distribution. Specifically, the η bin edges are defined to be:

$$\begin{aligned}\eta \in \{&-4.9, -4.0, -3.35, -3.16, -3.1, -3.0, -2.9, -2.8, -2.7, -2.6, -2.5,\\ &2.5, 2.6, 2.7, 2.8, 2.9, 3.0, 3.1, 3.35, 3.6, 4.0, 4.9\}\end{aligned} \tag{4.10}$$

In practice, for each cell in the slice a Gaussian distribution is fit to the Z mass spectrum of the cell. The mean of this fit, μ_{fit} is then used to derive a correction factor to the forward electron energy, α, which is defined to be:

$$\alpha = \frac{\mu_{fit}^2}{\text{median}(\mu_{fit})^2} - 1 \tag{4.11}$$

where median(μ_{fit}^2) refers to the median value of all fitted means within the η slice. The forward electron energy is then corrected as:

$$E_{corrected} = \frac{E_{uncorrected}}{1+\alpha} \tag{4.12}$$

As can be see in Figures 4.21, 4.22 and 4.23, once the uniformity correction is applied, the size of the modulations are significantly reduced in data and the energy response is, to a good approximation, uniform. Whilst the MVA already provides a mostly uniform energy response in Monte Carlo, it is still beneficial to apply a separate uniformity correction. It is also clear that the uniformity correction is of great importance in the FCal in data, where the non-MVA and uniformity corrected modulations are more than double the size of those in the EMEC due to the larger amounts of passive material in this region of the detector. Whilst the post-uniformity correction modulations in data in the FCal are still approximately 1.5 times larger than those in the EMEC, the resulting uniformity levels in Figures 4.21c, 4.22c and 4.23c are still perfectly adequate. Therefore, no further uniformity corrections are carried out for either the EMEC or the FCal.

15/16 Data	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	330	793	436
MVA + No Uniformity	310	753	402
MVA + Uniformity	103	123	108

(a) Data

15/16 Monte Carlo	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	130	149	113
MVA + No Uniformity	37	76	41
MVA + Uniformity	9	39	18

(b) Monte Carlo

Table 4.4: Tables of the average non-uniformity for each calibration variation in the EMEC, FCal and the combination for both data and Monte Carlo in the 2015/16 period.

The exact amount of non-uniformity present after each calibration step can be quantified by calculating the average non-uniformity. This can be done by fitting a straight line to the average invariant mass versus ϕ distribution and calculating the average absolute

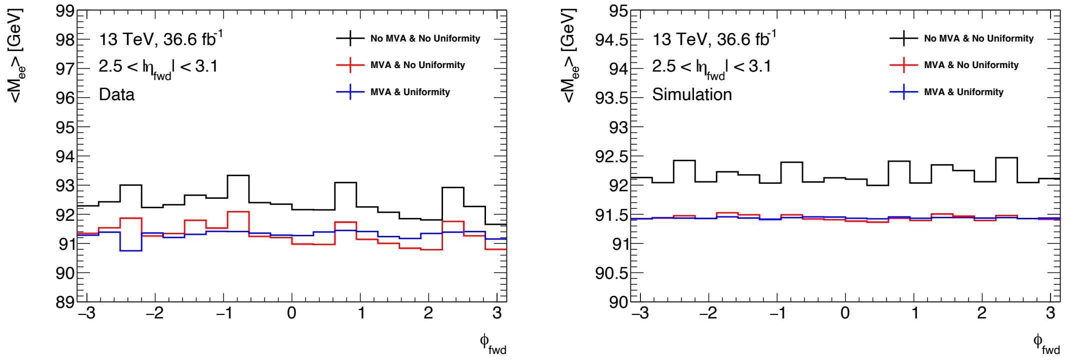

(a) Uniformity in data in the EMEC. (b) Uniformity in Monte Carlo in the EMEC.

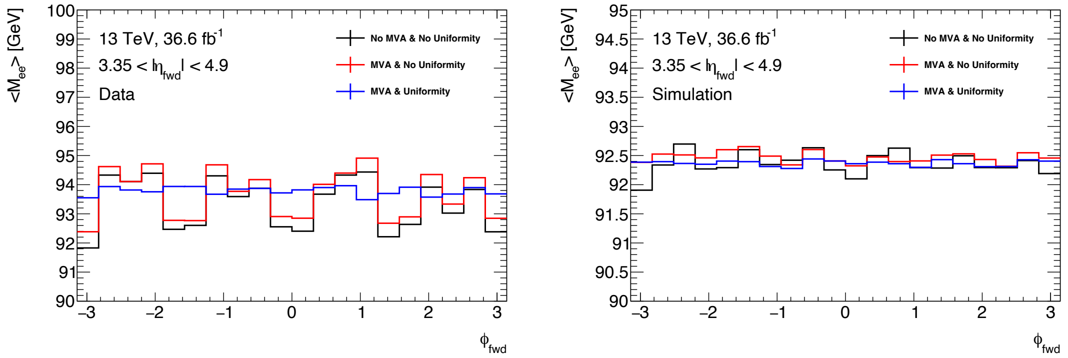

(c) Uniformity in data in the FCal. (d) Uniformity in Monte Carlo in the FCal.

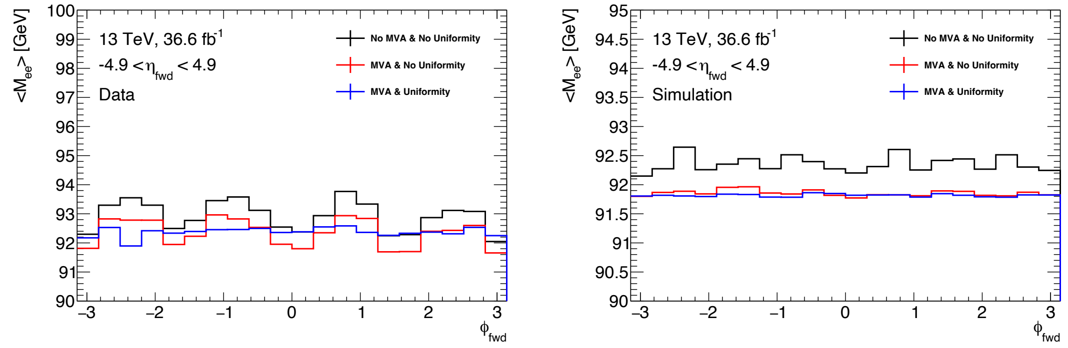

(e) Uniformity in data across all calorimeters. (f) Uniformity in Monte Carlo across all calorimeters.

Figure 4.21: Average invariant mass versus the ϕ coordinate of the forward electron in the EMEC, FCal and combination for the 2015/16 period. The three cases depicted correspond to no calibration, only the MVA calibration is applied and both the MVA and the uniformity correction are applied.

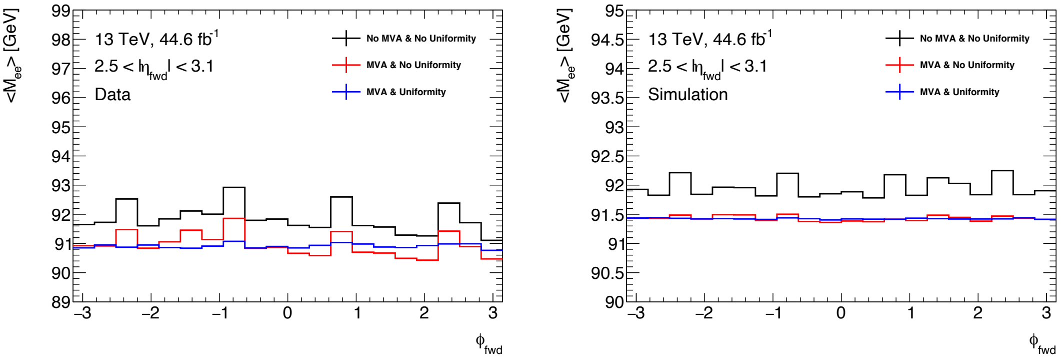

(a) Uniformity in data in the EMEC.

(b) Uniformity in Monte Carlo in the EMEC.

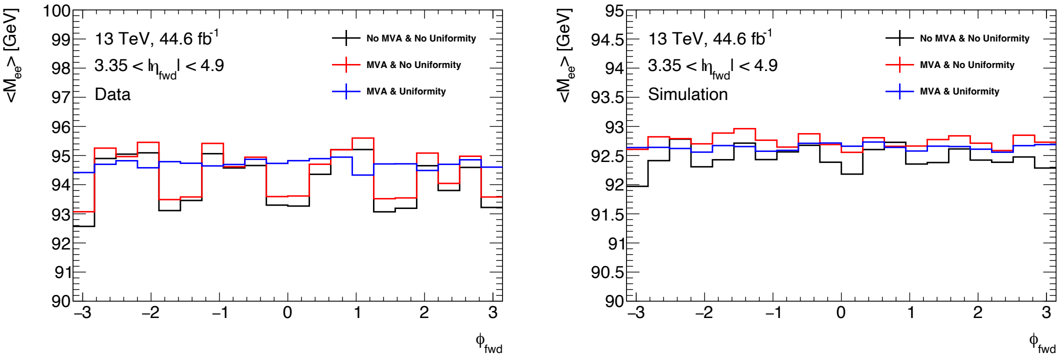

(c) Uniformity in data in the FCal.

(d) Uniformity in Monte Carlo in the FCal.

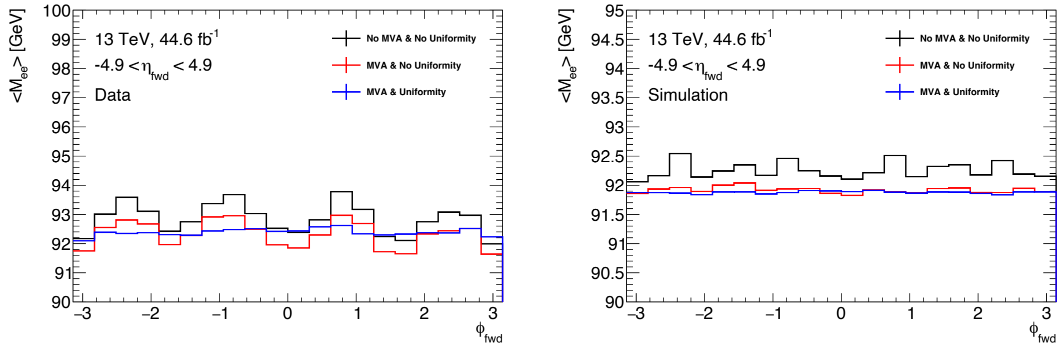

(e) Uniformity in data across all calorimeters

(f) Uniformity in Monte Carlo across all calorimeters

Figure 4.22: Average invariant mass versus the ϕ coordinate of the forward electron in the EMEC, FCal and combination for the 2017 period. The three cases depicted correspond to no calibration, only the MVA calibration is applied and both the MVA and the uniformity correction are applied.

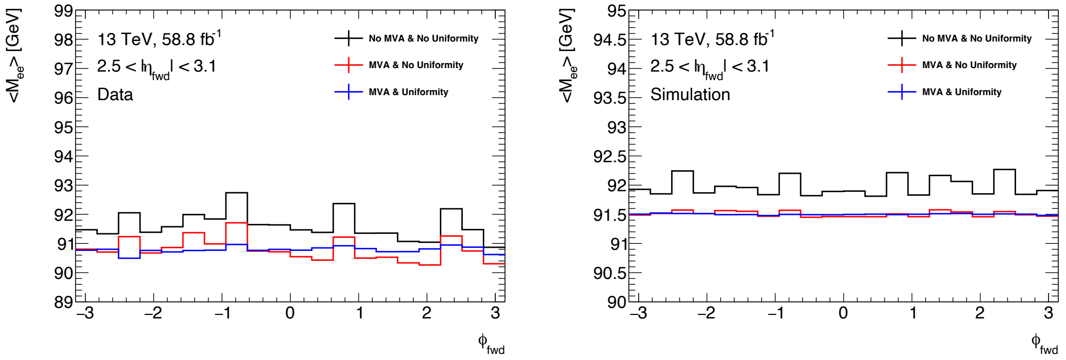

(a) Uniformity in data in the EMEC (b) Uniformity in Monte Carlo in the EMEC.

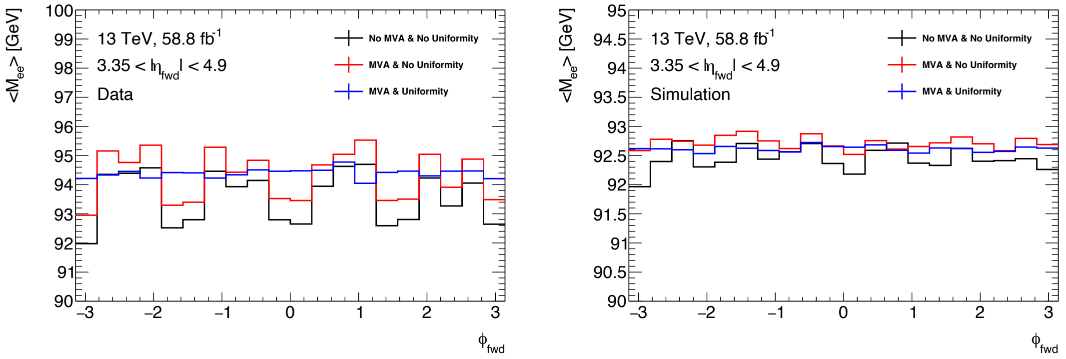

(c) Uniformity in data in the FCal. (d) Uniformity in Monte Carlo in the FCal.

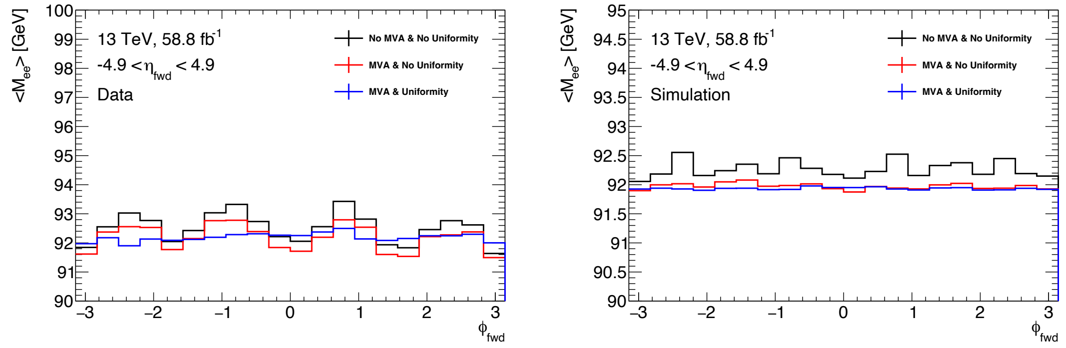

(e) Uniformity in data across all calorimeters. (f) Uniformity in Monte Carlo across all calorimeters

Figure 4.23: Average invariant mass versus the ϕ coordinate of the forward electron in the EMEC, FCal and combination for the 2018 period. The three cases depicted correspond to no calibration, only the MVA calibration is applied and both the MVA and the uniformity correction are applied.

17 Data	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	358	801	431
MVA + No Uniformity	317	761	379
MVA + Uniformity	61	117	91

(a) Data

17 Monte Carlo	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	130	152	116
MVA + No Uniformity	38	91	41
MVA + Uniformity	8	41	16

(b) Monte Carlo

Table 4.5: Tables of the average non-uniformity for each calibration variation in the EMEC, FCal and the combination for both data and Monte Carlo in the 2017 period.

18 Data	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	359	805	407
MVA + No Uniformity	310	764	366
MVA + Uniformity	75	118	110

(a) Data

18 Monte Carlo	Average Non-Uniformity [MeV]		
	EMEC	FCal	Combined
No MVA + No Uniformity	134	155	119
MVA + No Uniformity	40	85	41
MVA + Uniformity	8	38	16

(b) Monte Carlo

Table 4.6: Tables of the average non-uniformity for each calibration variation in the EMEC, FCal and the combination for both data and Monte Carlo in the 2018 period.

difference between the fit and the actual value in each ϕ bin. For each data taking period, these values displayed in Tables 4.4, 4.5 and 4.6 show that for the combination of the EMEC and FCal, the application of both the MVA and uniformity correction result in a decrease in non-uniformity of 75%, 79% and 73% (84%, 86% and 86%) in data (Monte Carlo) for 2015/16, 2017 and 2018 respectively, in comparison to not applying the MVA and uniformity correction. For the full Run 2 period, the average combined non-uniformity (accounting for the different luminosity of each period) was therefore 103 (16.5) MeV in data (Monte Carlo). In comparison to the Z mass of 91.19 GeV [72] i.e the energy scale at which the uniformity is being evaluated at, this means that the average non-uniformity is only a deviation of 0.1% in data and 0.02% in Monte Carlo.

In Situ Calibration

As with the uniformity correction, ideally the MVA would produce an identical Z mass spectrum in both data and Monte Carlo. However, the different energy responses for the two cases result in differing mass spectra. In general, Monte Carlo tends to underestimate the pileup noise which greatly broadens the electron energy resolution in the forward region. In contrast, the peak position in data does not correspond to the true Z peak value. These two major differences must, therefore, be reconciled before the CF channel can be used for physics analysis. This is therefore the task of in situ calibration; to smear the resolution in Monte Carlo and to shift the peak position in data.

In order to carry out these tasks, it is first necessary to build a template that accurately models the data. The template is formed through the convolution of some model function with the signal Monte Carlo, to which the background estimates[1] are added. It is important to note that the purpose of the model is to provide the necessary parameters to shift the data and smear the Monte Carlo and not to simply model either the data or Monte Carlo distributions. The template is then fit to the observed Z mass spectrum in the data, which will optimise the parameters of the model. After performing the fit, the mean value of the model, μ, will represent the difference between the peak position in data and Monte Carlo and can therefore be used as the shifting parameter. Since the peak position in Monte Carlo is much closer to the true Z peak, the shifting parameter is applied to the data rather than to the simulation. The Monte Carlo is then smeared by generating random data points from the model with the mean set to zero. Explicitly, the model is treated as a probability density function from which an x, y point is generated for each event in Monte Carlo. This x value is then used to perform the smearing. Mathematically these processes can be written as:

$$E_{Corrected} = \frac{1}{1+\mu} E_{Uncorrected} \qquad (4.13)$$

$$E_{Corrected} = (1+x)\, E_{Uncorrected} \qquad (4.14)$$

---

[1] Simulated top and electroweak and the data driven multijet estimate which will be rigorously defined in Sections 5.2 and 5.5 respectively.

Where Equation (4.13) describes the shifting of data with the mean of the fit, μ, and Equation (4.14) the smearing of the Monte Carlo with the x-coordinate of the generated data point from the fitted model.

In terms of which model should be used, it is well known that the Z resonance in a particle detector can be modelled by the convolution of the Breit-Wigner and Gaussian distributions. However, from a calibration standpoint it is better to use a model that is agnostic to the physics. Given that the forward region is highly impacted by the effects of pileup and bremsstrahlung, the Z boson cannot be modelled with a pure Gaussian distribution due to the resultant non-Gaussian energy response in the tails of the resonance. Therefore, the model considered is the so-called crystal ball (CB), which consists of a power law lower tail followed by a Gaussian core and upper tail. Explicitly it is defined to be:

$$f(m;\mu,\sigma,\alpha,n) = \begin{cases} A \times (B - \frac{m-\mu}{\sigma})^{-n} & \frac{m-\mu}{\sigma} < -\alpha \\ \exp\left(-\frac{(m-\mu)^2}{\sigma^2}\right) & \frac{m-\mu}{\sigma} > -\alpha \end{cases}$$

where m denotes a point on the invariant mass spectrum of the Z, μ is the mean of the distribution, σ the standard deviation, α the number of Gaussian standard deviations away from the mean that the Gaussian core becomes a power law for the lower tail, with n defining the decay of said tail. The coefficients A and B are then defined as:

$$A = \left(\frac{n}{|\alpha|}\right)^n \exp\left(\frac{-\alpha^2}{2}\right) \tag{4.15}$$

$$B = \frac{n}{|\alpha|} - |\alpha| \tag{4.16}$$

This crystal ball model can be further extended by replacing the Gaussian upper tail with a second power law distribution to give the so-called double sided crystal ball (DSCB) distribution. This DSCB model allows for a better flexibility in modelling the effects of pileup in the lower and upper tails whilst still accounting for the Gaussian broadened nature of the Z resonance at the expense of introducing more parameters that need to be fitted. Explicitly, it is defined to be:

$$f(m;\mu,\sigma,\alpha_L,\alpha_U,n_L,n_U) = \begin{cases} A_L \times (B_L - \frac{m-\mu}{\sigma})^{-n_L} & \frac{m-\mu}{\sigma} < -\alpha_L \\ \exp\left(-\frac{(m-\mu)^2}{\sigma^2}\right) & -\alpha_L < \frac{m-\mu}{\sigma} < \alpha_U \\ A_U \times (B_U - \frac{m-\mu}{\sigma})^{-n_U} & \frac{m-\mu}{\sigma} > \alpha_U \end{cases}$$

with the coefficients A and B replaced with separate values $A_{L,U}$ and $B_{L,U}$ for the lower and upper tails respectively but still defined by Equations (4.15) and (4.16). Correspondingly, the parameters α and n are defined separately for the lower and upper tails as $\alpha_{L,U}$ and $n_{L,U}$ respectively. All other parameters of the DSCB retain the same meaning as for the CB.

With the choice of model selected, the fitting process to derive the shift and smear parameters is then carried out separately for the data taking periods 2015/16, 2017 and 2018. The different filling schemes the LHC used to collide protons in each year resulted in different amounts of pileup and hence a non-uniform electron energy response across the full Run 2 dataset. Within each data taking period the fit is performed in a mass range of 66 - 150 GeV around the Z pole and in bins of η, where the edges are defined to be:

$$\begin{aligned}\eta \in \{&-4.9,-4.6,-4.3,-4.0,-3.35,-3.16,-3.1,-3.0,-2.9,-2.8,\\ &-2.7,-2.6,-2.5,2.5,2.6,2.7,2.8,2.9,3.0,3.1,3.35,3.6,4.0,4.3,4.6,4.9\}\end{aligned} \tag{4.17}$$

The results of the in situ calibration using both the CB and DSCB models can be seen in the EMEC, FCal and inclusively in Figures 4.24, 4.25 and 4.26. The CB model provides a good agreement between the data and expectation in the EMEC in 2015/16 and 2018, with the data to expectation agreement better than ±5% in the Z pole region of 80 - 100 GeV. Whilst the agreement is not flat and displays a clear ‘V’ shape in this region (therefore suggesting that some fine-tuning of the calibration is necessary), the overall size of this deviation from a flat agreement is small enough that it is not a problem. Outside of the pole region, the data to expectation agreement begins to deviate outside of the ±5% bounds but given the much lower amount of statistics in this region, the agreement that is displayed is still acceptable. However, the performance in the FCal is subpar with the CB model, with a large ‘W’ shape displayed in the Z pole region for each data taking period, as well as a significantly worsened agreement in both the lower and upper tails in comparison to the EMEC. Whilst this FCal performance can be attributed to the highly non-Gaussian nature of the upper and lower tails in this calorimeter which the CB cannot model correctly, it is unclear why the CB is unable to properly model the EMEC in 2017. Here a clear misalignment between the data and Monte Carlo distributions can be seen due to the steeply rising data to expectation agreement in the 80 - 100 GeV region, rather than an approximately flat line.

In comparison, the DSCB is able model both the EMEC for every year, with the data to expectation agreement comfortably within ±5% in the 80 - 100 GeV region albeit with the same ‘V’ shape as with the CB model. However, the main advantage of the DSCB is clearly in its modelling of the FCal, where it produces an approximately linear agreement between data and the expectation under the Z pole that varies only by ±2% for each year. The nominal model was therefore chosen to be the DSCB due to its reliability across both the EMEC and FCal and in each data taking period. The fitted value of each parameter for the nominal DSCB model is displayed in Figure 4.27.

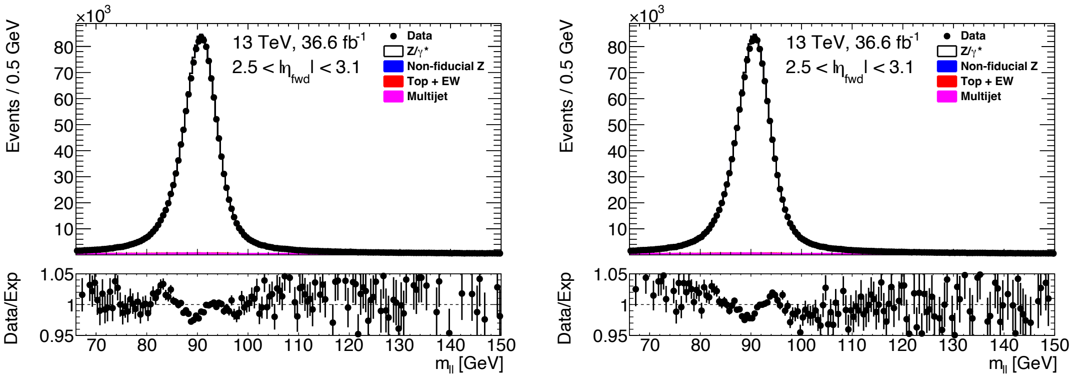

(a) EMEC with crystal ball template.

(b) EMEC with double sided crystal ball template.

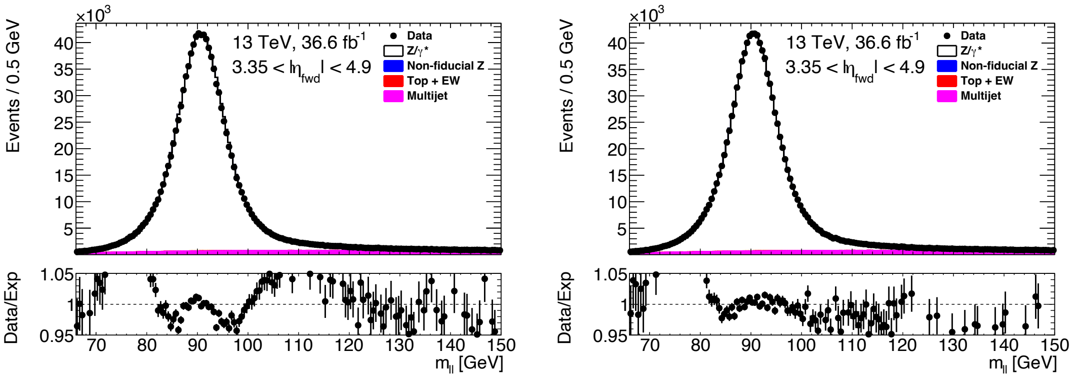

(c) FCal with crystal ball template.

(d) FCal with double sided crystal ball template.

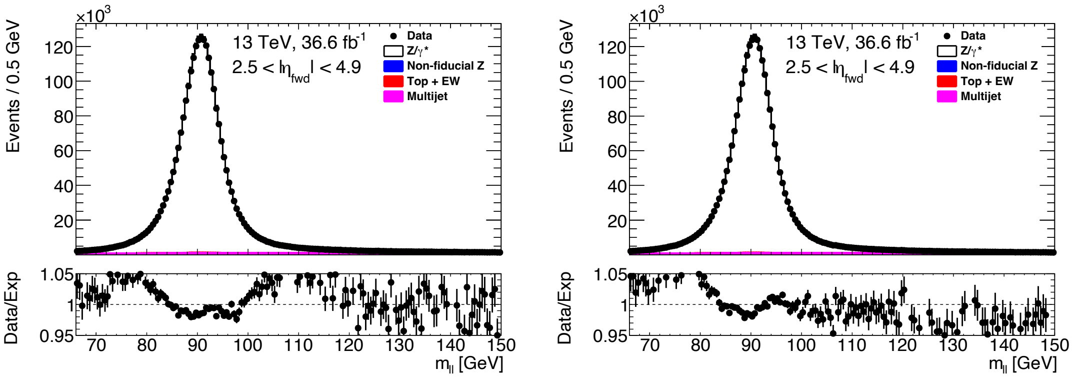

(e) Inclusive with crystal ball template.

(f) Inclusive with double sided crystal ball template.

Figure 4.24: The post-insitu calibration invariant mass distributions in the EMEC, FCal and inclusive combination for data and Monte Carlo using the 2015/16 dataset. The left column shows the result when a crystal ball model is used and the right a double sided crystal ball model. In all plots, the expectation is normalised to the data.

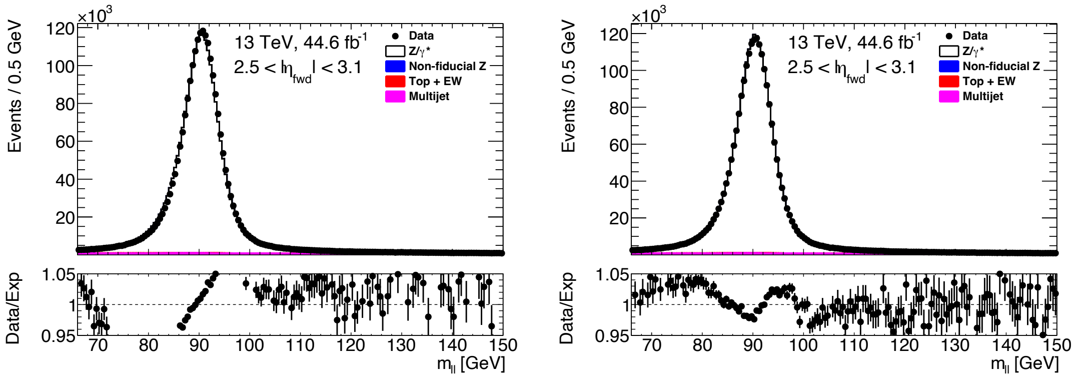

(a) EMEC with crystal ball template.

(b) EMEC with double sided crystal ball template.

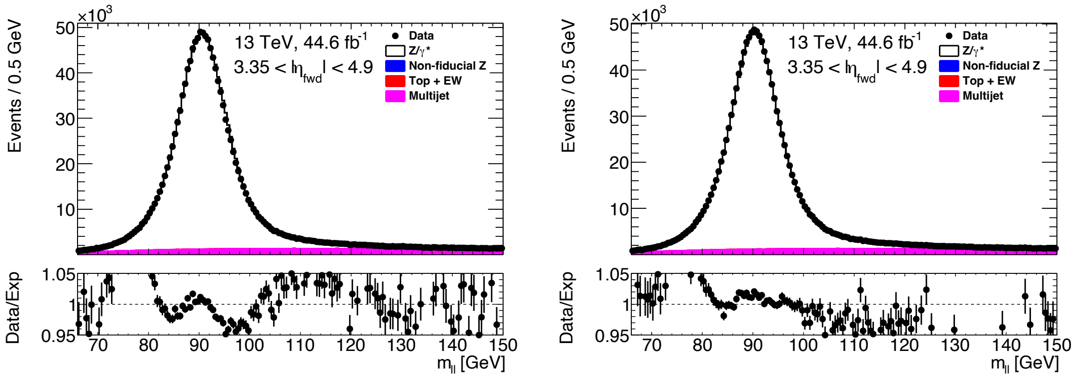

(c) FCal with crystal ball template.

(d) FCal with double sided crystal ball template.

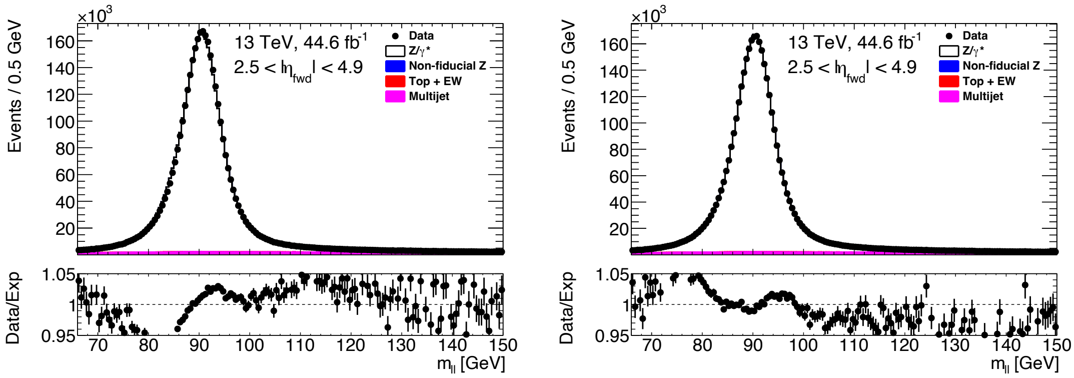

(e) Inclusive with crystal ball template.

(f) Inclusive with double sided crystal ball template.

Figure 4.25: The post-insitu calibration invariant mass distributions in the EMEC, FCal and inclusive combination for data and Monte Carlo using the 2017 dataset. The left column shows the result when a crystal ball model is used and the right a double sided crystal ball model. In all plots, the expectation is normalised to the data.

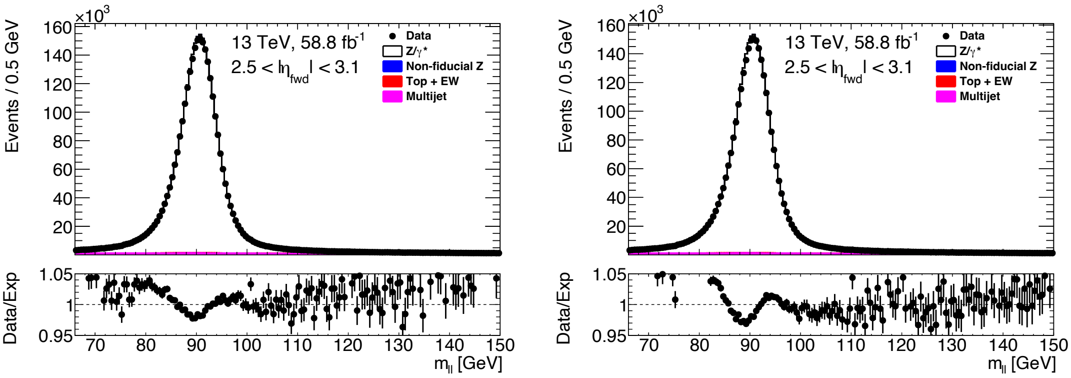

(a) EMEC with crystal ball template.

(b) EMEC with double sided crystal ball template.

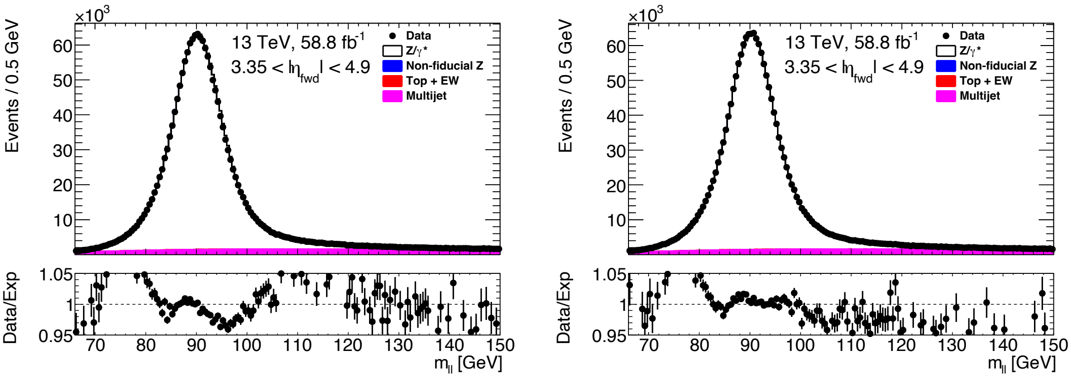

(c) FCal with crystal ball template.

(d) FCal with double sided crystal ball template.

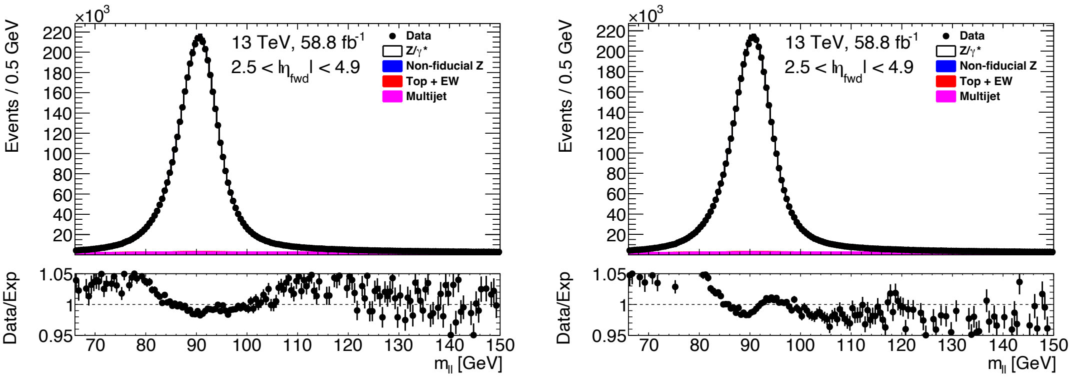

(e) Inclusive with crystal ball template.

(f) Inclusive with double sided crystal ball template.

Figure 4.26: The post-insitu calibration invariant mass distributions in the EMEC, FCal and inclusive combination for data and Monte Carlo using the 2018 dataset. The left column shows the result when a crystal ball model is used and the right a double sided crystal ball model. In all plots, the expectation is normalised to the data.

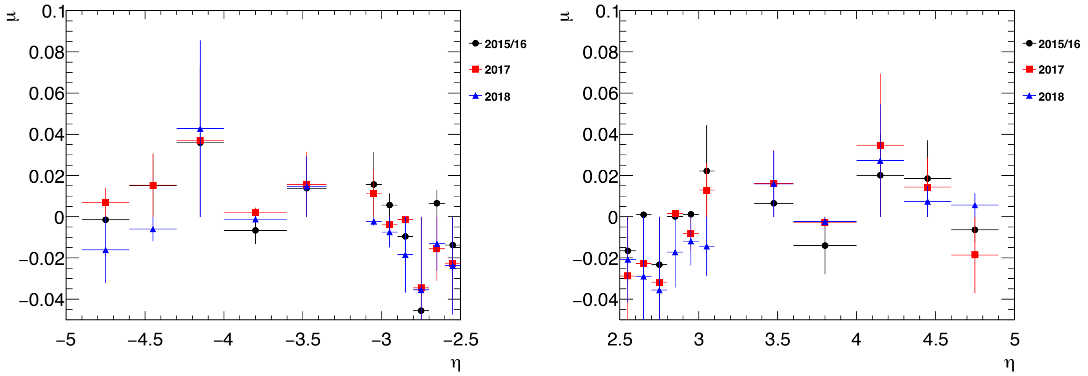

(a) Fitted μ for $\eta < 0$.

(b) Fitted μ for $\eta > 0$.

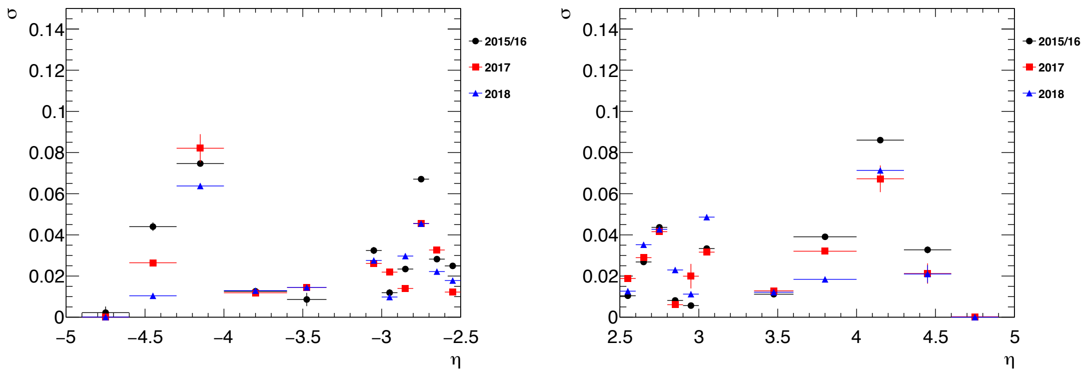

(c) Fitted σ for $\eta < 0$.

(d) Fitted σ for $\eta > 0$.

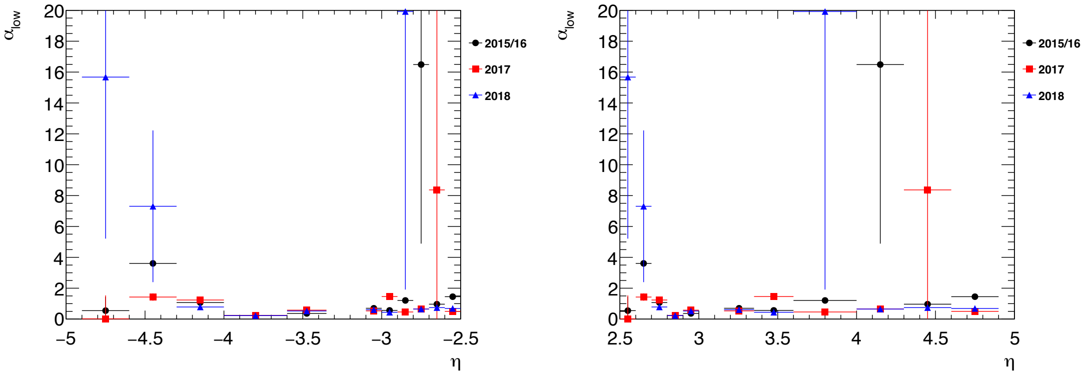

(e) Fitted α_L for $\eta < 0$.

(f) Fitted α_L for $\eta > 0$.

Figure 4.27: Fitted parameters of the double sided crystal ball model for each data taking period as a function of both negative (left column) and positive (right column) pseudorapidity.

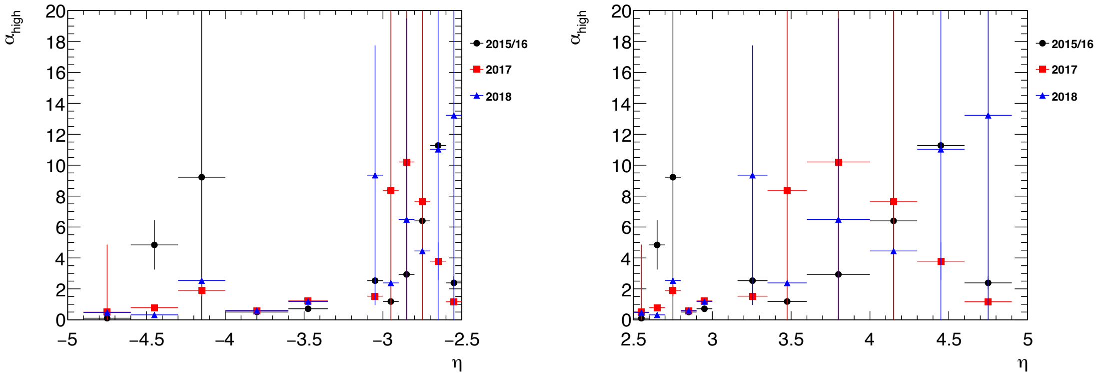

(g) Fitted α_U for $\eta < 0$.

(h) Fitted α_U for $\eta > 0$.

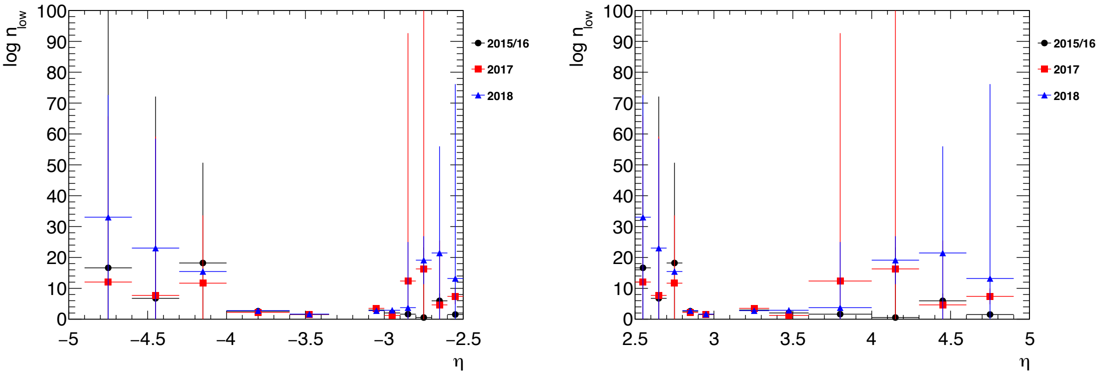

(i) Fitted $\log n_L$ for $\eta < 0$.

(j) Fitted $\log n_L$ for $\eta > 0$.

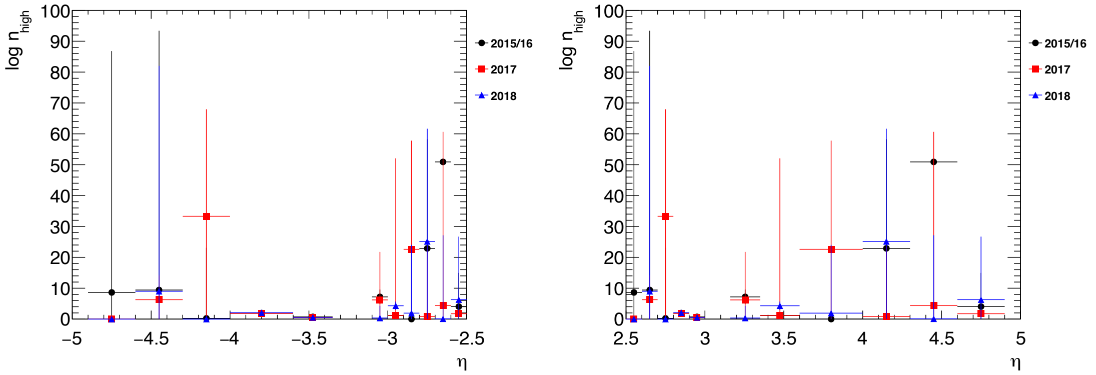

(k) Fitted $\log n_U$ for $\eta < 0$.

(l) Fitted $\log n_U$ for $\eta > 0$.

Figure 4.27: (continued) Fitted parameters of the double sided crystal ball model for each data taking period as a function of both negative (left column) and positive (right column) pseudorapidity.

From Figures 4.24, 4.25 and 4.26 it is clear that the choice of model has a significant effect on the post-calibration Z mass spectra. Many of the fitted parameters in Figure 4.27 also display large error bars, which if varied can also affect the shapes of the post-calibration mass spectra. Therefore, the CB model with nominal parameter values and the eigenvector decomposition of each DSCB parameter variation are treated as sources of systematic uncertainty to the in situ calibration.

The effect of these systematic variations on the invariant mass spectrum CF can be see in Figure 4.28, where in the 80 - 100 GeV invariant mass range the variations are of the level of 10%. Outside of this range, however, the variations quickly increase to 40% relative uncertainty in the lower tail for 2017 and 2018 (25% in 2015/16) and peak at approximately 100% in 2017 and 2018 (95% in 2015/16). Unsurprisingly, the calibration systematics are completely dominated by the variations of the DSCB model parameters, with the difference between the CB and DSCB models contributing a much more modest 5% maximum deviation from the nominal. This means that whilst the forward electrons are now calibrated and can be used for physics analysis, the large calibration uncertainties will most likely be the leading source of systematic uncertainty. These uncertainties could, however, be reduced in a future calibration version by making several adjustments to the calibration procedure:

- reducing the number of calibration bins of the FCal to overcome lower statistics and highly broadened lineshapes at high η.
- changing the fit range of the α and n parameters so that the range of values they span is decreased and to better ensure that the Gaussian to power law transition remains within the invariant mass range of the fit.
- fixing the value of certain fit parameters to reduce the total number of free parameters that must be fitted and simplify the model.

If these adjustments are fully investigated and optimised, then the parameter variations may be reduced to the level of the calibration model variation. If this is achieved, then it may also be worthwhile to introduce a further systematic uncertainty on the calibration mass range, since this is not included in Figure 4.28. Since the current calibration range of 66 - 150 GeV places a large emphasis on the tails of the Z resonance, using more restricted fit ranges that focus more on the Z boson peak would provide a robust test of the calibration model outside of its fit range.

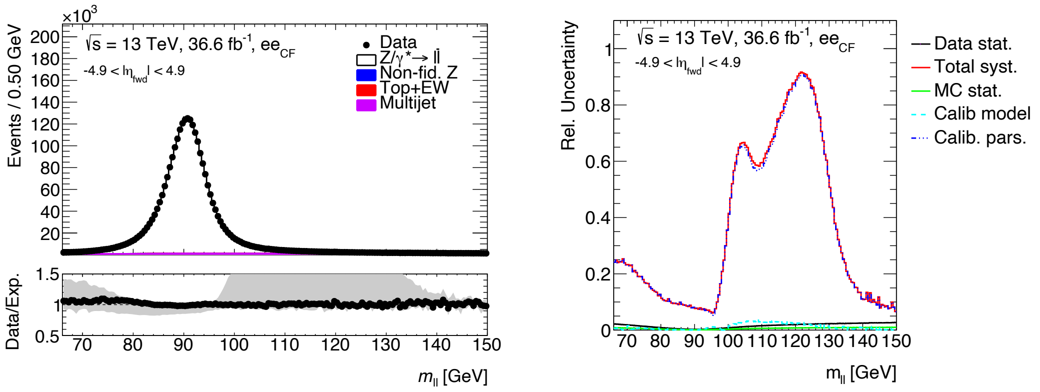

(a) Mass spectrum in 2015/16 (b) Breakdown of uncertainty in 2015/16

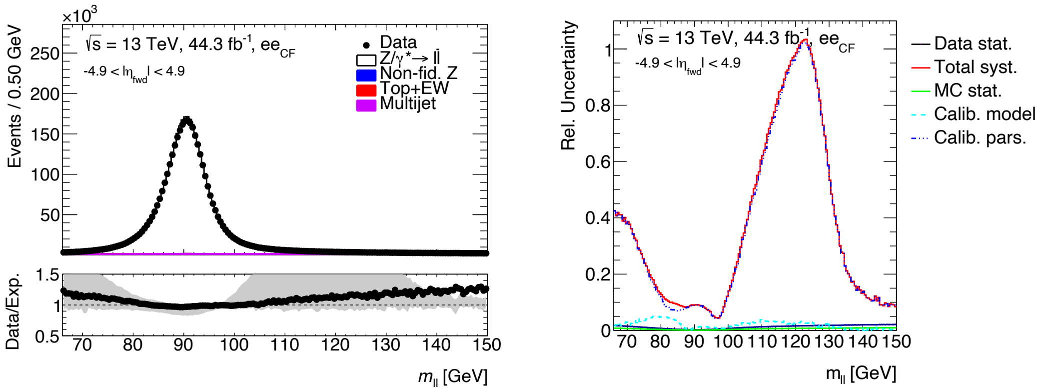

(c) Mass spectrum in 2017 (d) Breakdown of uncertainty in 2017

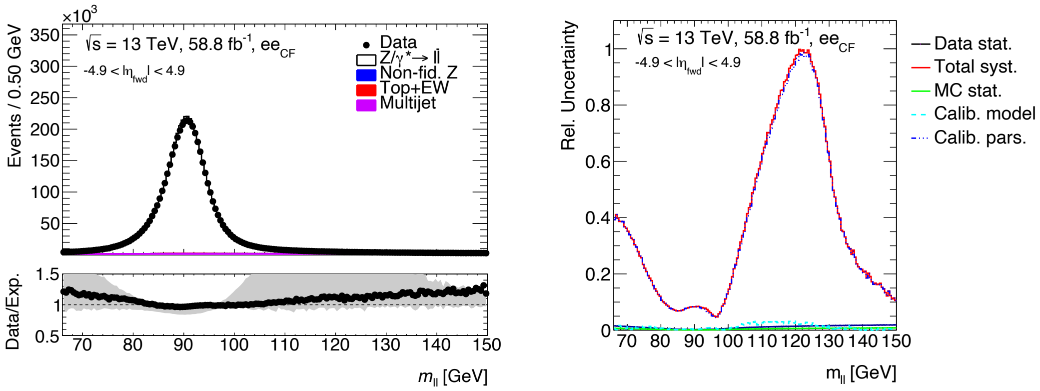

(e) Mass spectrum in 2018 (f) Breakdown of uncertainty in 2018

Figure 4.28: Pseudorapidity-inclusive Z mass spectrum for each data taking period (left column) and the corresponding symmetrised breakdown of systematic uncertainty (right column). Underneath the mass spectra, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from the calibration variations.

4.5 Muon Performance

The relatively long lifetime of the muon and, within the Standard Model, its single possible decay mode to electrons mean that it is unlikely muons will decay within ATLAS. The calorimeter is also ineffective at stopping muons and measuring their total energies since the muon can also only interact electroweakly and muon bremsstrahlung is heavily suppressed. This means that muons will only leave small energy deposits in both the electromagnetic and hadronic calorimeters. Therefore, in order to reconstruct and use muons for physics purposes, all of the necessary information must come from the tracks they leave.

4.5.1 Muon Reconstruction

With muon interactions in ATLAS spanning the Inner Detector, calorimeters and Muon Spectrometer (MS), there are multiple methods used to reconstruct candidate muons, each targeting a different type; combined, inside-out combined, muon-spectrometer extrapolated, segment-tagged and calorimeter-tagged muons [9]. Combined muons are formed by matching muon spectrometer tracks to tracks in the inner detector that also account for the small amounts of energy lost in the calorimeter. Conversely, inside out muons are first identified in the inner detector and their tracks are then extrapolated to the MS and if corresponding hits are found, then a full track hit will be performed. With the extended η coverage of the MS in comparison to the ID, it is possible for tracks to be identified in the MS that cannot be matched to any in the ID. These muons will have their path extrapolated from the MS towards interaction point and then classed as a muon-spectrometer extrapolated muon. If one segment in the MS is activated but does not form a full track, and an ID track can be extrapolated to this segment within a set of matching criteria, a muon will still be reconstructed and classified as segment-tagged. It is also possible to reconstruct calorimeter-tagged muons that leave no track in the MS using ID and calorimeter information. If an ID track is matched to the calorimeter signal of a minimum ionising particle, a muon will be reconstructed, whose properties are taken from the ID track. The combination of all five methods means that the chances of a muon not being reconstructed are very low, with an approximate 99% reconstruction efficiency achieved when applying the loose working point in the p_T range 20 - 100 GeV [9].

In order to differentiate between muons produced from the hard scattering and those produced from hadron decays, pileup interactions or cosmic rays, selection criteria are also used to associate reconstructed tracks to identified primary vertices. The transverse impact parameter, d_0 is defined to be the shortest distance in the transverse plane between the beamline and the point of closest approach of the muon track. The significance of the transverse impact parameter, $|d_0|/\sigma(d_0)$, is then required to be less than three to ensure a high degree of certainty that the muon track is close to the beam. A corresponding longitudinal impact parameter, z_0, is defined as the z coordinate at the point of closest approach of the primary vertex to the beamline. The shortest distance from the muon track to the primary vertex is then $z_0 \times \sin\theta$, where θ is the polar coordinate of the muon. In order to ensure the muon originated from the primary vertex, this distance is required to be $|z_0|\sin\theta < 0.5$ mm [9].

4.5.2 Muon Identification

Similar to electrons, it is also important to be able to classify the certainty with which a reconstructed muon is actually a real muon. However, unlike electron identification, which uses the likelihood method, muon identification primarily relies on a cut based approach [9]. This means that for a muon to pass a certain working point, it must pass a list of criteria rather than having a likelihood value larger than some threshold. The full set of muon identification working points and their respective criteria can be found in Reference [9].

4.5.3 Muon Calibration

Since the properties of muons within ATLAS can only be deduced from the tracks they leave behind, an accurate description of these tracks is imperative in order to measure the muon itself to a high degree of accuracy. In order to achieve this, it is necessary to make sure that any misalignments in both the Inner Detector and Muon Spectrometer are correctly accounted for.

As outlined by the ATLAS experiment [8], the alignment of the ID is influenced by movements of the individual detector subsystems during data taking that affect the reconstructed track positions. These movements are caused temperature fluctuations which in the short term correspond to the electrical load within ATLAS and longer term to the cycling of the power and cooling systems as well as the ramping of the magnetic fields. In addition, correlated movements of multiple detector components can also affect the momentum measurement of muons. In the barrel, the detector movements are typically of the order $\mathcal{O}(1\mu\text{m})$, increasing to $\mathcal{O}(10\mu\text{m})$ in the endcap. The detector subsystem movements are quantified by considering how the subsystem as a whole moves as well as considering the more granular behaviour of barrel and endcap modules and individual layers in the Pixel and SCT. Performing these studies show that the detector movements are typically of the order $\mathcal{O}(1\mu\text{m})$, increasing to $\mathcal{O}(10\mu\text{m})$ in the endcap. The behaviour of these movements as a function of time, temperature and instantaneous luminosity is well understood and hence corrected for in the track reconstruction algorithm. The detector components may also be deformed such that they are radially expanded or rotated with respect to the nominal position. Respectively, these deformations are known as length scale and sagitta biases with both affecting the measured momentum of a particle. For Run 2, these biases were found to be relatively small at 0.9×10^{-3} and 0.1 TeV^{-1} respectively [8].

In the MS, this is handled by the combination of an optical alignment system fitted into the spectrometer, that measures the positions of the individual components relative to each other, as well as to defined reference objects [43]. This allows for any thermal expansion or deformations of the system to be noticed and corrected for. By combining the information of the optical system with track information from special calibration runs where the toroid magnets are switched off, the design resolution of the MS can be achieved, namely that it is mostly $\Delta p_T/p_T = 4\%$ and rising to 10% at $p_T = 1$ TeV [43].

After performing the detector alignment procedures for both the ID and the MS, the remaining residual effects can still produce a non-negligible bias in the resultant muon momentum measurement [11]. To account for this, further corrections are therefore derived in a 48 × 48 $\eta - \phi$ grid which is granular enough to account for local residual effects but remains insensitive to statistical fluctuations. In data these corrections are derived using $Z \rightarrow \mu^+\mu^-$ decays, whereas the simulation also uses $J/\psi \rightarrow \mu^+\mu^-$ decays. The corrections are then validated using both $Z \rightarrow \mu^+\mu^-$ and $J/\psi \rightarrow \mu^+\mu^-$ channels as well as using $\Upsilon \rightarrow \mu^+\mu^-$ decays which are have not been used for calibration purposes and are therefore statistically independent. Unlike electron calibration, when correcting for simulated muons, the derived corrections will be used to update the simulation process itself which results in a better data to Monte Carlo agreement. After applying the corrections, the muon resolution is found to range from 1.8% - 3.0% and 2.3% - 3.4% with rising absolute pseudorapidity for J/ψ and Z decays respectively [11].

Chapter 5

Four-fold Drell-Yan Cross Section Measurement at 13 TeV

In order to measure the angular coefficients of the Drell-Yan cross section introduced in Chapter 2 a subset of the ATLAS full Run 2 dataset and a set of Monte Carlo simulated signal and background events are necessary. From these datasets a signal region can then be defined where the measurement will occur. After verification that the composition of the signal region is sensible the angular coefficients can be measured through the use of a profile likelihood fit. This profile likelihood fit will also consider the effects of numerous sources of systematic uncertainty in order to provide an accurate and precise value for the total uncertainty in each parameter. This chapter will therefore describe how the necessary datasets were chosen, how the signal region was defined, how the angular coefficients can be measured and finally a discussion of the results of the extraction.

5.1 Data

The scale of the 140 fb^{-1} ATLAS full Run 2 dataset means that there is an abundance of statistics available for analysis and thus should allow for the most precise measurement of the Drell-Yan cross section that has ever been achieved. However, since candidate $Z \rightarrow l^{+}l^{-}$ decays are only a subset of the overall dataset it would be an inefficient use of time to process every event that was recorded by ATLAS during Run 2 operations. Instead this analysis uses Derived Analysis Object Data (DAOD) files with the ATLAS Standard Model Derivation 4 (STDM4) i.e the analysis considers events that should only contain electrons or muons in the final state.

5.2 Simulation

In order to evaluate the effect of systematic uncertainties and ensure that the analysis has been carried out correctly, it is also necessary to have a simulated Monte Carlo dataset to accompany the data. This dataset should include both signal samples, which model the physics process of interest (in this case $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$) and also background samples, where the physics processes are not being studied directly, but whose final states can mimic those of the signal events. All of the Monte Carlo samples used for this analysis were produced centrally by the ATLAS collaboration, using the same STDM4 DAOD file format as with the data.

Signal Processes

Whilst the process of interest for this analysis is $pp \to Z \to l^+l^-$, the contribution from $pp \to \gamma^*$, where γ^* is an off-shell photon, cannot be ignored. The simulated signal, $Z/\gamma^* \to l^+l^-$, therefore encompasses both of these processes, with the truth level simulation provided by POWHEGBOX V1 [71, 59, 18, 75] using the CT10 NLO PDF [68] which is calculated at NLO in α_S. However, as previously discussed in Chapter 2, the simulation of an event in ATLAS follows several steps in order to properly simulate the full collision process. Therefore, the parton shower and hadronisation and underlying event in the detector are then simulated using PYTHIA version 8.186 [84]. The predictions from PYTHIA are then tuned using ATLAS data, with the parameter AZNLO [37], to improve the agreement of the simulation with the observed data. For the parton shower, the PDF set used was CTEQ6L1 was used [79]. For the simulation of final state radiation from the vector boson, the generator PHOTOS v3.52 was used [51]. Finally, the decays of any bottom or charm hadrons were handled by the EVTGEN generator v1.2.0 [80].

Background Processes

Since the detector signature of the signal processes is simply two leptons, there are many different background processes that have the same detector signature or can emulate it. Some of these backgrounds are reducible, in the sense that they can mainly be excluded from the signal region through the imposition of various kinematic and identification cuts. Other sources of background are, however, irreducible as they cannot be removed in significant quantities through the application of tighter cut criteria. For the Drell-Yan process, the reducible backgrounds are typically caused by low-energy jets being falsely reconstructed and classified as a lepton, whereas the irreducible backgrounds emerge from physics processes that produce real leptons but are falsely attributed as being from a Z decay. For this analysis, the main sources of irreducible background are non-fiducial Z events, electroweak processes. top quark processes, and photon induced processes.

In order to measure the Drell-Yan cross section in a region of the highest signal purity possible, a series of selection cuts are applied to the data and simulated Monte Carlo events as will be outlined in Section 5.3.

This means that the signal region does not correspond to the full phase space but instead to a fiducial volume that represents the phase space that can be accessed after the application of event selection criteria. Since the reconstructed lepton properties can differ from their truth level counterparts, it is also possible for events to migrate into the fiducial region that, in actuality, lie outside of it. Since they do not correspond to the signal region defined by the selection criteria, these 'non-fiducial' events are classed as a background. The amount of non-fiducial Z events in the signal region is estimated using signal Monte Carlo, by comparing the truth level properties of the decay leptons to their reconstructed values. Since signal region contamination by non-fiducial Z events is caused by the migration of events into the fiducial region its effects are concentrated at the edges of the fiducial phase space, becoming negligible directly under the Z peak.

Although the angular decomposition of the Drell-Yan cross section does hold for the τ lepton in theory, measuring this experimentally would lead to a much less precise result in comparison to the electron/muon final state. Since the τ can, however, decay leptonically, its contribution to the electron and muon signal regions must be included. Like the τ, diboson processes (WW, ZZ and WZ) can also produce real leptons or quarks that resemble a lepton in the detector. These final states can be falsely reconstructed as the decay of a single Z boson and so are treated as sources of background. For both the $Z \to \tau^+\tau^-$ and diboson processes, their contributions to the signal region are once again estimated through Monte Carlo simulation using the same generators and PDF sets as for the signal processes.

Alongside diboson events, fake Z detector signatures can also arise from processes involving a the top and antitop quarks. In the Standard Model, the top quark decays 99% of the time to a bottom quark via the emission of a W boson, meaning that a $t\bar{t}$ pair can act like a WW event in the detector. Additionally the top can be produced alongside a single W boson which can give the necessary lepton combinations to mimic the signal region. For processes involving the top quark, the combination of POWHEG-BOX and PYTHIA is again used to generate the Monte Carlo simulation but with the NNPDF23LO PDF [23] set used.

A much rarer source of background than the aforementioned, is that of lepton pair production from two photons. At Born level in QED, it is not possible for two photons to interact. However, through box diagrams it is possible for two photons to interact and produce a dilepton pair. Due to the lack of Born level photon-photon interactions, these processes are heavily suppressed and therefore do not have a large contribution to the signal region. However, for a precision analysis their effect must still be accounted for and so all possible photon-photon processes are described by three Monte Carlo samples using the LPair generator [26], HERWIG 7 [27, 31] with the MMHT2014lo68cl PDF set [68] and also PYTHIA 8 with the NNPDF23_QED PDF set [24].

5.3 Event Selection

In order to measure the Drell-Yan cross section as precisely as possible, the event selection must be appropriately chosen such that the signal region is as pure as possible with minimal amounts of background contamination. With the ATLAS detector, it is possible to reconstruct $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ events in three separate channels:

- Z decays to two central electrons (eeCC or CC)
- Z decays to one central and one forward electron (eeCF or CF)
- Z decays to two central muons (mmCC or MM)

For the purpose of this analysis only the CC and MM channels will be used to measure the Drell-Yan cross section but the CF channel could, in principle, also be used. The CF channel will, however, be utilised in a study on the extraction of the weak mixing angle in Chapter 6, where its event selection process will be outlined. Therefore, since there are two analysis channels to be studied - CC and MM - the event selection process is carried out separately for each channel in order to give the best possible response.

5.3.1 Triggers

The first stage in the event selection process is the application of triggers to give an initial dataset of candidate $Z \to l^+l^-$ events. For this analysis, the initial dataset therefore comprises events with either one or two possible lepton candidates through the application of single lepton and dilepton triggers respectively. Each trigger has identification and transverse momentum thresholds that must be met by the candidate lepton in order for the overall event to be recorded. Isolation requirements may also be applied by a trigger but this is not the case for every trigger that is used. For an event to be considered for the analysis, it only needs to pass one type of trigger.

For electrons the dilepton trigger requires that the two lepton candidates must pass loose identification and have a transverse momentum of at least 12 GeV for 2015 data and at least 17 GeV 2016 - 2018. Due to the limited output bandwidth from the HLT, it would not be possible to use these same thresholds for the single lepton trigger, hence the thresholds for the single lepton trigger are tighter. Correspondingly a candidate lepton must have a p_T of at least 24 (26) GeV and pass medium (tight) identification in 2015 (2016 - 2018).

In comparison to electrons, the detector signature of a muon is much more distinct and less likely to be imitated by a jet. Therefore, in order to achieve as high a trigger efficiency as possible, the transverse momentum threshold is lowered whilst still maintaining a high fake rejection efficiency. For the dilepton triggers this means that a candidate muon must have 10 (14) GeV in 2015 (2016 - 2018) or pass an asymmetric threshold of 18 (22) GeV and 8 GeV for the leading and subleading muon candidate respectively in 2015 (2016 - 2018). Similar to the electron case, the requirements on the single lepton trigger are tightened to 20 (26) GeV and passing loose (medium) identification in 2015 (2016 - 2018).

No matter the analysis channel, these trigger requirements are applied to both data and to Monte Carlo. Since the behaviour of the simulation differs from the data, any difference between the trigger efficiencies is corrected for through the application of a trigger scale factor as outlined in Section 4.3.

5.3.2 Data Quality

During data taking, the performance of the detector is constantly monitored. If any subsystem or part thereof is malfunctioning or the subsystem performance is suboptimal, this will be noted and filtered through to a Good Runs List (GRL) which contain all the individual lumiblocks where different detector subsystems were working. For this analysis, the data is taken from the standard GRLs with the additional requirement that the tracking systems and calorimeters were in an error free state. In addition, at least one primary vertex must also be reconstructed in the tracker for the event to be considered in the analysis.

5.3.3 Analysis Selections

The mentioned triggers, data quality and derivation selections that are applied to all ATLAS data and Monte Carlo simulation, provide a set of events which most likely contain leptons. However, this set of events does not necessarily comprise a good signal region, since there will still be large amounts of fakes within the selection and contributions from $\gamma^* \to l^+l^-$ that are not wanted. Therefore, the event selection process must continue at analysis level in order to measure the Drell-Yan cross section in a high purity region of $Z \to l^+l^-$ events, with minimal fakes.

Before the event sample can be purified, the set of leptons which are to be studied must first be established. Since the application of the trigger means that only at least one particle (or two for the dilepton triggers) matched the necessary criteria for the event to be recorded, there may be more than two particles present in the recorded event. When the trigger criteria are met, all detector activity is stored, including objects that did not pass the trigger requirements. Any surplus particles in an event are removed from the analysis by considering only the two leptons with the largest transverse momenta.

The effects of the trigger must also be taken into account when applying kinematic cuts to the leptons as well. Since the trigger efficiency changes greatly in the region near the threshold value and plateaus with increasing transverse momentum, a minimum p_T value is applied for an electron to be considered for analysis. This minimum value is, therefore, chosen to be sufficiently above the threshold to ensure that the trigger efficiency is in the plateau region. For the CC channel, this means that an electron must have a minimum p_T value of 22 GeV. Since the fake rate in the muon channel is much lower than for the electron channels, the minimum p_T value is also relaxed to 20 GeV.

In addition to the selection on the lepton transverse momentum, leptons that are reconstructed within certain parts of the detector are excluded from the analysis as well. For electrons this means cutting any reconstructed in the regions $1.37 < |\eta| < 1.52$ which corresponds to the transition region between the barrel and EMEC-OW. This transition region is characteriesed by larger amounts of passive material in comparison to other parts of the detector, thus degrading the quality of the energy response. The worsened electron performance in these regions would result in a less precise measured value for the Drell-Yan cross section which justifies its exclusion from the analysis. The pseudorapidity region for muons is defined to be $|\eta| < 2.4$ as this matches the extent of the muon trigger system. A further cut on the electron pseudorapidity of $|\eta| < 2.4$ is also applied so that the rapidity range spanned by the CC and MM channels will be the same.

In order to ensure that the leptons reconstructed are those from a primary vertex and reduce the effects of pileup, conditions on the impact parameters of the reconstructed tracks are also imposed. Namely, the absolute longitudinal impact parameter, z_0, must be less than 0.5 mm and the significance of the transverse impact parameter is less than five. For muons, the z_0 selection is the same as it is for electrons but the transverse impact parameter significance requirement is tightened, so that it has to be less than three.

As mentioned previously in Section 5.3.1, in order to pass the trigger some identification criterion must be met for one or both of the leptons. However, in order to improve the quality of the selected objects, the identification criteria are tightened further. For the CC and MM channels each lepton must pass at least the medium identification working point. In the MM channel, each muon must also pass the ‘fixed cut loose’ isolation working point.

The selection criteria defined so far mean that the surviving events are very likely to be high quality leptons and therefore suitable for analysis. However, the large mass range covered would result in increased amounts of signal region contamination from $\gamma^* \to l^+l^-$, were the event selection process to stop here. The invariant mass of the dilepton pair, m_{ll}, is therefore required to lie in the range $80 < m_{ll} < 102$ GeV around the Z pole.

With all application of all previous selection criteria, the signal region now has its full contingent of leptons that are suitable for the analysis. However, the purity of the signal region can still be improved further. The first improvement to the purity is by demanding that the two leptons should have opposite electric charges. This helps to eliminate backgrounds of a single real lepton plus a fake as well as backgrounds from same sign diboson events where both bosons decay leptonically. A final ‘truth matching’ step is also applied to the signal and background Monte Carlo, which involves matching the reconstructed objects to their truth level counterparts. This ensures that only ‘real’ leptons remain from the simulation and avoids the double counting of fakes, which must still be estimated.

Selection	Data	Z/γ^*	Top + EW Background
Initial	5.553×10^9	2.732×10^8	9.426×10^7
GRL	5.440×10^9	2.732×10^8	9.426×10^7
Data Quality	5.436×10^9	2.732×10^8	9.426×10^7
Two Leptons	1.382×10^9	1.499×10^8	2.942×10^7
Lepton Kinematics	9.250×10^7	8.554×10^8	1.534×10^6
Mass Window	7.904×10^7	7.691×10^7	3.416×10^5
Lepton ID	7.581×10^7	7.002×10^7	3.117×10^5
Opposite Sign	6.042×10^7	5.785×10^7	2.177×10^5
Truth Matching	6.042×10^7	5.785×10^7	2.164×10^5

Table 5.1: Number of events in the eeCC channel after the application of each selection criterion. The number of Monte Carlo estimated Z/γ^* and background events have been scaled to correspond to a luminosity of 140 fb^{-1}.

Selection	Data	Z/γ^*	Top + EW Background
Initial	5.553×10^9	2.733×10^8	9.448×10^7
GRL	5.440×10^9	2.733×10^8	9.448×10^7
Data Quality	5.435×10^9	2.733×10^8	9.448×10^7
Two Leptons	1.379×10^9	1.196×10^8	2.973×10^7
Lepton Kinematics	1.257×10^8	1.090×10^8	2.572×10^7
Mass Window	9.357×10^7	9.025×10^7	3.139×10^5
Lepton ID	8.407×10^7	8.161×10^7	2.780×10^5
Opposite Sign	8.407×10^7	8.161×10^7	2.757×10^5
Truth Matching	8.407×10^7	8.159×10^7	2.668×10^5

Table 5.2: Number of events in the MM channel after the application of each selection criterion. The number of Monte Carlo estimated Z/γ^* and background events have been scaled to correspond to a luminosity of 140 fb^{-1}.

5.4 Fiducial Volume and Acceptance

By following the event selection process of the previous selection, it is clear that it will not be possible to access the full phase space in which it is possible for $Z \to l^+l^-$ decays to occur. Instead, the analysis is carried out in what is known as the fiducial volume; the region of phase space that is accessible after the imposition of all analysis level selections and cuts. With truth level Monte Carlo it is possible to define the acceptance factor, A, which is the fraction of the fiducial phase space to the full phase space:

$$A = \frac{N_{fid}^{True}}{N_{All}^{True}} \tag{5.1}$$

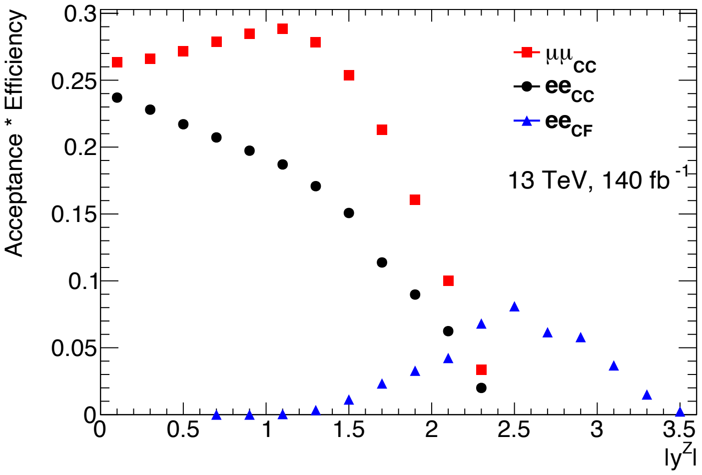

Figure 5.1: Acceptance × Efficiency for all three possible analysis channels as a function of the Z boson rapidity.

where N^{True}_{fid} is the number of truth level events within the 80 - 102 GeV mass range used for the analysis and N^{True}_{All} the total number of truth level events. However, it is also necessary to consider the effect of the efficiency as well, where the efficiency is the fraction of number of true events in the fiducial region that are also reconstructed ($N^{True}_{reco,fid}$) to N^{True}_{fid}. The product of acceptance and efficiency therefore describes the total amount of the full phase space that is covered at reconstruction level and is simply the ratio of $N^{True}_{reco,fid}$ to N^{True}_{All} i.e

$$Acceptance \times Efficiency = \frac{N^{True}_{fid}}{N^{True}_{All}} \times \frac{N^{True}_{reco,fid}}{N^{True}_{fid}} = \frac{N^{True}_{reco,fid}}{N^{True}_{All}} \tag{5.2}$$

In Figure 5.1 this quantity is displayed as a function of the rapidity for each of the three analysis channels, where clear differences can be seen. In general, since muons have a much more distinct detector signature than an electron, they are reconstructed with a much higher efficiency than electrons. They also do not suffer from a degraded energy response in the transition region between the barrel and endcap calorimeters, which means that this region is included in the analysis and correspondingly increases the acceptance. With the imposition of the $|\eta| < 2.4$ cut, the acceptance × efficiency correspondingly falls to zero at $|y^Z| = 2.4$. In comparison, the kinematics of the CF channel mean that the $|y^Z| < 1.2$ is forbidden but this is compensated by an increase in the Z rapidity phase space up to 3.6. The acceptance × efficiency is also much lower in the CF channel than the other two analysis channels due to the lower efficiencies associated with forward electrons, as well as the tighter identification requirements necessary to reduce the large number of fakes.

5.5 Multijet Estimate

Since many physics processes are well described through Monte Carlo modelling, the amount of these backgrounds in the signal region can be estimated by simply counting the number of luminosity-weighted Monte Carlo events that pass the selection criteria. However, semileptonic decays of heavy quarks, the misidentification of hadrons as leptons and photon conversions to electrons are all processes that are not well modelled by Monte Carlo but are still present in the data. Collectively these types of events are known as ‘multijet’ events and in order to estimate their contribution to the signal region, data-driven methods must be employed. In the case of this analysis, the fake-factor method was chosen to estimate the amount of multijet [12].

Using the fake factor method, it has been verified that the amount of multijet in the MM channel is less than 0.1% of the signal region. Therefore, the multijet estimate needs to only be carried out for the CF and CC channels.

5.5.1 Fake Factor Method

The aim of the fake factor method is to estimate the amount of multijet in the signal region through measuring the amount of multijet in well defined control regions and then multiplying it with a transfer factor (the fake factor) that extrapolates the result to the signal region. Therefore, the fake factor method is a two step procedure: firstly estimating the fake factor and secondly calculating the amount of multijet events in the signal region.

In order to define the fake factor, two event regions must be defined: a ‘tight’ region that is mostly signal with little multijet and an orthogonal ‘loose’ region which is enriched in multijet. Here tight and loose do not necessarily correspond to the tight and loose identification working points introduced in Section 4.2 but are generic labels used to describe the method. The resultant combination of the tight and loose regions is then known as the baseline selection. Within this baseline selection two efficiencies must be defined, ε_r which is the fraction of real leptons passing the tight selection and ε_f, the fraction of fakes passing the tight selection.

CF Channel

The simplest case in a dilepton final state is that where one lepton can considered can always be considered real and the other may be a fake. This is exactly the case in the CF channel where it is much more likely that the forward electron will turn out to be a fake than the central electron due to the lack of tracking information. Whilst the CF channel is not being used to measure the Drell-Yan cross section, the fake factor method can be used to estimate the amount of multijet background in the signal region used for the measurement of the weak mixing angle.

The first quantity that must be considered is the total number of fakes in the loose region, N_f^l, which will simply be the number of events in the loose region, N^l, with the number of real events, N_r^l, subtracted:

$$N_f^l = N^l - N_r^l \tag{5.3}$$

N_f^l can also be determined by considering the fake efficiency, ε_f. Since ε_f is the proportion of fakes passing the tight selection, then $(1 - \varepsilon_f)$ will be the proportion of fakes passing loose:

$$N_f^l = (1 - \varepsilon_f) N_f^b \tag{5.4}$$

where N_f^b is the total number of fakes in the baseline selection, which can be directly calculated when combining Equations (5.3) and (5.4):

$$N_f^b = \frac{1}{1 - \varepsilon_f}(N^l - N_r^l) \tag{5.5}$$

The corresponding number of fakes in the tight region will then simply be N_f^b multiplied by the fake efficiency. This relation can then be combined with Equation (5.5) to yield an expression for the number of fakes in the tight selection as a function of N^l and N_r^l:

$$\begin{aligned} N_f^t &= \varepsilon_f N_f^b \\ &= \frac{\varepsilon_f}{1 - \varepsilon_f}(N^l - N_r^l) \end{aligned} \tag{5.6}$$

Or equivalently:

$$N_f^t = F(N^l - N_r^l) \tag{5.7}$$

where the fake factor, F, is defined to be:

$$F = \frac{\varepsilon_f}{1 - \varepsilon_f} \tag{5.8}$$

Whilst Equations (5.7) and (5.8) are valid for any final state where a single lepton may be a fake, these terms must be translated so that they make sense in the context of the measurement of the angular coefficients. For this, four regions must be defined:

- a control region, approximately where the analysis is to be carried out, IDOn. This region includes all forward electrons that pass the tight identification working point and the reconstructed Z mass is between $80 < m_{ee} < 100$ GeV (i.e on pole), except for forward electrons with $|\eta| > 4.0$ where the mass cut is relaxed to $80 < m_{ee} < 150$ GeV due to the broadening of the Z boson peak in this region.
- another control region, where the forward electron fails the very very loose working point [1] but passes the on pole requirement, nIDOn.
- a further control region, where the forward electron passes tight identification but fails the on pole requirement, IDOff.
- a final control region, where the forward electron fails both the very very loose working point and also the on pole requirement, nIDOff.

From these regions, the IDOn and nIDOn control regions can be identified as the tight and loose regions that were previously defined when deriving the expressions for the number of fakes in the tight region and the fake factor. By considering the definition of the fake efficiency, it can be written as the ratio of multijet events in the signal region, $N_{\mathrm{MJ}}^{\mathrm{IDOn}}$ to the total amount of multijet in the baseline selection, $N_{\mathrm{MJ}}^{\mathrm{b}}$. Correspondingly, $(1-\varepsilon_f)$ will become the ratio of multijet events in the nIDOn region to the total amount of multijet in the baseline selection.

$$\varepsilon_f = \frac{N_{\mathrm{MJ}}^{\mathrm{IDOn}}}{N_{\mathrm{MJ}}^{\mathrm{b}}} \tag{5.9}$$

$$1 - \varepsilon_f = \frac{N_{\mathrm{MJ}}^{\mathrm{nIDOn}}}{N_{\mathrm{MJ}}^{\mathrm{b}}} \tag{5.10}$$

$$F \equiv \frac{\varepsilon_f}{1-\varepsilon_f} = \frac{N_{\mathrm{MJ}}^{\mathrm{IDOn}}}{N_{\mathrm{MJ}}^{\mathrm{nIDOn}}} \tag{5.11}$$

If it is now assumed that the ratio of multijet in the IDOn region to the amount of multijet in the nIDOn region is independent of the invariant mass range chosen[2], then the fake factor will be independent of the choice of tight and loose selections, i.e:

$$\frac{N_{\mathrm{MJ}}^{\mathrm{IDOn}}}{N_{\mathrm{MJ}}^{\mathrm{nIDOn}}} = \frac{N_{\mathrm{MJ}}^{\mathrm{IDOff}}}{N_{\mathrm{MJ}}^{\mathrm{nIDOff}}} = F \tag{5.12}$$

In order to calculate the fake factor, it is therefore necessary to know the amount of multijet in the control region. For the purposes of this analysis, anything that is not a

---

[1] The very very loose (VVL) working point is defined such that the signal efficiency is 99.9%. An 'electron' that fails this working point is therefore almost certainly a fake.

[2] A reasonable assumption to make since the underlying physics processes that create multijet events are uncorrelated with Z boson production.

real electron is classified as multijet, meaning that, assuming the amount of real electron contamination in the control region is small, for a given control region the amount of multijet is:

$$N^{\text{Region}}_{\text{MJ}} = N^{\text{Region}}_{\text{data}} - \hat{\mu} \cdot N^{\text{Region}}_{\text{MC}} \tag{5.13}$$

where $\hat{\mu}$ is a factor to account for any residual normalisation effects that have not been dealt with the application of weights and scale factors to the simulated Monte Carlo events. By considering the equality on the left hand side of Equation (5.12), subtracting one side and rearranging, it can be shown that:

$$A\hat{\mu}^2 + B\hat{\mu}^2 + C = 0 \tag{5.14}$$

where:

$$\begin{aligned}
A &= N^{\text{nIDOn}}_{\text{MC}} \cdot N^{\text{IDOff}}_{\text{MC}} - N^{\text{SR}}_{\text{MC}} \cdot N^{\text{nIDOff}}_{\text{MC}} \\
B &= N^{\text{SR}}_{\text{data}} \cdot N^{\text{nIDOff}}_{\text{MC}} - N^{\text{nIDOn}}_{\text{data}} \cdot N^{\text{IDOff}}_{\text{MC}} + N^{\text{SR}}_{\text{MC}} \cdot N^{\text{nIDOff}}_{\text{data}} - N^{\text{nIDOn}}_{\text{MC}} \cdot N^{\text{IDOff}}_{\text{data}} \\
C &= N^{\text{nIDOn}}_{\text{data}} \cdot N^{\text{IDOff}}_{\text{data}} - N^{\text{SR}}_{\text{data}} \cdot N^{\text{nIDOff}}_{\text{data}}
\end{aligned} \tag{5.15}$$

The quadratic equation can then be used to calculate the value of $\hat{\mu}$ from Equation (5.14) and hence the fake factor defined by Equation (5.11) can also be calculated. When the fake factor is calculated a statistical error on its value can also be derived, meaning that the a systematic uncertainty can also be placed on the amount of multijet background that will eventually be estimated.

With the fake factor now calculated, the next step is to estimate the amount of multijet in the signal region. To do this Equation (5.7) is used, with the use of the analysis signal region as the tight selection and a final control region for the loose selection, CR. CR is defined as the region where the forward electron fails to pass the tight ID requirement. Translating the $N^l - N^l_r$ term from Equation (5.7) to the CR then reads as:

$$N^{\text{CR}}_{\text{data}} - \mu \cdot N^{\text{CR}}_{\text{MC}} = N^{\text{CR}}_{\text{MJ}} \tag{5.16}$$

where μ plays the same normalisation role as $\hat{\mu}$ in Equation (5.13) but must be recalculated to take into account any normalisation differences from the change in definition of control regions and is hence renamed to emphasise this difference. Explicitly, μ can be written as:

$$\mu = \frac{N^{\text{SR}}_{\text{data}} - FN^{\text{CR}}_{\text{data}}}{N^{\text{SR}}_{\text{MC}} - FN^{\text{CR}}_{\text{MC}}} \tag{5.17}$$

After recalculating μ, the multijet in the signal region can finally be calculated as:

$$N_{\mathrm{MJ}}^{\mathrm{SR}} = F N_{\mathrm{MJ}}^{\mathrm{CR}} \tag{5.18}$$

CC Channel

In the CC channel, both electrons are just as likely to be fakes as the other, thus complicating the multijet estimate. The fake factor method, can however, still be used by considering combinations of real and fake electrons in the same general baseline, loose and tight selections as for the CF multijet estimate. The baseline sample will therefore contain three separate sources of fake events:

- Both leptons are fake: $N_{ff}^{b}(FF)$
- The first lepton is real and the second is fake: $N_{rf}^{b}(PF)$
- The first lepton is fake and the second is real: $N_{fr}^{b}(FP)$

The total number of fakes in the tight selection is therefore:

$$N_{f}^{t} = \varepsilon_{r,1}\,\varepsilon_{f,2}\,N_{rf}^{b} + \varepsilon_{f,1}\,\varepsilon_{r,2}\,N_{f,r}^{b} + \varepsilon_{f,1}\,\varepsilon_{f,2}\,N_{ff}^{b} \tag{5.19}$$

Where $\varepsilon_{r1/2}$ is the efficiency of real leptons one and two passing the tight selection and $\varepsilon_{f1/2}$ the efficiency of fake leptons one and two passing the tight selection. As with the CF channel, it is possible to deduce expressions for N_{rf}^{b}, N_{fr}^{b} and N_{ff}^{b} which allows for the total number of fakes in the tight selection to be calculated. If first the number of events where lepton one is real and passes tight and lepton two is fake and passes loose is considered, then:

$$N_{rf}^{tl} = N^{tl} - N_{rr}^{tl} = \varepsilon_{r,1}(1-\varepsilon_{f,2})N_{rf}^{b} \tag{5.20}$$

which when rearranged reads as:

$$\varepsilon_{r,1}N_{rf}^{b} = \frac{1}{1-\varepsilon_{f,2}}(N^{tl} - N_{rr}^{tl}) \tag{5.21}$$

Similarly, it can also be shown that:

$$\varepsilon_{r,2}N_{fr}^{b} = \frac{1}{1-\varepsilon_{f,1}}(N^{lt} - N_{rr}^{lt}) \tag{5.22}$$

The number of events where both leptons are fake and loose can also be written as:

$$N_{ff}^{ll} = N^{ll} - N_{rr}^{ll} = (1-\varepsilon_{f,1})(1-\varepsilon_{f,2})N_{ff}^{b} \tag{5.23}$$

Correspondingly, the number of events in the baseline sample, where both leptons are fake will be:

$$N^{b}_{ff} = \frac{1}{(1-\varepsilon_{f,1})(1-\varepsilon_{f,2})}(N^{ll} - N^{ll}_{rr}) \tag{5.24}$$

By substituting Equations (5.21), (5.22) and (5.24) back into Equation (5.19) , the number of fakes in the tight selection is therefore:

$$N^{t}_{f} = \frac{\varepsilon_{f,2}}{1-\varepsilon_{f,2}}(N^{tl} - N^{tl}_{rr}) + \frac{\varepsilon_{f,1}}{1-\varepsilon_{f,1}}(N^{lt} - N^{lt}_{rr}) + \frac{\varepsilon_{f}, 1}{1-\varepsilon_{f,1}}\frac{\varepsilon_{f,2}}{1-\varepsilon_{f,2}}(N^{ll} - N^{ll}_{rr}) \tag{5.25}$$

Equation (5.25) can then be simplified by defining fake factors F_1 and F_2:

$$F_1 = \frac{\varepsilon_{f,1}}{1-\varepsilon_{f,1}} \text{ and } F_2 = \frac{\varepsilon_{f,2}}{1-\varepsilon_{f,2}} \tag{5.26}$$

The final expression for the number of fakes in the tight selection then becomes:

$$N^{t}_{f} = (N^{tl} - N^{tl}_{rr})F_2 + (N^{lt} - N^{lt}_{rr})F_1 + (N^{ll} - N^{ll}_{rr})F_1F_2 \tag{5.27}$$

Similarly to the CF channel, this theoretical treatment of the multijet background must be translated to the analysis level. First, in order to define the baseline sample, events with one lepton passing and another failing the single lepton trigger are used, with events passing solely the dilepton trigger not considered. Whilst this introduces a bias into the calculation of the fake factors, the likelihood that one or both of the electrons passing a dilepton trigger is fake, is much smaller than that of a lepton failing a single lepton trigger being fake. The case of events passing the dilepton trigger as being fake can therefore be ignored.

The baseline sample is further split into events where the leading lepton passes the single lepton trigger and events where the leading lepton fails the single lepton trigger. This is done so that individual fake factors for both cases can be calculated, which correspond to F_2 and F_1 respectively. Within this baseline selection, the tight selection is then electrons that pass the medium identification requirement and the dilepton invariant mass is within the range $75 < m_{ee} < 125$ GeV. The loose selection then contains electrons that find themselves outside of this mass range, that fail the medium identification working point and fail tracking and calorimetric isolation criteria, namely:

- in a cone of $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.2$, the ratio of the p_T of all tracks within this cone to the electron p_T should be less than 0.14

- in a cone of $\Delta R = 0.2$ around the electromagnetic calorimeter cluster, the ratio of the transverse energy, $E_T = \sqrt{m^2 + p_T^2}$, of all deposits within the cone to the electron transverse energy, should be less than 0.2.

As for the CF channel, the fake factors are calculated using the four control regions: IDOn, nIDOn, IDOff and nIDOff. Explicitly, F_1 and F_2 are calculated using the off pole control regions in Equation (5.12). It is also assumed that there is no correlation between the identification of the two leptons so that each version of the control region is orthogonal from the other. The $\hat{\mu}$ normalisation constant that is necessary to calculate the fake factor is again calculated according to Equation (5.14).

As in the CF channel, the CC channel fake factors are calculated in bins of η and p_T in order to produce a multijet estimate that can account for any kinematic dependence of both signal and fakes. The fake factors corresponding to this binning can be seen in Figure 5.2. The statistical uncertainty on these fake factors is again considered as a systematic uncertainty when measuring the angular coefficients.

With the fake factors F_1 and F_2 defined, then the total amount of multijet in the signal region can now be estimated by using Equation (5.27) with the translation that:

$$\begin{aligned} N_f^t &= N_{\mathrm{MJ}}^{\mathrm{SR}} \\ N^{tl} - N_{rr}^{tl} &= N_{\mathrm{MJ}}^{\mathrm{PF}} \\ N^{lt} - N_{rr}^{lt} &= N_{\mathrm{MJ}}^{\mathrm{FP}} \\ N^{ll} - N_{rr}^{ll} &= N_{\mathrm{MJ}}^{\mathrm{FF}} \end{aligned} \tag{5.28}$$

where the regions are defined as:

- SR: the signal region where the extraction of the angular coefficients is carried out.
- PF: a control region where the lepton passing the single lepton trigger also passes the medium ID requirement. In addition the lepton failing the single lepton trigger must fail medium ID and also fail the tracking and calorimetric isolation requirements previously introduced.
- FP: a control region where the lepton passing the single lepton trigger fails the medium ID requirement and the tracking and calorimetric isolation requirements. In addition the lepton failing the single lepton trigger must pass medium ID.
- FF: a control region where both leptons fail medium ID and also fail the tracking and calorimetric isolation requirements.

After defining each region, Equation (5.27) can similarly be rewritten as an expression for the total amount of multijet in the signal region:

$$N_{\mathrm{MJ}}^{\mathrm{SR}} = F_2\, N_{\mathrm{MJ}}^{\mathrm{PF}} + F_1\, N_{\mathrm{MJ}}^{\mathrm{FP}} + F_1\, F_2\, N_{\mathrm{MJ}}^{\mathrm{FF}} \tag{5.29}$$

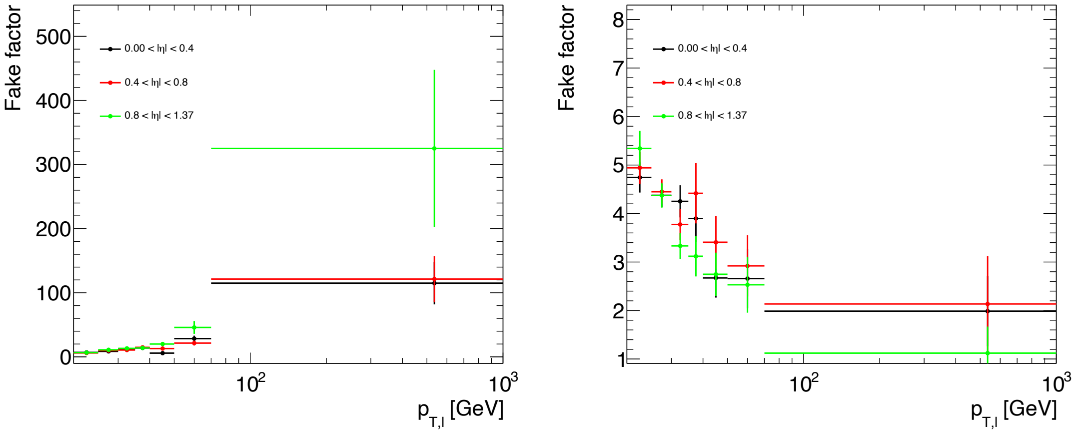

(a) Fake factors in the barrel region for electrons passing the single lepton trigger.

(b) Fake factors in the barrel region for electrons failing the single lepton trigger.

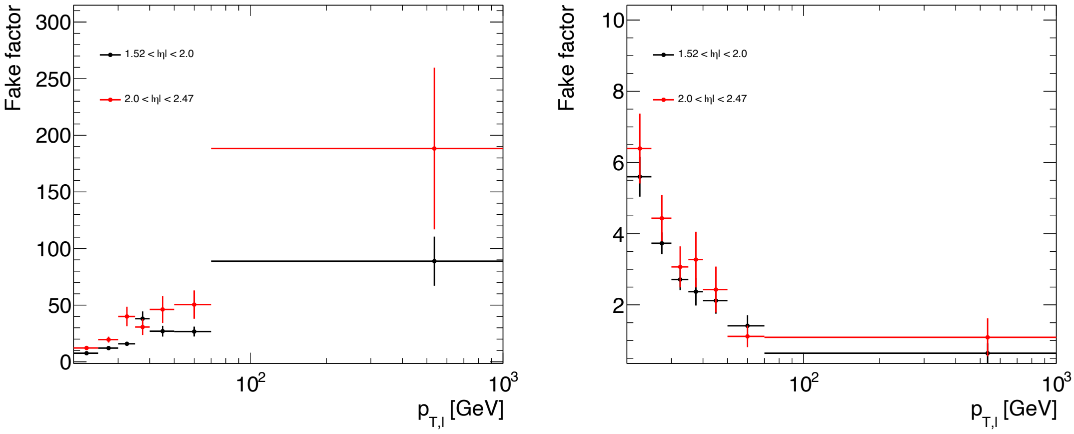

(c) Fake factors in the endcap region for electrons passing the single lepton trigger.

(d) Fake factors in the endcap region for electrons failing the single lepton trigger.

Figure 5.2: Fake factors with their statistical uncertainty for the central electron channel as a function of the electron transverse momentum split by electrons passing the single lepton trigger (left column) and those failing the single lepton trigger (right). The top row displays the fake factors in the barrel region and the bottom row those in the endcap.

As in the CF channel, the amount of multijet in each control region is calculated by Equation (5.16).

When calculating the amount of multijet in each control region, Equation (5.16) cannot be used for each control region since it does not take into account any duplication of multijet in the FF region that is also counted in the PF and FP regions. Therefore the total amount of multijet in each region must be calculated as:

$$\begin{aligned} N_{\mathrm{MJ}}^{\mathrm{FF}} &= N_{\mathrm{data}}^{\mathrm{FF}} - \mu \cdot N_{\mathrm{MC}}^{\mathrm{FF}} \\ N_{\mathrm{MJ}}^{\mathrm{PF}} &= N_{\mathrm{data}}^{\mathrm{PF}} - \mu \cdot N_{\mathrm{MC}}^{\mathrm{PF}} - N_{\mathrm{MJ}}^{\mathrm{FF}} \\ N_{\mathrm{MJ}}^{\mathrm{FP}} &= N_{\mathrm{data}}^{\mathrm{FP}} - \mu \cdot N_{\mathrm{MC}}^{\mathrm{FP}} - N_{\mathrm{MJ}}^{\mathrm{FF}} \end{aligned} \tag{5.30}$$

where any overlap between the regions (and thus double counting of multijet events) is accounted for by subtracting the amount of multijet in the FF region from the PF and FP regions. Similarly, the μ normalisation parameter in Equation (5.30) must also take into account any duplication of multijet events and so is defined as:

$$\mu = \frac{N_{\mathrm{data}}^{\mathrm{SR}} - F_2\, N_{\mathrm{data}}^{\mathrm{PF}} - F_1\, N_{\mathrm{data}}^{\mathrm{FP}} + F_1\, F_2 N_{\mathrm{data}}^{\mathrm{FF}}}{N_{\mathrm{MC}}^{\mathrm{SR}} - F_2\, N_{\mathrm{MC}}^{\mathrm{PF}} - F_1\, N_{\mathrm{MC}}^{\mathrm{FP}} + F_1\, F_2 N_{\mathrm{MC}}^{\mathrm{FF}}} \tag{5.31}$$

The invariant mass distribution in each control region and the resultant amount of multijet in the signal region can be seen in Figure 5.3. Both the FF and FP regions show a flat distribution versus m_{ee} meaning there is no signal contamination here and that the regions have been properly defined under the assumptions of the fake factor method, i.e that the loose regions should be enriched in multijet with little to no real contamination. There is, however, some clear signal contamination in the PF region as the Z boson resonance can be seen. However, since this amount of contamination does not dominate over the multijet in this region, then it is still permissible. The ratio of the data to the expectation in each control region plot is also exactly one by construction due to the explicit normalisation to the data when μ is calculated. When the amount of prompt electrons is no longer normalised to the data, the data to expectation ratio begins to vary as can be seen in the signal region plot.

5.6 Systematic Uncertainties

As currently defined, the nominal dataset that is produced by the event selection process does not take into account any uncertainty on the various theoretical and experimental effects that are necessary to correctly weight the Monte Carlo and calibrate the leptons used for the analysis. Any change to the weighting or calibration is known as a systematic uncertainty. Beginning with the Monte Carlo generation process, the only source of theoretical uncertainty considered by this analysis is that of the parton distribution

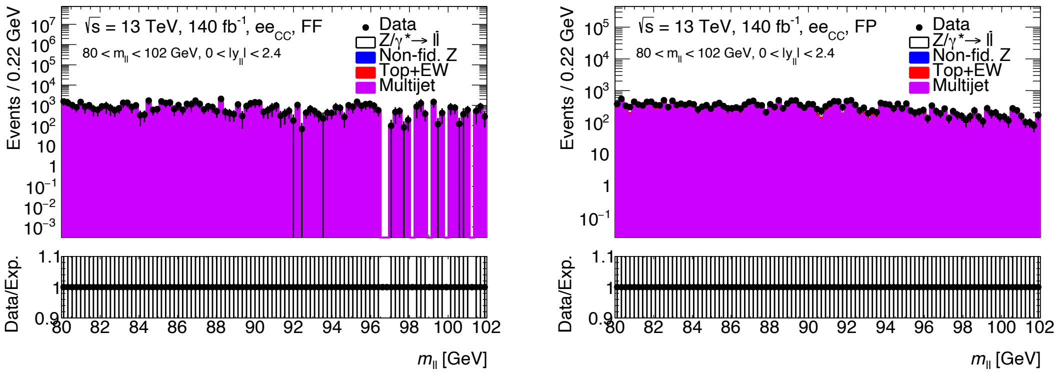

(a) FF region.

(b) FP region.

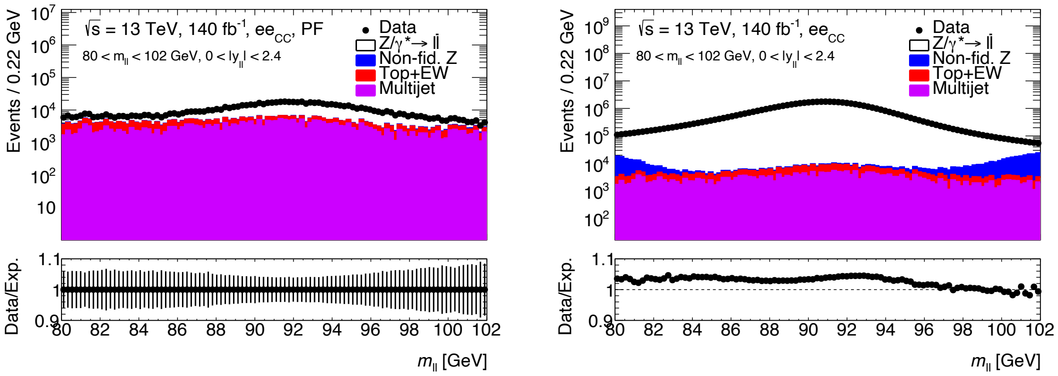

(c) PF region.

(d) Signal region.

Figure 5.3: The m_{ee} distributions in the three control regions (FF, FP and PF) used to estimate the multijet and the resultant distribution in the signal region. The data to expectation ratio contains statistical errors and those from the normalisation. By construction the data to expectation ratio is exactly one in the FF, FP and PF regions.

functions. As was explained in Section 2.3.3, it is not possible to analytically calculate the PDFs and their value must be extracted experimentally with a corresponding uncertainty. This uncertainty therefore affects the modelling of the $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ processes and hence the simulated Monte Carlo datasets. The values for the uncertainties in the PDF sets used for this analysis are taken from the LHAPDF tool [30]. Whilst it would be possible to consider other theoretical uncertainties such as those from the choice of factorisation and renormalisation scale, choice of event generator, parton shower model and electroweak and QED corrections, it was established that their effects are minimal by the previous ATLAS measurement of the angular coefficients [6]. Therefore, these effects are neglected for this analysis in order to avoid adding unnecessary systematic uncertainties to the measurement.

When the simulated Monte Carlo is weighted to the amount of data recorded by ATLAS, the first experimental uncertainty is consequently that of the integrated luminosity. For the full Run 2 dataset, this was measured to be 140 fb^{-1} $\pm$ 0.87% [5], meaning there is also a ±0.87% uncertainty applied to the Monte Carlo normalisation weight. All other sources of experimental systematic uncertainty arise from the calibration and efficiency corrections to Monte Carlo. Since the process of reconstructing and calibrating electrons and muons are fundamentally different, these systematic uncertainties must be split by analysis channel.

Unlike the identification, reconstruction and trigger variations, the systematic uncertainties on the calibration do not arise from the variation of different scale factors. Instead, the calibration nuisance parameters are split into two categories; those that affect the resolution and those that affect the energy scale at which the electron is reconstructed. The resolution is heavily influenced by the amount of material an electron must travel through before it showers in the calorimeter as well as any residual smearing from the in situ calibration process. Hence most of the nuisance parameters on the electron resolution are variations covering the amount of material in different areas of the detector an electron may encounter. Similarly the electron energy scale is also affected by material distributions but also by the high voltage system, the topocluster threshold values used for reconstruction and also the scale adjustments from the in situ calibration.

For the electron channel, the full set of systematic uncertainties that cover the calibration process as well as the trigger, reconstruction and identification efficiencies are provided centrally by the e/γ combined performance group [13, 14]. For the muon channel, the systematic uncertainties on the calibration process, trigger, reconstruction, track to vertex association (TTVA) and isolation efficiencies as well as the sagitta corrections are provided by the muon combined performance group [11, 9].

The last source of systematic uncertainty is that on the number of background events in the signal region. Since the amount of background is specified wholly through the event selection cuts, the systematic uncertainty on the background is also estimated at analysis level. To be conservative, a flat ±20% variation is placed on the amount of top and electroweak background that enters into the signal region. A further systematic is applied on the choice of generator used to simulate the background as well. As was discussed in Section 5.5.1, the statistical uncertainty on the fake factors used to estimate the amount of multijet in the central electron channel is also used as a systematic uncertainty.

5.7 Composition of the Signal Region

After the application of all the analysis cuts and estimating the number of fake events, the modelling of the Monte Carlo in the resultant signal region must be validated.This is done by analysing the control plots of the process in question, namely the kinematics of the selected events and the quantities that necessary to extract the angular coefficients. For the kinematics, this means looking at plots for the individual leptons ($\eta_{lead/sub}$, $p_T^{lead/sub}$) as well as those covering the dilepton, i.e boson, properties (m_{ll}, p_T^{ll}, y_{ll}, ϕ_{ll}). Since the measurement of the angular coefficients is defined by the Collins-Soper reference frame, it is also necessary to look at the corresponding angular distributions $\cos\theta_{CS}$ and ϕ_{CS}.

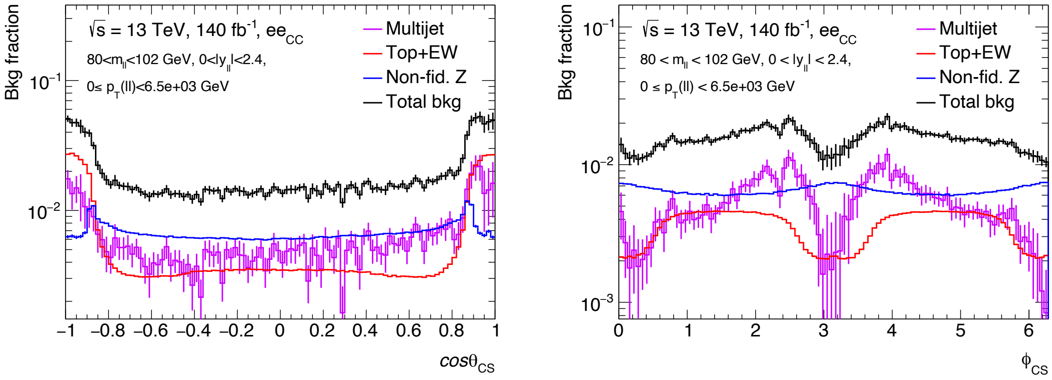

(a) Background composition of $\cos\theta_{CS}$.

(b) Background composition of ϕ_{CS}.

Figure 5.4: Fraction of the total number of events in the signal region that corresponds to each source of background events as a function of $\cos\theta_{CS}$ and ϕ_{CS}.

With the imposition of the analysis level selection cuts, the signal region that is created is a highly pure sample of $Z \to l^+l^-$ events. The fractional contribution of the total background, as well as its individual components, is shown as a function of $\cos\theta_{CS}$ and ϕ_{CS} in Figure 5.4 for the electron channel. This shows that the total background is of the order of 2% for the electron channel, with similar amounts of each background type contributing to the total. With the overall background levels established, the distribution of the individual variables can now be analysed.

5.7.1 Dilepton Variables

The simplest and, arguably, one of the most important checks to make using the control plots, is that the events labelled in Tables 5.1 and 5.2, are indeed from $Z \to l^+l^-$ decays.

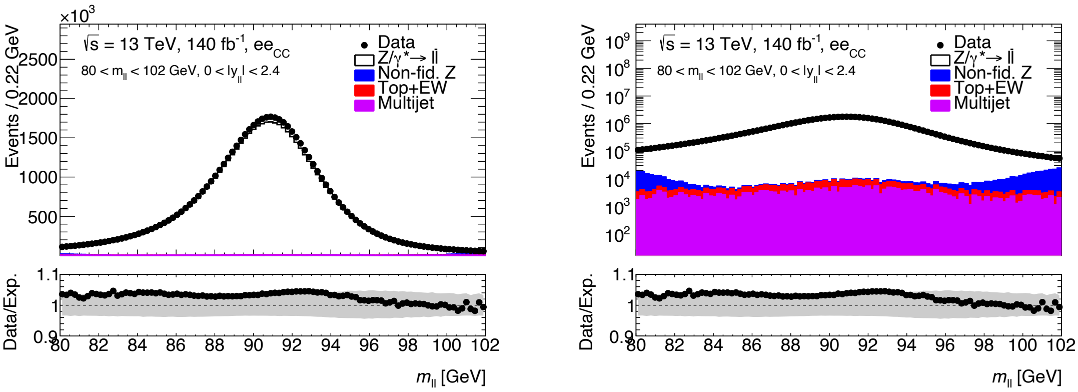

(a) m_{ee} distribution.

(b) Logarithmically scaled m_{ee} distribution.

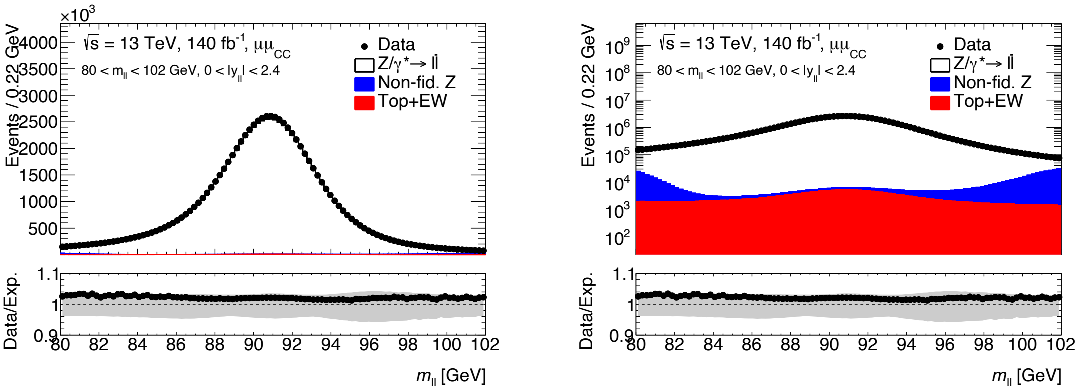

(c) $m_{\mu\mu}$ distribution.

(d) Logarithmically scaled $m_{\mu\mu}$ distribution.

Figure 5.5: Signal region inclusive control plots of the dilepton invariant mass distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

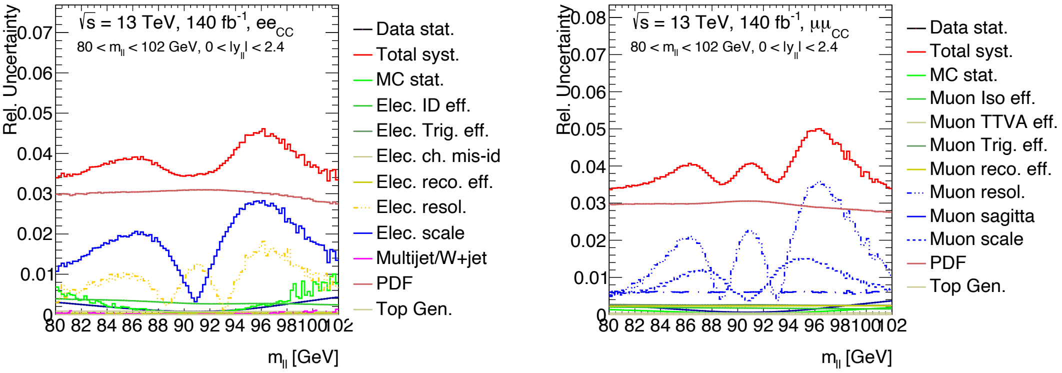

(a) m_{ee} uncertainty breakdown.

(b) $m_{\mu\mu}$ uncertainty breakdown.

Figure 5.6: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the m_{ll} distribution for the inclusive signal region in both the electron and muon channels. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

In Figure 5.5, the invariant mass of the dilepton pair is displayed for both the electron and muon channels. In both channels a clear peak at 91 GeV with two decaying tails can be observed; the hallmark of the Z boson resonance. This therefore confirms that the signal region has been correctly specified and the analysis can proceed. Figures 5.5b and 5.5d display the invariant mass spectrum using a logarithmic scale, revealing that the background contribution is typically an order of magnitude lower than the signal, which highlights the effects the effects of the cleaning cuts applied that make the signal region as pure as possible. For both channels, the agreement between the data and the expectation shows that there is approximately 5% more data than is expected. However, since the Monte Carlo is only generated at NLO in α_S it does not take into account any higher order QCD corrections that are inherently part of the data. This discrepancy between the data and expectation is therefore to be expected.

With the signal region having been confirmed as populated with Z bosons, its momentum spectrum must also be checked. Since the Z is produced at the LHC through the collision of two partons, any imbalance in their momenta or a collision of a heavy flavour with a lighter flavour, will result in the produced Z also having a non-zero (transverse) momentum. If there is also emission of QCD ISR during the collision, the Z will recoil and thus be provided with a further source of momentum. The production of Z momentum is, therefore, dependent on the effects of non-perturbative QCD, as well as all orders of perturbative QCD. These effects are, however, not well modelled by the signal Monte Carlo which was calculated only to NLO in α_S. As shown for the muon channel in Figures 5.7e and 5.7f, this causes the data to expectation agreement to begin to worsen from approximately 70 GeV onwards and quickly becomes a more than 20% disagreement

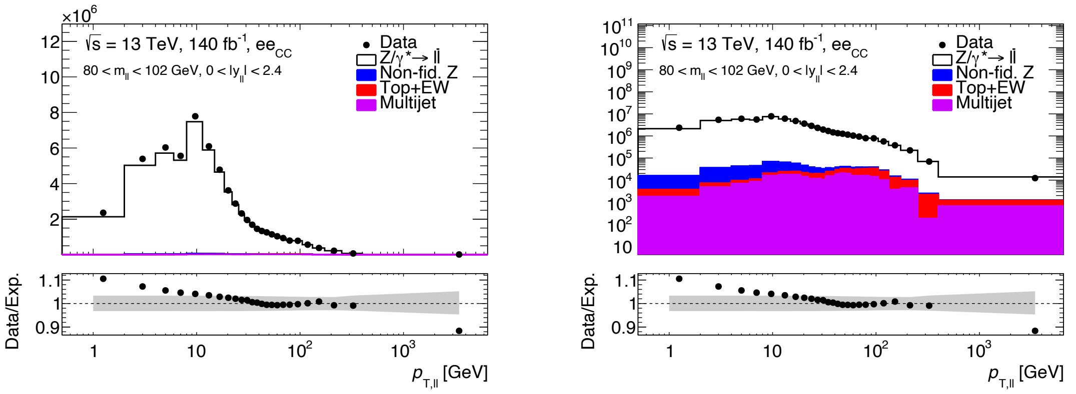

(a) p_T^{ee} distribution.

(b) Logarithmically scaled p_T^{ee} distribution.

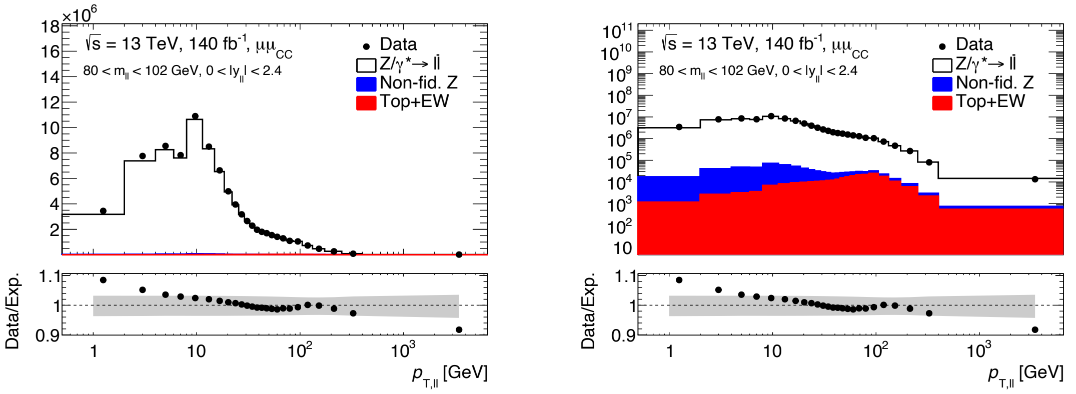

(c) $p_T^{\mu\mu}$ distribution.

(d) Logarithmically scaled $p_T^{\mu\mu}$.

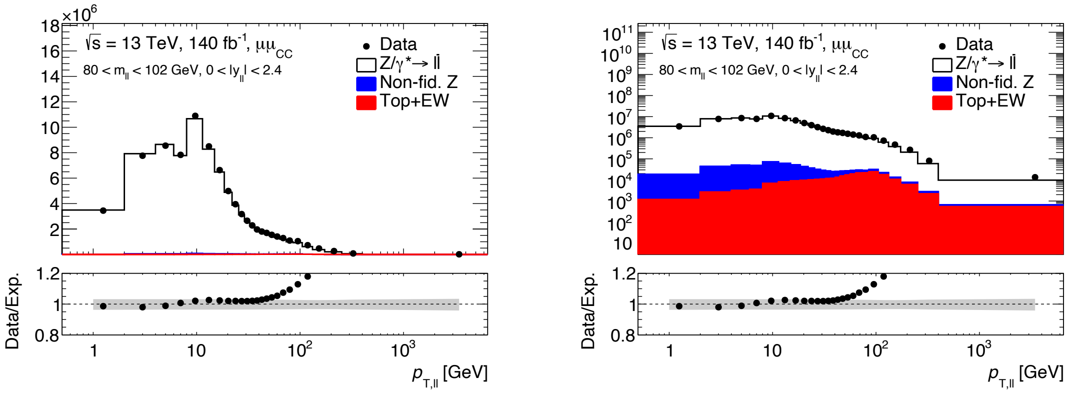

(e) $p_T^{\mu\mu}$ distribution without reweighting.

(f) Logarithmically scaled $p_T^{\mu\mu}$ distribution without reweighting.

Figure 5.7: Signal region inclusive control plots of the dilepton transverse momentum distribution for both the electron and muon channels. Additionally, the effect of not applying the transverse momentum reweighting is displayed for the muon channel. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

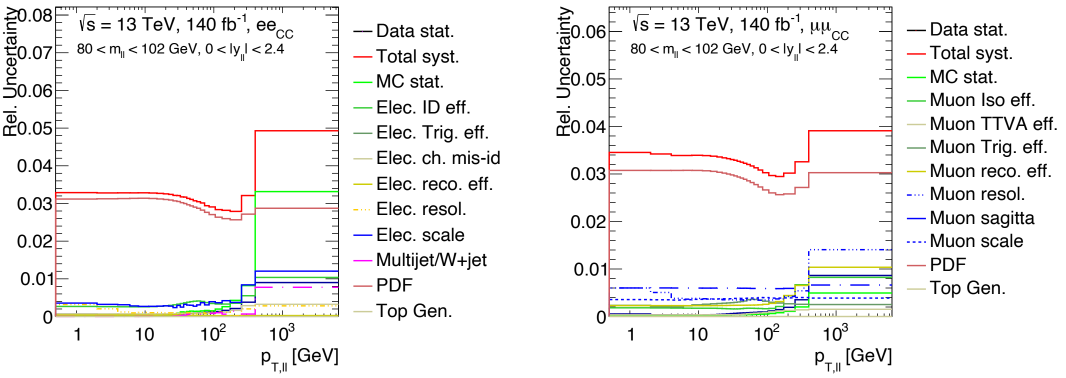

(a) p_T^{ee} uncertainty breakdown.

(b) $p_T^{\mu\mu}$ uncertainty breakdown.

Figure 5.8: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the $p_{T,ll}$ distribution for the inclusive signal region in both the electron and muon channels. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

between the data and the expectation. Therefore, in order to provide a more accurate description of the expectation, the Monte Carlo is reweighted to the predictions of DY-Turbo [32], which provides predictions for the transverse momentum spectrum of the Drell-Yan vector boson at N^{3}LO with resummation to N^{4}LL accuracy. The result of this is a dilepton transverse momentum distribution that accurately matches the data from approximately 10 to several hundred GeV of dilepton transverse momentum as can be seen in Figures 5.7a - 5.7d. At very low and very high transverse momentum the effects of non-perturbative and higher order perturbative QCD (i.e higher than the order used for DYTurbo predictions) become more apparent as the agreement between the data and expectation worsen to ±10%. However, as these regions are more sparsely populated than the bulk of the distribution, the overall agreement of the transverse momentum distribution is very good.

Although the PDFs contribute to the Z boson's transverse momentum, their greatest effect on Z production is in the rapidity distribution. In order to produce a Z boson at high rapidity, one of the colliding partons must have a correspondingly high Björken x. Since the PDFs become very small at high Björken x, the probability of producing a Z at high rapidity is correspondingly greatly diminished. Therefore, the Z boson rapidity spectrum should display a peak at zero rapidity (where the PDFs are higher), which decreases as rapidity increases. This is exactly what is seen in both the dielectron and dimuon rapidity distributions in Figure 5.9, thus confirming the expected behaviour. Similarly to the invariant mass spectrum, the slight disagreement between the data and expectation is again due to the lack of higher order QCD corrections in the Monte Carlo samples.

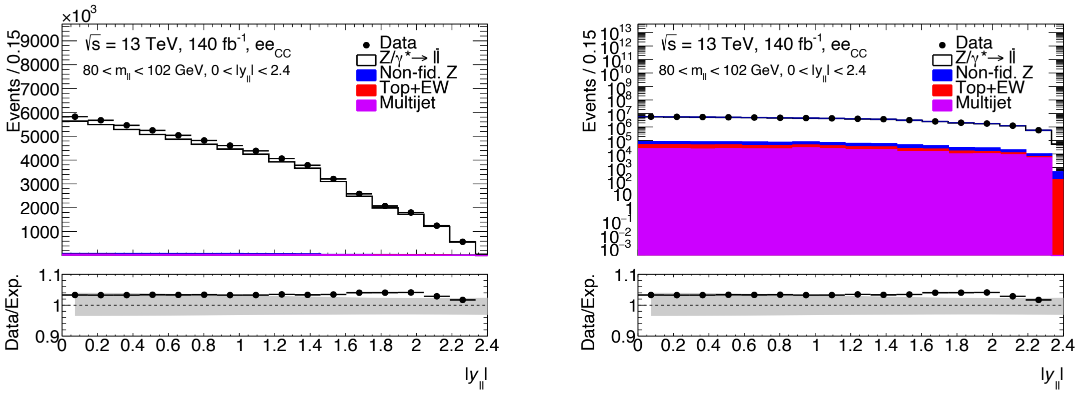

(a) y_{ee} distribution. (b) Logarithmically scaled y_{ee} distribution.

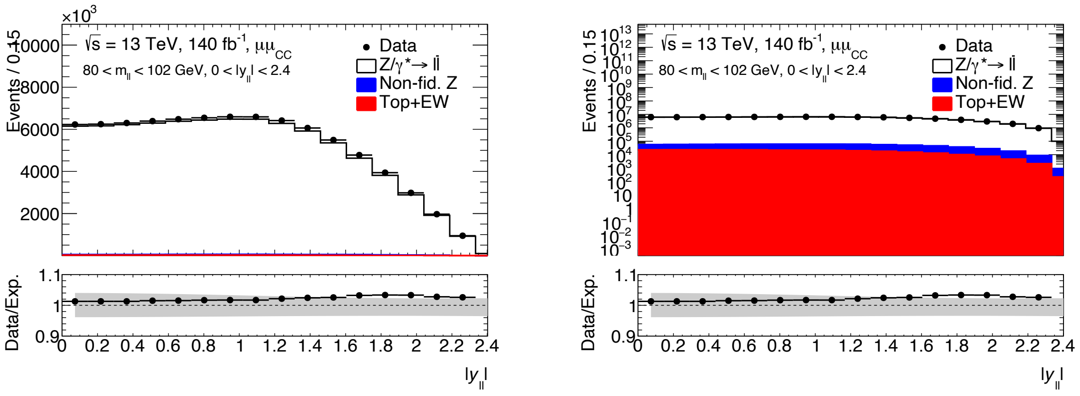

(c) $y_{\mu\mu}$ distribution. (d) Logarithmically scaled $y_{\mu\mu}$ distribution.

Figure 5.9: Signal region inclusive control plots of the dilepton rapidity distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

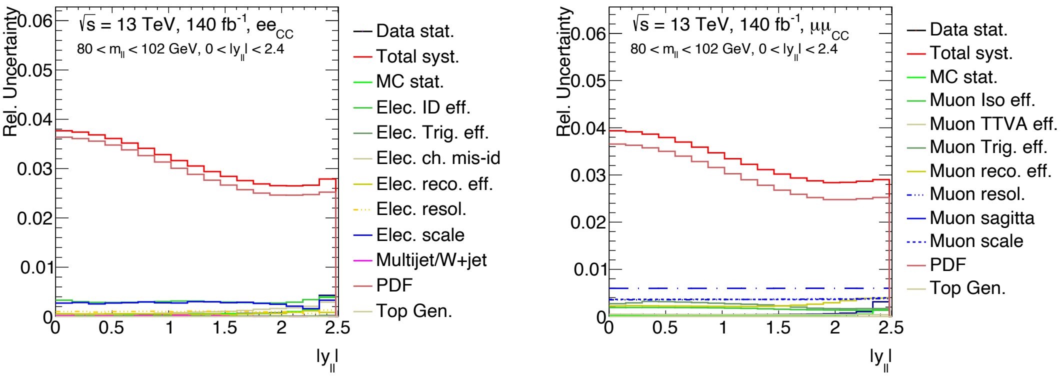

(a) y_{ee} uncertainty breakdown.

(b) $y_{\mu\mu}$ uncertainty breakdown.

Figure 5.10: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the y_{ll} distribution for the inclusive signal region in both the electron and muon channels.

Unlike the rapidity distribution, the PDFs have no similar effect on the azimuthal angle of Z production, meaning that there should be no preferred azimuthal angle of the Z. As shown in Figure 5.11, this is exactly the case for both the electron and muon channels. Whilst there are some small statistical fluctuations visible in the agreement of the data to the expectation for all the plots in Figure 5.11, these fluctuations remain within the uncertainty range.

No matter which dilepton variable is considered, the most dominant uncertainty in the data to expectation agreement plots is that of the PDF variations as is seen in Figures 5.6, 5.8, 5.10 and 5.12. However, depending on the variable, different sources of uncertainty can also become important. For the invariant mass breakdown in Figure 5.6, in both channels, the second and third largest sources of uncertainty are from the resolution and scale of the underlying leptons respectively. These scale and resolutions effects arise from the imperfect alignment of the Z boson mass spectra in data and Monte Carlo. For the Z transverse momentum in Figure 5.8, the last bin shows the Monte Carlo statistical uncertainty becoming higher than that of the PDFs for the electron channel but not in the muon channel. This is most likely an artefact of the randomness of the Monte Carlo generation process, in the sense that simply not enough events were generated in this region of phase space for the electrons but were for the muons.

5.7.2 Single Lepton Variables

The dilepton quantities are all reconstructed from the measured values of the individual leptons, meaning that these underlying distributions must also be checked to ensure that they individually make sense. The pseudorapidity distributions of both the leading and subleading leptons for the electron and muon channels in Figures 5.13 and 5.14 respec-

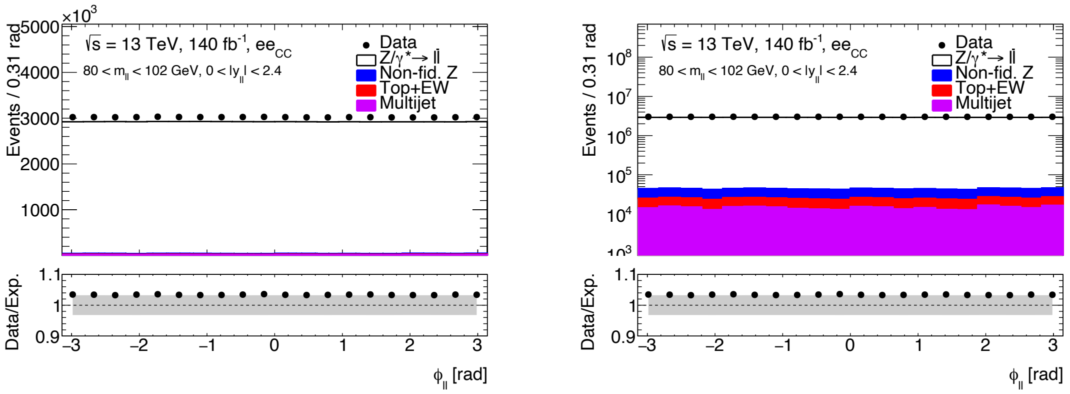

(a) ϕ_{ee} distribution. (b) Logarithmically scaled ϕ_{ee} distribution.

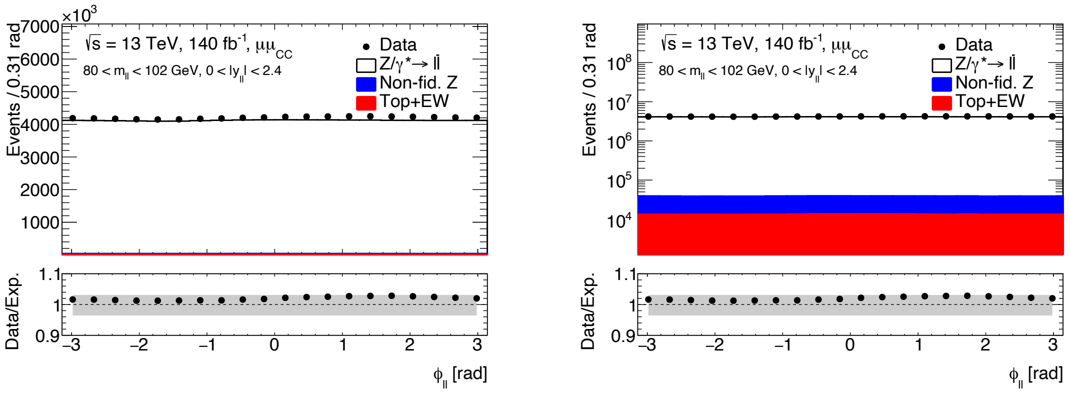

(c) $\phi_{\mu\mu}$ distribution. (d) Logarithmically scaled $\phi_{\mu\mu}$ distribution.

Figure 5.11: Signal region inclusive control plots of the dilepton azimuthal distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

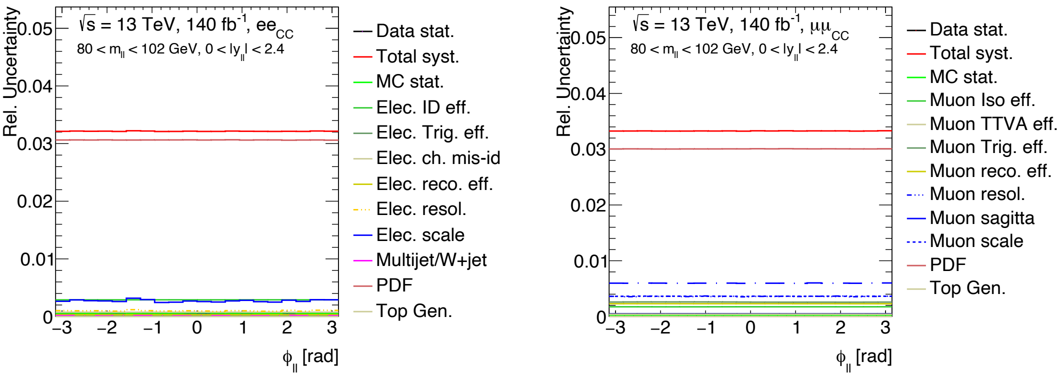

(a) ϕ_{ee} uncertainty breakdown.

(b) $\phi_{\mu\mu}$ uncertainty breakdown.

Figure 5.12: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the ϕ_{ll} distribution for the inclusive signal region in both the electron and muon channels. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

tively, show an approximately symmetric distribution of events around a pseudorapidity of zero, which then slightly increases, before beginning to fall with increasing magnitude of the pseudorapidity. The slowly falling number of events as the pseudorapidity increases is to be expected as the number of events at high dilepton rapidity also decreases, as is seen in Figure 5.9. Of further note is the large drop in the number of events in the region $1.37 < |\eta| < 1.52$ for both the leading and subleading electrons, as can be seen in Figures 5.13a, 5.13b, 5.14a and 5.14b. This corresponds to the transition region between the barrel and endcap calorimeters and is thus cut from the analysis for the electron channel. Since the passage of muons through this region is much less impacted by passive material, this region is not cut for the muon channel. Similarly, the small dip in events at $\eta = 0$ is caused by small gaps between the two sides of the the Inner Detector and Muon Spectrometer. Since the barrel calorimeter is a single object, and therefore without any gap at $\eta = 0$, the dip in events is much more pronounced in the muon channel than for the electron channel.

With the positional aspects of each lepton considered, their kinematics must also be examined. For both the leading and subleading leptons the transverse momentum plots in Figures 5.16 and 5.17 show a rapidly increasing number of events, with a sharp peak just above 45 GeV, before again falling rapidly. As is shown in Figure 5.9, the modal event in the signal region is a Z produced with a very low rapidity meaning that the longitudinal momentum of the Z is correspondingly very small. Figure 5.7 also shows a large number of events with low amounts of Z transverse momentum. Therefore, the decay signature of these events will be when the total energy of the Z will be transferred equally to the two leptons and will be carried almost exclusively as transverse

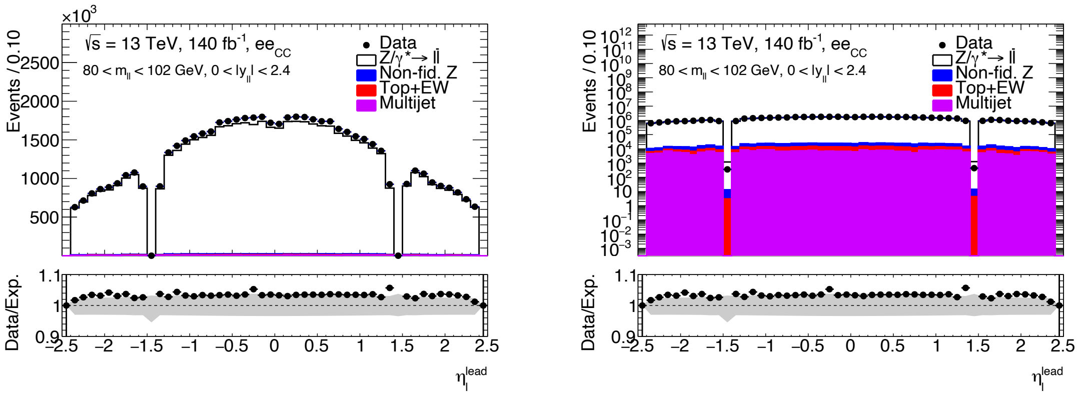

(a) η_e^{lead} distribution. (b) Logarithmically scaled η_e^{lead} distribution.

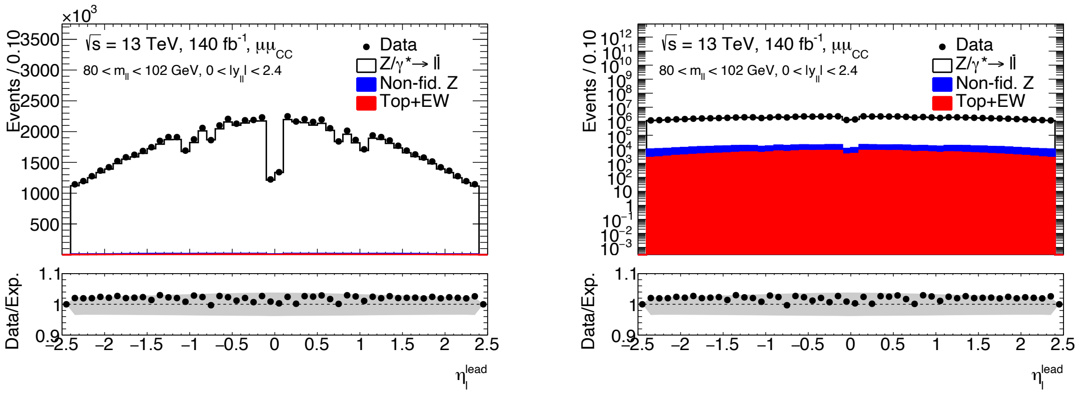

(c) η_μ^{lead} distribution. (d) Logarithmically scaled η_μ^{lead} distribution.

Figure 5.13: Signal region inclusive control plots for the leading momentum lepton pseudorapdity distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

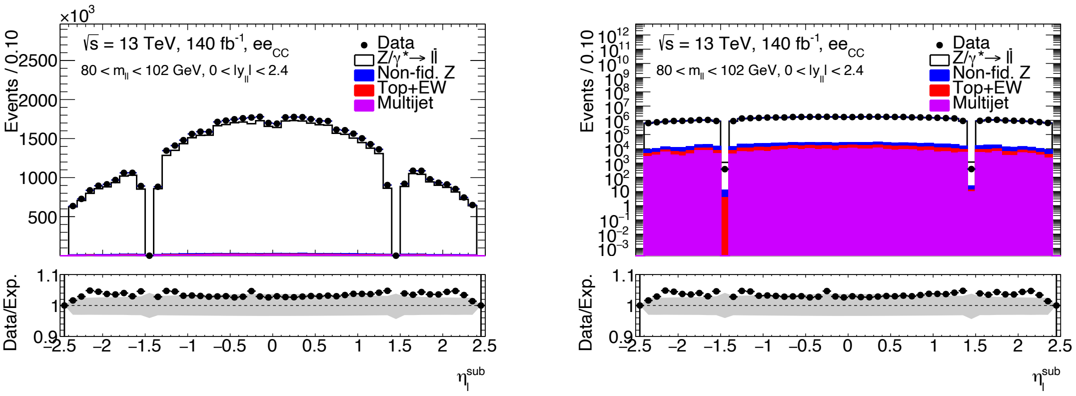

(a) η_e^{sub} distribution.

(b) Logarithmically scaled η_e^{sub} distribution.

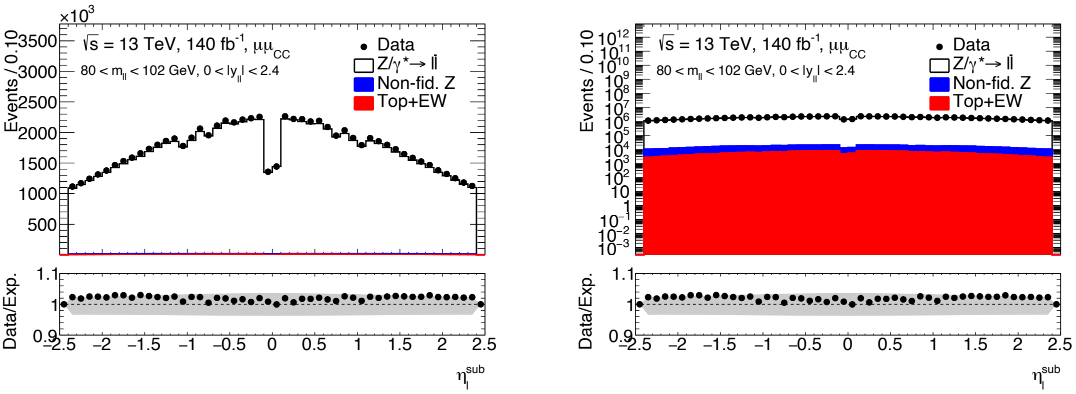

(c) η_μ^{sub} distribution.

(d) Logarithmically scaled η_{mu}^{sub} distribution.

Figure 5.14: Signal region inclusive control plots for the subleading momentum lepton pseudorapdity distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

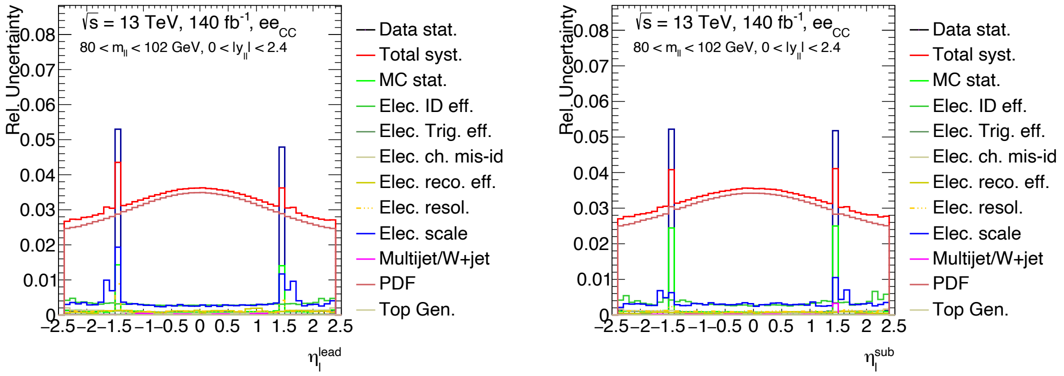

(a) η_e^{lead} uncertainty breakdown.

(b) η_e^{sub} uncertainty breakdown.

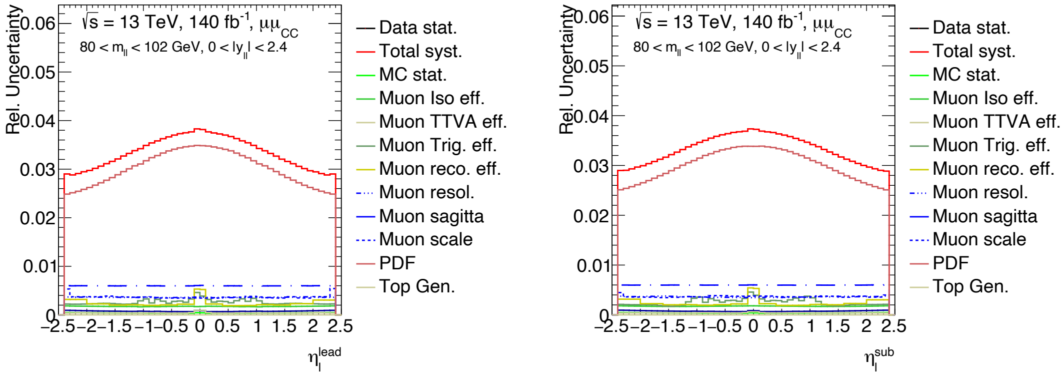

(c) η_μ^{lead} uncertainty breakdown.

(d) η_μ^{sub} uncertainty breakdown.

Figure 5.15: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the η distribution for the inclusive signal region. The top row shows the breakdown for the leading and subleading central electrons, whereas the bottom row displays the same distributions for the leading and subleading muons. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

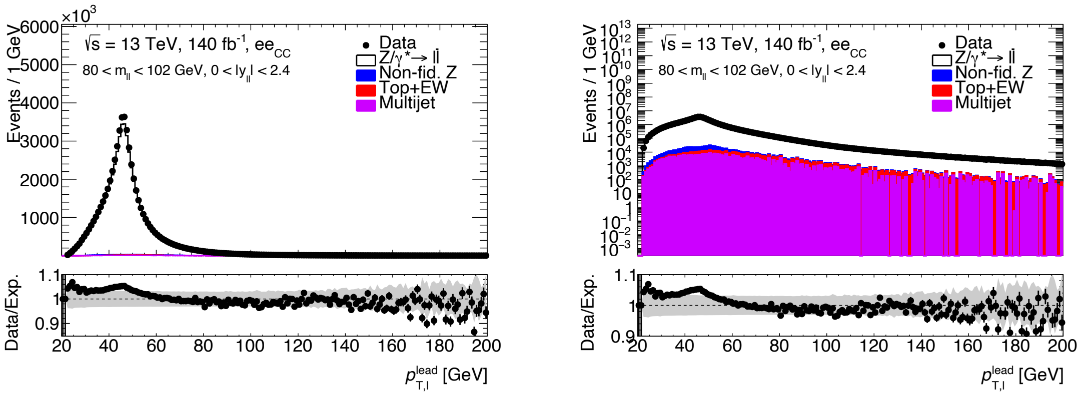

(a) $p_{T,e}^{\mathrm{lead}}$ distribution. (b) Logarithmically scaled $p_{T,e}^{\mathrm{lead}}$ distribution.

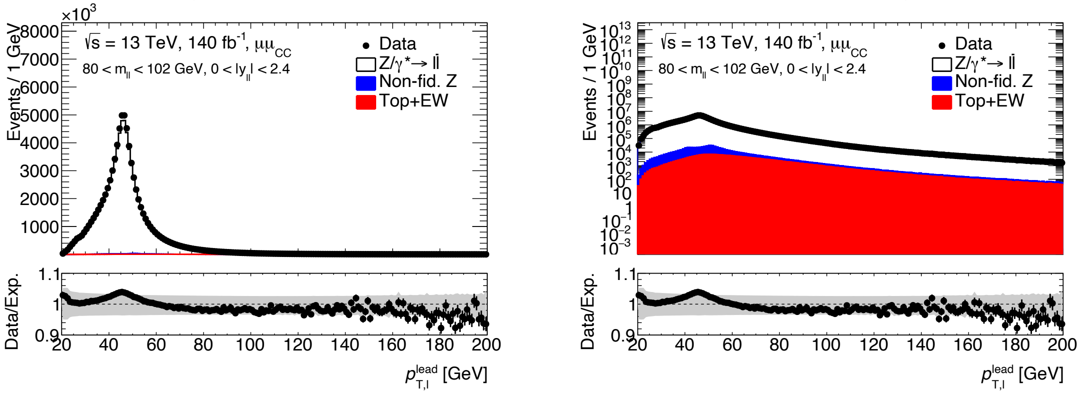

(c) $p_{T,\mu}^{\mathrm{lead}}$ distribution. (d) Logarithmically scaled $p_{T,\mu}^{\mathrm{lead}}$ distribution.

Figure 5.16: Signal region inclusive control plots for the leading momentum lepton transverse momentum distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

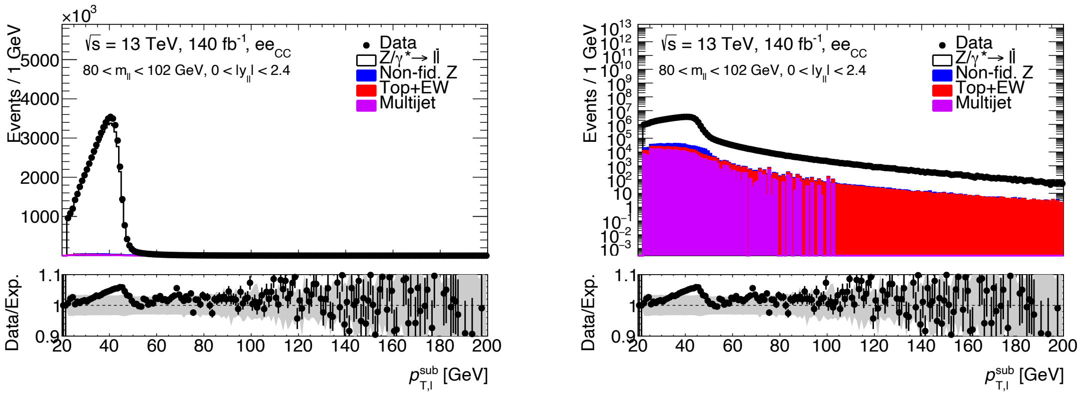

(a) $p_{T,e}^{\mathrm{sub}}$ distribution.   (b) Logarithmically scaled $p_{T,e}^{\mathrm{sub}}$ distribution.

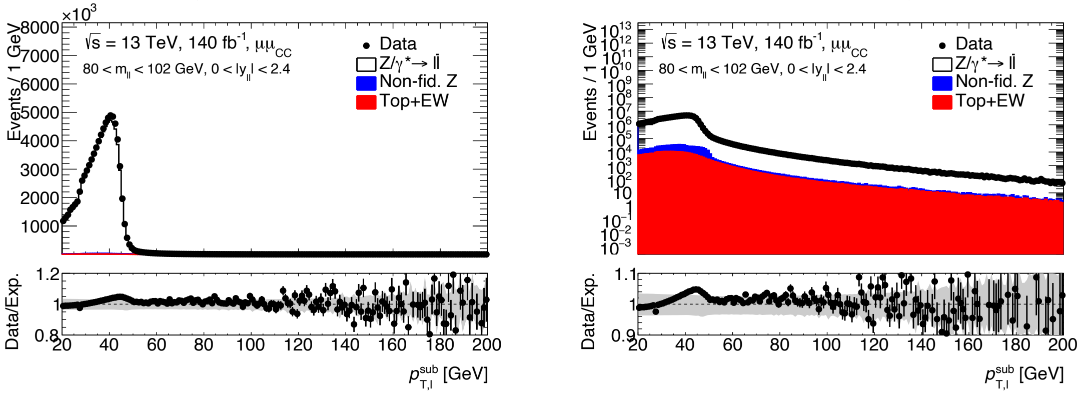

(c) $p_{T,\mu}^{\mathrm{sub}}$ distribution.   (d) Logarithmically scaled $p_{T,\mu}^{\mathrm{sub}}$ distribution.

Figure 5.17: Signal region inclusive control plots for the subleading momentum lepton transverse momentum distribution for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

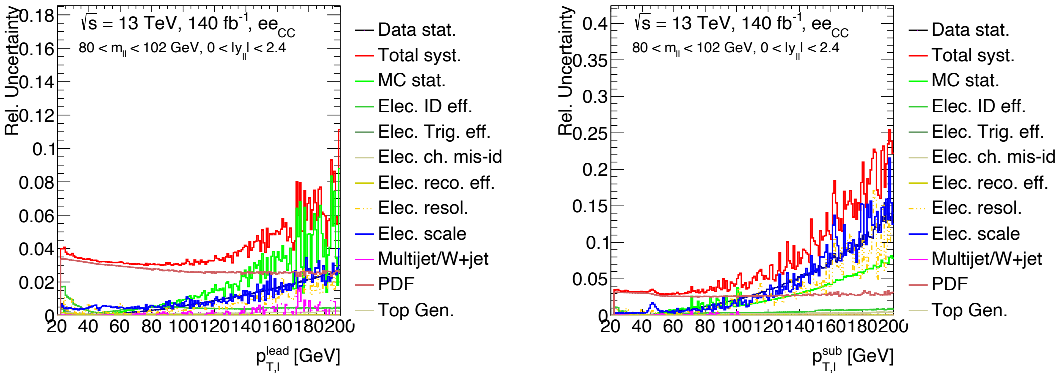

(a) $p_{T,e}^{\text{lead}}$ uncertainty breakdown.

(b) $p_{T,e}^{\text{sub}}$ uncertainty breakdown.

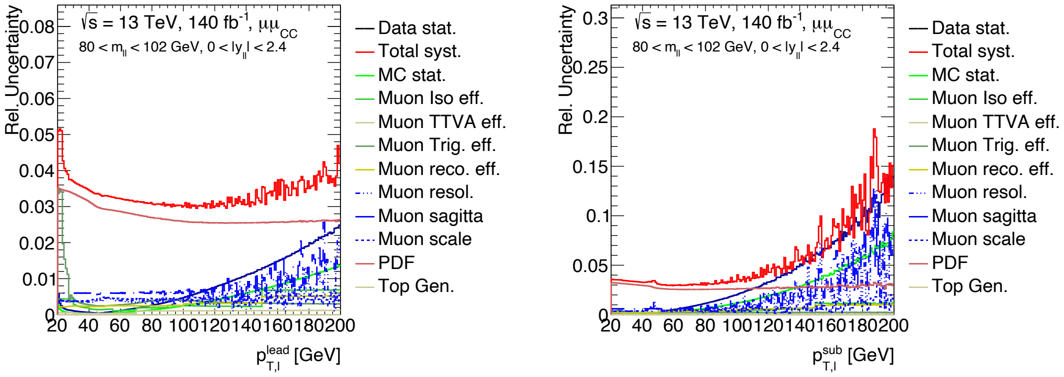

(c) $p_{T,\mu}^{\text{lead}}$ uncertainty breakdown.

(d) $p_{T,\mu}^{\text{sub}}$ uncertainty breakdown.

Figure 5.18: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the $p_{T,lepton}$ distribution for the inclusive signal region. The top row shows the breakdown for the leading and subleading central electrons, whereas the bottom row displays the same distributions for the leading and subleading muons. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

momentum, i.e the transverse momentum should peak at approximately half the of the Z mass. As can be seen in Figures 5.16 and 5.17, this is exactly the case, albeit the subleading transverse momentum is much more asymmetric than the leading transverse momentum by definition. Whilst it may appear concerning that Figures 5.16b and 5.17b show the multijet contribution disappearing in several bins (and indeed, completely at high p_T for the subleading electron), this is due to the data driven nature of the multijet estimate, rather than interesting physics occurring. If the procedure outlined in Section 5.5.1 fails and predicts an unphysical negative amount of multijet in a given region, then it is set to zero to circumvent this issue.

Since it was shown that for the dilepton variables, the by far dominant source of uncertainty was from the PDF variations, this dominance will naturally filter through to the single lepton uncertainty breakdowns as is shown in Figures 5.15 and 5.18. However, for both leading and subleading electrons as well as the subleading muons, the effects of low statistics become dominant at high transverse momentum (approximately 120 GeV for the subleading and 140 GeV for the leading leptons). The effects of the calibration also become important for the subleading leptons, since these are derived from p_T inclusive measurements they do not fully account for the high transverse momentum phase space and therefore result in larger variations. It is also surprising that the muon trigger efficiency systematic spikes to 3.5% relative uncertainty at low p_T of the leading muon in Figure 5.18c. Whilst this systematic does return to a more reasonable value by 23 GeV, more investigation is still required to see where this issue comes from. However, as it is still confined to a relatively small region of phase space, the extraction of the angular coefficients can still continue under the caveat that the values derived from the muon channel do not represent the final measurement.

5.7.3 Collins-Soper Variables

As introduced in Section 2.4.1, Equations (2.26) and (2.27) show that the Collins-Soper variables are explicitly dependent upon the single lepton and dilepton kinematics. Therefore the different p_T and η selections applied to electrons and muons will result in different distributions for the CS frame variables, although the underlying physics remains the same. Exactly this behaviour is seen in Figures 5.19 and 5.20, where the broad characteristics of the CS frame variables are the same across both channels but with some small differences. The overall agreement between the data and the expectation is again very good, with variations on the level of 4% which is again mainly due to the limitations of the Monte Carlo simulation. Figure 5.21 also shows that again the systematic uncertainty is driven by an approximate 3% uncertainty due to PDF modelling, meaning that all experimental effects are well under control.

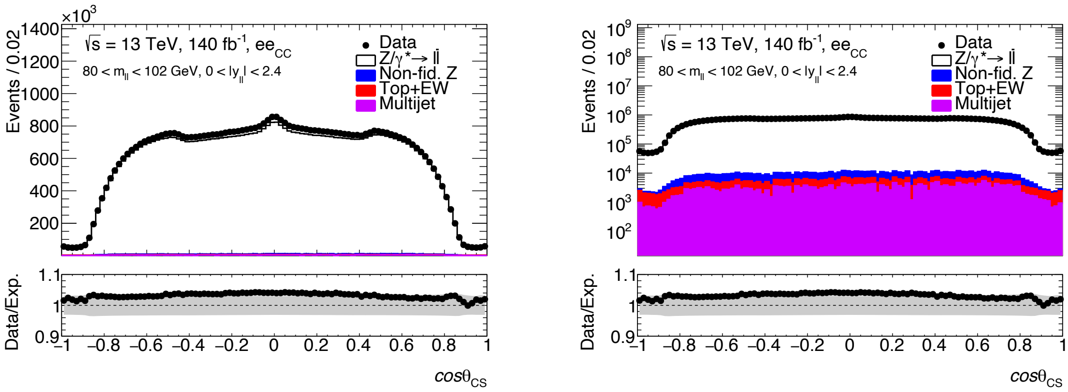

(a) $\cos\theta_{CS,e}$ distribution.

(b) Logarithmically scaled $\cos\theta_{CS,e}$ distribution.

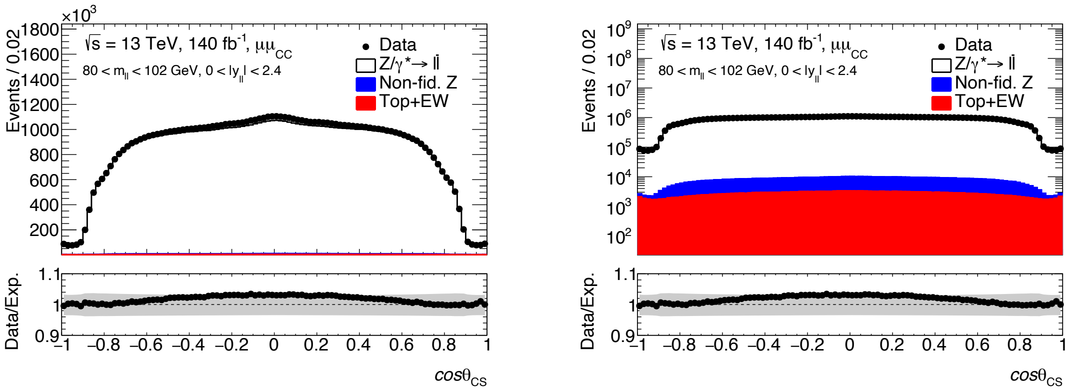

(c) $\cos\theta_{CS,\mu}$ distribution.

(d) Logarithmically scaled $\cos\theta_{CS,\mu}$ distribution.

Figure 5.19: Signal region inclusive control plots of the polar angle distribution in the Collins-Soper frame for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

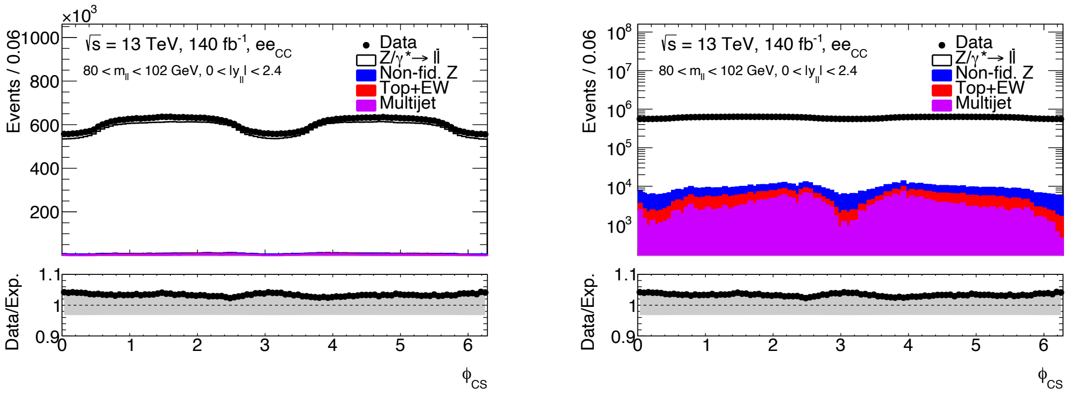

(a) $\phi_{CS,e}$ distribution.

(b) Logarithmically scaled $\phi_{CS,e}$ distribution.

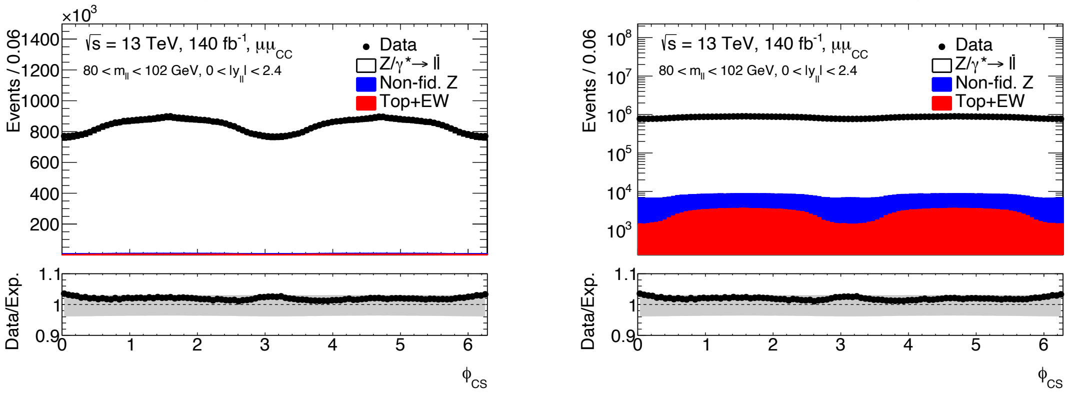

(c) $\phi_{CS,\mu}$ distribution

(d) Logarithmically scaled $\phi_{CS,\mu}$ distribution.

Figure 5.20: Signal region inclusive control plots of the azimuthal angle distribution in the Collins-Soper frame for both the electron and muon channels. The left column shows the distribution nominally and the right column logarithmically to better highlight the magnitude of the background components in each bin. Underneath each distribution, the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) is also displayed. The grey band in this ratio represents the combined uncertainty from all the systematic variations, including also the uncertainty in the luminosity.

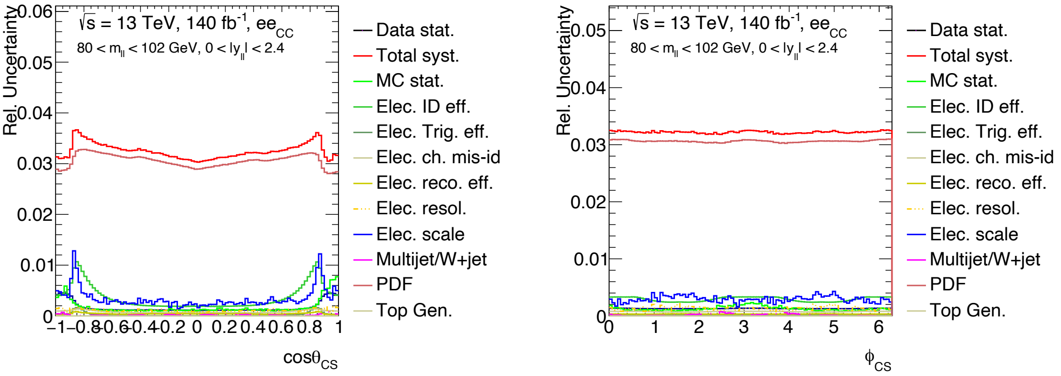

(a) $\cos\theta_{CS,e}$ uncertainty breakdown.

(b) $\phi_{CS,e}$ uncertainty breakdown.

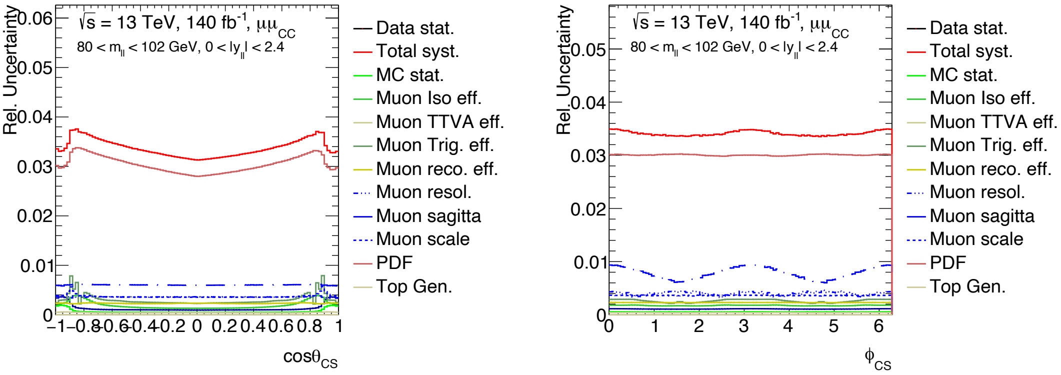

(c) $\cos\theta_{CS,\mu}$ uncertainty breakdown.

(d) $\phi_{CS,\mu}$ uncertainty breakdown.

Figure 5.21: Symmetrised breakdown of all sources of systematic uncertainty in the agreement of data to the expectation (Monte Carlo signal, background and fake estimate) in the Collins-Soper distributions for the inclusive signal region. The top row shows the breakdown for the central electron channel, whereas the bottom row displays the same distributions for the muon channel. The uncertainty in the luminosity is disregarded from these plots since it affects only the overall normalisation.

5.8 Analysis Strategy

5.8.1 $\mathbf{A_i}$ Decomposition

In Chapter 2, the theoretical motivation to study the neutral current Drell-Yan process was established, namely as a method to study both the $q\bar{q} \to Z$ production as well as the $Z \to l^+l^-$ decay itself. Whilst the lepton pair could be taken from any of the three lepton generations, for the purposes of this thesis the lepton pair will be considered to either be an e^+/e^- or μ^+/μ^- pair. By considering the Feynman rules and with the correct parameterisation, it was also motivated in Chapter 2 that the four-fold differential cross section can be decomposed as shown in Equation (5.32) (neglecting any mass dependence since the analysis is inclusive in the region $80 < m_{ll} < 102$ GeV around the Z pole).

$$\begin{aligned}\frac{d\sigma}{dy^Z\, dp_T^Z\, d\phi\, d\cos\theta_{CS}} &= \frac{3}{16\pi}\frac{d\sigma^{U+L}}{dy^Z\, dp_T^Z}\{(1+\cos^2\theta_{CS}) + \frac{1}{2}A_0(1-3\cos^2\theta_{CS}) \\ +A_1(\sin 2\theta_{CS}\cos\phi_{CS}) &+ \frac{1}{2}A_2\sin^2\theta_{CS}\cos 2\phi_{CS} + A_3\sin\theta_{CS}\cos\phi_{CS} + A_4\cos\theta_{CS} \\ &+A_5\sin^2\theta_{CS}\sin 2\phi_{CS} + A_6\sin 2\theta_{CS}\sin\phi_{CS} + A_7\sin\theta_{CS}\sin\phi_{CS}\}\end{aligned} \tag{5.32}$$

In Equation (5.32), all quantities are defined at truth level (denoted with the superscript Z). Correspondingly, at reconstruction level, the transverse momentum (p_T^{ll}), rapidity (y^{ll}) of the Z can be reconstructed from its decay products. From these lab-frame reconstructed quantities, the equivalent Collins-Soper variables can be calculated by boosting from the lab-frame to the Collins-Soper frame. This then allows the angular coefficients and the unpolarised cross section to be extracted from the data using the same method as was used by ATLAS in the corresponding 8 TeV measurement [6]. However, unlike the 8 TeV analysis, this measurement uses a variable sized binning in both the transverse momentum and rapidity of the Z. Specifically the bin edges, at both truth and reconstruction levels, used for this analysis are:

$$\begin{aligned}p_T^{ll}, p_T^Z \in \{&0, 2.0, 4.0, 6.0, 11.4, 14.9, 18.5, 22.0, 25.5, 29.0, 32.6, \\ &36.4, 40.4, 44.9, 50.2, 56.4, 63.9, 73.4, 85.4, 105.0, \\ &132.0, 173.0, 253.0, 400.0, 6500.0\}\ \text{GeV}\end{aligned} \tag{5.33}$$

$$|y^{ll}|, |y^Z| \in \{0, 0.4, 0.8, 1.2, 1.6, 2.0, 2.4\} \tag{5.34}$$

The Collins-Soper variables are also split into eight uniform bins in the range [-1,1] and $[-\pi, \pi]$ for $\cos\theta_{CS}$ and ϕ_{CS} respectively.

5.8.2 Template Method

By inspection, Equation (5.32) can be seen to have the form $\frac{d\sigma}{dy^Z\, dp_T^Z\, d\phi\, d\cos\theta_{CS}} \propto P_8 + \sum_{i=0}^{7} A_i P_i$, where the A_i are the angular coefficients as already defined and the P_i a set of nine analytic polynomials. Since these polynomials directly linked to the angular coefficients, they shall be written explicitly as:

$$
\begin{aligned}
P_0 &= \frac{1}{2}(1 - 3\cos^2\theta) & P_1 &= \sin\theta\cos\phi & P_2 &= \sin\theta\sin\phi \\
P_3 &= \sin 2\theta\cos\phi & P_4 &= \cos\theta & P_5 &= \sin\theta\sin\phi \\
P_6 &= \sin^2\theta\cos 2\phi & P_7 &= \sin^2\theta\sin 2\phi & P_8 &= 1 + \cos^2\theta
\end{aligned}
\tag{5.35}
$$

where the angles θ and ϕ are defined in the Collins-Soper frame i.e they are θ_{CS} and ϕ_{CS} with the CS subscript dropped for brevity. Since these polynomials are fully analytic, then their shapes in the $\cos\theta/\phi$ plane can also be easily calculated and displayed, as is shown in Figure 5.22. If the polynomial is then paired with its corresponding angular coefficient, the product is then known as a template.

To construct the templates at truth level, the angular coefficients must themselves be known. Whilst they are not used as an input variable for Monte Carlo generators, their values can still be extracted through the exploitation of the orthogonality of the harmonic polynomials. The moment of a polynomial, P, can be defined as:

$$
\langle P(\cos\theta, \phi)\rangle = \frac{\int P(\cos\theta, \phi) d\sigma(\cos\theta, \phi) d\cos\theta d\phi}{\int d\sigma(\cos\theta, \phi) d\cos\theta d\phi}
\tag{5.36}
$$

When the moment of each polynomial in Equation (5.35) is calculated, the result depends only on a single angular coefficient.

$$
\begin{aligned}
\langle P_0\rangle &= \frac{3}{20}(A_0 - \frac{2}{3}) & \langle P_1\rangle &= \frac{1}{4}A_3 & \langle P_2\rangle &= \frac{1}{5}A_6 \\
\langle P_3\rangle &= \frac{1}{5}A_1 & \langle P_4\rangle &= \frac{1}{4}A_4 & \langle P_5\rangle &= \frac{1}{4}A_7 \\
\langle P_6\rangle &= \frac{1}{10}A_2 & \langle P_7\rangle &= \frac{1}{5}A_5 & &
\end{aligned}
\tag{5.37}
$$

From the moments presented in Equation (5.37), a reference value of angular coefficient i can be extracted in p_T^Z, y^Z, i.e truth, bin j, A_{ij}^{ref}. Whilst these reference values of the angular coefficients of course do not correspond to detector level results, it is still crucial that these reference values are calculated. In order to extract the angular coefficients in the full phase space from reconstruction level data, it is necessary that all the detector

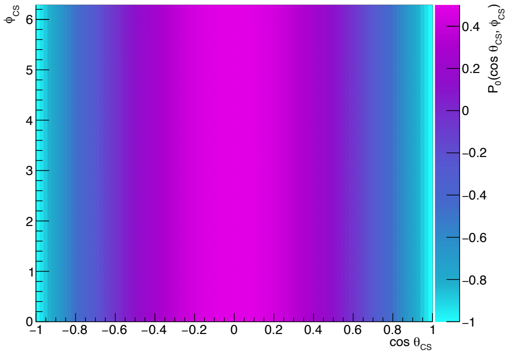

(a) P_0

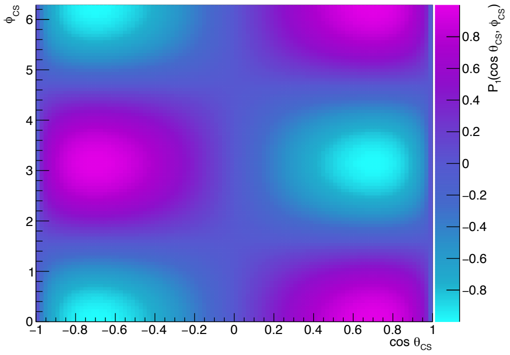

(b) P_1

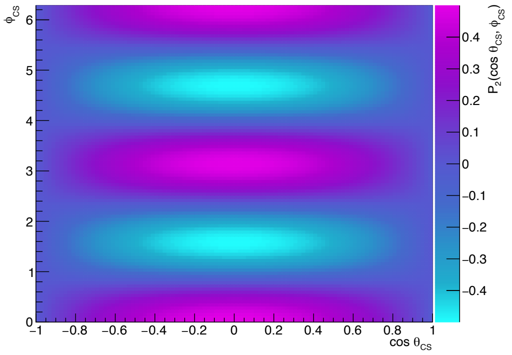

(c) P_2

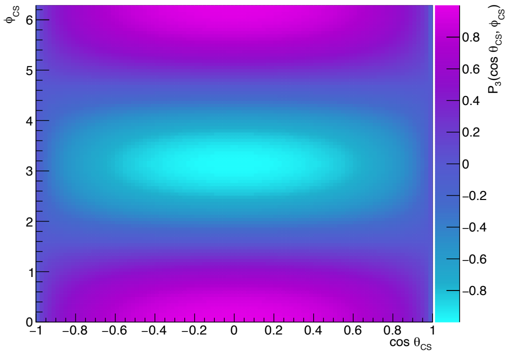

(d) P_3

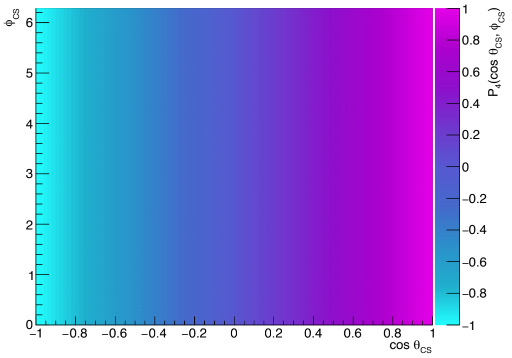

(e) P_4

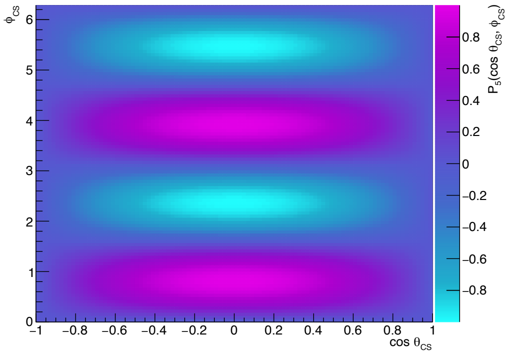

(f) P_5

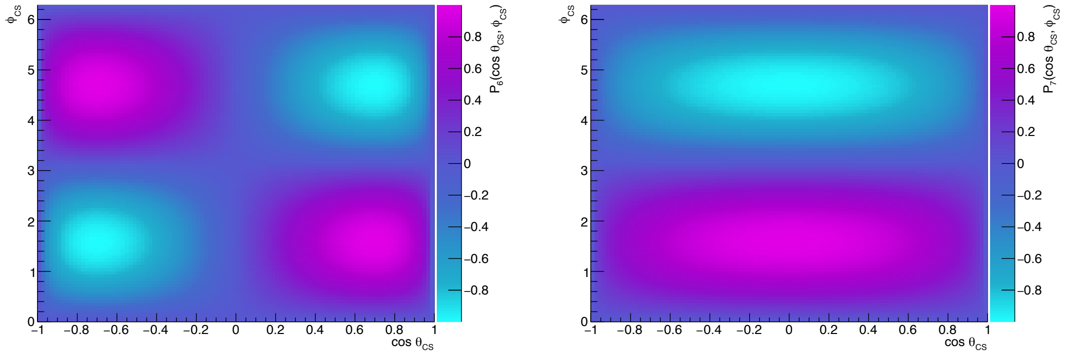

(g) P_6

(h) P_7

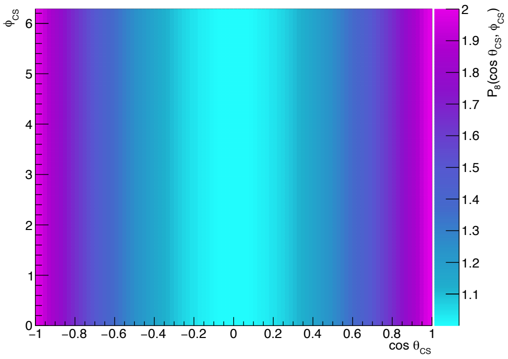

(i) P_8

Figure 5.22: Analytical polynomials P_0 to P_8 as functions of $\cos\theta_{CS}$ and ϕ_{CS} as defined by Equation (5.35).

effects are contained within the P_i i.e the polynomials must be folded from truth level to reconstruction level.

If a general truth level variable, t, with an underlying distribution $f(t)$, is considered, it will be reconstructed as variable r, with the underlying distribution $g(r)$. This reconstruction level distribution is then related to the truth level, through the convolution of $f(t)$ with the probability of reconstructing variable r, given its true value t, $p(r|t)$:

$$g(r) = \int f(t)p(r|t)dt \tag{5.38}$$

Since $p(r|t)$ is cannot be derived analytically, the Monte Carlo simulation is instead used as a proxy for this quantity, $p^{MC}(r|t)$. The overall Monte Carlo event weight, $w_{evt}(r,t)$ contains truth information in the generator weight but also reconstruction information through the application of scale factors, making it a good candidate to estimate $p(r|t)$. If the event weight is then divided by the expected cross section, it will then become a probability distribution function, which defines $p^{MC}(r|t)$:

$$p_j^{MC}(r|t) = \frac{w_{evt}(r,t)}{f_j(\cos\theta^Z_{CS}, \phi^Z_{CS})} \text{ where} \tag{5.39}$$

$$f_j(\cos\theta^Z_{CS}, \phi^Z_{CS}) = \sigma_j \left[(1+\cos^2\theta_{CS})_j + \sum_{i=0}^{7} A_{ij}^{ref} P_{ij}(\cos\theta^Z_{CS}, \phi^Z_{CS}) \right] \tag{5.40}$$

where j denotes a specific p_T^Z, y^Z bin and all other symbols are as previously defined. By construction p^{MC} is, to a good approximation, agnostic to the Monte Carlo generator used since the division of f_j should remove any generator dependent effects in w_{evt}. This, therefore, guarantees generality in the extraction of the angular coefficients which shouldn't be dependent on how the Monte Carlo is produced.

With a suitable proxy for $p(r|t)$ now defined, the truth level polynomials P_{ij} can be folded to a reconstruction level template t_{ij}^n:

$$t_{ij}^n = \sum_{events\,\in (p_T^Z, y^Z)_j (p_T^{ll}, y^{ll})_n} P_{ij}(\cos\theta^Z_{CS}, \phi^Z_{CS}) \cdot p_j^{MC}(r|t) \tag{5.41}$$

where n is the index of the p_T^{ll}, y^{ll} reconstruction level bin that the event falls into and all other variables are as previously defined. As shown in Figure 5.23, when the t_{ij}^n are constructed the broad characteristics of their truth level counterparts in Figure 5.22 are retained but the information is much more coarse due to the binning. The simulated detector smearing and efficiency effects are also visible due to the non-uniform value distribution in the templates.

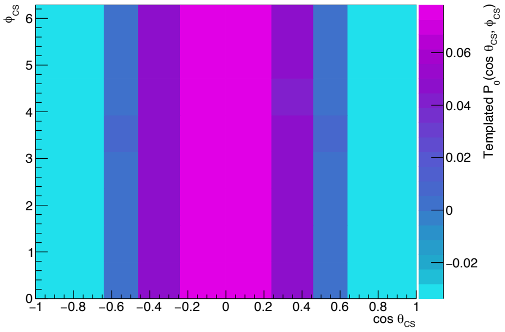

(a) Folded P_0.

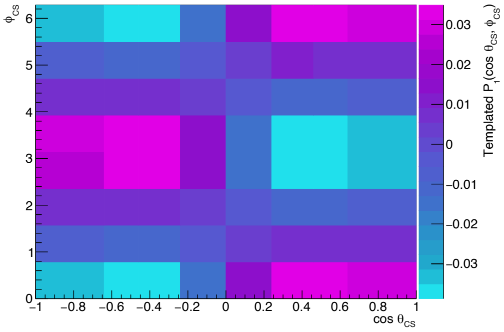

(b) Folded P_1.

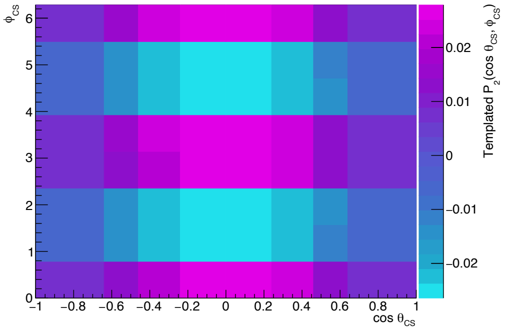

(c) Folded P_2.

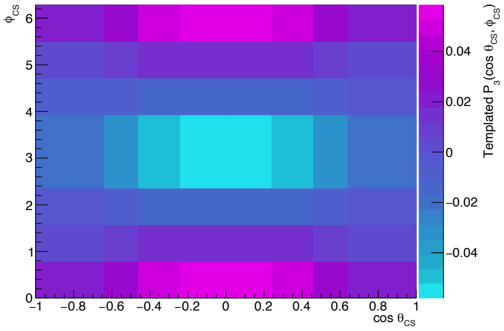

(d) Folded P_3.

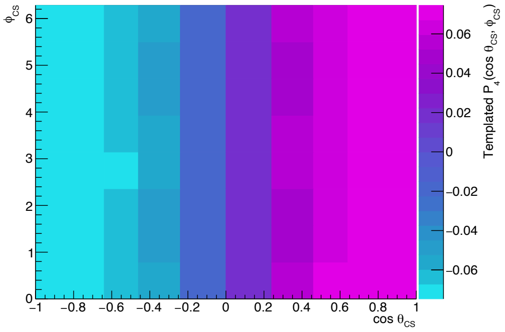

(e) Folded P_4.

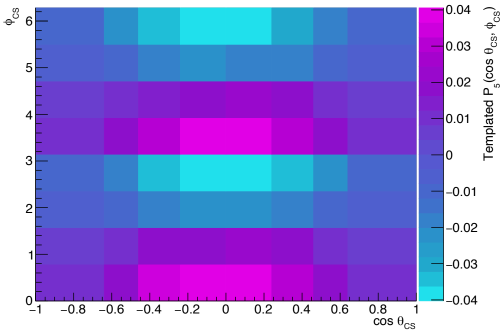

(f) Folded P_5.

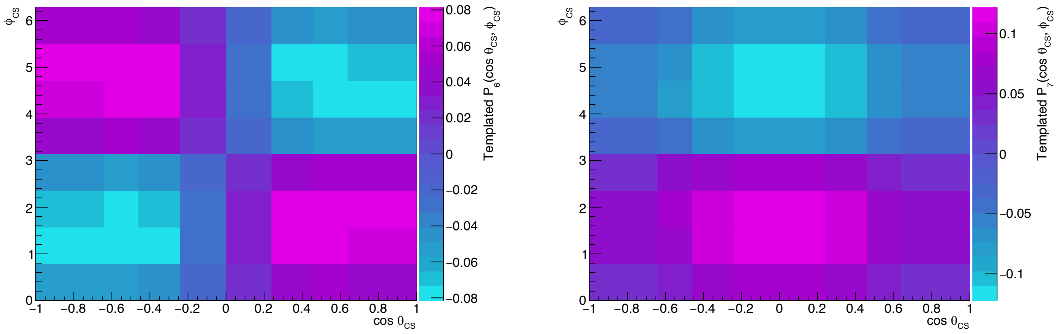

(g) Folded P_6.

(h) Folded P_7.

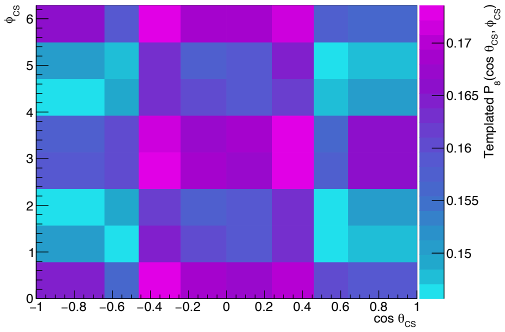

(i) Folded P_8.

Figure 5.23: Folded polynominals of P_0 to P_8 as a function of $\cos\theta_{CS}$ and ϕ_{CS} for the first p_T^Z, y^Z, p_T^{ll}, y^{ll} bin i.e the templates t_{00}^0 to t_{80}^0 as defined by Equation (5.41).

As currently presented, the t^n_{ij} represent only the templated signal polynomials and do not take into account the effect of background events that are also present in the signal region. To account for this, separate background templates are created for the top and electroweak, non-fiducial, and multijet backgrounds. These background templates are normalised to the product of their respective cross section and the integrated luminosity for the simulated backgrounds.

5.8.3 Statistical Methods

With the templates defined, the total number of expected events at reconstruction level in bin n will be given by summing over the expected number of events from all the truth level bins and adding the contribution from the background:

$$N^n_{\text{exp}}(A,\sigma) = \sum_{j=1}^{N_{Z\,bins}} \sigma_j \cdot L \cdot \left[t^n_{8,j} + \sum_{i=0}^{7} A_{ij} \cdot t^n_{ij} \right] + \sum_{Bkgs} T^n_B \tag{5.42}$$

where A is the full set of angular coefficients and σ the unpolarised cross section. As written Equation 5.42 considers only the nominal case and therefore does not include the effect of any systematics which change the number of expected events. In order to get the best estimate of the overall uncertainty, all of the sources of systematic uncertainty outlined in Section 5.6 must be included into Equation 5.42.

The effect of systematic uncertainties are encoded in individual nuisance parameters denoted as β and γ whose combination is denoted as $\theta = \{\beta, \gamma\}$. The β type nuisance parameters each describe a source of uncertainty that can affect both the overall shape and normalisation of the template distribution. These types of uncertainty are constrained with a unit Gaussian distribution $G(0|\beta^m, 1)$, where $\beta^m = \pm 1\sigma$ of the parameter under consideration such that the mean is interpolated between the nominal template i.e mean of zero, and systematic templates i.e $\beta^m = \pm 1\sigma$. In the muon channel there are 62 total β type nuisance parameters, whereas for the electron channel there are 867. The much larger number of β^m in the electron channel is driven by the binning of the identification scale factors, which is much more fine than that used in the muon channel.

In comparison to the β type nuisance parameters, the γ type of affect only the overall normalisation of the template and do not affect the shape. For the extraction of the angular coefficients, this means that each individual γ^n represents the statistical limitation of the generated Monte Carlo dataset in a given analysis bin, n. Since the total number of Monte Carlo events in bin n is an effective number of events made from the weighting of the total number of generated events to the observed luminosity and detector efficiencies, it does not necessarily represent the total number of expected events by the Standard Model. This means that the γ^n can be constrained by considering the Poissonian probability of observing $N^n_{\text{MC,eff}}$ given an expectation of $\gamma N^n_{\text{MC,eff}}$, i.e $P(N^n_{\text{MC,eff}}|\gamma N^n_{\text{MC,eff}})$. Given the 25 p^{ll}_T, 6 y^{ll}, 8 $\cos\theta_{CS}$ and 8ϕ_{CS} bins, the total number of bins is $N_{\text{bins}} = 25 \times 6 \times 8 \times 8 = 9600$. Correspondingly, there are therefore 9600 γ^n nuisance parameters that need to be constrained.

With the systematics defined, their effect on the number of expected events can now be incorporated into Equation (5.42). The β nuisance parameters with their effect on the calibration and efficiencies of the leptons will have a corresponding effect on the signal and background templates, meaning that all templates will become a function of β. Each reconstruction bin must also be multiplied by the corresponding γ parameter to incorporate the power of the Monte Carlo statistics. The expected number of events in reconstruction bin n can then be written as:

$$N^n_{exp}(A, \sigma, \theta) = \left\{ \sum_{j=1}^{N_{Z\,bins}} \sigma_j \cdot L \cdot \left[t^n_{8,j}(\beta) + \sum_{i=0}^{7} A_{ij} \cdot t^n_{ij}(\beta) \right] + \sum_{Bkgs} T^n_B(\beta) \right\} \cdot \gamma^n \tag{5.43}$$

In order to finally fit the angular coefficients and unpolarised cross section, a likelihood function, $\mathcal{L}$, must be constructed that takes into account both the nominal case and the systematic effects. This likelihood is taken as the product of the probability of observing N events given a certain number of expected events and the previously defined constraint of the γ^n nuisance parameters, in a single reconstruction level bin n. The product of these probabilities is then multiplied with the same probability from all the other reconstruction level bins and multiplied with the Gaussian constraints from the β nuisance parameters. Mathematically, this can be expressed as:

$$\mathcal{L}(A, \sigma, \theta) = \prod_{n}^{N_{\text{reco bins}}} \left\{ P(N^n_{\text{obs}} | N^n_{\text{exp}}(A, \sigma, \theta)) P(N^n_{\text{MC, eff}} | \gamma^n N^n_{\text{MC, eff}}) \right\} \cdot \prod_{m}^{M} G(0|\beta^m, 1) \tag{5.44}$$

Explicitly, $N_{\text{reco bins}}$ includes every $p_T ll$, y^{ll}, $\cos\theta_{CS}$ and ϕ_{CS} bin meaning that in total $N_{\text{reco bins}} = 25 \times 6 \times 8 \times 8 = 9600$.

Maximum Likelihood Fitting

From Equation (5.44), it would be possible to fit the angular coefficients and unpolarised cross section by simply using an iterative process to find the values that would maximise the value of the likelihood function. However, by the Neyman-Pearson lemma [73] the most powerful statistical test is instead the likelihood ratio, Λ [50]. In the case of this analysis, the likelihood ratio can be defined as the ratio of the likelihood of a single angular coefficient (or the unpolarised cross section) in truth bin j given the nuisance parameters θ, to the global likelihood of all angular coefficients and the unpolarised cross section.

$$\Lambda(A_{ij},\sigma_j) = \frac{\mathcal{L}(\hat{A}(A_{ij}),\hat{\sigma_j}(\sigma_j),\hat{\theta}(A_{ij},\sigma_j))}{\mathcal{L}(\hat{A},\hat{\sigma},\hat{\theta})} \tag{5.45}$$

Instead of maximising Λ to fit all the parameters of interest, the test statistic $q_{A_{ij}} = -2\log\Lambda(A_{ij})$ is instead minimised using the Newton minimisation algorithm [74] within the FTL framework that was developed specifically for the analysis [21]. By Wilk's theorem [88], for very large datasets the distribution q asymptotically approaches the χ^2 distribution with which for this analysis is satisfied by the unprecedented size of the ATLAS Run 2 dataset. This asymptotic behaviour of q allows for the uncertainties in the fit values to be easily computed since $q = q_{min} \pm 1$ will correspond to the 1-σ confidence intervals of the χ^2 distribution. By performing the fit in this way, the final fit values of the angular coefficients and the unpolarised cross section will be, from a statistics standpoint, their optimal values with the smallest possible uncertainty.

Goodness of Fit

As described in the previous section, the fit process cares only that the minimisation procedure converges to provide a value of the parameter that is being evaluated and does not provide any indication as to the quality of the fit itself. Instead, this is evaluated through Pearson's χ^2 test [76]. The χ^2 function is defined as [50]:

$$\chi^2 = \sum_{i=1}^{n} \frac{(x_i - \mu_i)^2}{\sigma_i^2} \tag{5.46}$$

where the i represent the individual observed data points, x_i the fit value, μ_i the model value i.e the expectation value and σ_i the uncertainty in the fitted value. It is important to note that here a data point refers to a single fit parameter in a single bin, rather than a sum over all events recorded by ATLAS.

By itself, the χ^2 value is not necessarily a good metric to measure the quality of the fit, since its value depends on the total number of parameters being fit, which for this analysis is the 8 angular coefficients plus the unpolarised cross section in 25 p_T^Z bins and 6 y^Z bins meaning there are 1350 individual parameters in the central electron and muon channels. Therefore, it is better to consider the value χ^2/N_{dof}, where N_{dof} is the number of degrees of freedom within the fit:

$$N_{\text{dof}} = N_{\text{bins}} - N_{\text{fit parameters}} \tag{5.47}$$

where N_{bins} is the total number of bins in p_T^Z, y^Z, $\cos\theta_{CS}$ and ϕ_{CS} and $N_{\text{fit parameters}}$ the total number of parameters that have been fit as previously discussed. Considering the 25 p_T^Z bins, 6 y^Z, 8 $\cos\theta_{CS}$ and 8 ϕ_{CS} bins, $N_{\text{dof}} = 8250$ in the central electron and muon channels. If the fit has converged well, then the corresponding χ^2/N_{dof} should be approximately one. If the value is, however, much larger than this, then the fit has likely either not converged or the model is not an appropriate description of the data. Conversely if $\chi^2/N_{\mathrm{dof}} \approx 0$, then it is likely that overfitting has occurred since the data is almost perfectly represented by the model. Since the data also includes statistical fluctuations, no model should be able to reproduce it with 100% accuracy.

Through the χ^2/N_{dof} value, the global goodness of fit can be estimated. However, this value does not indicate the effect of the individual nuisance parameters on the fit. Therefore, for a given nuisance parameter, θ, associated with a source of systematic uncertainty defined in Section 5.6, its corresponding pull is defined as:

$$p = \frac{\hat{\theta} - \theta_0}{\sigma_\theta} \tag{5.48}$$

where $\hat{\theta}$ is the best fit value obtained from minimising the likelihood ratio for that nuisance parameter, θ_0 is the expected value and σ_θ the uncertainty in the nuisance parameter. Since the underlying distribution of the nuisance parameters is considered to be a normally distributed Gaussian, the uncertainty in the pull should correspondingly also be one. If this is not the case and $\Delta p < 1$, then the nuisance parameter is said to be constrained by the data. If a nuisance parameter is constrained, then the corresponding systematic uncertainty is also smaller in size.

Assuming that the Monte Carlo modelling of the system is correct, means that the observed value should not lie too far from the expected value, i.e that the numerator of Equation (5.48) should be approximately zero. If the central value of the pull function differs greatly from this, the nuisance parameter is said to be pulled. A pulled nuisance parameter therefore means that there is some aspect of the data that is not modelled correctly by the underlying Monte Carlo simulation.

Considering the pulls on all systematic uncertainties is therefore another way of judging the overall quality of the fit and the underlying model. If many of the nuisance parameters differ by more than one sigma, or greater, from the expectation or are greatly constrained (or some combination of both), then the underlying model may have to be revisited in order to have an acceptable level of systematic uncertainty in the resultant measurement. Since the full Run 2 ATLAS dataset used for this measurement is more than five years old, the underlying models used for the systematic variations have been very thoroughly tested and refined. This means that for this measurement, the pulls should show no significant deviations from the expectation value and also not be strongly constrained.

Asimov Dataset

Whilst the quality of the fit can be determined by looking at the observed data, if any problems are observed that must be corrected additional conscious or unconscious biases may be introduced into the analysis. Instead the fit process can be performed by using an

‘Asimov dataset’. This dataset replaces the observed number of events with the expected number of events, meaning that the values it returns will always equal the expectation value. This has the effect that the fit values is exactly equal to those calculated using the moments method since the folding of the templates to reconstruction level mean that the angular coefficients are effectively being extracted at truth level. However, unlike the moments method the Asimov dataset also allows for the effect of systematic uncertainties on the parameters of interest to be gauged.

When using the Asimov dataset, any problems in the fitting process will also show up when the pull plots for each nuisance parameter are created. By definition, the Asimov pull plots should have a central value of zero and the uncertainty equal to one. Therefore, if these two properties are not observed (modulo some small variation), a problem has occurred somewhere in the data preparation or fitting process. The Asimov dataset therefore acts as a validation dataset to ensure that there are no problems with the fitting process itself, whilst still remaining blinded to the actual observed data.

5.9 Single Lepton Channel Results

5.9.1 Electron Channel Pull Plots

The combination of the electron and muon channels will naturally give the most precise measurement of the angular coefficients. Before this fit is performed it is important to validate each channel individually in order to spot any potential problems. For the electron channel the overall fit result returned a $\chi^2/N_{\mathrm{dof}} = 11183.3/8250$ suggesting that a good agreement between the fitted values and the expectation values. However, the χ^2 value does not indicate if there were any issues with specific systematic groups in the fit. As mentioned in Section 5.8.3, the pull plots provide an important check on the quality of the fit and should therefore be assessed before unblinding the results of the fit in each channel. For the electron channel, the full set of pull plots can be seen in Appendix A with some examples presented in the following. The nuisance parameters pertaining to the electron charge identification and identification efficiencies in Figure 5.24 show that the data is described very well by the underlying model as the vast majority of nuisance parameters almost exactly match the expectation value. For those that show some deviation from the expectation, almost all of these are within $\pm 1\sigma$ and show very little constraints. Only the parameters EL_EFF_ID_CorrUncertaintyNP7 and EL_CHARGEID_SYStotal show pulls of approximately 1.7σ and -1.1σ respectively. However, since these are the only two identification efficiency nuisance parameters that are pulled by more than one sigma, it is more likely that the model is simply not as accurate as it could be in this single region of phase space, rather than a problem with the fit.

The case is similar for the trigger and reconstruction efficiencies in Figure 5.25, since most of these pulls lie well within one sigma of the expectation. Only the nuisance parameters EL_EFF_Trigger_SIMPLIFIED_UncorrUncertaintyNP5 and EL_EFF_Reco_CorrUncertaintyNP2 lie outside of the one sigma range, with the greatest of these deviations being the latter with a pull of 1.28σ from the expectation. However, as for the identification efficiency parameters, since the overall number of nuisance parameters that are pulled in this systematic group is small in comparison to the total, the overall fit results should not be adversely affected.

Since the underlying material models are highly unlikely to exactly match the material distributions of the physical detector, it would be expected that there should be some spread away from zero in the corresponding calibration pull plots. In Figure 5.26, this is exactly what is seen, with only a single nuisance parameter outside of the $\pm 1\sigma$ range (EG_SCALE_G4). Whilst the nuisance parameters on the resolution and the in situ scale systematic may be slightly conservative since they are constrained, the overall behaviour of the calibration nuisance parameters means that they do not pose a problem to the extraction of the angular coefficients.

In contrast to the other systematic groups displayed so far, where the vast majority of the nuisance parameters agree very well or at least within 1σ of the expectation value, the fake factor variations display very large pulls. Indeed, 13 of the 30 variations have a pull greater than $\pm 2\sigma$, which indicates this is more than just a statistical fluctuation. However, as is shown in the background composition plots in Figure 5.4, the average amount of multijet in the signal region is very small, typically less than 1% when viewed inclusively and depending on which $\cos\theta_{CS}/\phi_{CS}$ bin is being analysed. Indeed, the multijet accounts for only 0.5% of events in the signal region when integrating over all analysis bins. Correspondingly, the systematic uncertainty on the angular coefficients resulting from the variation of the fake factors will still be low, even if thirteen of the nuisance parameters are pulled by an amount greater than 2σ. Therefore, the current non-ideal state of these nuisance parameters can be overlooked with the caveat that they will be revisited in the future.

The final set of nuisance parameters that must be considered are those from the PDF variations. These are implemented as variations to the PDF weight that is applied to every simulated event through the LHAPDF6 [30] tool. In Figure 5.28 the corresponding pull plots for these nuisance parameters are shown to agree well with the pre-fit expectation, with no parameters outside of the $\pm 1\sigma$ range. Therefore, after checking the quality of the fit through the pull plots there are no major issues with the electron channel.

5.9.2 Muon Channel Pull Plots

Since the fit using the muon channel is performed separately to the electron channel, it is also necessary to check the pull plots for this channel before unblinding the results. However, the muon channel use a simpler systematic model that focusses on broader regions of phase space in comparison to the electron channel where better results are achieved with a finer binning of scale factors. This therefore means that the number of nuisance parameters in the muon channel (covering the same broad areas as the electron channel)

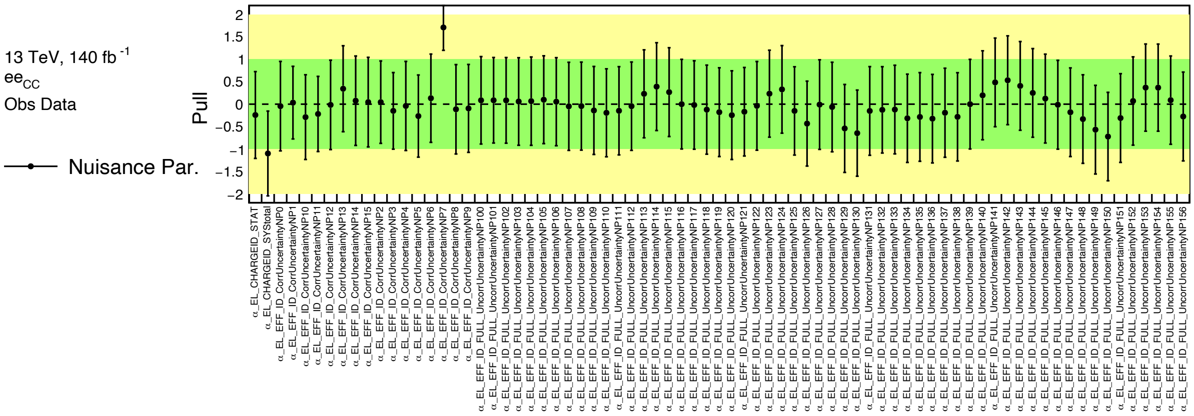

Figure 5.24: Pull plots for the electron (charge) identification efficiency variations.

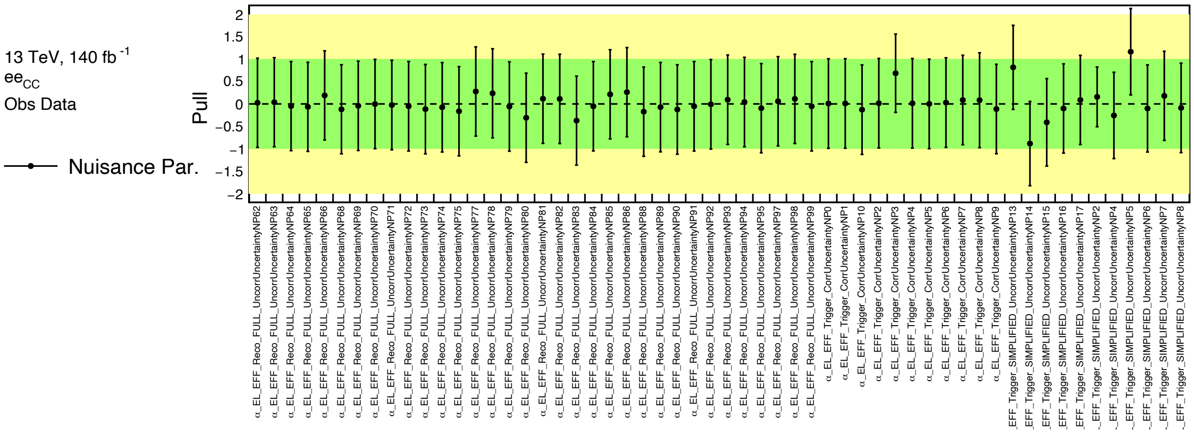

Figure 5.25: Pull plots for the electron reconstruction and trigger efficiency variations.

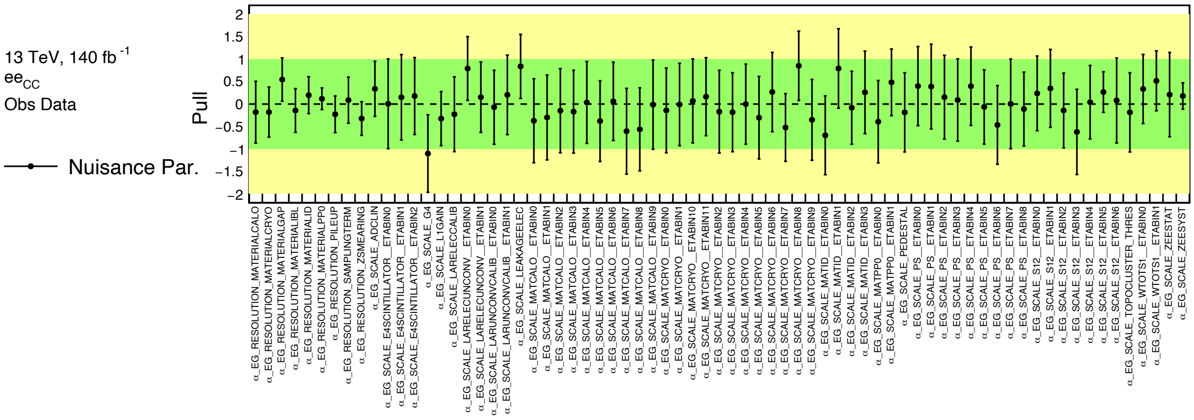

Figure 5.26: Pull plots for the electron calibration variations.

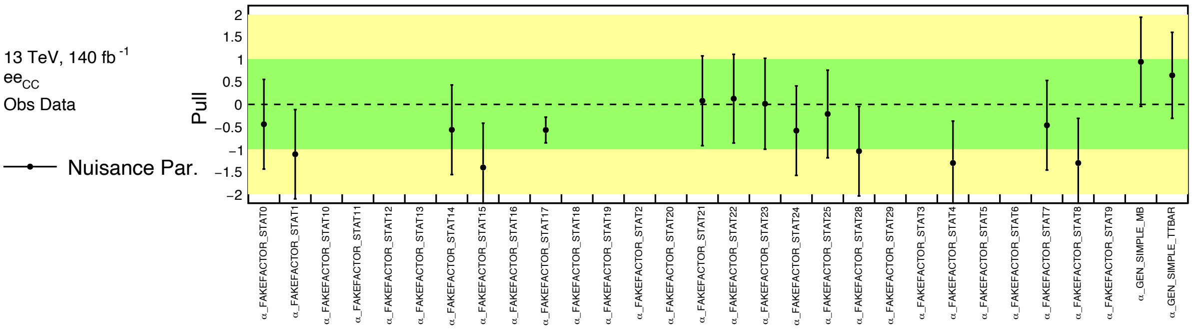

Figure 5.27: Pull plots for the electron fake factor and background Monte Carlo generator variations.

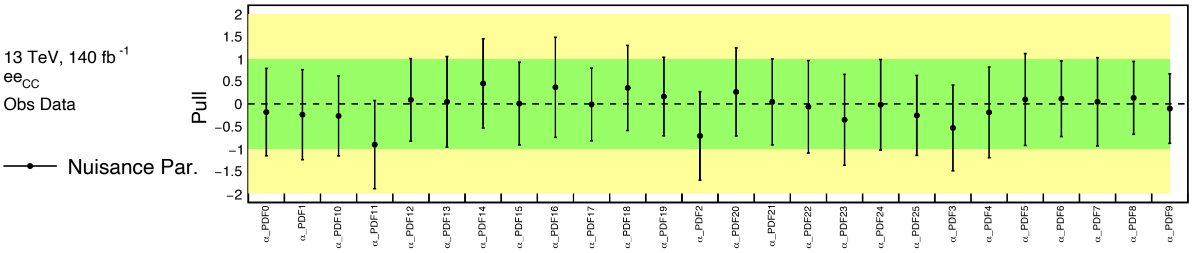

Figure 5.28: Pull plots for the electron PDF variations.

is significantly reduced in comparison to the electron channel. The full set of muon nuisance parameter pull plots are provided in Appendix B . Since $\chi^2/N_{\mathrm{dof}} = 8584/8250$ for the muon channel fit, it is unsurprising that the vast majority of the muon nuisance parameters agree well with the pre-fit exception. However, Figure 5.29 shows that the reconstruction efficiency nuisance parameters MUON_EFF_RECO_STAT and MUON_EFF_RECO_TRUTH display a pull value of -1.8σ which is significantly outside of the $\pm 1\sigma$ range. Since the muon nuisance parameters cover much more broad regions of phase space than their electron counterparts, a pull of this size cannot be as easily overlooked as it can be in the electron channel. This suggests that these nuisance parameters require further investigation to determine whether this is a problem of the fit or with the application of the nuisance parameters themselves. Similarly, the trigger nuisance parameter MUON_EFF_TrigStatUncertainty is pulled and constrained to $-2.7 \pm 0.5\sigma$ and MUON_EFF_TrigSystUncertainty to $0.6 \pm 0.2\sigma$. Therefore, with these nuisance parameters displaying behaviour that is far from the expectation, the muon results cannot be regarded as finalised. It is still, however, valid to look at the the unblinded result plots in the muon channel albeit with the caveat that they are subject to further improvement.

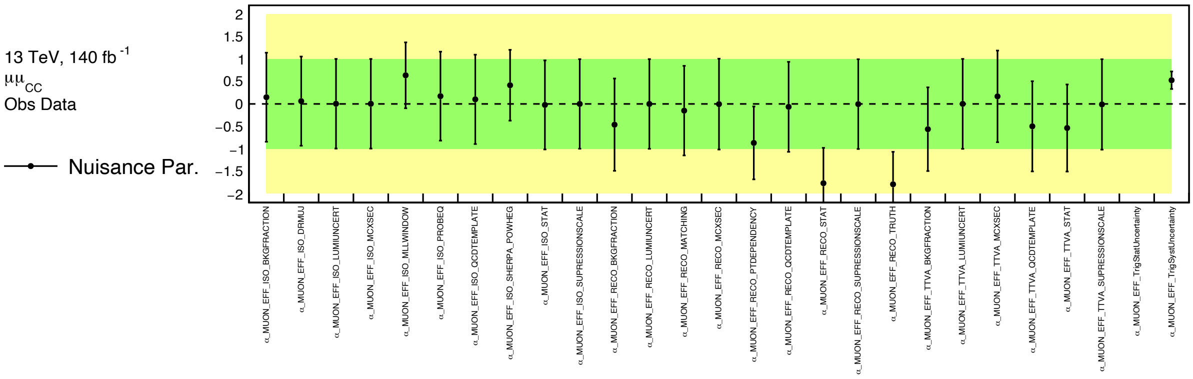

Figure 5.29: Pull plots for the muon efficiency variations.

Breakdown of the Uncertainty

Each of the systematic groups presented in Sections 5.9.1 and 5.9.2 will represent a source of systematic uncertainty in the measurement of each angular coefficient. Therefore, Figure 5.30 shows the uncertainty breakdown in the electron channel by systematic group, including the uncertainty from the data and Monte Carlo statistics, for each angular coefficient at low and high rapidity. From Figure 5.30 it is clear that the largest source of uncertainty is the data statistics, even at low rapidity where the number of events is at its most plentiful. Whilst it may be surprising that this analysis is statistically limited given the large size of the ATLAS full Run 2 dataset, the Drell-Yan cross section has also never been measured four-fold differentially before. It is, therefore, due to this very high level of granularity (particularly in the dilepton transverse momentum) that the analysis bins remain statistically limited. Similarly, the leading systematic for each angular coefficient is the Monte Carlo statistics, which can only be improved further by generating a significantly larger number of Monte Carlo events. After the Monte Carlo statistics, the next two leading systematics (of approximately the same significance) are the electron resolution and scale variations. Given the dependence of $\cos\theta_{CS}$ and ϕ_{CS} on the lepton momentum, it is not surprising that these variations have a, relatively, significant impact on the uncertainty in the angular coefficients. Indeed, given enough statistics in both data and Monte Carlo, these would have been the leading systematics which underscores the importance of having as accurate an electron calibration model as possible.

Even though the analysis is statistically limited, the level of precision that is, however, reached is remarkable. Figure 5.30q shows that the relative uncertainty in the unpolarised cross section at low rapidity remains below the percent level for dilepton transverse momenta below 125 GeV. This means that in this single analysis bin, the level of precision reached is greater than the entire rapidity integrated measurement performed by ATLAS using the Run 1 dataset. Whilst this precision is statistically limited below 50 GeV of p_T^{ll}, it becomes systematically limited above this point due to the increasing significance of the background and Monte Carlo generator background ('Top Gen' in each plot) uncertainty. At these high energy levels, the physics processes described by the simulated backgrounds become on shell and hence result in larger amounts of background contamination in the signal region which affect the precision of the overall fit result. It could however be argued that the flat 20% uncertainty applied to the level of simulated background is too conservative and if it were to be relaxed slightly, the precision would correspondingly improve. However, even with this level of uncertainty in the amount of background, the overall precision level is still lower than one percent, except from in the most extreme regions of phase space at low rapidity. Unlike the angular coefficients, the electron scale and resolution variations do not have as large an effect when considering the uncertainty in the unpolarised cross section. Instead, the electron identification efficiency variations become the next most important variation after the Monte Carlo statistics at the per mille level before the background uncertainty becomes dominant at 25 GeV of p_T^{ll}.

As is to be expected, Figure 5.30r shows that with increasing rapidity the overall precision in the unpolarised cross section worsens due to a decreasing amount of statistics. The leads to the precision worsening to between 3 and 9% of the measured cross section when combining the statistical and systematic uncertainties. At high rapidity, the identification efficiency, PDF and calibration variation all become approximately equivalent but remaining below the percent level and far below the Monte Carlo statistics uncertainty.

Similarly to the electron channel, the uncertainties in the muon channel are broken down by systematic group in Figure 5.31. As is the case for the electron channel, the uncertainties in the angular coefficients are once again statistically limited at both

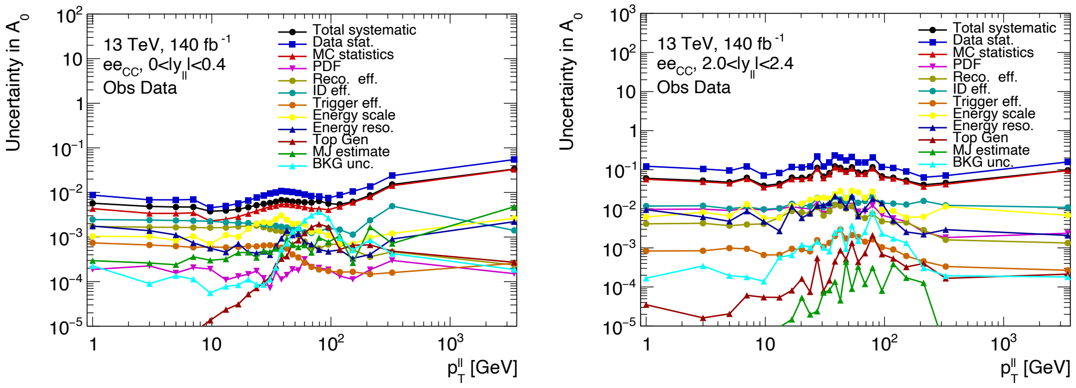

(a) Uncertainty in A_0 at low rapidity. (b) Uncertainty in A_0 at high rapidity.

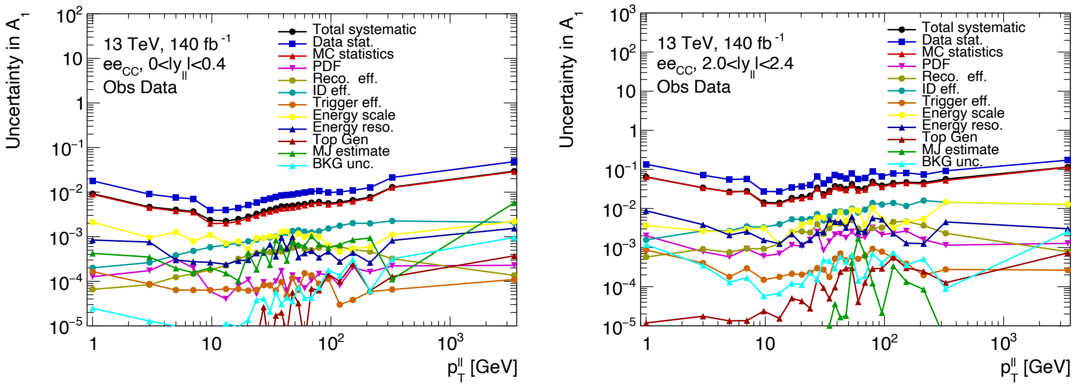

(c) Uncertainty in A_1 at low rapidity. (d) Uncertainty in A_1 at high rapidity.

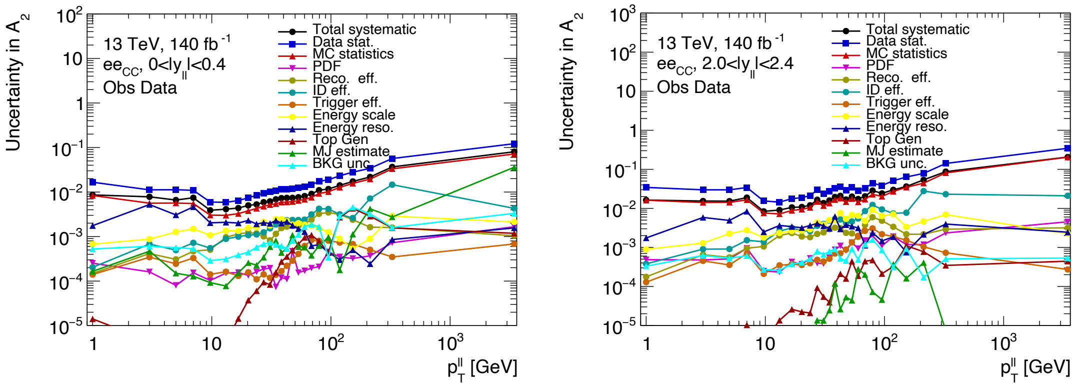

(e) Uncertainty in A_2 at low rapidity. (f) Uncertainty in A_2 at high rapidity.

Figure 5.30: The uncertainty in each angular coefficient broken down by systematic group in the central electron channel. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

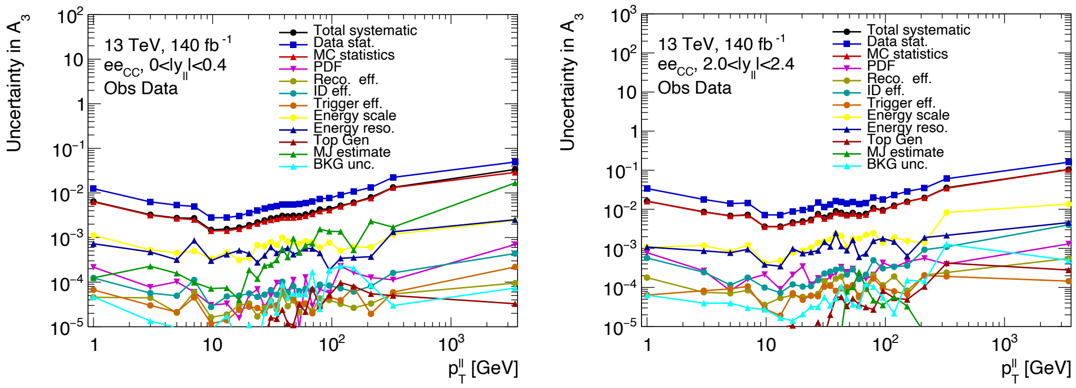

(g) Uncertainty in A_3 at low rapidity.

(h) Uncertainty in A_3 at high rapidity.

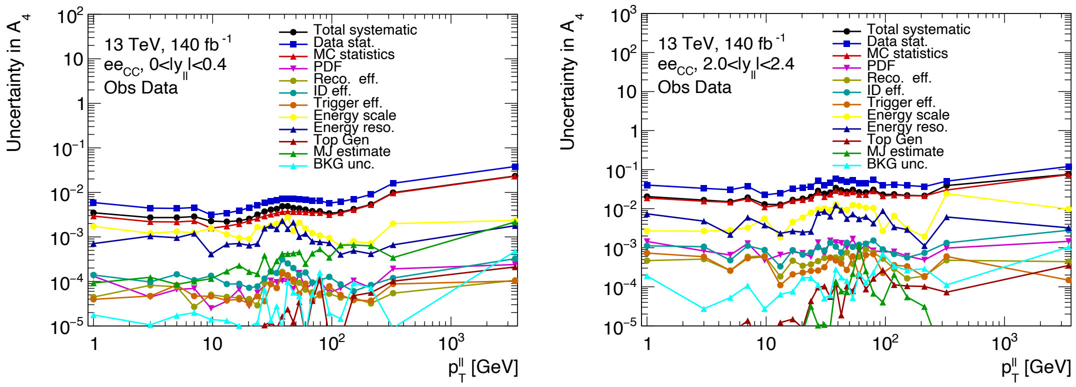

(i) Uncertainty in A_4 at low rapidity.

(j) Uncertainty in A_4 at high rapidity.

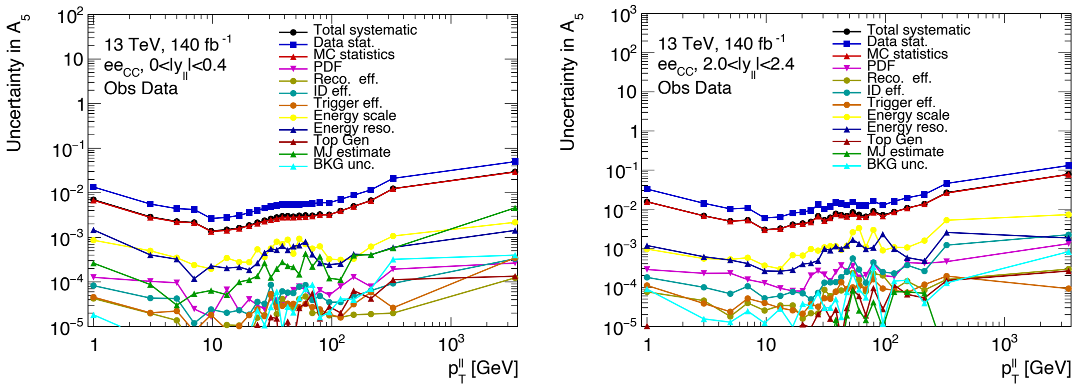

(k) Uncertainty in A_5 at low rapidity.

(l) Uncertainty in A_5 at high rapidity.

Figure 5.30: (continued) The uncertainty in each angular coefficient broken down by systematic group in the central electron channel. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

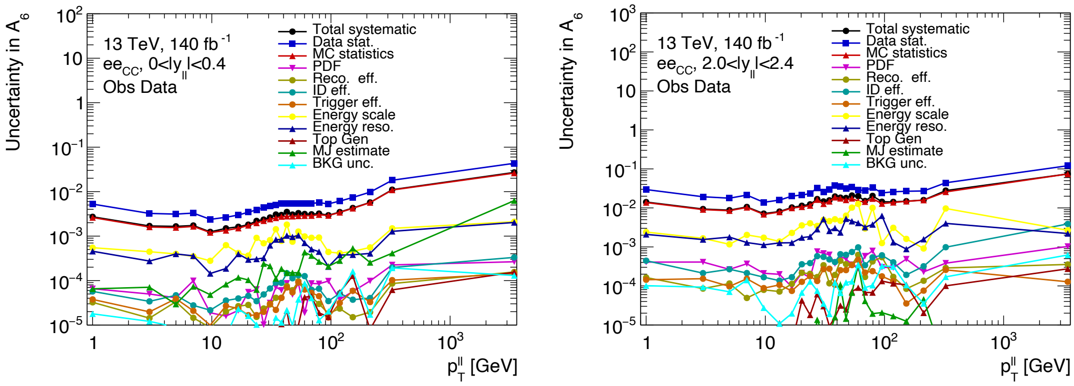

(m) Uncertainty in A_6 at low rapidity.

(n) Uncertainty in A_6 at high rapidity.

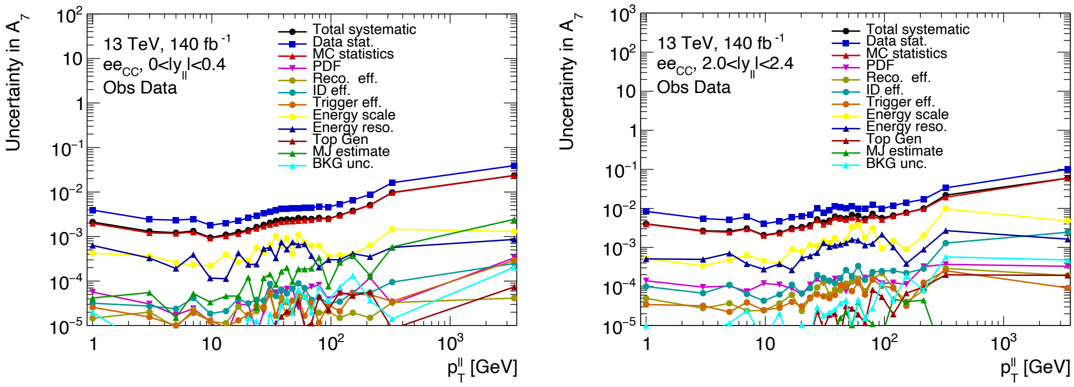

(o) Uncertainty in A_7 at low rapidity.

(p) Uncertainty in A_7 at high rapidity.

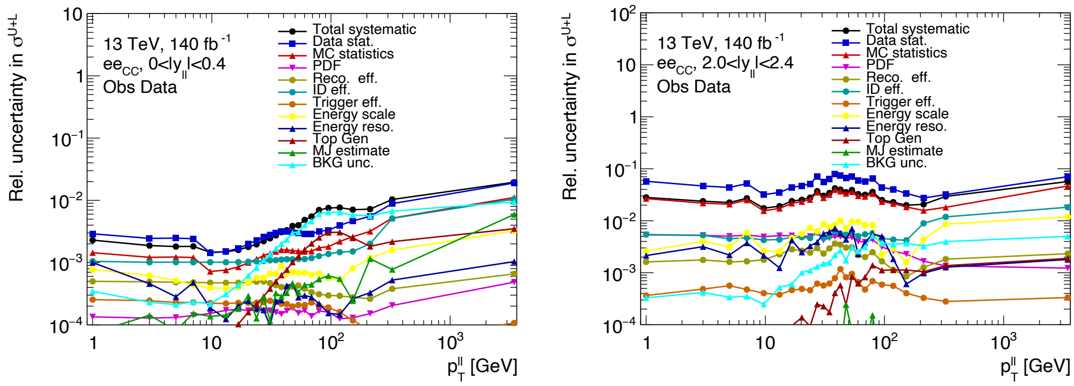

(q) Relative uncertainty in σ^{U+L} at low rapidity.

(r) Relative uncertainty in σ^{U+L} at high rapidity.

Figure 5.30: (continued) The uncertainty in each angular coefficient broken down by systematic group in the central electron channel. For the unpolarised cross section (σ^{U+L}) the relative uncertainty is shown in order to gauge the level of precision achieved by the analysis. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

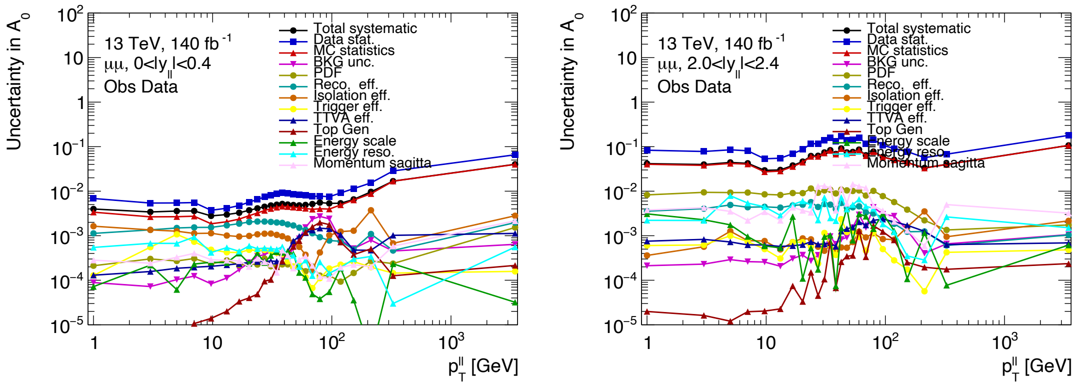

(a) Uncertainty in A_0 at low rapidity.

(b) Uncertainty in A_0 at high rapidity.

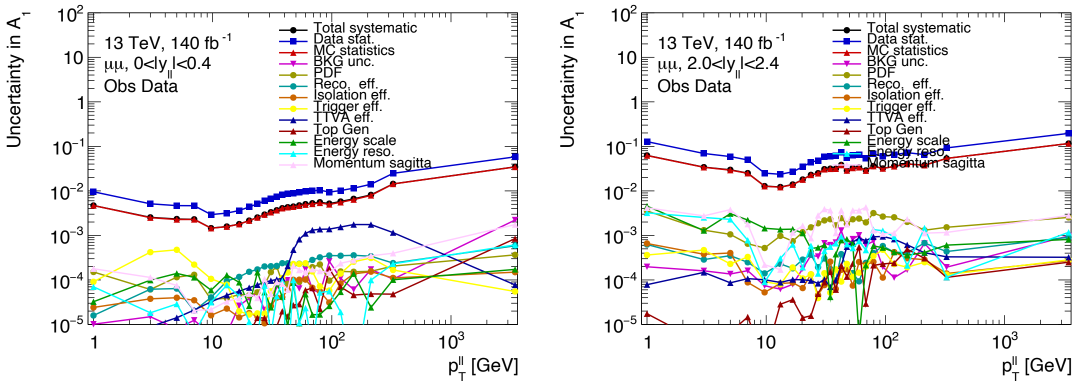

(c) Uncertainty in A_1 at low rapidity.

(d) Uncertainty in A_1 at high rapidity.

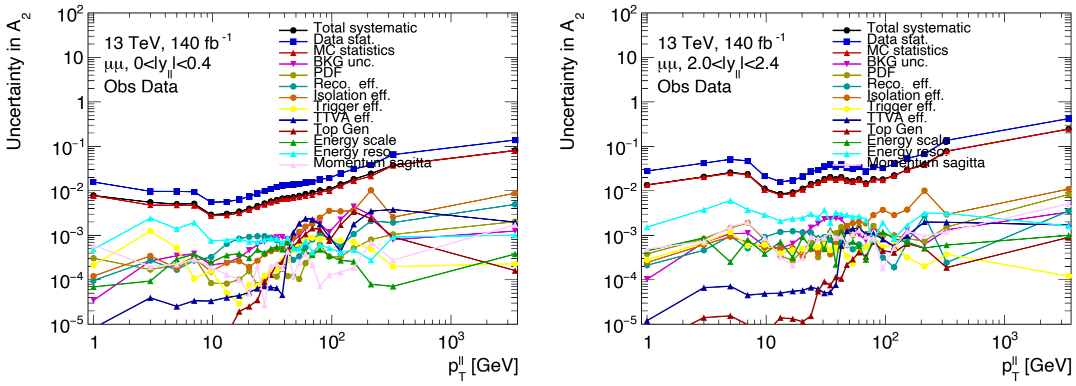

(e) Uncertainty in A_2 at low rapidity.

(f) Uncertainty in A_2 at high rapidity.

Figure 5.31: The uncertainty in each angular coefficient broken down by systematic group in the central muon channel. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

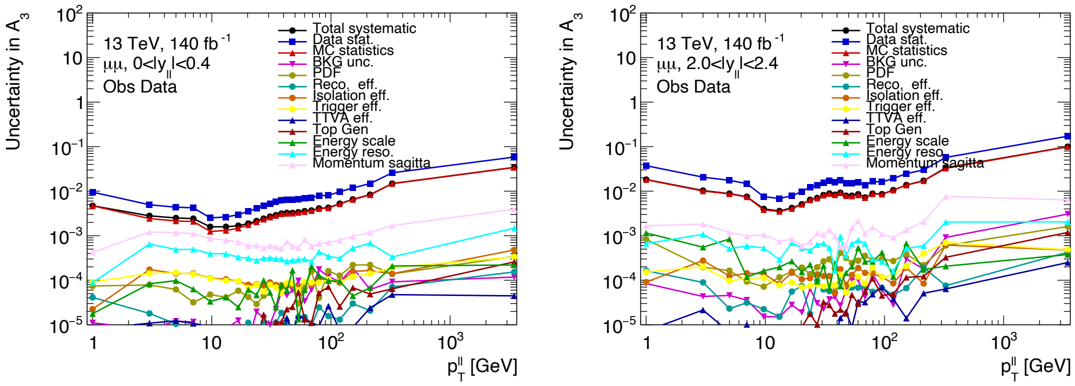

(g) Uncertainty in A_3 at low rapidity. (h) Uncertainty in A_3 at high rapidity.

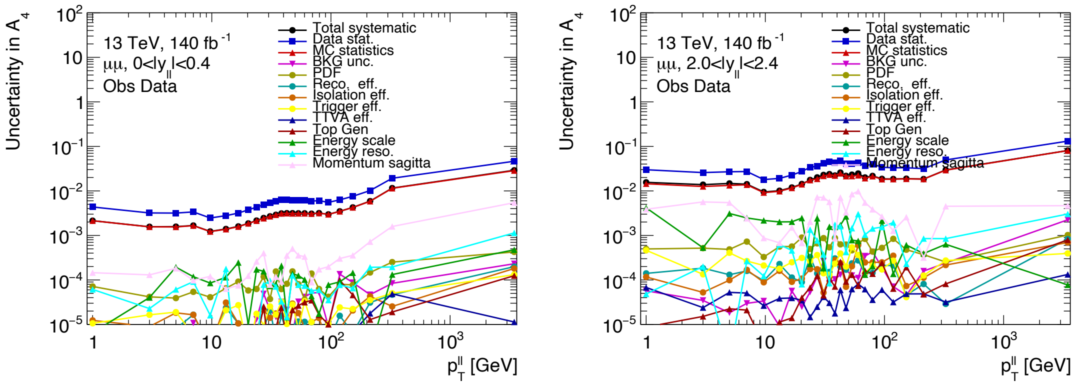

(i) Uncertainty in A_4 at low rapidity. (j) Uncertainty in A_4 at high rapidity.

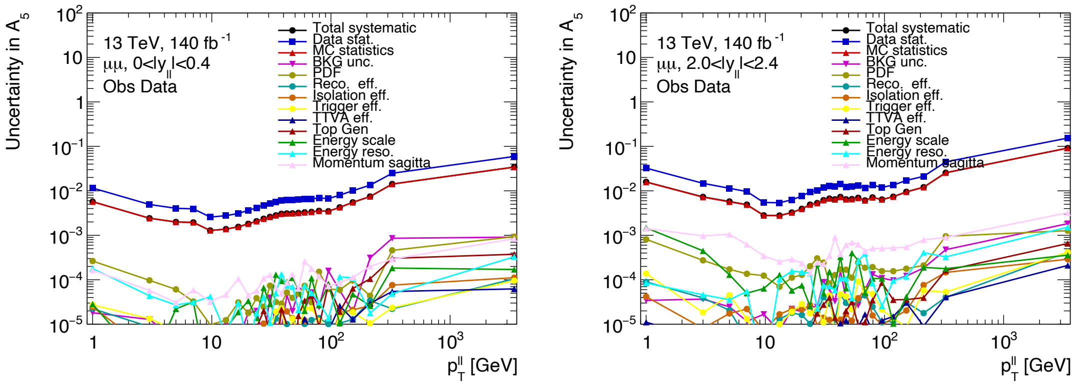

(k) Uncertainty in A_5 at low rapidity. (l) Uncertainty in A_5 at high rapidity.

Figure 5.31: (continued) The uncertainty in each angular coefficient broken down by systematic group in the central muon channel. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

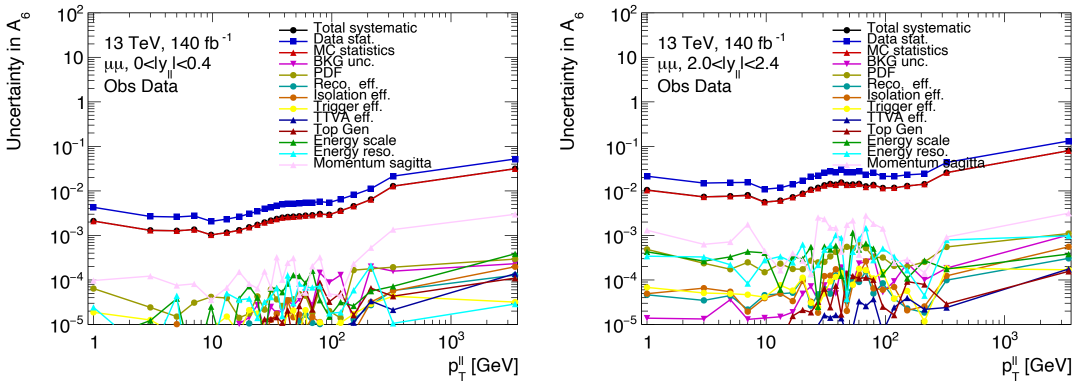

(m) Uncertainty in A_6 at low rapidity.

(n) Uncertainty in A_6 at high rapidity.

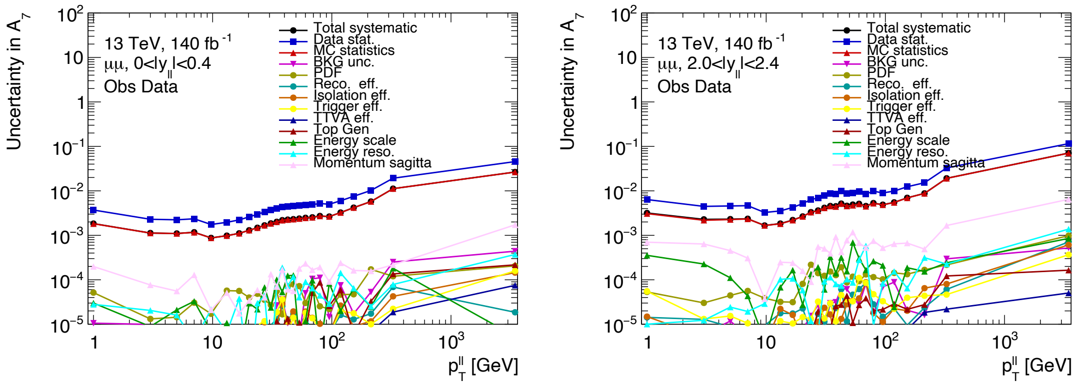

(o) Uncertainty in A_7 at low rapidity.

(p) Uncertainty in A_7 at high rapidity.

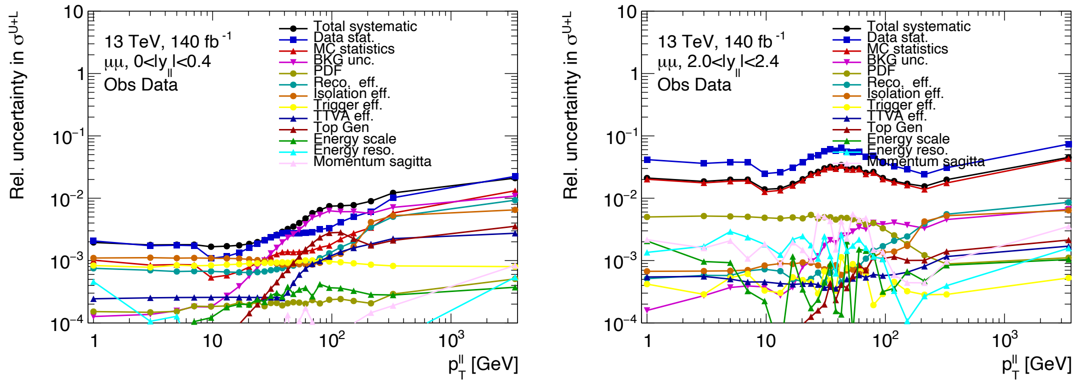

(q) Relative uncertainty in σ^{U+L} at low rapidity.

(r) Relative uncertainty in σ^{U+L} at high rapidity.

Figure 5.31: (continued) The uncertainty in each angular coefficient broken down by systematic group in the central muon channel. For the unpolarised cross section (σ^{U+L}) the relative uncertainty is shown in order to gauge the level of precision achieved by the analysis. The left column shows the uncertainty in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

low and high rapidity. Whilst the statistical uncertainty is slightly better in the muon channel than it is in the electron channel, they are both of approximately the same order of magnitude, as would be expected from lepton flavour universality. This difference in statistical sensitivity is simply due to the greater acceptance in the muon channel from the lack of a cut in the $1.37 < |\eta| < 1.52$ pseudorapidity region. The systematics are once again also limited by the amount of Monte Carlo statistics, meaning that there can be no significant improvement to the results in this channel without a much greater Monte Carlo simulation campaign. After accounting for statistical effects in both data and Monte Carlo, the other systematic groups are typically all of the same order of magnitude, with no systematic clearly dominating. It is only for A_3 to A_7 that the uncertainty from the momentum sagitta bias becomes a clearer next to leading systematic at both low and high rapidity.

With the extremely high level of precision achieved in the relative uncertainty in the unpolarised cross section in the electron channel, it is no surprise that this precision is replicated in the muon channel, as evidence by Figures 5.31q and 5.31r. However, unlike the electron channel, at low rapidity in the muon channel, the systematic uncertainty matches that of the statistical uncertainty below 40 GeV of dilepton transverse momentum instead of being purely statistically limited. Above this point, the systematic error begins to dominate over the statistical effects due to the increasing importance of the background and Monte Carlo generator systematics, which begin to peak around 100 GeV as in the electron channel. The overall precision is, however, at a level of approximately 0.2% below 30 GeV of p_T^{ll} and remains below 1% below 200 GeV, which accounts for the vast majority of the phase space probed at low rapidity. Since it is known that there are issues with the trigger efficiency systematics from the pull plots that must be fixed, it may also be possible to reduce the associated error, meaning that at low dilepton transverse momentum the muon channel may become statistically limited instead of the current balance between statistics and systematics.

As is to be expected, the balance between the statistical and systematical uncertainty is broken at high rapidity due to the decrease in statistics in this region, meaning that the precision in the unpolarised cross section worsens and is limited by statistics. The systematic uncertainty also becomes dominated by the Monte Carlo statistics, with all other systematic groups at least one order of magnitude lower. As in the electron channel, the total level of precision varies between 3-9% across the measured dilepton transverse momentum range when combining the statistical and systematic uncertainties.

5.9.3 Compatibility of the Electron and Muon Channels

Figures 5.30 and 5.31 show that both the electron and muon channels are statistically limited, rather than systematically. Therefore the simplest way to reduce the overall uncertainty will be to combine the results from both channels, reducing the statistical uncertainty by a factor of approximately $1/\sqrt{2}$. However, in order to do this the electron and muon channels must be compatible with each other, i.e they must obey lepton flavour universality and provide the same central values in the measured angular coefficients.

To rigorously test the compatibility of the two channels, the likelihood is modified in such a way that it will instead fit the delta between the two channels for each coefficient, rather than the angular coefficient itself. The delta is defined to be the proportional difference between the electron and muon channel with respect to the electron channel, i.e for a given angular coefficient:

$$\Delta A_i = \frac{A_{i,ee} - A_{i,MM}}{A_{i,ee}} \tag{5.49}$$

By using this method to test the compatibility of the two channels, all of the delta values fitted will properly account for all sources of systematic uncertainty across both channels, as well as the correlations between the variables. This therefore allows for a much more robust measure of the compatibility in comparison to simply subtracting the results of the individual single channel fits. The result of this fit gives a $\chi^2/N_{\mathrm{dof}} = 23697/16500$ which suggests it has converged well and its results are trustworthy.

In Figure 5.32, the deltas defined by Equation (5.49) can be seen for all of the angular coefficients and the unpolarised cross section at low and high rapidity. At low rapidity, the angular coefficients are mostly clustered around zero, with some variations on the level of $\pm 2\%$ from the central value in the electron channel. In almost all bins, these variations are within $\pm 2\sigma$ of zero meaning that the two channels show no significant evidence of having different values. At high rapidity the delta values remain scattered around zero. However, the scale of these variations becomes larger, with $\pm 10\%$ of the electron value typical for most angular coefficients. As with the single channel results the lower number of events results in a higher statistical error which compensates for the larger variance from zero. This means that at high rapidity, zero still remains with $\pm 2\sigma$ in almost all bins and therefore shows no significant evidence of deviation from lepton flavour universality.

For the unpolarised cross section at low rapidity, the delta values show the electron channel has an approximately 1% greater magnitude than the muon channel below 100 GeV of dilepton transverse momentum. Above this point the muon channel has a larger value for the unpolarised cross section. This suggests that there may be some systematic effect between the two channels that isn't properly being taken into account which is causing this disagreement, given it is highly unlikely to be caused by violation of lepton flavour universality. Since there is a known issue with the muon SF systematics, as shown in Figure 5.29, it may be that fixing fixing the application of these scale factors may restore the compatibility of the two channels. However, since most of the data points agree with zero within $\pm 3\sigma$ then this apparent 1% fluctuation may still be a statistical effect. Whilst this is still a large deviation from zero, it is still not so large that it can be considered as irrefutable evidence of disagreement between the electron and muon channels. In contrast, the high rapidity unpolarised cross section is much more similar to the angular coefficients than its low rapidity counterpart. Here, the delta values vary around $\pm 10\%$ from the electron value, with a preference for the negative values.

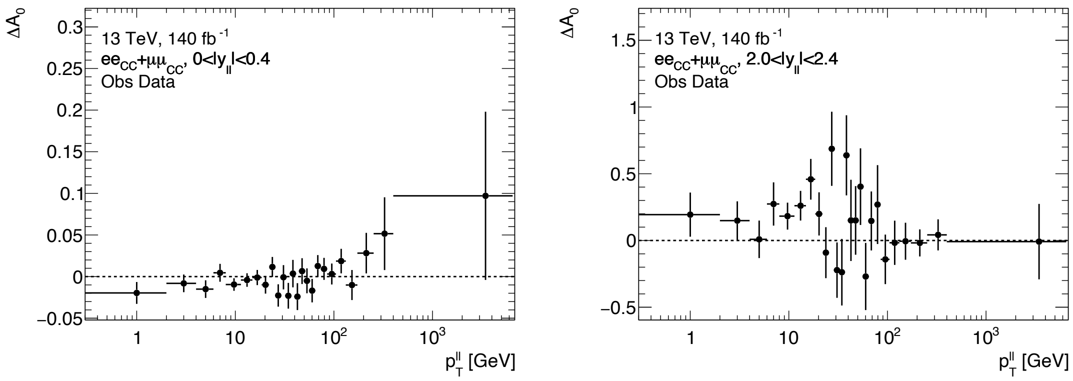

(a) ΔA_0 at low rapidity.

(b) ΔA_0 at high rapidity.

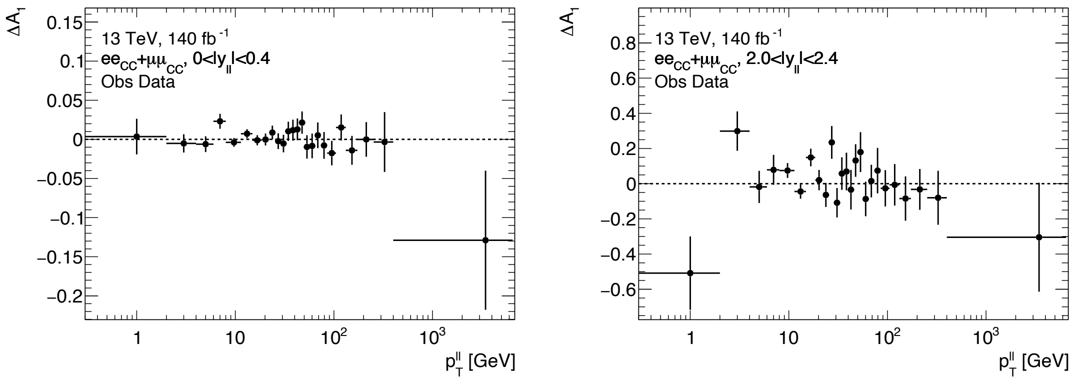

(c) ΔA_1 at low rapidity.

(d) ΔA_1 at high rapidity.

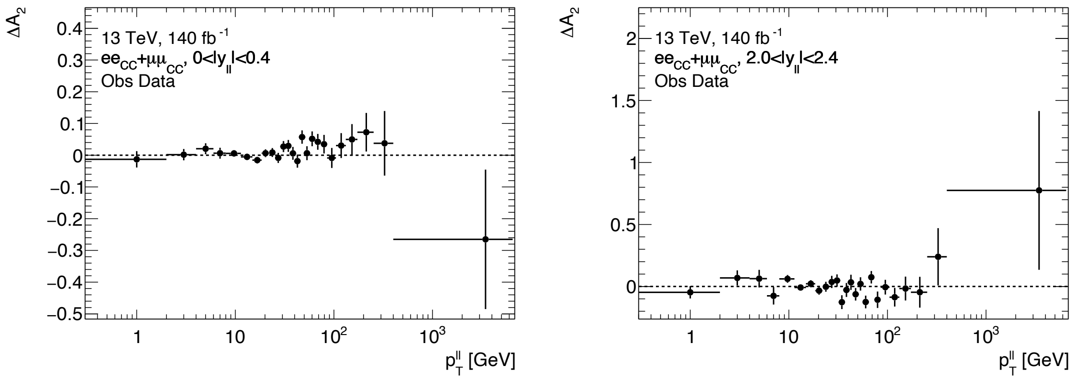

(e) ΔA_2 at low rapidity.

(f) ΔA_2 at high rapidity.

Figure 5.32: Fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The left column shows the the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic effects.

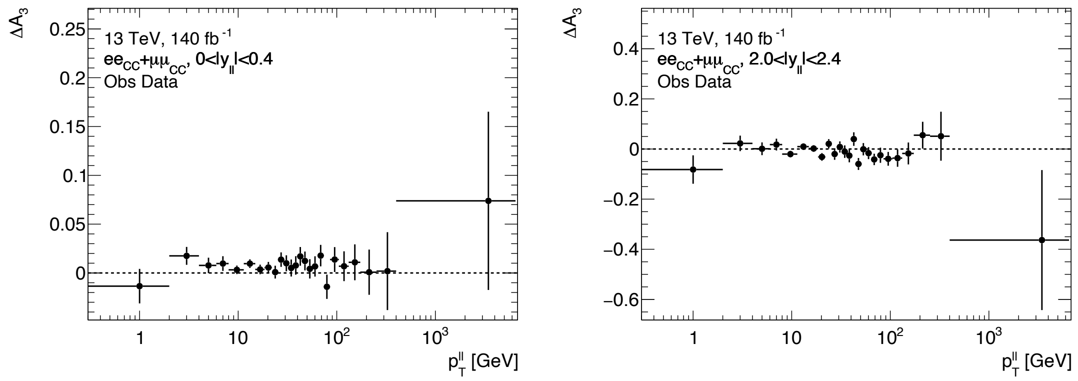

(g) ΔA_3 at low rapidity.

(h) ΔA_3 at high rapidity.

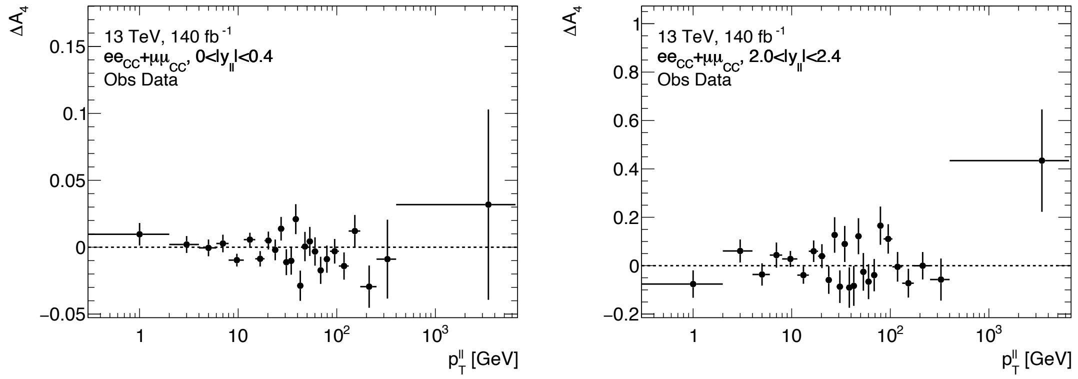

(i) ΔA_4 at low rapidity.

(j) ΔA_4 at high rapidity.

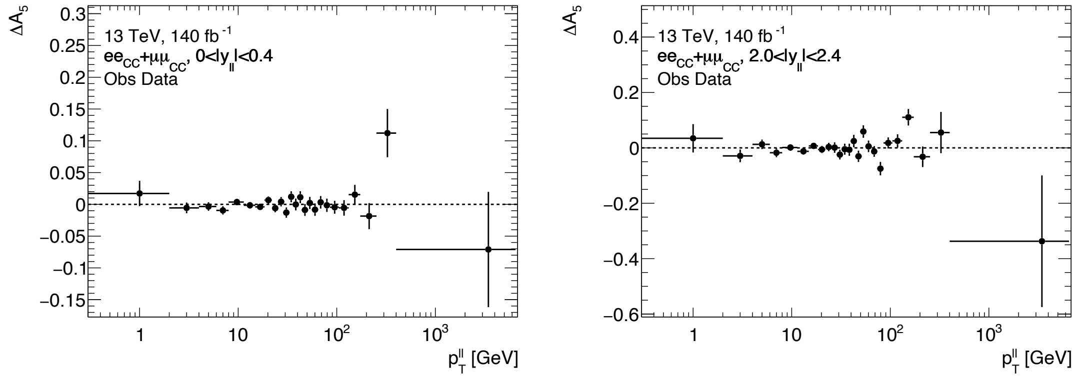

(k) ΔA_5 at low rapidity.

(l) ΔA_5 at high rapidity.

Figure 5.32: (continued) Fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The left column shows the the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic effects.

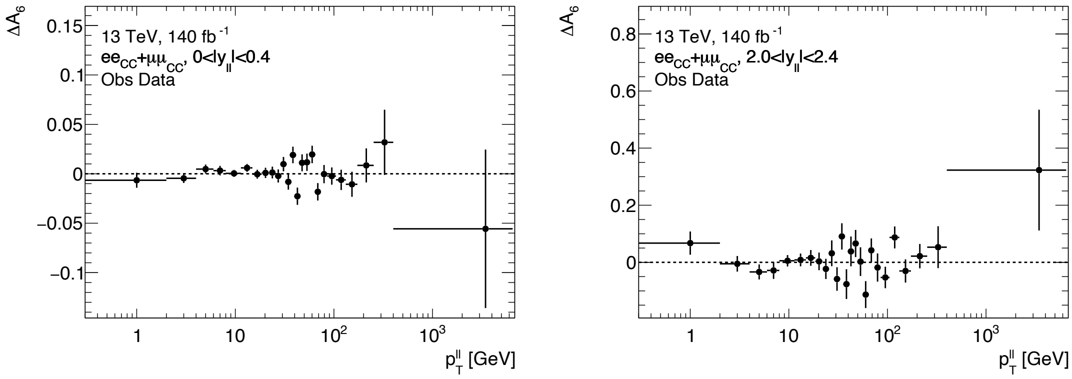

(m) ΔA_6 at low rapidity.

(n) ΔA_6 at high rapidity.

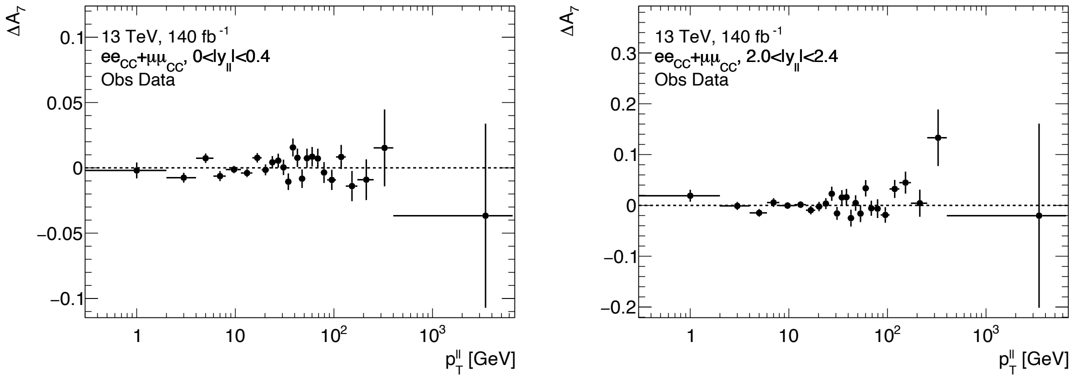

(o) ΔA_7 at low rapidity.

(p) ΔA_7 at high rapidity.

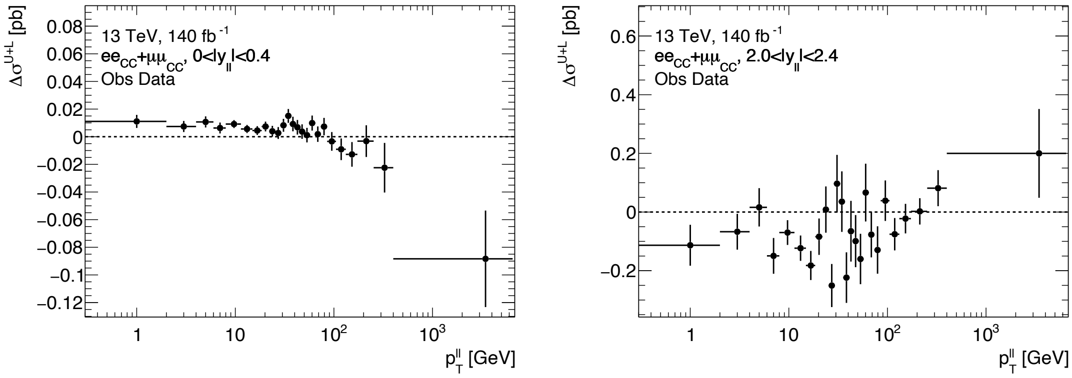

(q) $\Delta\sigma^{U+L}$ at low rapidity.

(r) $\Delta\sigma^{U+L}$ at high rapidity.

Figure 5.32: (continued) Fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The left column shows the the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic effects.

The majority of the high rapidity delta values are also compatible with zero within $\pm 2\sigma$ just as with the angular coefficients. This therefore means that in this region, the electron and muon channels are in good agreement. A full set of agreement plots for each angular coefficient and the unpolarised cross section in every rapidity bin can be found in Appendix C.

Given the size of the error bars at high rapidity in Figure 5.32, it is reasonable to consider whether it would be possible to improve them and have the best possible sensitivity to the agreement between the two channels. However, as shown in Figure 5.33 the total uncertainty in the delta values are dominated by the statistical uncertainty, except for the low rapidity unpolarised cross section delta. This therefore means that improving the systematic uncertainty will not substantially improve the sensitivity to the agreement between the electrons and muons. Furthermore, any improvement in the systematic uncertainty would do little to increase the number of standard deviations necessary for the delta values to agree with zero and thus increase any possible evidence of lepton flavour universality violation. Therefore, given the current good agreement between both the electron and muon channels, the two channels can be considered to be compatible.

5.9.4 Angular Coefficients Results Plots

Assuming lepton flavour universality, the Z boson coupling to electrons should be exactly the same as to muons. Thus the distribution of the angular coefficients should be exactly the same, modulo statistical fluctuations. Therefore to show to be able to properly compare the results of the two single channel fits, in Figure 5.34 they are plotted for each angular coefficient at low and high rapidity as a function of the transverse momentum of the Z. The exact agreement between the two channels is approximately what is seen at both low and high rapidity. Where there are divergences between the two channels (mainly at high p_T^Z where there are not a lot of statistics), the difference is typically within error bounds of the two channels, meaning that there is a generally good agreement of the two channels.

The shapes of the angular coefficients also show that not every helicity state contributes equally to the overall Drell-Yan cross section. It is clear that at both low and high rapidity A_0 and A_2 contribute the largest amount to the cross section whilst most of the other angular coefficients are centred around zero and have a much smaller contribution. At low rapidity, there is only a very small contribution to the overall cross section from A_4 and from A_3 at high p_T^{ll} only. At high rapidity A_1, A_3, A_4 and A_6 have an appreciably larger magnitude than at low rapidity, however, A_0 and A_2 continue to have much larger values than the other coefficients. Since A_4 is essentially a proxy for the forwards-backwards asymmetry, which is known to increase as a function of rapidity, it is therefore to be expected that A_4 has also increased with rapidity. As A_3 also contributes to the forwards-backwards asymmetry, this same reasoning explains the slight increase in the magnitude of A_3 at high rapidity as well.

Of particular interest is the departure of A_2 from zero in the first p_T^{ll} bin at both low

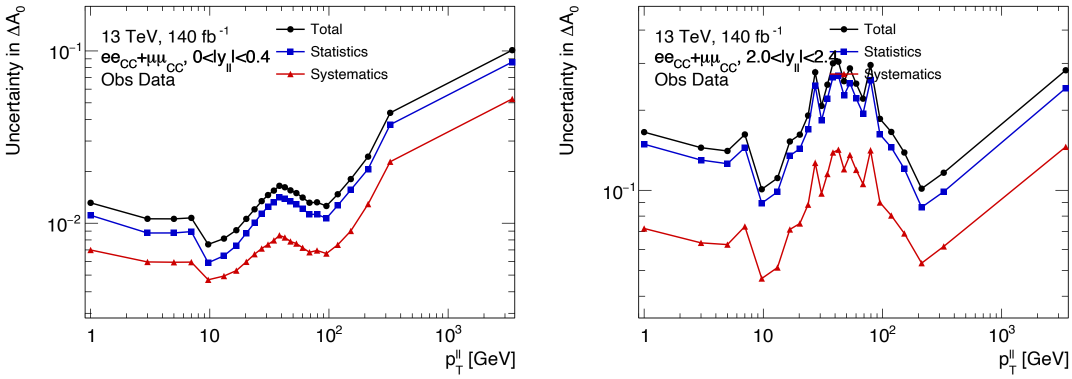

(a) Uncertainty in ΔA_0 at low rapidity. (b) Uncertainty in ΔA_0 at high rapidity.

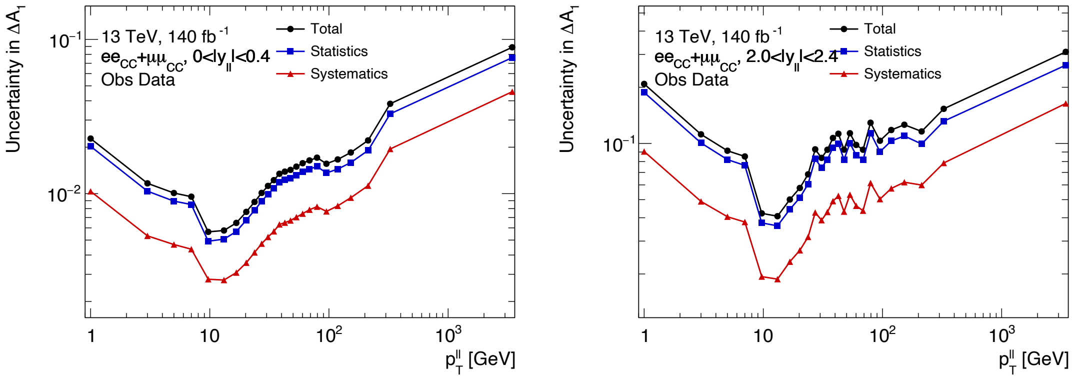

(c) Uncertainty in ΔA_1 at low rapidity. (d) Uncertainty in ΔA_1 at high rapidity.

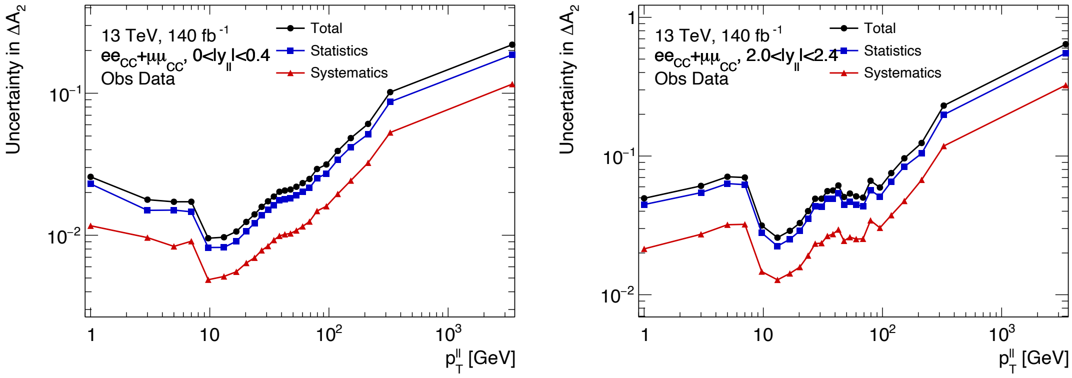

(e) Uncertainty in ΔA_2 at low rapidity. (f) Uncertainty in ΔA_2 at high rapidity.

Figure 5.33: The total uncertainty in the fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and broken down into its statistical and systematic components and as a function of the dilepton transverse momentum. The left column shows the uncertainty lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

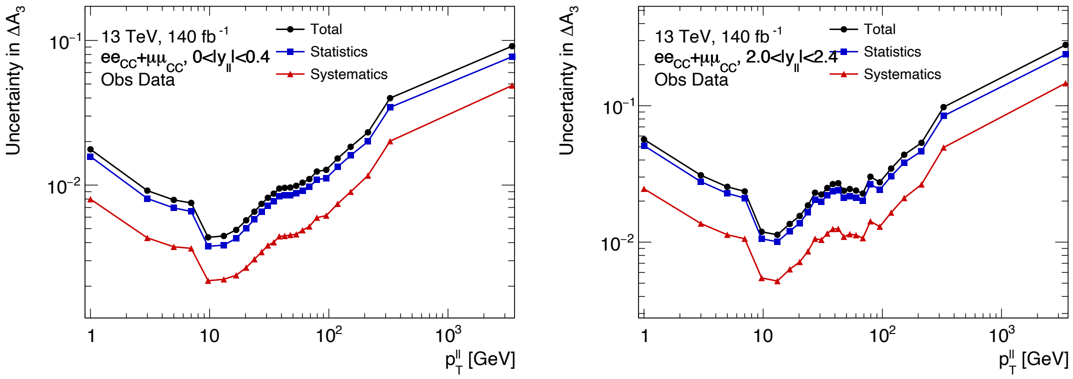

(g) Uncertainty in ΔA_3 at low rapidity. (h) Uncertainty in ΔA_3 at high rapidity.

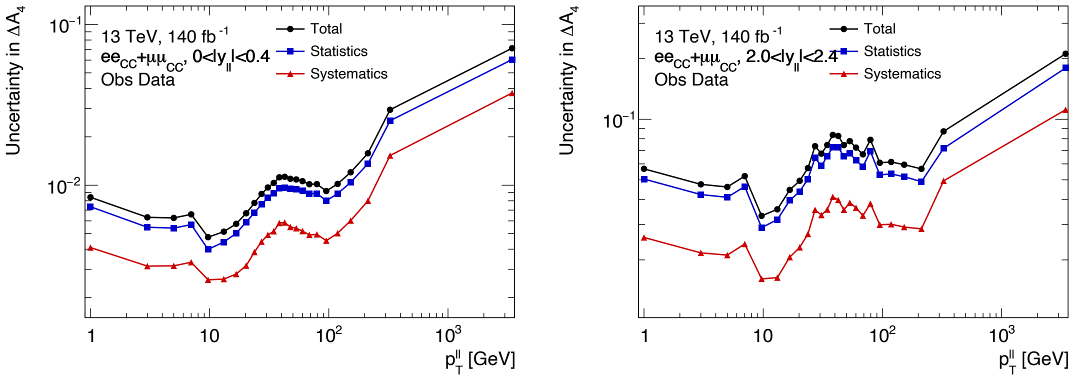

(i) Uncertainty in ΔA_4 at low rapidity. (j) Uncertainty in ΔA_4 at high rapidity.

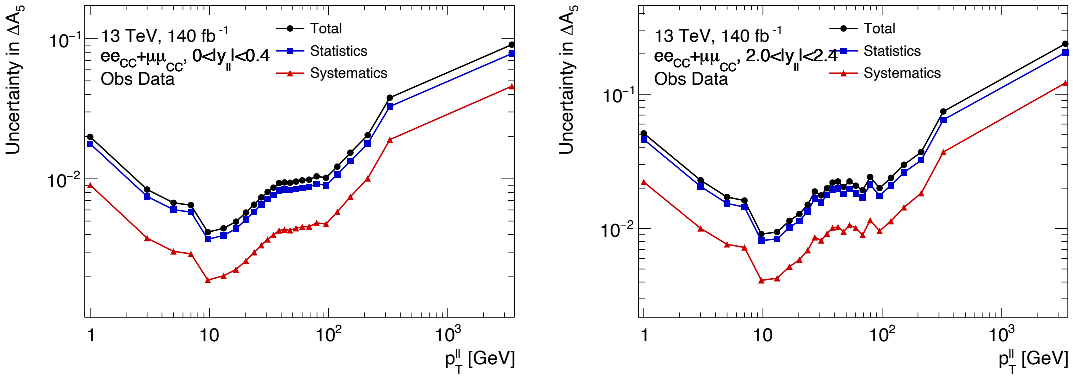

(k) Uncertainty in ΔA_5 at low rapidity. (l) Uncertainty in ΔA_5 at high rapidity.

Figure 5.33: (continued) The total uncertainty in the fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and broken down into its statistical and systematic components and as a function of the dilepton transverse momentum. The left column shows the uncertainty lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

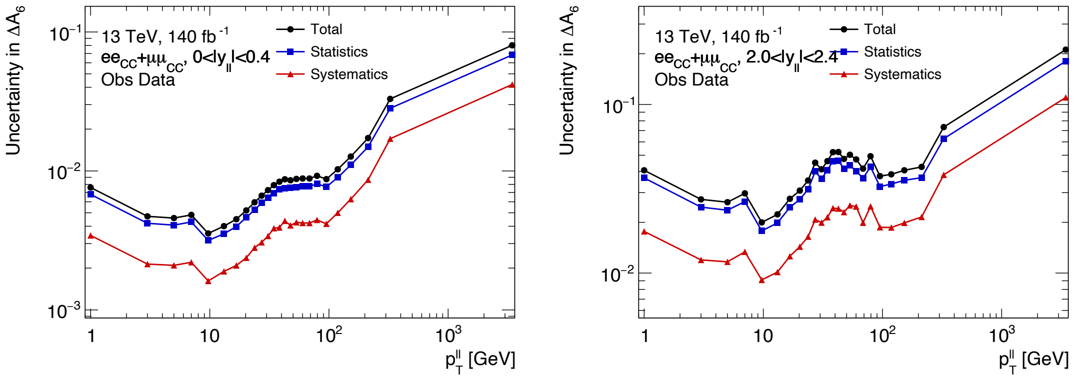

(m) Uncertainty in ΔA_6 at low rapidity. (n) Uncertainty in ΔA_6 at high rapidity.

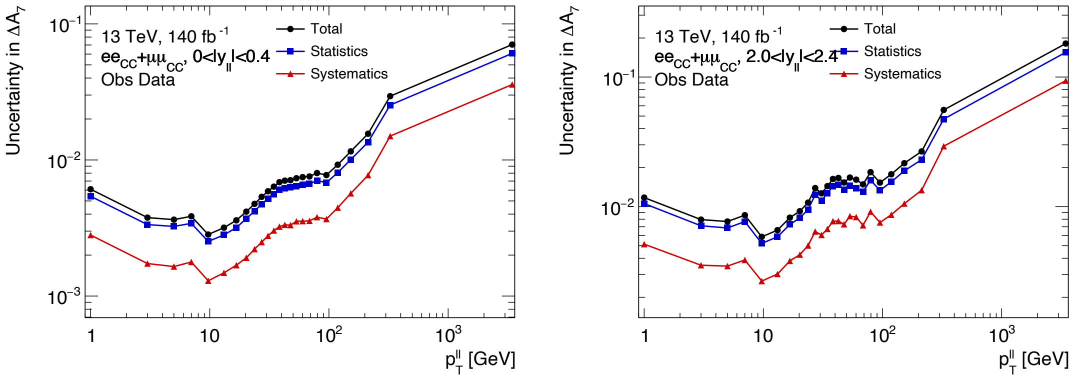

(o) Uncertainty in ΔA_7 at low rapidity. (p) Uncertainty in ΔA_7 at high rapidity.

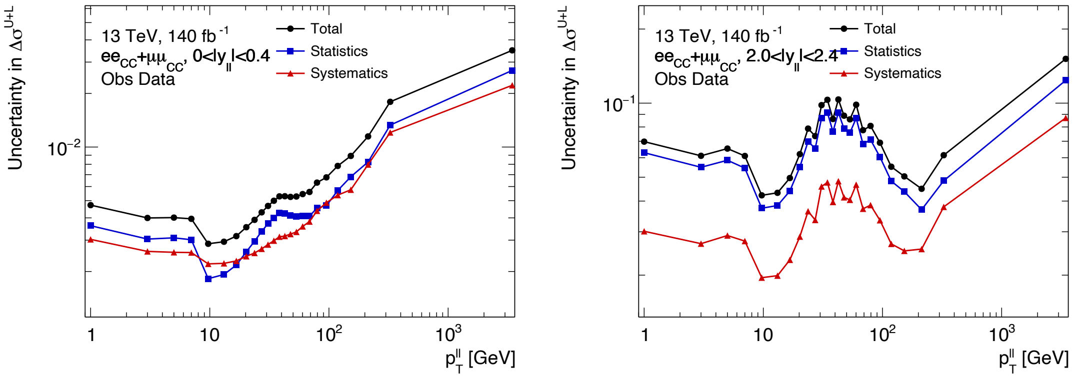

(q) Uncertainty in $\Delta\sigma^{U+L}$ at low rapidity. (r) Uncertainty in $\Delta\sigma^{U+L}$ at high rapidity.

Figure 5.33: (continued) The total uncertainty in the fitted difference in the angular coefficients and unpolarised cross section between the electron and muon channels, as a fraction of the electron channel result and broken down into its statistical and systematic components and as a function of the dilepton transverse momentum. The left column shows the uncertainty lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$).

and high rapidity, where theoretically it should be exactly zero. Similar evidence for a non-zero A_2 value at low transverse momentum has been found previously by ATLAS at 8 TeV [6] and by the LHCb experiment at 13 TeV [15], thus adding credence to the hypothesis that this is more than just a statistical fluctuation. If this were to be confirmed as being more than a 5σ difference from zero, then this could be considered to be evidence for new physics. However, this is very much not the case currently and will require further study in the future with larger datasets.

The unpolarised cross section in Figures 5.34q and 5.34r show an increasing distribution that peaks in the $8 < p_T^{ll} < 11.4$ GeV before beginning to fall as the dilepton transverse momentum increases. The dip in the 6 - 8 GeV bin can be explained as an artefact of the bin widths chosen for the extraction to be performed in, rather than some unexplained physics effect. Whilst the shape of the distribution is same at both low and high rapidity, the agreement between the electron and muon channels is visibly better at low rapidity than it is at high rapidity. However, this difference between the channels is much more likely to be due to statistical fluctuations from the lower number of events in this region rather than a violation of lepton flavour universality.

The plots in Figure 5.34 also display the Asimov results with the corresponding Monte Carlo statistical error, which represents what the angular coefficients are expected to look like from the simulation. The Asimov results show an overall good agreement with the observed data in the electron and muon channels by broadly following the trends established by the data. In general this agreement is, however, not as good at high p_T^{ll}, particularly for A_0 and A_2. This is, however, simply an effect of comparing a Monte Carlo simulation produced at NLO in α_S to the data which contains all perturbative orders and non-perturbative effects of QCD. Whilst it also may seem concerning that the Asimov value of A_0 is unphysically less than zero at low p_T^{ll}, this is due to a well known limitation of the POWHEG generator. Therefore, these Asimov values can be disregarded in this region.

5.9.5 Combination of the Electron and Muon Channels

With the compatibility of the electron and muon channels having been verified, the next logical step is to combine the results of the two channels. This will result in the central values being averaged across the two channels and the statistical uncertainty being reduced by a factor of approximately 1.4 due to the doubling of available statistics. The results of the combination in the same high and low rapidity bins as for the single channel results can be seen in Figure 5.35. As is to be expected, the combination plots in Figure 5.35 share the same characteristics as their single channel counterparts. In addition to the values extracted from the observed data, Figure 5.35 also depicts fixed order NNLO DYTurbo predictions of each angular coefficient. Since these predictions are one order higher in QCD than the NLO POWHEG Asimov values in Figure 5.34, they are a better match to the data with some small deviations at high p_T^Z values. This acts as a useful validation of the predictions of DYTurbo which are of great importance when measuring other parameters of the Standard Model such as the mass of the W boson.

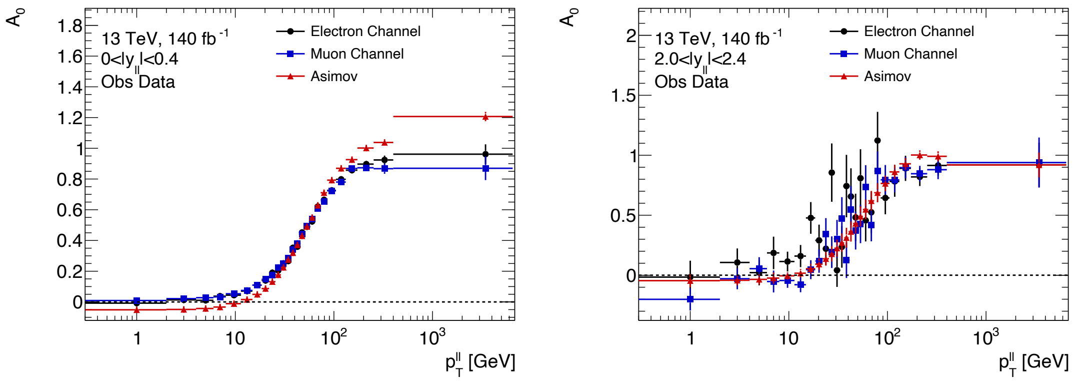

(a) A_0 at low rapidity. (b) A_0 at high rapidity.

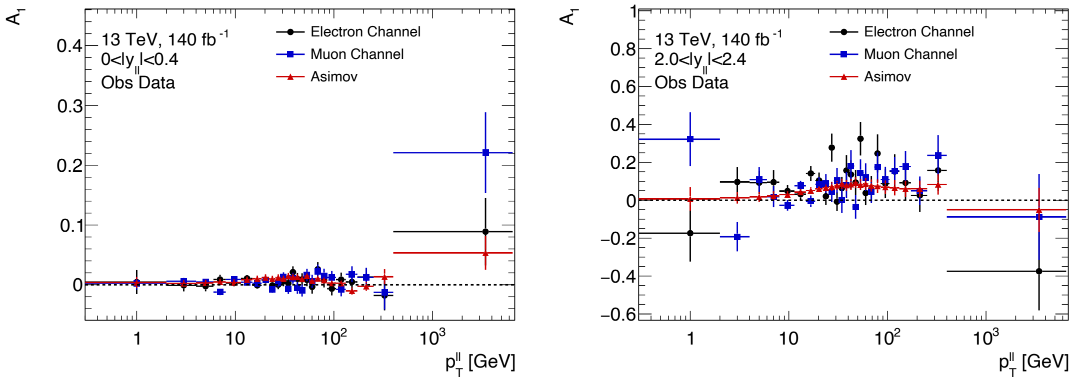

(c) A_1 at low rapidity. (d) A_1 at high rapidity.

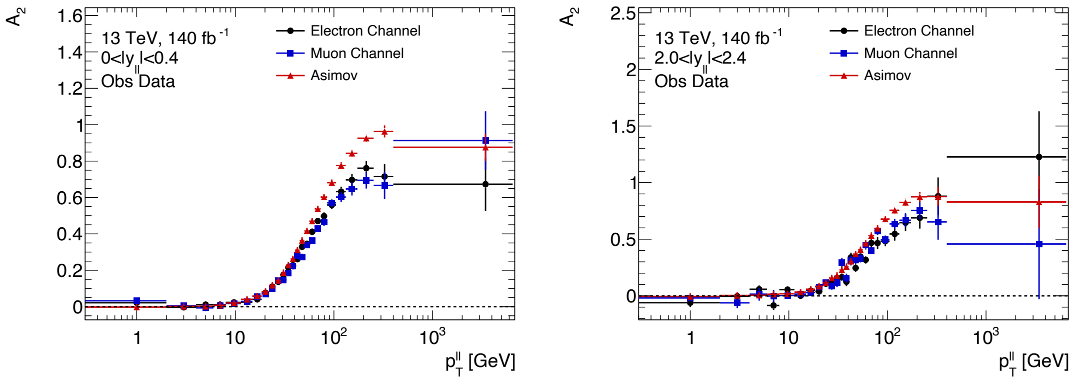

(e) A_2 at low rapidity. (f) A_2 at high rapidity.

Figure 5.34: The angular coefficients of the Z boson measured in observed data for the electron (black points) and muon (blue points) channels as a function of the dilepton transverse momentum. The left column shows the coefficients in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic uncertainties in each channel.

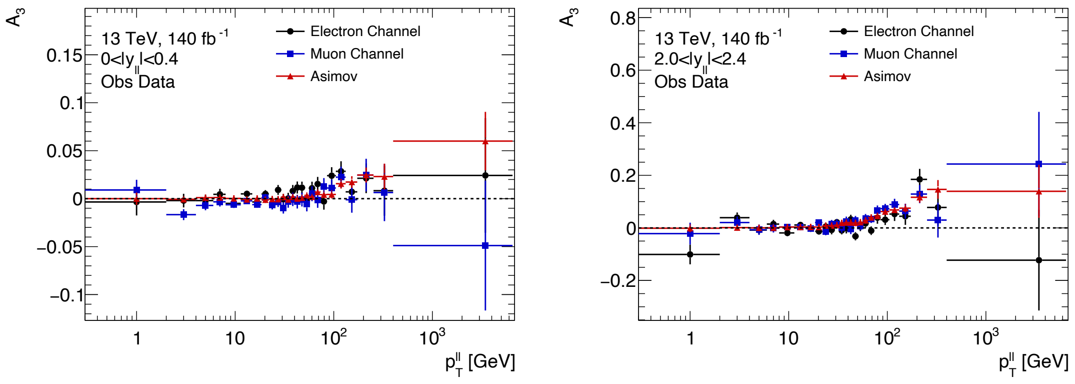

(g) A_3 at low rapidity.

(h) A_3 at high rapidity.

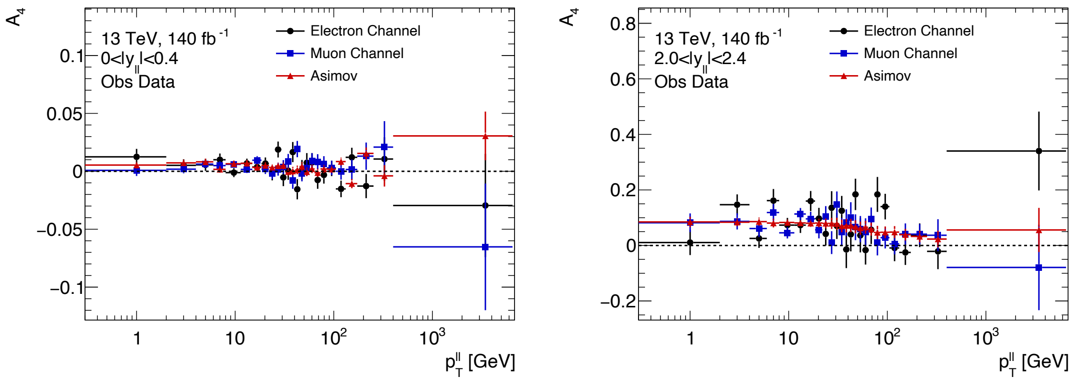

(i) A_4 at low rapidity.

(j) A_4 at high rapidity.

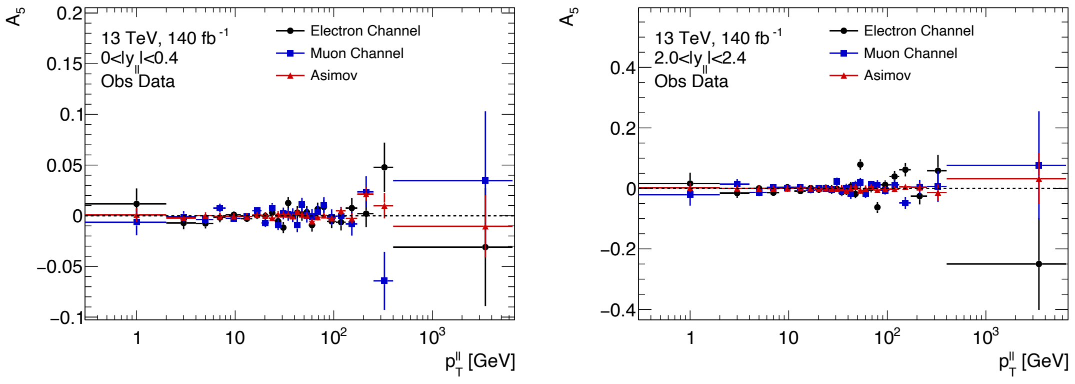

(k) A_5 at low rapidity.

(l) A_5 at high rapidity.

Figure 5.34: (continued) The angular coefficients of the Z boson measured in observed data for the electron (black points) and muon (blue points) channels as a function of the dilepton transverse momentum. The left column shows the coefficients in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic uncertainties in each channel.

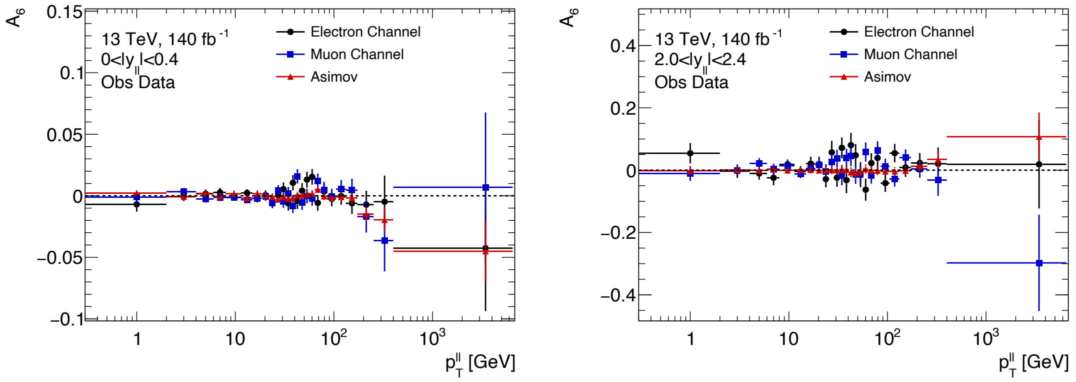

(m) A_6 at low rapidity.

(n) A_6 at high rapidity.

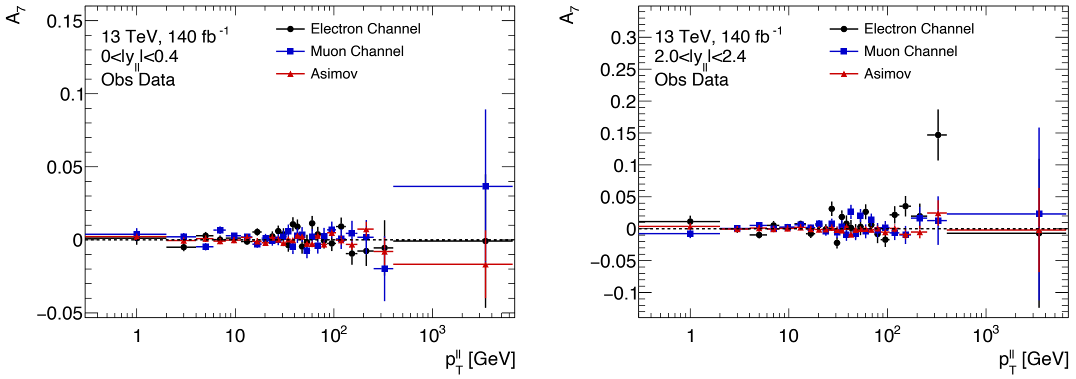

(o) A_7 at low rapidity.

(p) A_7 at high rapidity.

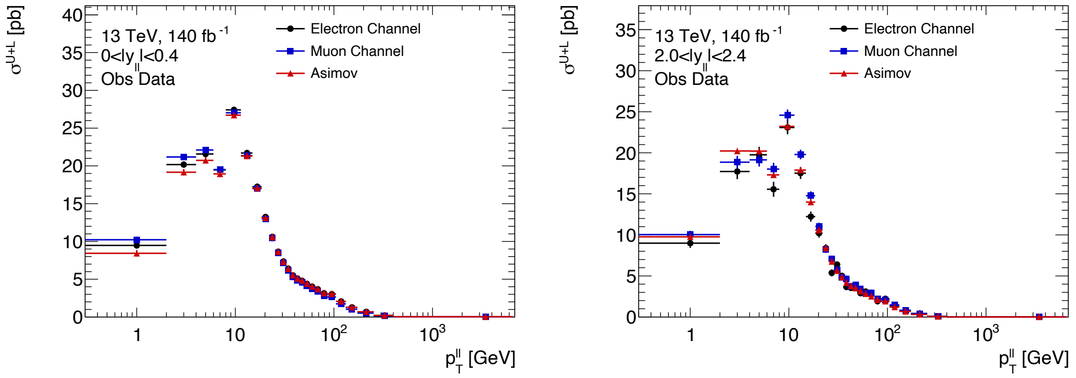

(q) σ^{U+L} at low rapidity.

(r) σ^{U+L} at high rapidity.

Figure 5.34: (continued) The angular coefficients of the Z boson measured in observed data for the electron (black points) and muon (blue points) channels as a function of the dilepton transverse momentum. The left column shows the coefficients in the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.0 < |Y_{ll}| < 2.4$). The error bars correspond to the combination of statistical and systematic uncertainties in each channel.

The full set of combination results for every rapidity bin can be found in Appendix D where the same characteristics of the single channel fits can be observed.

As is expected the statistical component of the uncertainties is approximately 1.4 times smaller than the individual channels as can be seen in Figure 5.36. This improves the the precision in the unpolarised cross section to approximately 2 per mille for $p_T^{ll} < 15$ GeV before beginning to increase but still remaining below 1% until a dilepton transverse momentum of 250 GeV. This increase in uncertainty is once again due to the background and Monte Carlo generator systematics peaking and plateauing in this region, even when not explicitly depicted. Even with this large reduction in the statistical uncertainty, it is still the dominant source of uncertainty for the angular coefficients and unpolarised cross section at low momentum. Whilst the eventual addition of the central-forward channel will improve the statistical uncertainty in the region $1.2 < |Y_{ll}| < 2.4$, this will not result in a tripling of the total statistics since the acceptance and efficiencies associated with the central-forward channel are much lower than their central-central counterparts. Therefore, with the current analysis binning, the best estimate of the angular coefficients will only come with a new data taking campaign where the amount of data recorded is significantly greater than that of Run 2 of the LHC.

Lam - Tung Relation

Following on from the combination results, it is also possible to test the Lam-Tung relation [69], i.e that $A_0 - A_2 = 0$. This is done by simply taking the fitted values of A_0 and A_2 from the combination result and subtracting them. The uncertainty can then be calculated by considering the usual error propagation rules and taking into account the effect of the covariance between A_0 and A_2. Whilst the Lam-Tung relation is expected to be zero, this is only at leading order in perturbation theory. When QCD vertices are considered in the Drell-Yan process, the Lam-Tung relation will become non-zero. Figure 5.37 shows that at 13 TeV, $A_0 - A_2$ clearly becomes non-zero and incompatible with zero. Therefore, as expected, the Lam-Tung relation is violated. The magnitude of the violation is approximately the same across all rapidity bins (excluding the last bin which lacks sensitivity due to limited statistics) meaning that the violation does not depend on rapidity and is purely a momentum related effect. At low dilepton transverse momentum values below 10 GeV, the magnitude of the violation is very small and compatible with zero. The violation then begins to increase before peaking at approximately 0.2. This behaviour is to be expected since the dilepton transverse momentum is directly caused by QCD. At low p_T^{ll} values there is little QCD activity and hence the Lam-Tung relation remains unviolated. However, as the QCD activity increases at larger momentum values, the magnitude of the violation also increases.

5.10 Outlook

The results presented in this chapter represent not only the highest precision measurement of the angular coefficients of the Z boson to date, but also the first four-fold

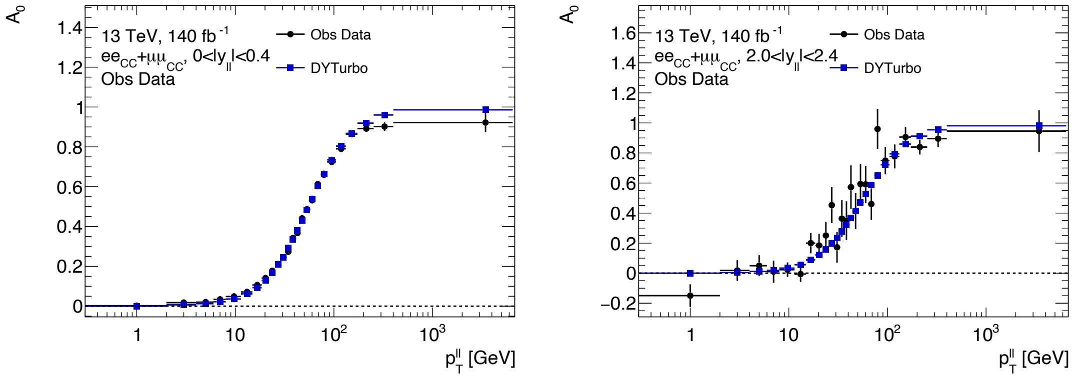

(a) A_0 at low rapidity.

(b) A_0 at high rapidity.

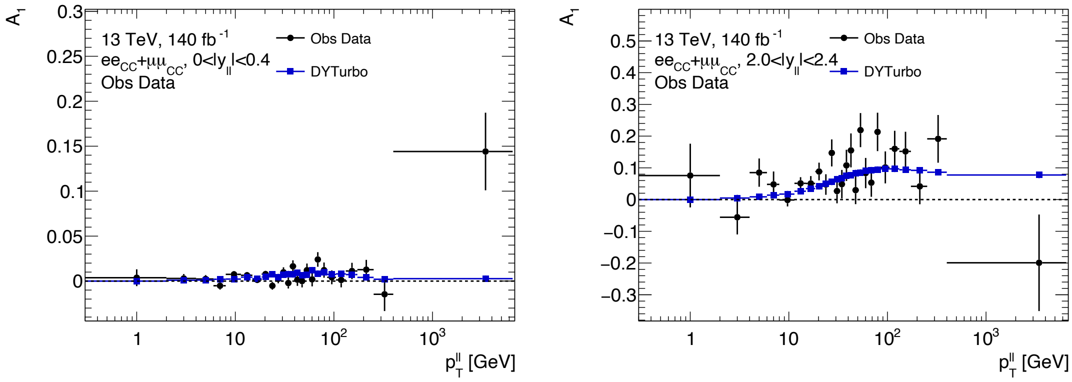

(c) A_1 at low rapidity.

(d) A_1 at high rapidity.

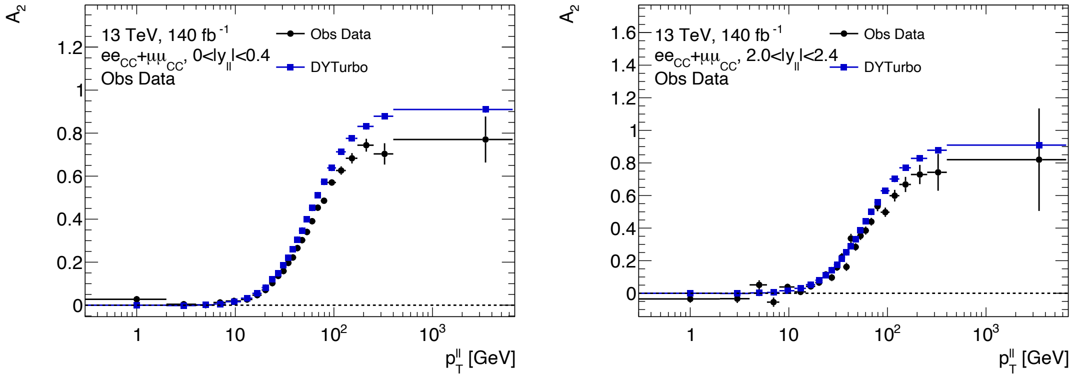

(e) A_2 at low rapidity.

(f) A_2 at high rapidity.

Figure 5.35: The angular coefficients and unpolarised cross section for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

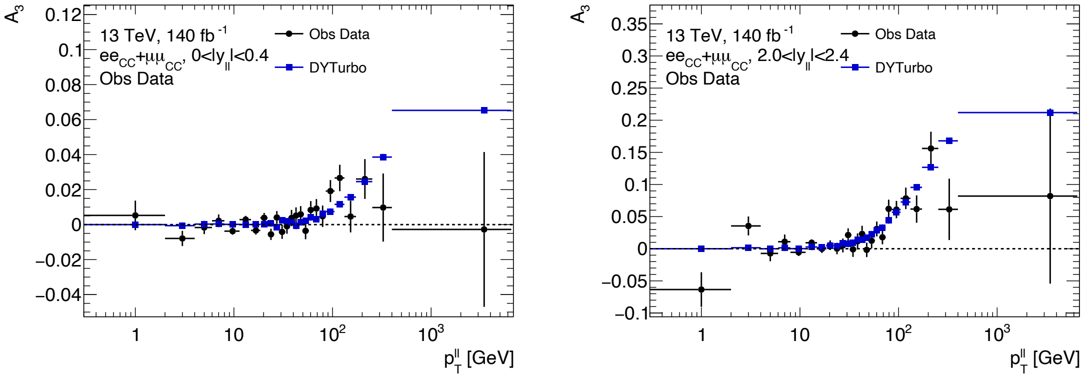

(g) A_3 at low rapidity.

(h) A_3 at high rapidity.

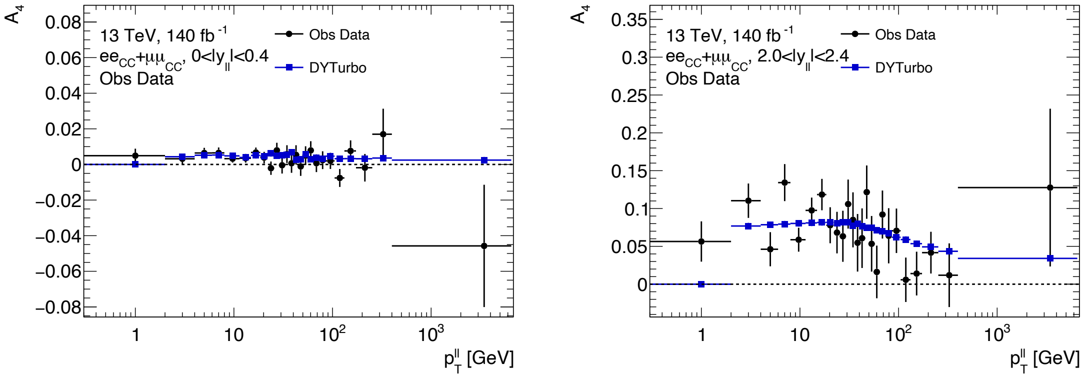

(i) A_4 at low rapidity.

(j) A_4 at high rapidity.

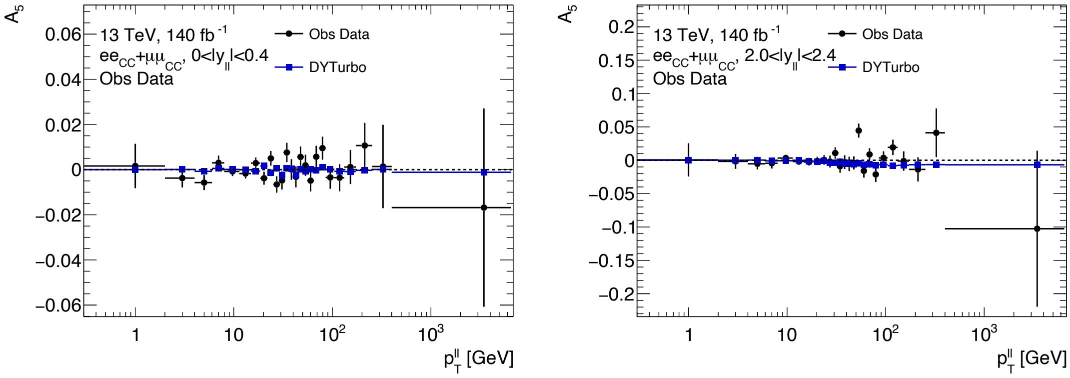

(k) A_5 at low rapidity.

(l) A_5 at high rapidity.

Figure 5.35: (continued) The angular coefficients and unpolarised cross section for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

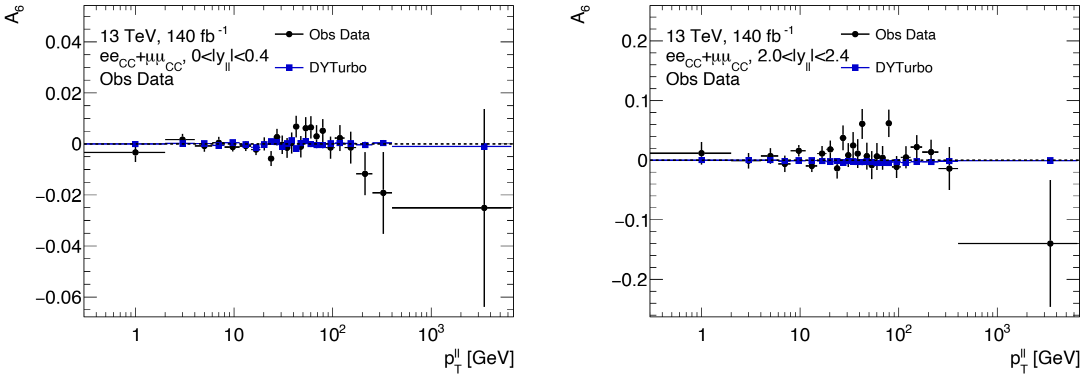

(m) A_6 at low rapidity.

(n) A_6 at high rapidity.

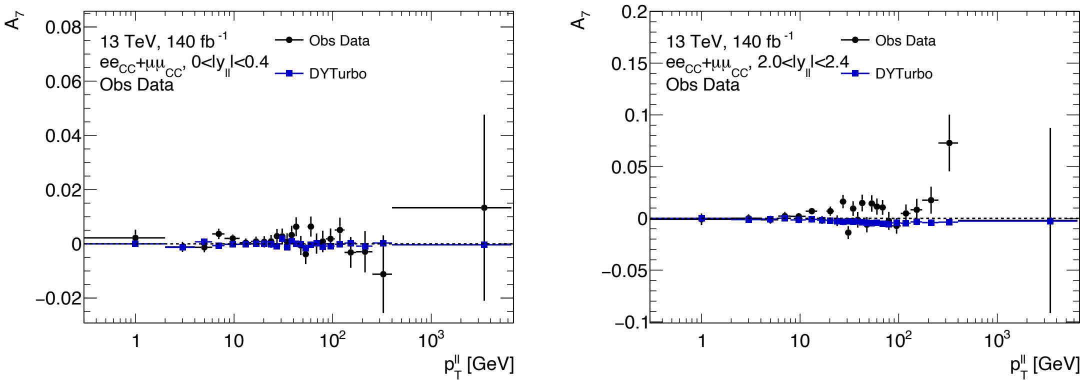

(o) A_7 at low rapidity.

(p) A_7 at high rapidity.

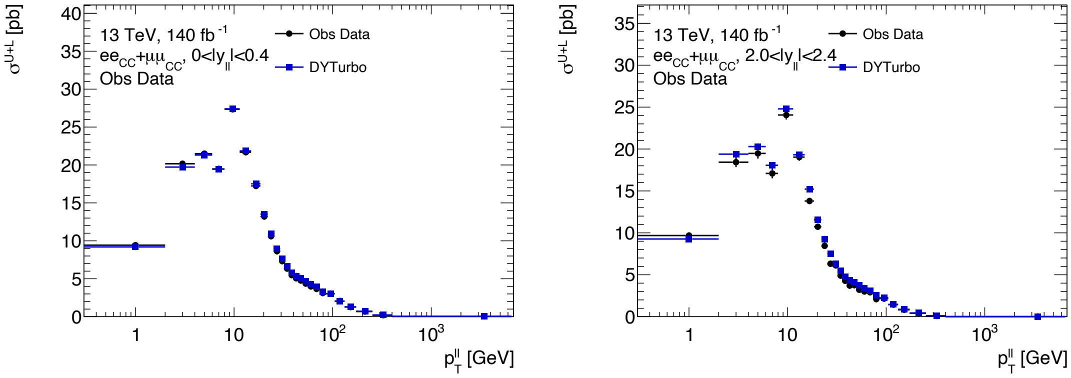

(q) σ^{U+L} at low rapidity.

(r) σ^{U+L} at high rapidity.

Figure 5.35: (continued) The angular coefficients and unpolarised cross section for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

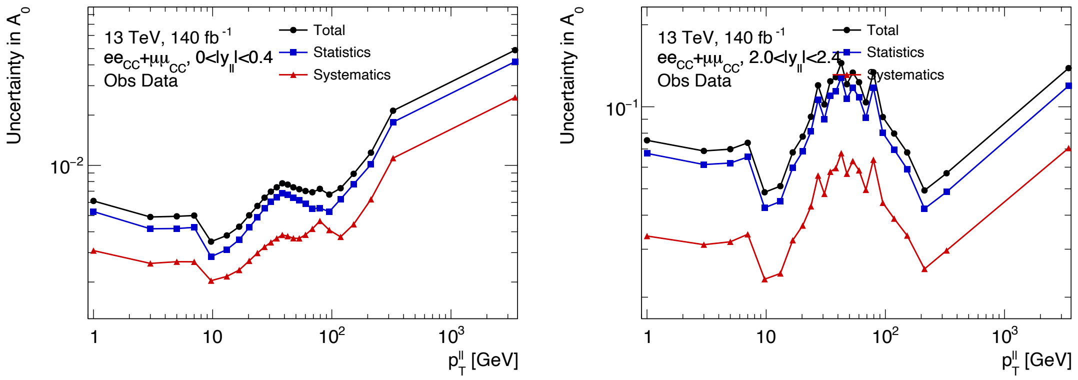

(a) Uncertainty in A_0 at low rapidity. (b) Uncertainty in A_0 at high rapidity.

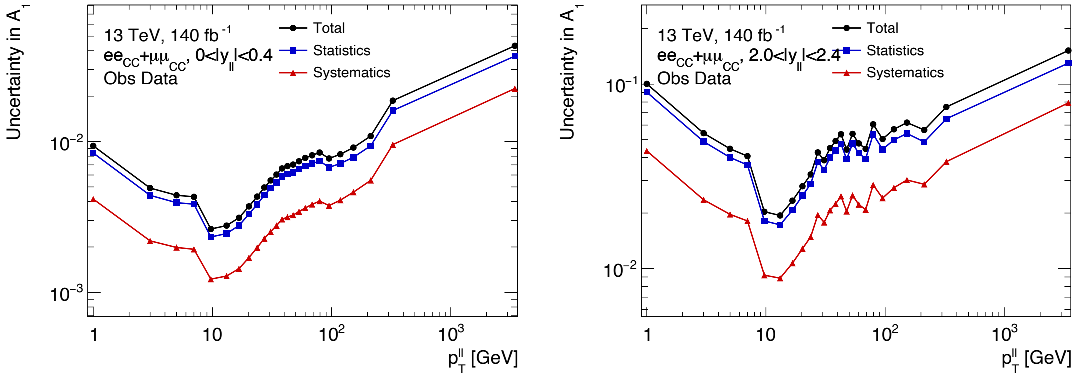

(c) Uncertainty in A_1 at low rapidity. (d) Uncertainty in A_1 at high rapidity.

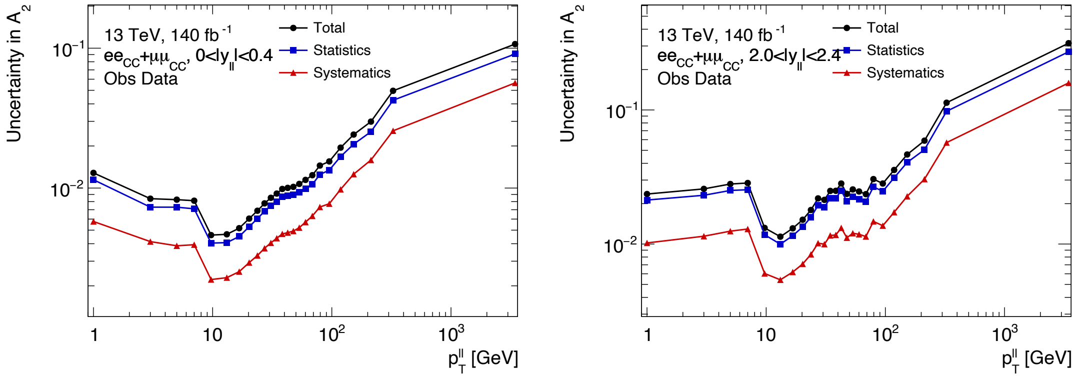

(e) Uncertainty in A_2 at low rapidity. (f) Uncertainty in A_2 at high rapidity.

Figure 5.36: The uncertainty in the angular coefficients and unpolarised cross section for the combined electron and muon channels as a function of the dilepton transverse momentum. The uncertainty is broken down into its total statistical and systematic components for brevity. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

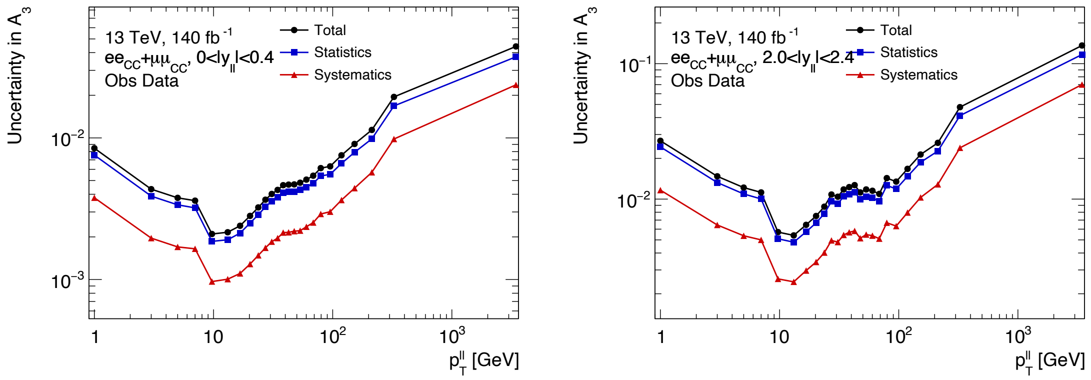

(g) Uncertainty in A_3 at low rapidity.

(h) Uncertainty in A_3 at high rapidity.

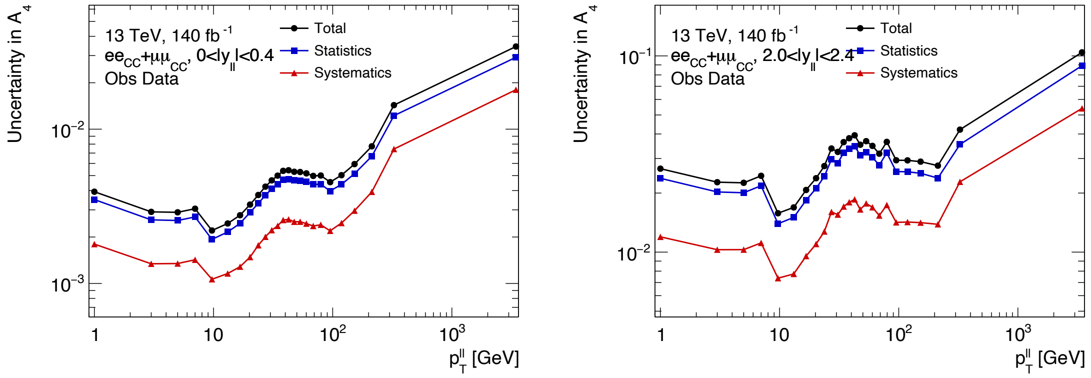

(i) Uncertainty in A_4 at low rapidity.

(j) Uncertainty in A_4 at high rapidity.

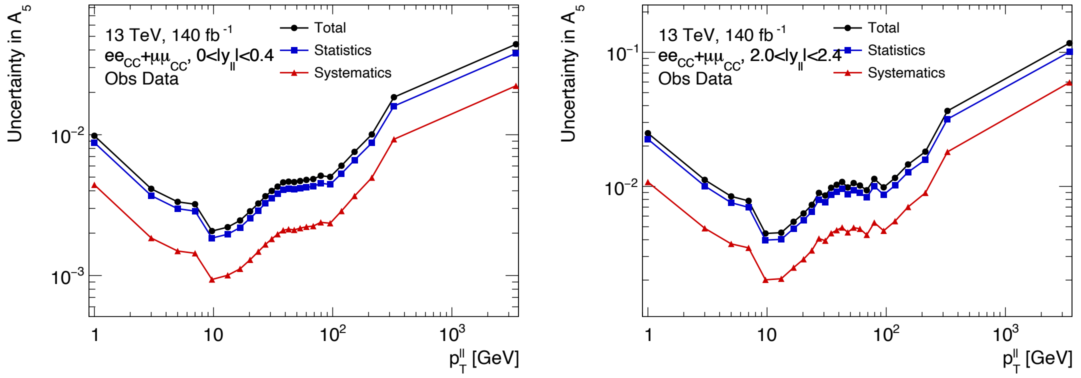

(k) Uncertainty in A_5 at low rapidity.

(l) Uncertainty in A_5 at high rapidity.

Figure 5.36: (continued) The uncertainty in the angular coefficients and unpolarised cross section for the combined electron and muon channels as a function of the dilepton transverse momentum. The uncertainty is broken down into its total statistical and systematic components for brevity. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

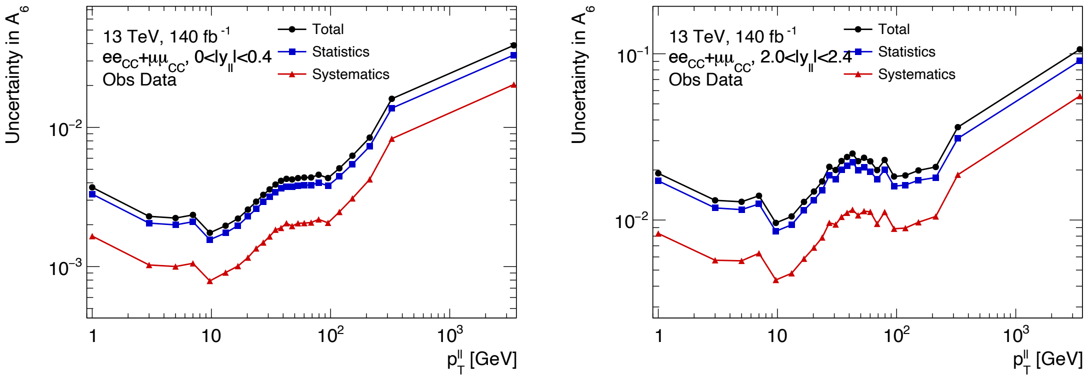

(m) Uncertainty in A_6 at low rapidity.

(n) Uncertainty in A_6 at high rapidity.

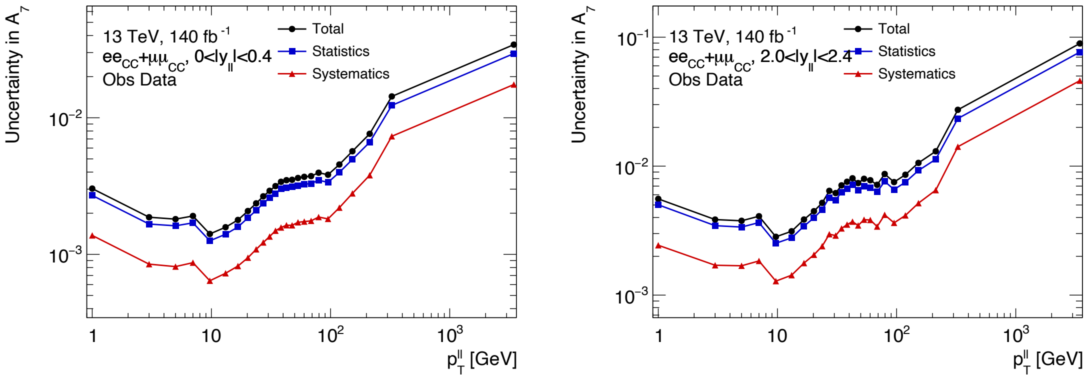

(o) Uncertainty in A_7 at low rapidity.

(p) Uncertainty in A_7 at high rapidity.

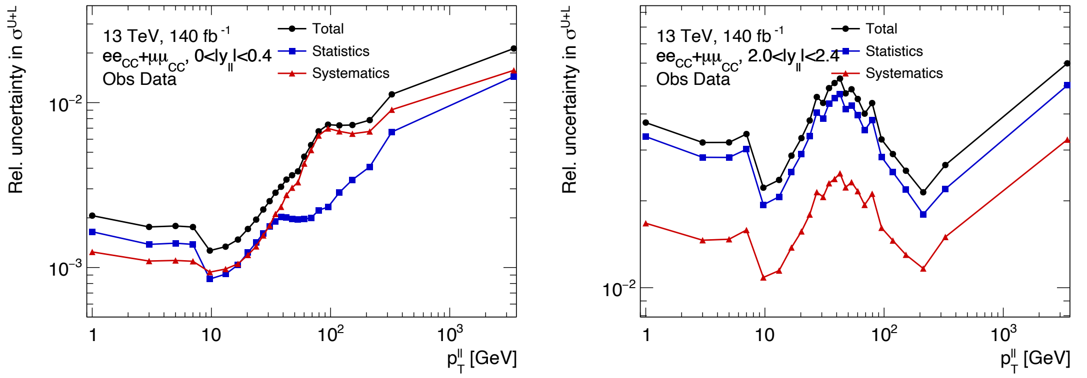

(q) Relative uncertainty in σ^{U+L} at low rapidity.

(r) Relative uncertainty in σ^{U+L} at high rapidity.

Figure 5.36: (continued) The uncertainty in the angular coefficients and unpolarised cross section for the combined electron and muon channels as a function of the dilepton transverse momentum. The uncertainty is broken down into its total statistical and systematic components for brevity. The left column shows the lowest rapidity bin ($0 < |Y_{ll}| < 0.4$) and the right column at the highest central rapidity ($2.4 < |Y_{ll}| < 2.4$).

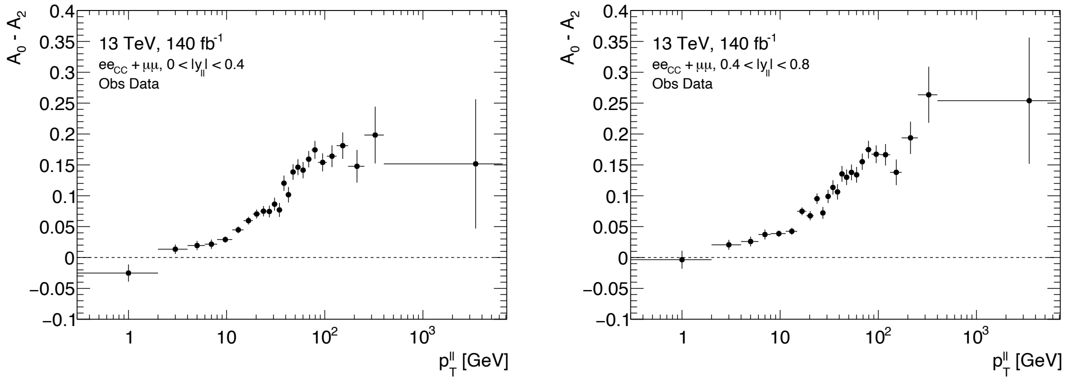

(a) $0 < |Y_{ll}| < 0.4$ (b) $0.4 < |Y_{ll}| < 0.8$

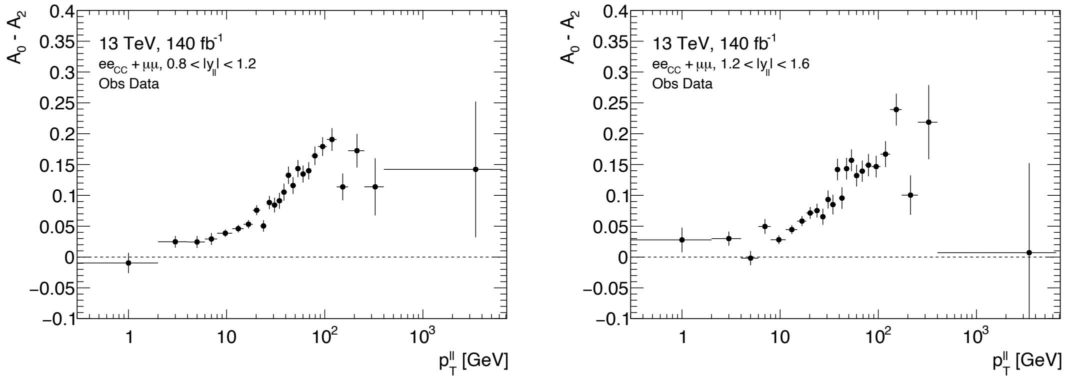

(c) $0.8 < |Y_{ll}| < 1.2$ (d) $1.2 < |Y_{ll}| < 1.6$

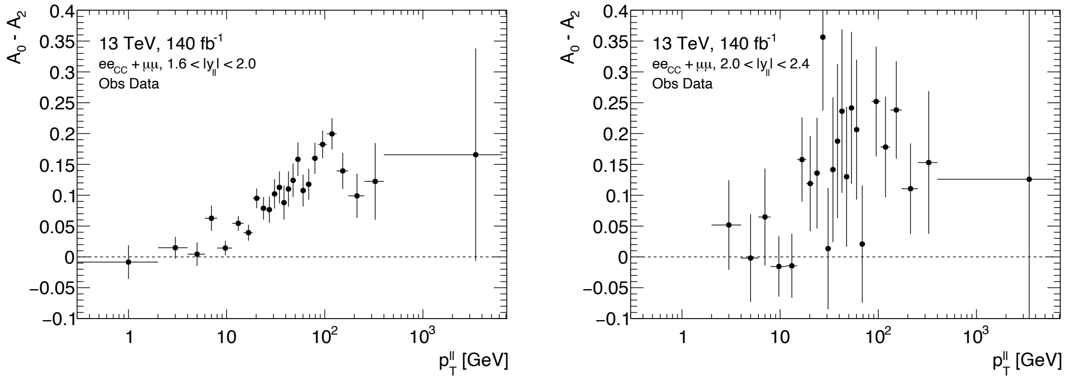

(e) $1.6 < |Y_{ll}| < 2.0$ (f) $2.0 < |Y_{ll}| < 2.4$

Figure 5.37: The Lam-Tung ($A_0 - A_2$) relation for the combination of the electron and muon channels as a function of the dilepton transverse momentum in the observed data. The error bars correspond to the combination of statistical and systematic uncertainties including correlations between A_0 and A_2 in each bin but neglecting any correlations between analysis bins.

differential measurement of the Drell-Yan cross section through the extraction of the coefficients in bins of dilepton rapidity and transverse momentum. They improve greatly on the previous ATLAS measurement of the coefficients at 8 TeV not only due to the reduced uncertainties but also through the much higher granularity nature of the results. These results will therefore aid the study of perturbative QCD and improve the modelling of the Z boson production at hadron colliders and hence improve the background modelling in searches for new physics.

However, these results do not constitute the final measurement. First, the analysis in the muon channel must be revisited in order to solve the issues that are seen with the isolation and trigger efficiency systematics. With this complete, the electron and muon channels will represent the currently best possible precision on the angular coefficients in this region. Whilst fixing these systematics may improve the corresponding uncertainty slightly, the uncertainty in the angular coefficients will still remain (data and Monte Carlo) statistically limited. Therefore, if this analysis were to be repeated at the end of Run 3 of the LHC, assuming the same extraction procedure, $300fb^{-1}$ of collected data and neglecting Z production cross section effects from the higher centre of mass energy ($\sqrt{s} = 13.6$ TeV for Run 3, in comparison to the $\sqrt{s} = 13$ TeV of Run 2), the statistical uncertainty could be reduced by a factor of 1.5. With the planned high luminosity upgrade to the LHC, the total amount of data recorded by ATLAS may reach 3000 fb^{-1}, at which point the statistical uncertainty would fall by a factor of 4.6 in comparison to this measurement. Correspondingly for any hypothetical Run 3 or high luminosity version of this analysis, the corresponding amount of Monte Carlo will also have to be greatly increased so that this does not become the leading source of uncertainty.

However, whilst Run 3 of the LHC is still ongoing and the high luminosity data taking is in the future, the measurement can still be improved using the Run 2 dataset. Namely, the rapidity range used to extract the coefficients can be extended to $|Y_{ll}| < 3.6$ through the central-forward electron channel. This will allow for a greater precision to be reached in the region $1.2 < |Y_{ll}| < 2.4$ since more events will be available to decrease the statistical uncertainty, whilst also improving the measurement of the region $2.4 < |Y_{ll}| < 3.6$ that was last probed by ATLAS at 8 TeV. The inclusion of the central-forward channel will also be crucial when measuring the effective weak mixing angle due to the vastly greater sensitivity in this region, a sensitivity study of which will be presented in the next chapter.

Chapter 6

Extraction of the Weak Mixing Angle

6.1 The Weak Mixing Angle

As introduced in Chapter 2, the weak mixing angle is one of the fundamental parameters of the Standard Model, which allows for the mixing of the W_μ and B_μ fields into the physically observed photon and Z boson. Furthermore, the weak mixing angle can also be used to constrain the values of the masses of the $W^\pm$ and Z bosons. Since the effects of new physics may couple to the Standard Model electroweak gauge bosons, measuring the weak mixing angle as precisely as possible will help to identify the effects of any new physics, or lack thereof. This chapter will present an Asimov sensitivity study on the extraction weak mixing angle at 13 TeV in both the central and central-forward regions from the A_4 parameter previously measured in Chapter 5.

6.1.1 Maximising Sensitivity to A_4

In Equation (2.46) the effective weak mixing angle is explicitly linked to the axial and vector couplings to the leptons. The axial and vector couplings are also linked to the forwards-backwards asymmetry of the Z boson through Equation (2.43). Since A_{FB} is, at leading order, simply a rescaling of the A_4 coefficient through Equation (2.42), the weak mixing angle can be directly extracted from a measurement of A_4. Whilst the A_3 parameter also contributes to the parity violation exhibited by the Z boson, as measured in Chapter 5 and shown in Figure 5.35 this angular coefficient becomes non-zero only at high Z boson transverse momentum values. Since this region of phase space is much more sparsely populated than the lower momentum region, its contribution to the weak mixing angle is minimal in comparison to A_4. Therefore, in this study the extraction of the weak mixing angle is carried out using only the A_4 parameter.

In order to maximise the sensitivity to the weak mixing angle, the value of A_4 must be maximised. In Chapter 5 it was found that in the mass region 80 - 102 GeV the value of A_4 increases as a function of the Z rapidity. As was explained in Section 5.7.1, at high

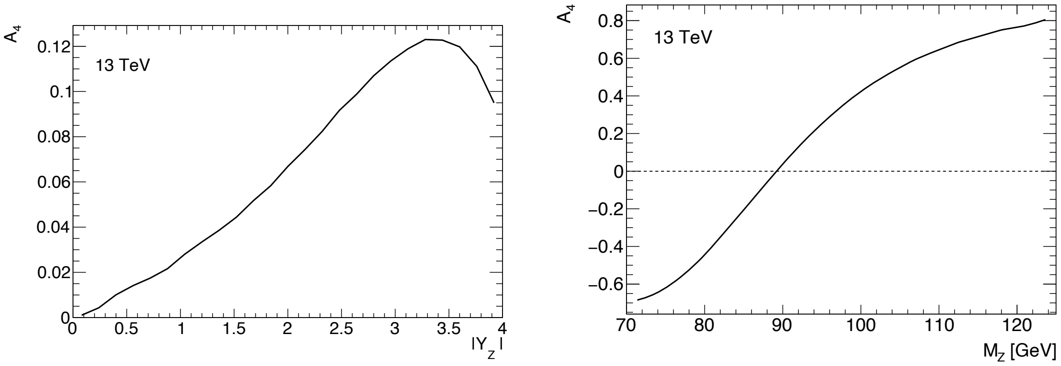

(a) A_4 versus Y_Z (b) A_4 versus M_Z

Figure 6.1: Predicted values of A_4 from Monte Carlo simulated 13 TeV $Z \rightarrow e^+e^-$ events using the POWHEG generator as a function of the Z boson rapidity and invariant mass. The presented values are calculated using the moments method described in Section 5.8.2.

rapidity one of the colliding partons has a high Björken x value and is therefore more likely to be a sea quark than a valence quark. As the forwards backwards asymmetry is enhanced in sea quarks, A_4 grows correspondingly larger with increasing rapidity as can be seen in Figure 6.1a. In addition, the interference between the virtual photon and Z boson mediated processes result in a large variance of A_4 outside of the Z pole region as is shown by Figure 6.1b.

By considering these points, the binning used to measure A_4 must be changed from that used in Chapter 5 to measure the angular coefficients. Since it can be seen in Figure 5.35 that A_4 is approximately constant versus the Z transverse momentum, the statistical sensitivity is maximised by extracting A_4 in a single inclusive p_T^Z bin. In order to increase the rapidity region that can be accessed, the CF channel that was introduced in Section 5.3 is also included in addition to the CC and MM channels that were used to measure the full set of angular coefficients. Two invariant mass sidebands above and below the Z pole region are also introduced. Specifically the bin edges used are:

$$p_T^{ll}, p_T^Z \in \{0, 4000\}\,\mathrm{GeV} \tag{6.1}$$

$$|y^{ll}|, |y^Z| \in \{0, 0.4, 0.8, 1.2, 1.6, 2.0, 2.4, 2.8, 3.2, 3.6\} \tag{6.2}$$

$$m_{ll}, m_Z \in \{70, 80, 100, 125\} \tag{6.3}$$

For the CC and MM channels the rapidity binning begins at 0 and ends at 2.4, whereas the kinematics of the CF channel require that it starts at 1.2 and ends at 3.6.

Since the work on the CF channel is not as progressed as in the central channels, the following sections describing the extraction of A_4 and the weak mixing angle will use an Asimov dataset as outlined in Section 5.8.3. This will allow for an estimate of the sensitivity to the weak mixing angle at 13 TeV to be determined, whilst remaining unbiased towards the observed data.

6.2 Central Forward Channel

In Section 5.3, the central-forward (CF) channel was introduced as a decay channel of the Z boson to one central electron and to one forward electron. As is shown in Figure 5.1, the much lower acceptance and efficiencies associated with forward electrons mean that there are significantly less events available to be studied in this region than in the central channels. However, as outlined in Section 6.1.1 the vastly improved sensitivity to the weak mixing angle that is afforded by the higher rapidity range covered by the CF channel means that the effects of lower statistics are mostly cancelled out. Nevertheless, as with the central channels not every event in the CF channel is suitable for analysis and an event selection process must be followed. This process is similar to that outlined in Section 5.3 but with several differences.

The first difference is that only single lepton triggers are used to select events in the CF channel. Since forward electrons are found outside of the tracking acceptance defined in Section 3.2.2, it is impossible to distinguish between forward electrons and forward photons. In addition there is a much larger number of fake forward electrons in comparison to fake central electrons. For both of these reasons it would be impractical to trigger solely on objects in the EMEC and FCal. Therefore only forward electrons that are recorded with a central electron that fulfils the criteria for the single lepton triggers are considered for analysis.

As a direct consequence of this increased fake rate for forward electrons, the tight identification criterion is also applied in order to minimise the total amount of fakes. Tight identification is also applied to the central electrons to in the CF pair in order to match the identification criteria that was applied at trigger level. A further consequence of the trigger requirements is that the minimum p_T of a central electron is increased to 27 GeV. For forward electrons the minimum p_T is set to 20 GeV in order to maximise statistics but minimise the amount of multijet events from low energy jets.

The pseudorapidity selection on the central electron is also changed to $|\eta| < 2.47$ excluding the transition region between the barrel and endcap calorimeters in order to maximise the statistics at the highest possible dilepton rapidities. The pseudorapidity for forward electrons is defined to cover $2.5 < |\eta| < 4.9$, with the transition region region between the EMEC and FCal of $3.1 < |\eta| < 3.35$ excluded due to the larger amounts of passive material there.

In the central channels, both leptons are required to have opposite electric charges and Monte Carlo events are required to be truth matched. These two selection criteria are,

Selection	Data	Z/γ^*	Top + EW Background
Initial	5.553×10^9	2.090×10^8	3.433×10^7
GRL	5.440×10^9	2.090×10^8	3.433×10^7
Data Quality	5.435×10^9	2.090×10^8	3.433×10^7
Two Leptons	1.276×10^8	2.751×10^7	7.476×10^6
Mass Window	9.427×10^7	2.086×10^7	1.865×10^6
Lepton Kinematics	8.331×10^7	2.047×10^7	1.730×10^6
Lepton ID	1.364×10^7	1.169×10^7	8.178×10^4

Table 6.1: Cutflow showing the effect of applying each selection step to the data and Monte Carlo simulated signal and background events. The signal and backgrounds have been scaled to correspond to a luminosity of 140 fb^{-1}.

however, not applied in the CF channel since it is not possible to directly measure the charge of a forward electron without tracking information and since the truth level information was missing for forward electrons. The effect of the application of each CF selection step can be seen in Table 6.1.

The number of signal Z/γ^* and top and electroweak background events in the CF channel are estimated using simulated Monte Carlo predictions as for the CC and MM channels. The amount of multijet background is estimated using the fake factor method as described in Section 5.5.1.

6.2.1 Control Plots

As with the central channels used to measure the angular coefficients of the Drell-Yan cross section, it is important to ensure that the CF signal region created after the imposition of the analysis level cuts in Section 6.2 is sensible. As is to be expected, the invariant mass spectrum that is displayed in Figure 6.2 shows a clear resonance around the Z peak. The ratio of the data to the expectation is also within ±5% of a perfect agreement across the m_{ee} range considered showing the effectiveness of both the new forward electron calibration and forward identification scale factors that have been developed. As was to be expected from Section 4.4.2, outside of the direct Z pole region the systematic uncertainty in this agreement increases greatly due to the imperfection of the forward electron calibration.

Due to the kinematics of the CF region, the dilepton rapidity shown in Figure 6.3 behaves slightly differently than it does in the CC region by beginning to rise, before peaking and falling again. Nevertheless, this behaviour is well modelled by the Monte Carlo since the data to expectation ratio is within ±10% of perfect agreement until a rapidity of 3.2. Above this point the agreement worsens as the simulations predicts more signal than there is data. This is most likely because the Monte Carlo is only simulated at NLO in QCD yet higher order effects become more prominent at higher rapidity values. This suggests that it may be beneficial to reweight the Monte Carlo to higher order

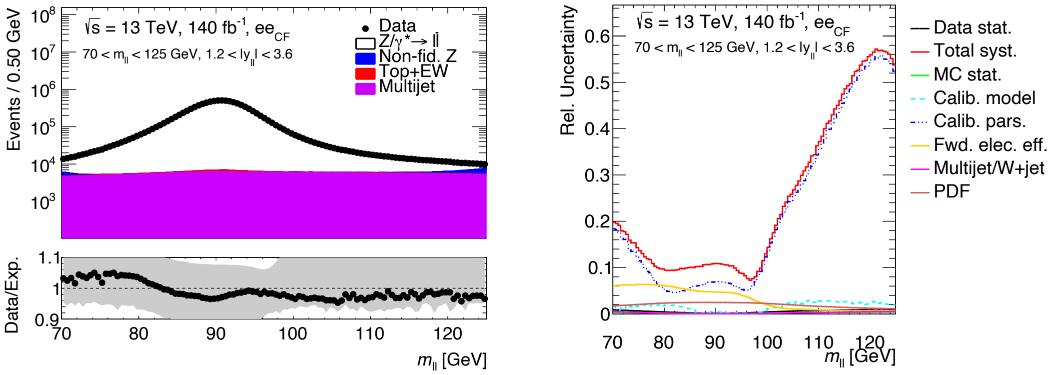

(a) Invariant mass distribution.

(b) Symmeterised breakdown of uncertainty in M_{ll}.

Figure 6.2: The invariant mass distribution control plot with the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) in the CF channel.The grey band in this ratio represents the combined uncertainty from all the systematic variations, along with the uncertainty in the luminosity. The breakdown of this uncertainty is shown in the right column for each systematic type considered, excluding the uncertainty in the luminosity.

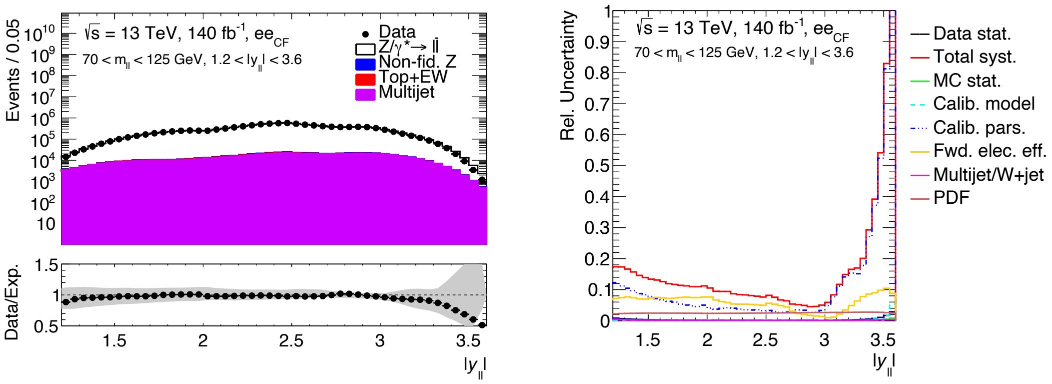

(a) Dielectron rapidity distribution.

(b) Symmeterised breakdown of uncertainty in y_{ll}.

Figure 6.3: The dielectron rapidity distribution control plot with the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) in the CF channel.The grey band in this ratio represents the combined uncertainty from all the systematic variations, along with the uncertainty in the luminosity. The breakdown of this uncertainty is shown in the right column for each systematic type considered, excluding the uncertainty in the luminosity.

QCD predictions as a function of the dilepton rapidity in order to improve the data to expectation ratio. In any case, this high rapidity region has a much larger associated systematic uncertainty that covers this difference due to the imprecision of the forward electron calibration in the FCal.

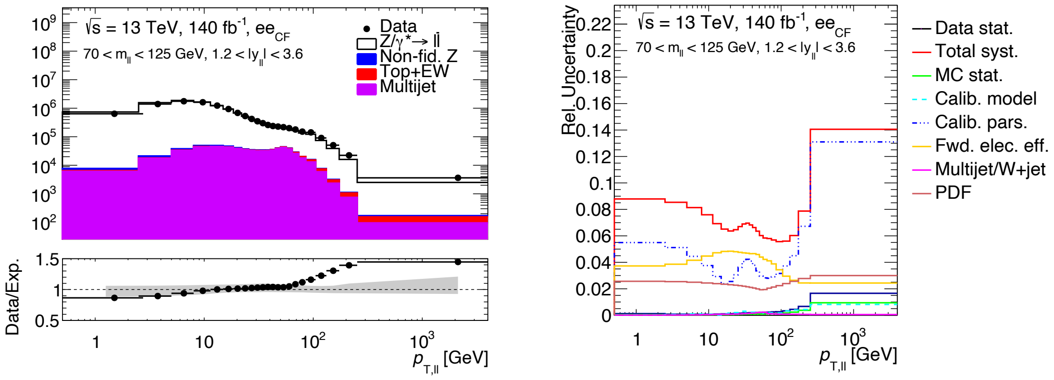

(a) Dielectron transverse momentum distribution.

(b) Symmeterised breakdown of uncertainty in p_T^{ll}.

Figure 6.4: The dielectron transverse momentum distribution control plot with the ratio of the data to the expectation (Monte Carlo signal, background and fake estimate) in the CF channel.The grey band in this ratio represents the combined uncertainty from all the systematic variations, along with the uncertainty in the luminosity. The breakdown of this uncertainty is shown in the right column for each systematic type considered, excluding the uncertainty in the luminosity.

Whilst a reweighting of the dilepton transverse momentum to higher order predictions of DYTurbo is applied to the Monte Carlo, as can be seen in Figure 6.4 it is clearly insufficient in the CF channel. Here the data to expectation ratio shows a disagreement of up to 50% which is not the case for the central channels. In order to improve the data to expectation ratio, the reweighting that is applied can be binned in the dilepton rapidity as is done for the central channels. As work on the CF channel is also still ongoing, this agreement will be improved at a future point.

Since it is impossible to distinguish between forward electrons and forward photons and the lack of tracking vetoes mean it is more likely for a low energy jet to be classified as a forward electron, it is to be expected that there are large levels of multijet background in the CF channel. This is indeed what is seen in Figure 6.5, where the multijet contribution dominates the total background fraction. In the 80 - 100 GeV mass region the background makes up approximately 5% of the total signal region since the Z boson is resonant in this region. Outside of this region the total background fraction increases greatly and varies between 10 - 50% due to increasing levels of both multijet events and non-fiducial Z bosons.

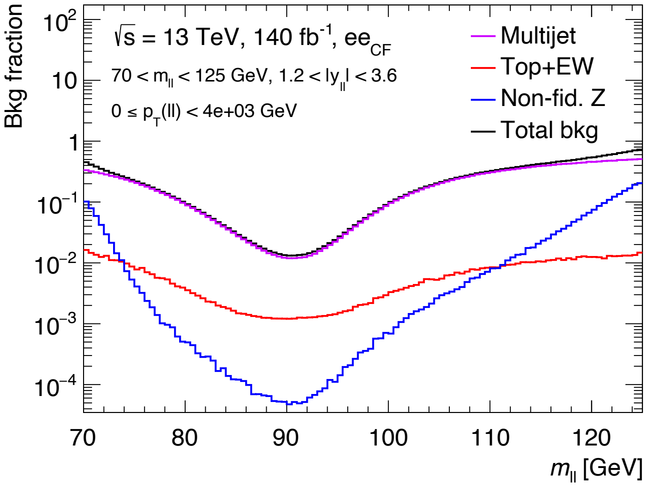

Figure 6.5: Fraction of the signal region that corresponds to each source of background events as a function of the dilepton invariant mass in the CF channel.

6.3 A_4 Results

For each of the analysis bins defined in Section 6.1.1, the same procedure as was used for measuring the angular coefficients in Chapter 5 is again used for the extraction of the A_4 coefficient in all three analysis channels. In order to present a conservative estimate on the uncertainty of A_4, all of the angular coefficients and unpolarised cross section are simulataneously fit alongside A_4 although their values are not relevant for the extraction of the weak mixing angle[1].

6.3.1 Central Channels

The Asimov A_4 results using the binning described in Section 6.1.1 can be seen in Figure 6.6. In all M_{ll} regions the value of A_4 increases as a function of rapidity, with the magnitude of A_4 larger outside of the Z boson pole in both sidebands as is to be expected. In all but two analysis bins A_4 is statistically rather than systematically limited. Given the small difference between the total uncertainty and the statistical uncertainty, the values of A_4 in the central channels are as accurate as they can be and cannot be improved further. This therefore means that the sensitivity to the weak mixing angle has also been maximised for this dataset and choice of selection criteria.

6.3.2 Central Forward Results

In Figure 6.7 the expected A_4 values follow the same trends as those in the Figure 6.6 but with a larger magnitude as is to be expected from the higher rapidity regions be-

---

[1] The value of the unpolarised cross section is an exception to this statement due to its power to constrain PDF uncertainties which will be further discussed in Section 6.5.

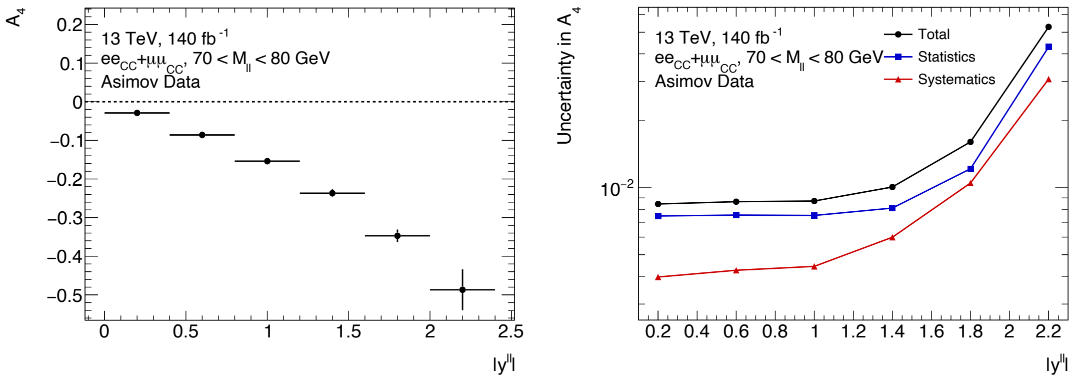

(a) A_4 for $70 < M_{ll} < 80$ GeV (b) Uncertainty in A_4 for $70 < M_{ll} < 80$ GeV

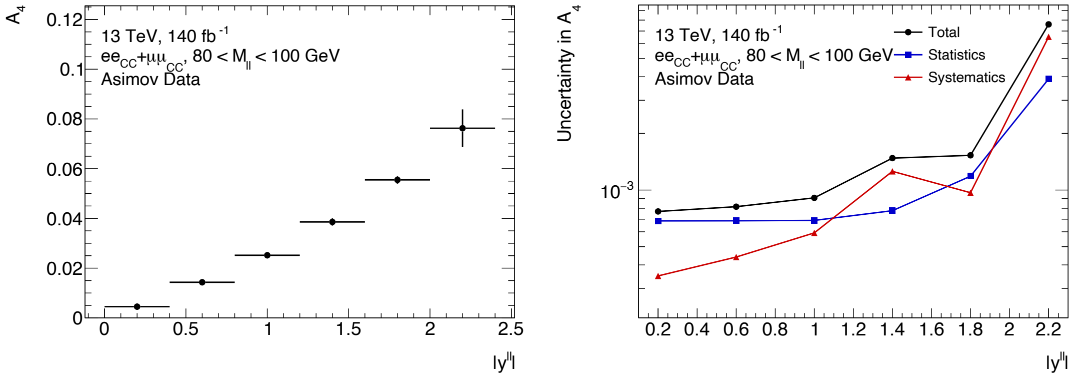

(c) A_4 for $80 < M_{ll} < 100$ GeV (d) Uncertainty in A_4 for $80 < M_{ll} < 100$ GeV

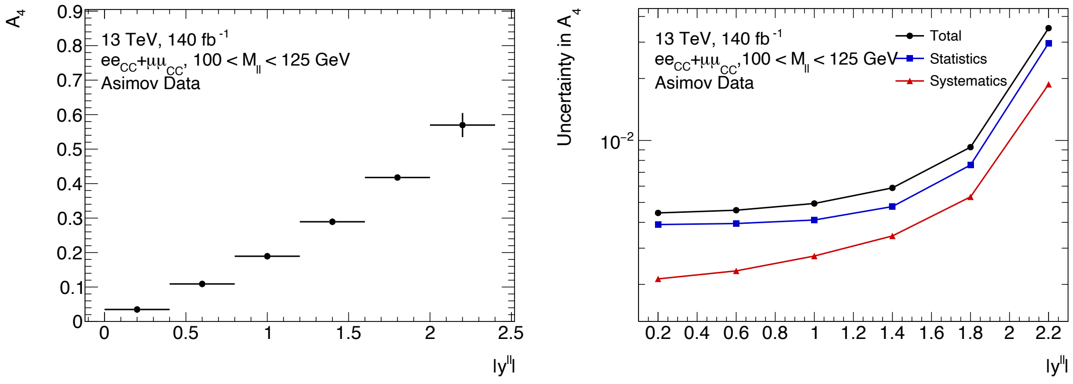

(e) A_4 for $100 < M_{ll} < 125$ GeV (f) Uncertainty in A_4 for $100 < M_{ll} < 125$ GeV

Figure 6.6: Plots of the Asimov value of A_4 (left column) and the corresponding breakdown of statistical and systematic uncertainty (right column) for the combination of the central electron and muon channels. The plots are displayed as a function of the dilepton rapidity in each invariant mass bin.

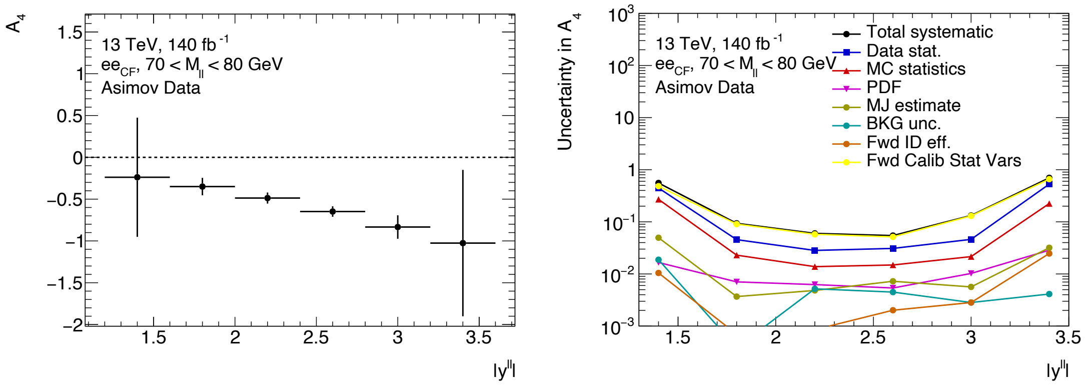

(a) A_4 for $70 < M_{ll} < 125$ GeV (b) Uncertainty in A_4 for $70 < M_{ll} < 80$ GeV

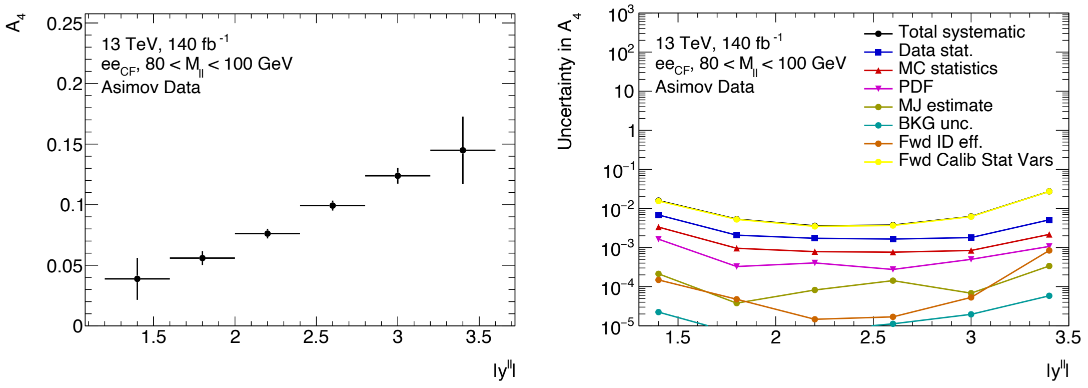

(c) A_4 for $80 < M_{ll} < 100$ GeV (d) Uncertainty in A_4 for $80 < M_{ll} < 100$ GeV

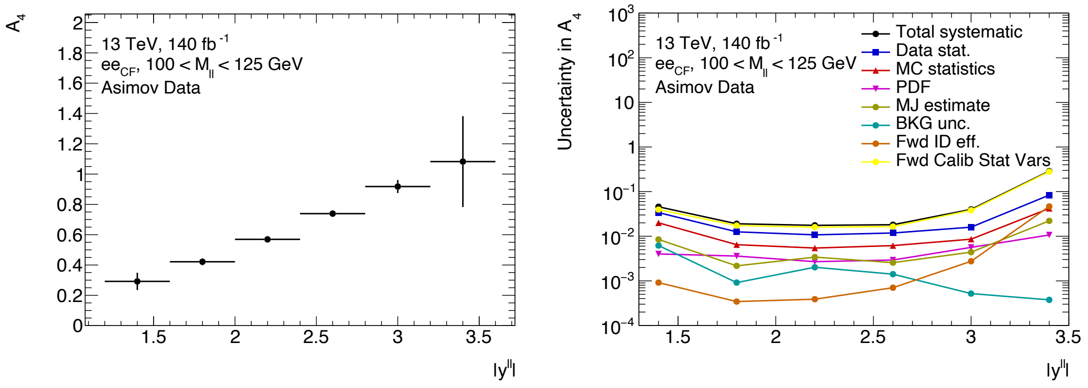

(e) A_4 for $100 < M_{ll} < 125$ GeV (f) Uncertainty in A_4 for $100 < M_{ll} < 125$ GeV

Figure 6.7: Plots of the Asimov value of A_4 (left column) and the corresponding breakdown of statistical and systematic uncertainty (right column) for the central forward channel. The plots are displayed as a function of the dilepton rapidity in each invariant mass bin.

ing probed. However, the uncertainties on each value is larger for every analysis bin in comparison to the central values. Whilst this is largely driven by the lower amount of statistics in the CF channel, the leading (systematic) uncertainty is the variation of the forward electron calibration parameters. As was outlined in Section 4.4.2, the calibration process on the forward electrons could be refined further which would have the effect of constraining its associated systematic uncertainty. The PDF variations are then the second most important systematic variation in the $80 < M_{ll} < 100$ GeV region and are of approximately the same order of magnitude as the multijet background systematic in the mass sidebands. The multijet uncertainty could also be improved further if the fake factors were calculated for the full Run 2 period rather than the individual data taking years as they currently are. This would have the effect of reducing the statistical uncertainty on each fake factor and hence the associated systematic uncertainty. However, since the multijet uncertainty is dwarfed by that from the forward electron calibration any improvements here would be minimal.

6.3.3 Combination of Central and Central Forward Channels

After analysing the central and central forward channels separately, the result from the combination of all three channels can be seen in Figure 6.8. As is to be expected the statistical sensitivity is improved where the central and central forward channels overlap i.e in the range $1.2 < |Y_{ll}| < 2.4$. This means that that only the highest rapidity region in each invariant mass bin is expected to have a significant uncertainty. However, since this bin is systematically limited by the forward electron calibration, the final sensitivity when evaluating using the observed data could be improved further.

6.4 Extraction Method

In order to extract the weak mixing angle from the A_4 values predicted in Section 6.3, the A_4 term in the likelihood described by Equation (5.44) is modified. Instead of fitting A_4, it is replaced by a linear model of the weak mixing angle i.e:

$$A_{4,j} = a_j(\theta) \cdot \sin^2 \theta^l_{eff} + b_j(\theta) \tag{6.4}$$

where the a and b parameters are fit in analysis bin j following the same method described in Section 5.8.3. The a and b parameters will also be affected by the same systematic variations as for the measurement of the angular coefficients and hence they are both functions of the nuisance parameters, θ. By rearranging Equation (6.4), a value for the weak mixing angle can then be estimated in each analysis bin by using the separately fitted values of A_4, a and b. A final value for $\sin^2 \theta^l_{eff}$ can then be calculated by taking the weighted average of its respective value in each analysis bin. However, since this study presents the Asimov results on the extraction of the weak mixing angle, the central value itself is irrelevant. Instead, this work will assume a central value that corresponds to that of the current most precise measurement of the weak mixing angle.

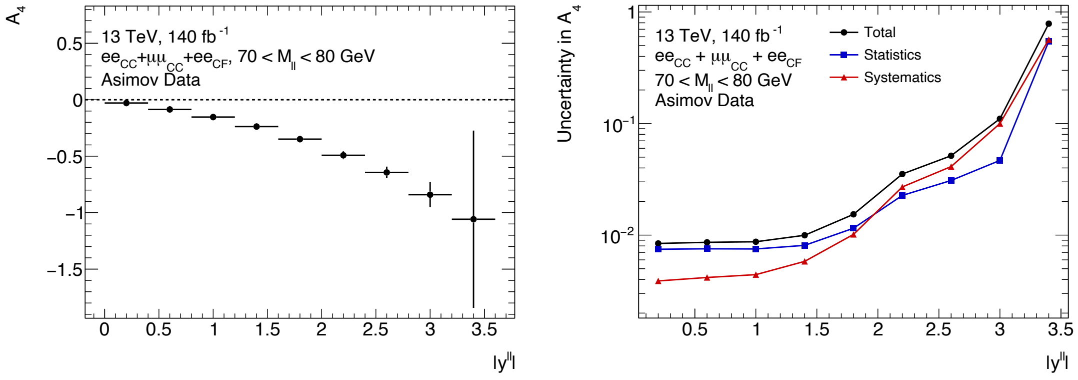

(a) A_4 for $70 < M_{ll} < 100$ GeV (b) Uncertainty in A_4 for $70 < M_{ll} < 80$ GeV

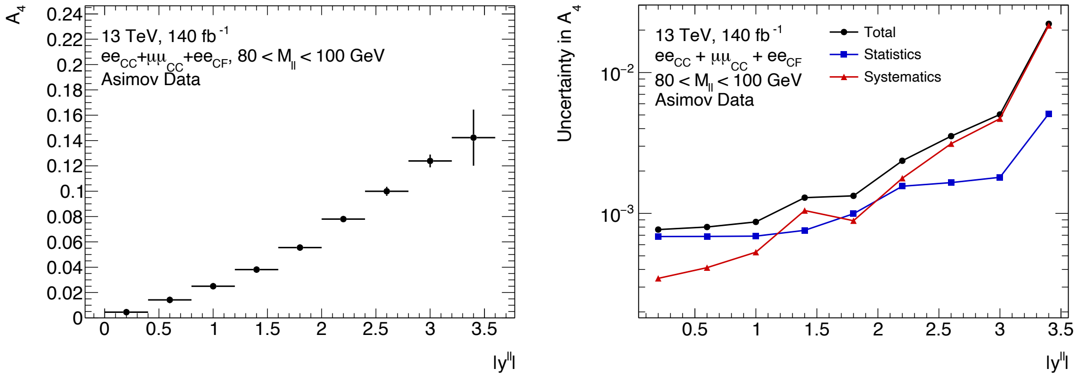

(c) A_4 for $80 < M_{ll} < 100$ GeV (d) Uncertainty in A_4 for $80 < M_{ll} < 10$ GeV

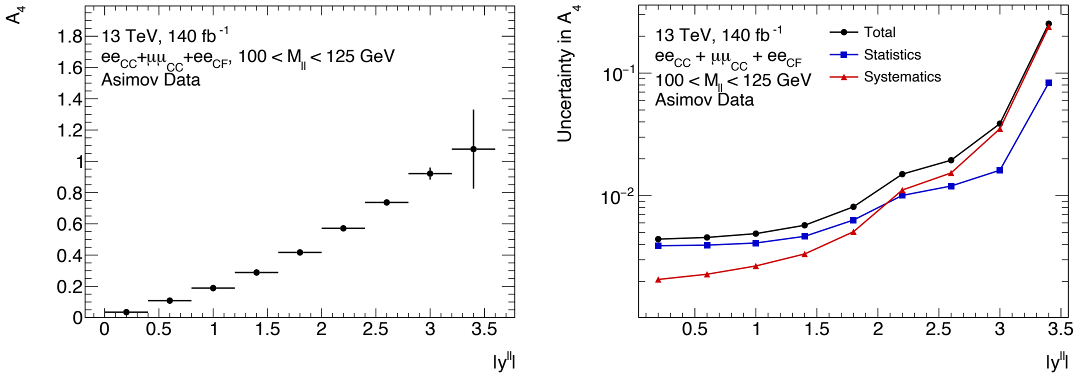

(e) A_4 for $100 < M_{ll} < 125$ GeV (f) Uncertainty in A_4 for $100 < M_{ll} < 125$ GeV

Figure 6.8: Plots of the Asimov value of A_4 (left column) and the corresponding breakdown of statistical and systematic uncertainty (right column) for the combination of the central electron and muon channels and central forward electron channel. The plots are displayed as a function of the dilepton rapidity in each invariant mass bin.

The final uncertainty in the weak mixing angle is estimated using the xFitter framework [16] to perform the statistical analysis. This statistical analysis works by following the method of Pinto et al [78] to propagate the decomposed uncertainties of A_4 (shown in Section 6.3) and the unpolarised cross section into an uncertainty in the weak mixing angle. The extraction of the weak mixing angle is then performed at leading order in QCD and incorporates the effects of PDF profiling in order to minimise the theoretical uncertainty caused by the PDFs.

6.5 Weak Mixing Angle Results

Uncertainty	Δ eeCC (10^{-5})	Δ MM (10^{-5})	Δ eeCF (10^{-5})
Stat	44	34	36
MC Background	4	4	1
Central Electron Calibration	21	-	-
Muon Calibration	-	13	-
Forward Calibration	-	-	18
Central Lepton Efficiency	5	2	-
Forward ID Efficiency	-	-	5
Multijet Background	2	-	2
PDF (exp)	7	8	2
PDF (theory)	26	23	22
MC Stat	-	-	-
Total	56	44	46

Table 6.2: Breakdown of statistical and systematic uncertainty in the weak mixing angle for each analysis channel when only the A_4 parameter is used for fitting. All uncertainties are quoted in units of 10^{-5}.

The results presented in Table 6.2 show the result of extracting the weak mixing angle using only the A_4 parameter in each of the three analysis channels using the NNPDF 4.0 PDF set [25]. In each channel, the theoretical uncertainty caused by varying the PDFs at truth level is the largest source of uncertainty after the statistical uncertainty. This uncertainty can, however, be constrained by updating the weight of a replica PDF set using some new data that is itself sensitive to the PDFs [28]. Since the PDFs affect the distribution of the unpolarised cross section as a function of invariant mass and rapidity the unpolarised cross section is used to update the PDF weights. Therefore, the unpolarised cross section is concurrently fit with the reparameterised A_4 in order to provide an accurate updating of the PDF weights. When this is done, the PDF theory uncertainty in the CF channel falls by 10% from 22×10^{-5} to 20×10^{-5} and by similar amounts in the CC and MM channels.

The electron energy calibration also represents a relatively large source of uncertainty at 21×10^{-5} and 18×10^{-5} for the central and forward electrons respectively. For the CF

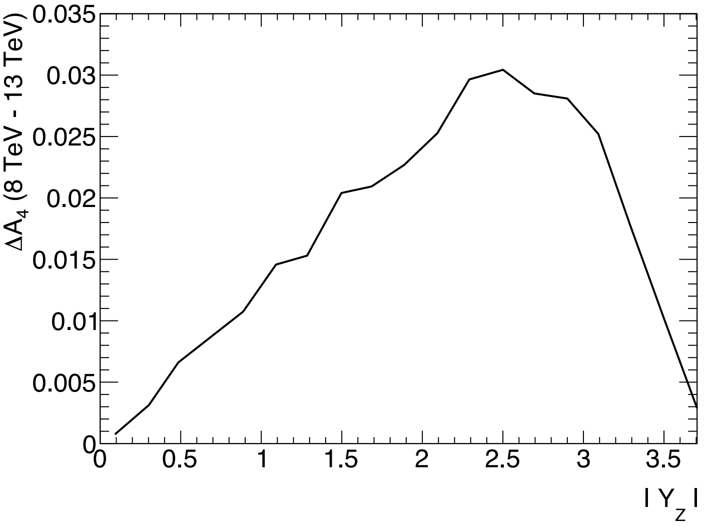

Figure 6.9: The difference between A_4 at 8 and 13 TeV as a function of the Z rapidity. The values used are taken from DYTurbo predictions.

channel it is expected that this number will be lowered in the future with an improved calibration following the outlined steps in Section 4.4.2. However, even if the forward calibration were to be perfected and have a corresponding uncertainty of zero, then the total uncertainty in the CF channel would only fall by 9% to 42×10^{-5}. Overall, this means that the sensitivity in the CF channel has not yet been maximised and can still be improved further.

The statistical uncertainty in the CF channel of 36×10^{-5} is also higher than the corresponding 8 TeV measurement of 21×10^{-5} [41] despite a seven-fold increase in statistics from 8 to 13 TeV. As can be seen in Figure 6.9, the CF channel is particularly susceptible to the effects of dilution since the difference in the DYTurbo predicted A_4 values is larger here than it is in the CC region. In addition, the higher pileup conditions of Run 2 also play a significant role. During 8 TeV data taking, there were on average 21 interactions per bunch crossing [45] in comparison to 34 during the Run 2 data taking [46]. This increased number of interactions caused the amount of pileup noise to increase. Since the cell sizes of the EMEC-IW and FCal are larger than those in the EMB and EMEC-OW this increase in pileup noise cannot be accounted for in the forward region as well as it can be centrally. This therefore means that the resolution of the Z boson is worsened and the statistical sensitivity is reduced. As described in Section 4.4.2, whilst the MVA calibration is partially able to offset this effect it is unable to replicate the performance of the central region.

In comparison to the PDF and calibration uncertainties, all other sources of uncertainty are relatively small. This means that any further optimisation of these systematic groups would only have a very small effect on the final uncertainty that can be achieved.

Uncertainty	Uncertainty (10^{-5})
MC Background	3
CC + MM Calibration	5
CF Calibration	6
Efficiency	2
Multijet	1
PDF (Experimental)	5
PDF(Theory)	14
Total Systematic	18
Statistical	22
Total	28

Table 6.3: Breakdown of statistical and systematic uncertainty in the weak mixing angle for the combination of the two central lepton channels and central forward electron channel when both the A_4 and σ^{U+L} parameters are used for fitting. All uncertainties are quoted in units of 10^{-5}.

After considering the uncertainty in each individual channel, when the three channels are combined and the fit is carried out using the unpolarised cross section the resultant uncertainties are as displayed in Table 6.3. The total estimated uncertainty in this combination of 28×10^{-5} would be an extremely precise result. If it were to be replicated when looking at observed data, it would be the most precise measurement of the weak mixing angle ever at a hadron collider.

When these uncertainties are compared to other measurements of the weak mixing angle that are summarised in Figure 6.10, this uncertainty would be 9.7% better than the preliminary 13 TeV measurement of the weak mixing angle by the CMS experiment [65] and 22% better than the previous preliminary ATLAS 8 TeV measurement of the weak mixing angle [41]. It would also slightly improve upon the precision achieved in the combination result from the Large Electron Positron collider (LEP) and Stanford Large Detector (SLD) when measuring the weak mixing angle from the forwards-backwards asymmetry of bottom quarks [17].

When the three channels are combined, and assuming a central value of the current most precise measurement of the weak mixing angle, the final estimate of the weak mixing angle is then:

$$\sin^2\theta^l_{eff} = 0.23152 \pm 0.00022\,(\text{stat}) \pm 0.00010\,(\text{systematic}) \pm 0.00014\,(\text{theory}) \tag{6.5}$$

$$\sin^2\theta^l_{eff} = 0.23152 \pm 0.00028 \tag{6.6}$$

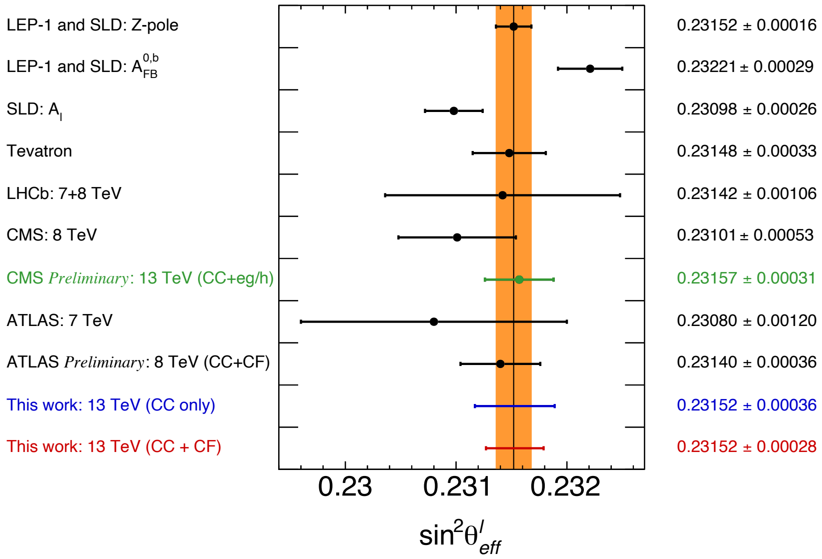

Figure 6.10: An overview of various measurements of the weak mixing angle from lepton colliders, the Tevatron and the LHC.

6.6 Outlook

Although the predicted uncertainty in the weak mixing angle is impressive, it is still be possible to improve it further. Since the dominant uncertainty in both the central electron and muon channels is the statistical uncertainty in the data, if the high mass sideband upper limit is increased from 125 to 150 GeV then the statistical uncertainty should correspondingly decrease. However, given that the effects of dilution at such high invariant mass values will weaken the effectiveness of the gained events, the effect on the statistical uncertainty will likely be minimal. Whilst the lepton calibration systematics also provide a relatively large contribution to the overall uncertainty, it is unlikely that the central electron and muon calibration processes will be further optimised for the Run 2 data.

In comparison, improvements to the central forward channel are much more feasible. If the improvements to the cailbration process that are outlined in Section 4.4.2 are implemented then this will likely reduce the large calibration uncertainties that are particularly visible in the high mass sideband in Figure 6.2. The statistical error could also be improved if the resolution of the Z mass spectrum were also to be improved and brought closer to 8 TeV levels. However, this would only be possible through a much improved MVA calibration which is unlikely to be possible given the constraints imposed by the previously mentioned imperfect modelling of the pileup noise.

After combining the three analysis channels at 13 TeV, the next logical step would then be to combine the 13 TeV results with the 8 TeV results. It is estimated that if this is done the uncertainty could be as low as 24×10^{-5}, which would be a further 14% improvement in the uncertainty on the combined 13 TeV result. This forecast is, however, much more speculative given that the ATLAS 8 TeV result is yet to be finalised and the 13 TeV measurement must still be optimised. Nevertheless, it is truly spectacular that such levels of precision are within the realms of possibility from a hadron collider and are a testament to the extremely efficient operation of the LHC and also to the physicists who have worked to keep the systematic uncertainties in check.

Chapter 7

Conclusion

When the Higgs boson was discovered in 2012, it would have been simple to declare the Standard Model as complete and move on to looking solely for BSM phenomena after measuring the properties of the new Higgs particle. However, instead of withering on the vine, the study of the Standard Model remains alive and well. Indeed, it is flourishing as new Standard Model processes have been measured for the first time and well known processes and parameters have been measured at ever increasing levels of precision. It is highly likely that signs of new physics will manifest themselves as small departures from the theoretical values of Standard Model given the lack of direct evidence for new physics provided by BSM searches. It is with this backdrop that this thesis has presented a $\sqrt{s} = 13$ TeV measurement of the Drell-Yan cross section through its decomposition to 8 angular coefficients and an unpolarised cross section. Through the A_4 coefficient it has also been possible to perform a study on the sensitivity to the weak mixing angle, one of the fundamental parameters of the Standard Model.

The measurement of all of the angular coefficients and unpolarised cross section has been presented differentially in the rapidity and transverse momentum in both the electron and muon channels as well as the combination of the two channels. This measurement reaches a very high level of precision by exploiting the 140 fb^{-1} of the ATLAS Run 2 dataset. Despite the large amount of events this provides for analysis purposes, the statistical uncertainty in the data still remains the leading source of uncertainty in this measurement. The total systematic uncertainty is also dominated by statistical uncertainty of the finite number of simulated Monte Carlo events which shows that all other sources of systematic uncertainty are well understood and controlled for. After verifying that the electron and muon channels were compatible with each other, the angular coefficients and unpolarised cross section were measured in the combination of the two channels. In this combination measurement, the precision in the unpolarised cross section ranges from approximately 0.2% at low rapidity to approximately 3% at high rapidity which highlights the very high levels of precision that have been reached. In addition, the combination results are compared to the Standard Model predictions for the angular coefficients from the DYTurbo package where a good agreement is observed.

In order to maximise the effectiveness of the sensitivity study to the weak mixing angle, a further analysis channel was considered alongside the two central electron and muon channels, that of the central forward electron channel. In order to use the forward electrons detected by ATLAS for physics analysis purposes, an entirely new forward electron calibration chain was developed. The foundation of this calibration chain was a multivariate algorithm composed of an ensemble of boosted decision trees which was trained to correct the reconstructed energy of an electron towards its true energy. This resulted in large improvements to both the electron energy scale and resolution across the forward phase space which naturally filtered through to the Z boson mass spectrum in the central-forward channel. Some fine tuning of the calibration was necessary through the application of a uniformity correction before an in situ calibration was derived. This in situ calibration had the effect of matching the simulated mass spectrum of the Z boson to that in data to a very precise level of within ±5% across the full Run 2 period. Whilst the systematic uncertainties were found to be large, improvements to the process have been identified which should reduce these uncertainties greatly.

Once the forward electron calibration work was completed, the A_4 coefficient could be measured in the combination of the central electron and muon and central forward electron channels. By doing this both the accessible rapidity region could be extended to $|Y_{ll}| < 3.6$ which had the effect of maximising the sensitivity to the weak mixing angle. By translating the uncertainty in A_4 to an uncertainty on the weak mixing angle, the full Run 2 ATLAS measurement of the weak mixing angle may reach the level of 28×10^{-5} in its level of precision. If this value were to be reached, or indeed improved further, it would be a spectacularly precise result that would become competitive with results achieved at lepton colliders. It would also completely disprove the old adage that hadron colliders are only useful for discovering new physics whilst lepton colliders are necessary for precision measurements and cement the LHC as both a discovery machine and a precision machine.

A | Electron Channel Pull Plots

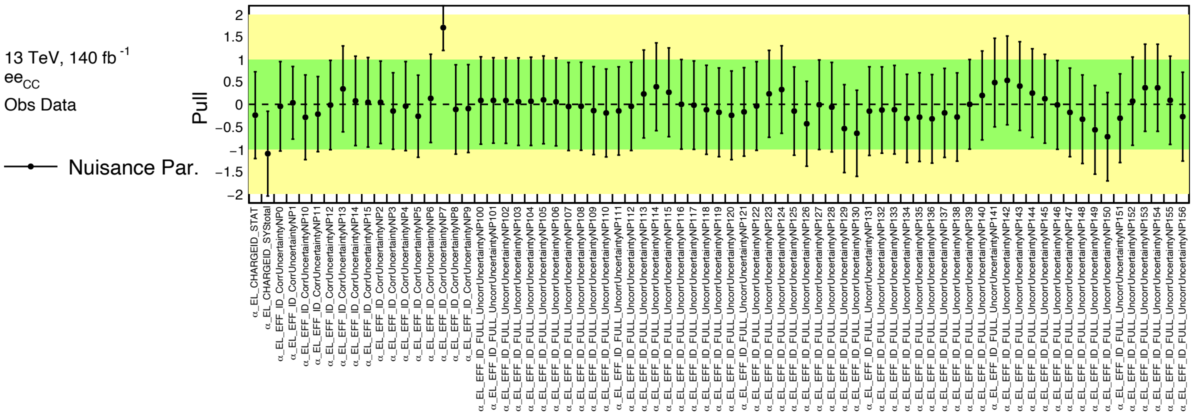

Figure 1: Pull plots for the electron (charge) identification efficiency variations.

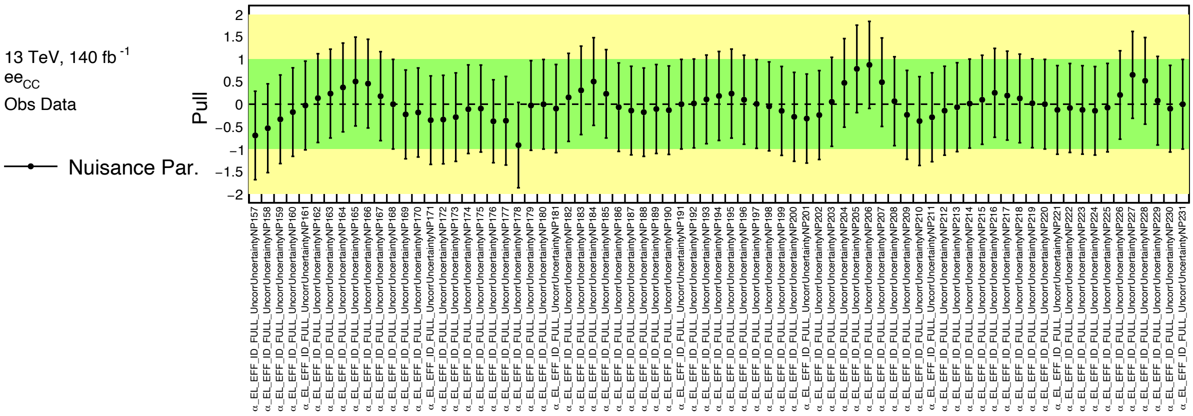

Figure 2: Pull plots for the electron identification efficiency variations.

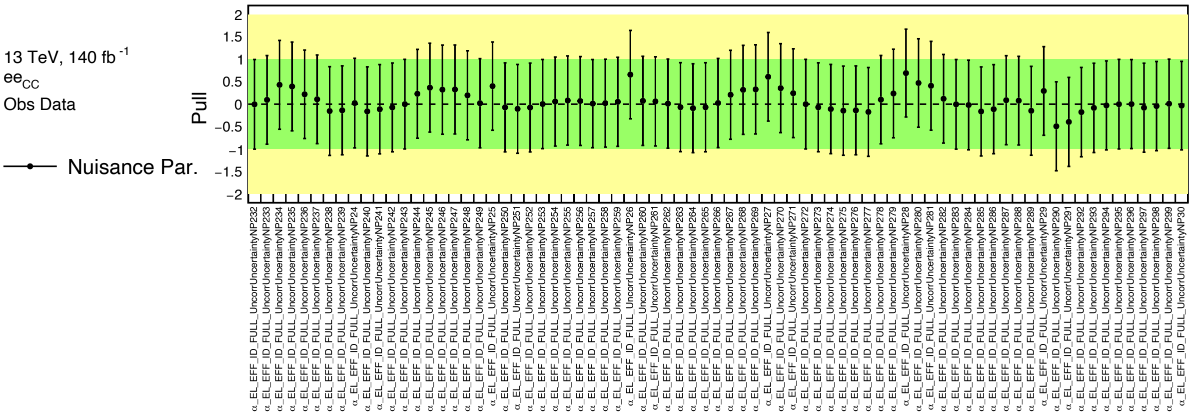

Figure 3: Pull plots for the electron identification efficiency variations.

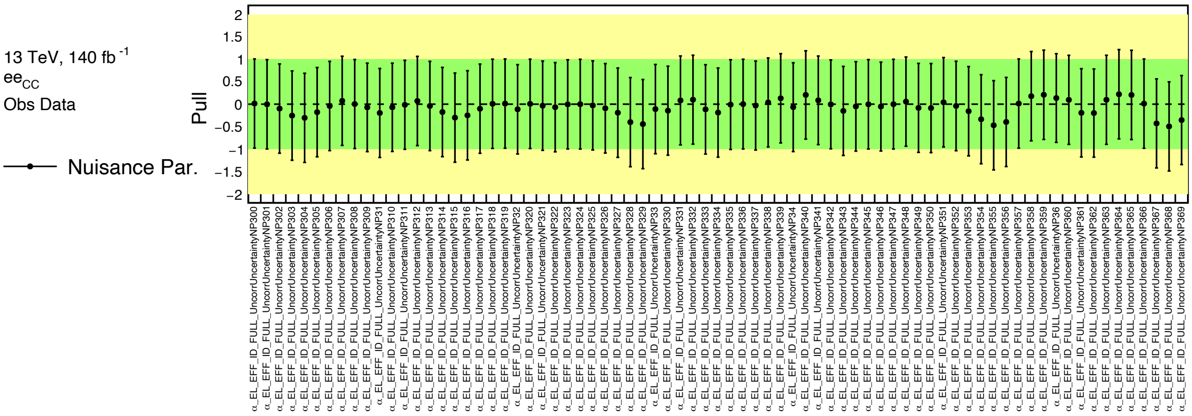

Figure 4: Pull plots for the electron identification efficiency variations.

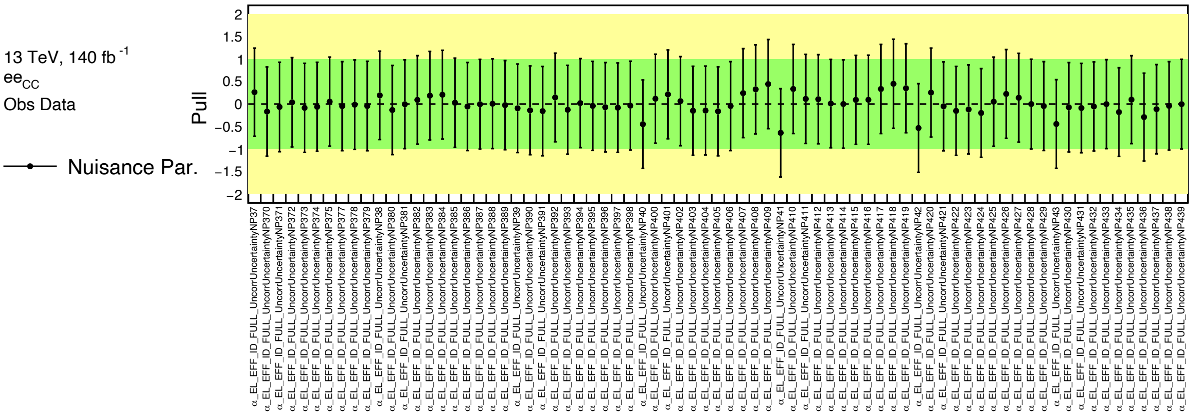

Figure 5: Pull plots for the electron identification efficiency variations.

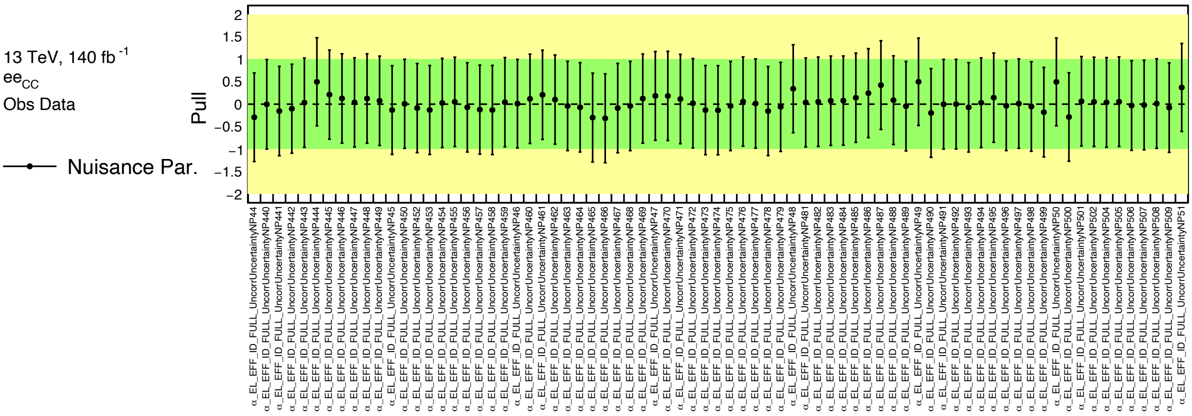

Figure 6: Pull plots for the electron identification efficiency variations.

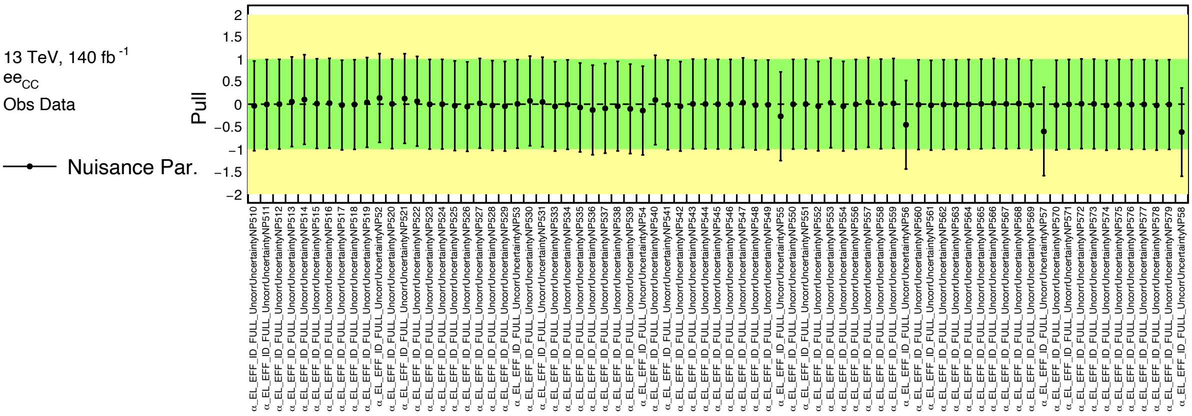

Figure 7: Pull plots for the electron identification efficiency variations.

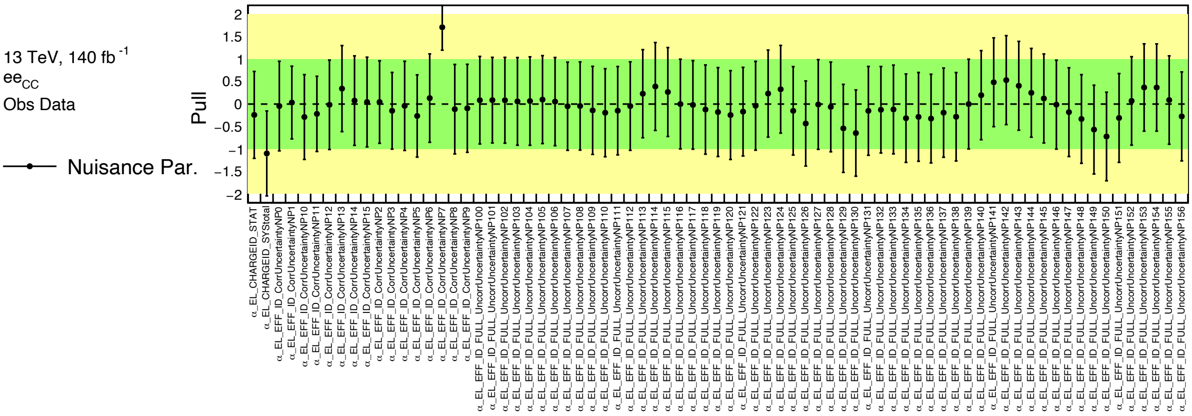

Figure 8: Pull plots for the electron identification efficiency variations.

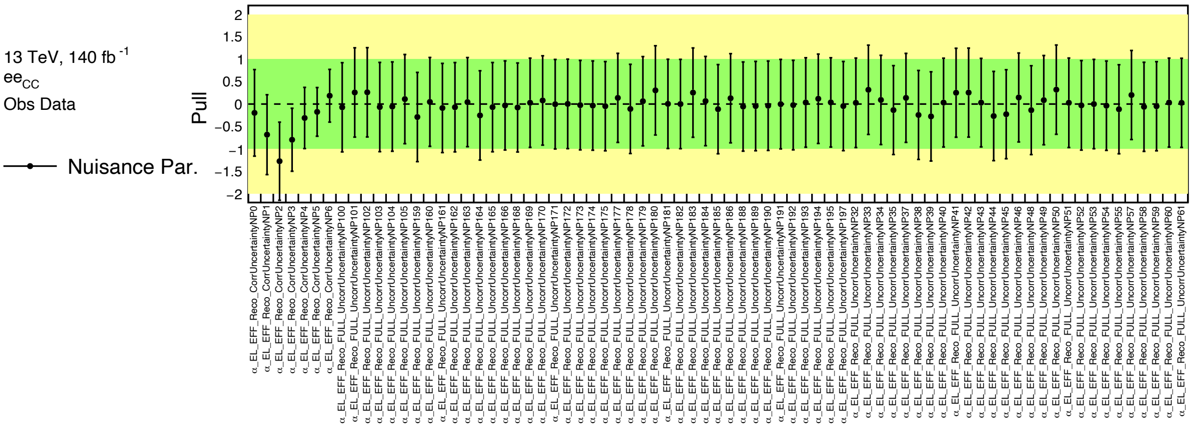

Figure 9: Pull plots for the electron reconstruction efficiency variations.

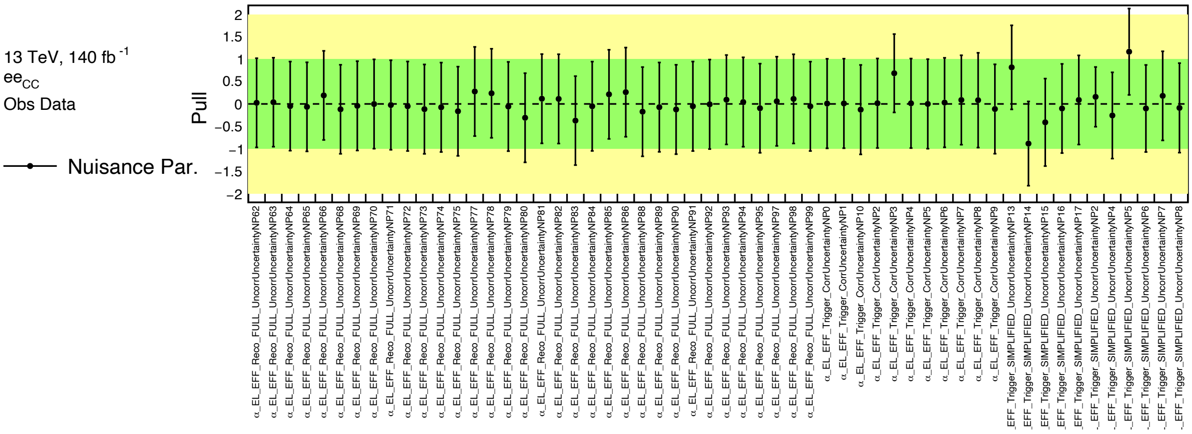

Figure 10: Pull plots for the electron reconstruction and trigger efficiency variations.

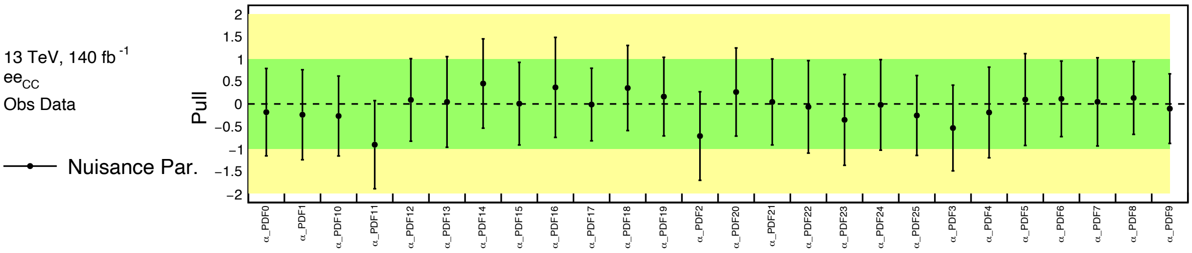

Figure 11: Pull plots for the electron PDF variations.

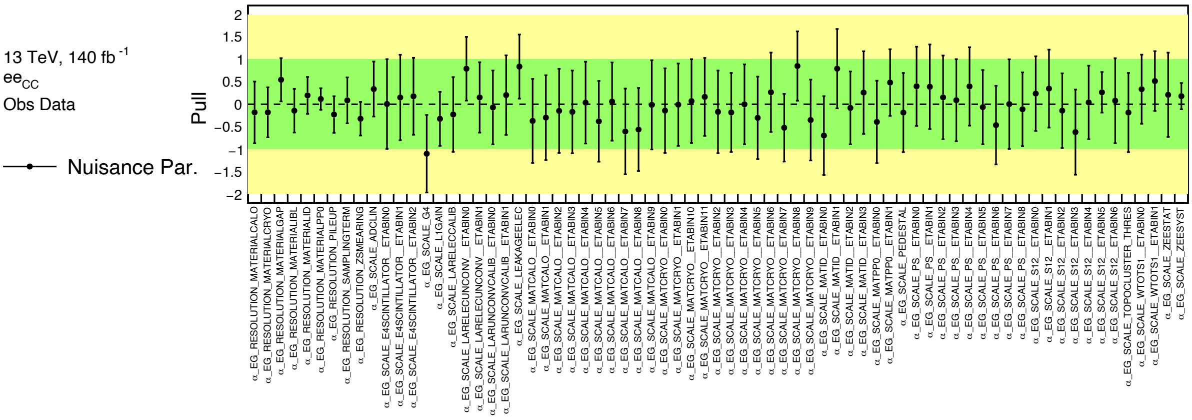

Figure 12: Pull plots for the electron calibration variations.

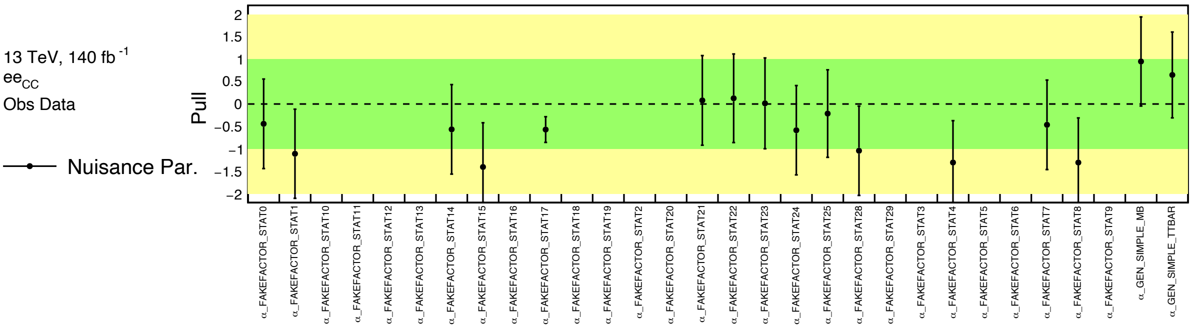

Figure 13: Pull plots for the electron fake factor and background Monte Carlo generator variations.

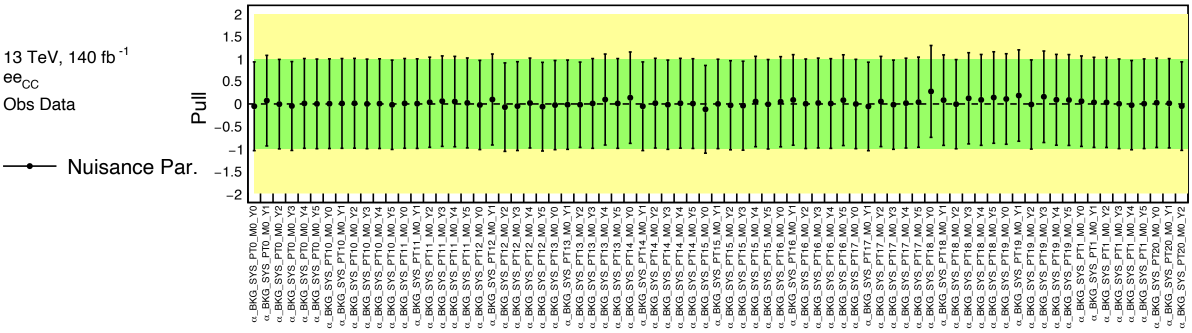

Figure 14: Pull plots for the electron Monte Carlo background variations.

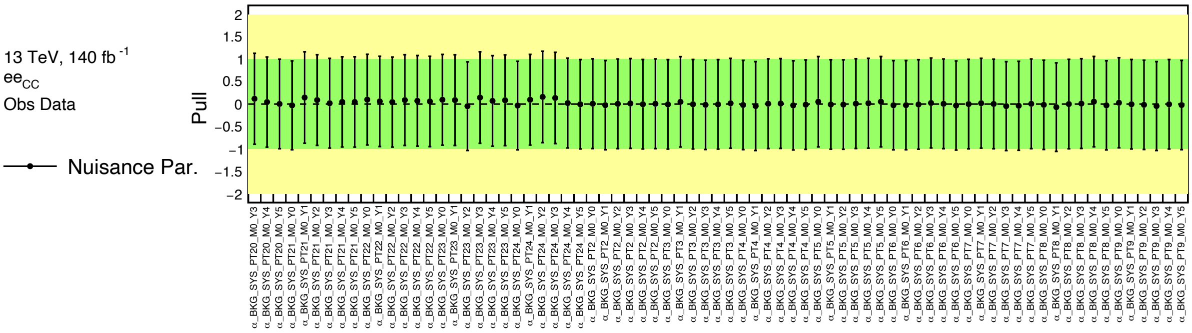

Figure 15: Pull plots for the electron Monte Carlo background variations.

B | Muon Channel Pull Plots

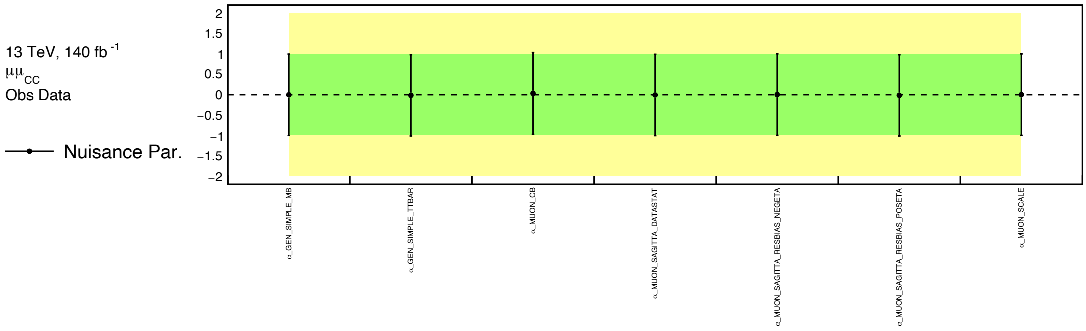

Figure 16: Pull plots for the calibration variations.

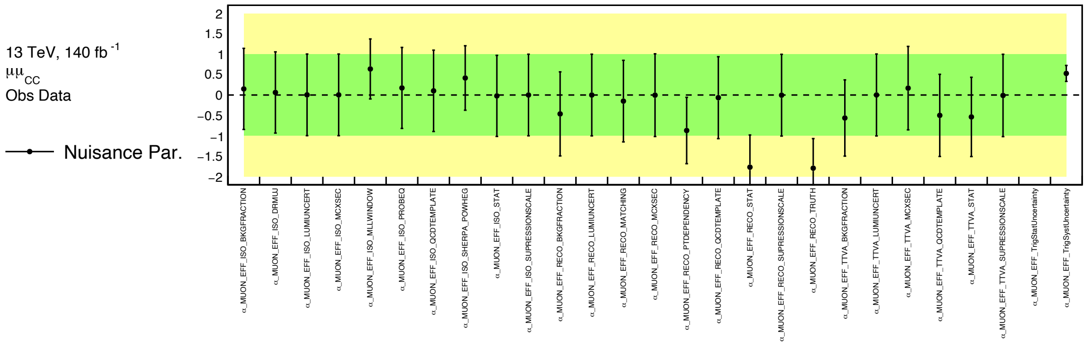

Figure 17: Pull plots for all efficiency variations.

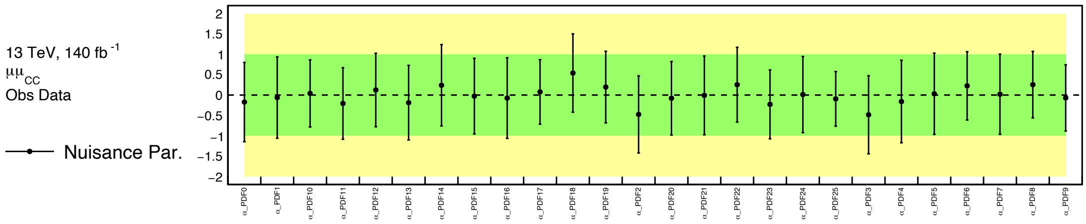

Figure 18: Pull plots for the PDF variations.

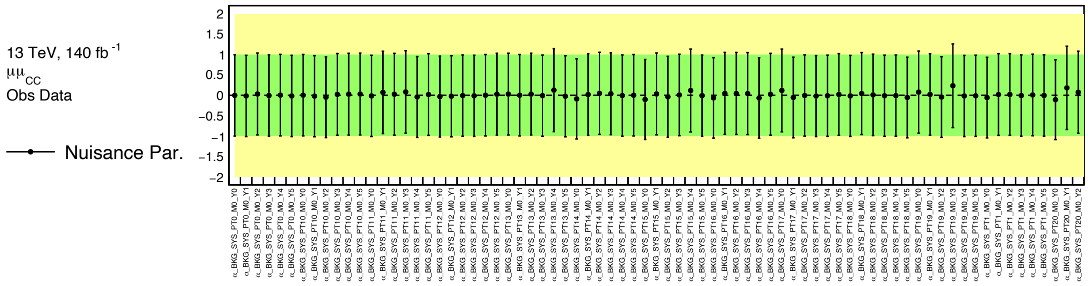

Figure 19: Pull plots for the background variations.

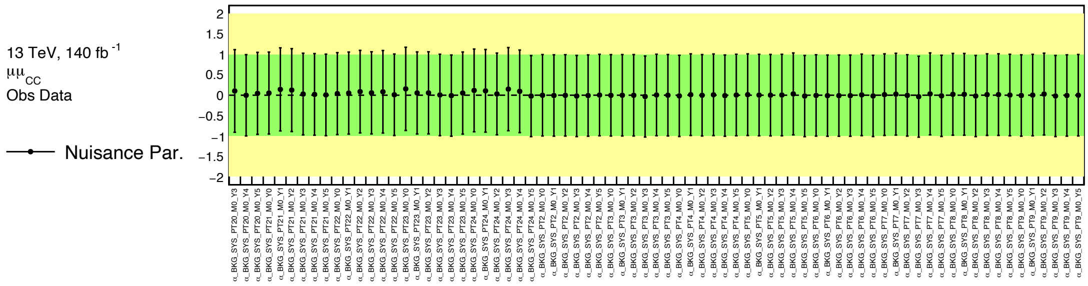

Figure 20: Pull plots for the background variations..

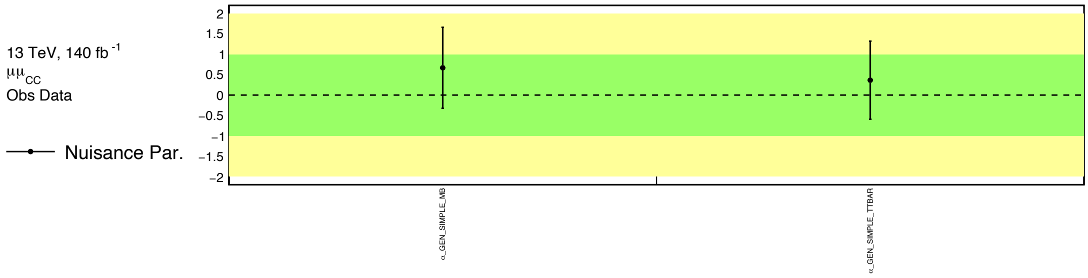

Figure 21: Pull plots for the generator variations.

C | Compatibility of the Electron and Muon Channels

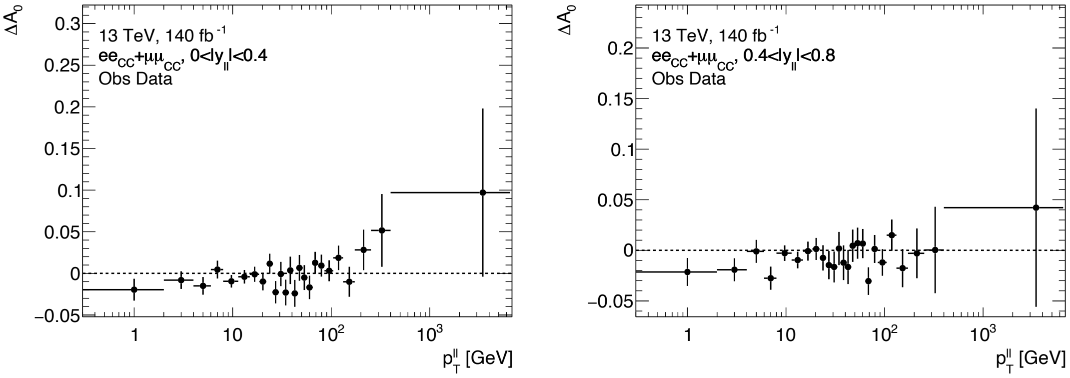

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

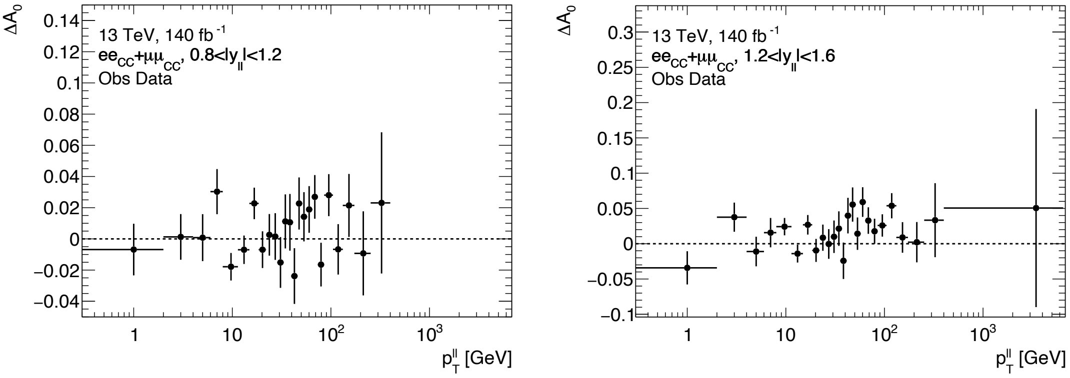

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

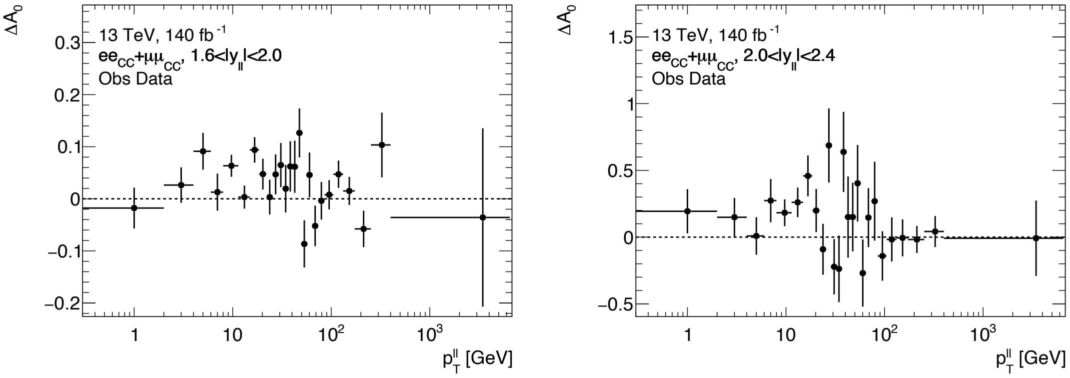

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 22: Fitted difference in A_0 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

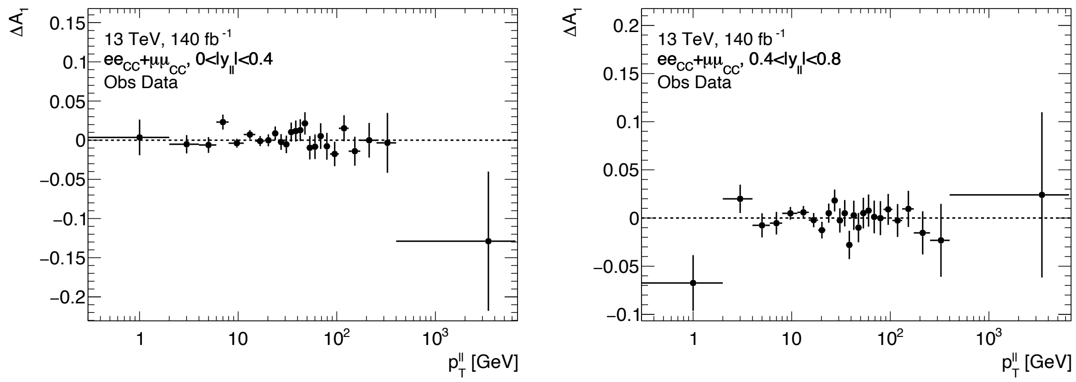

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

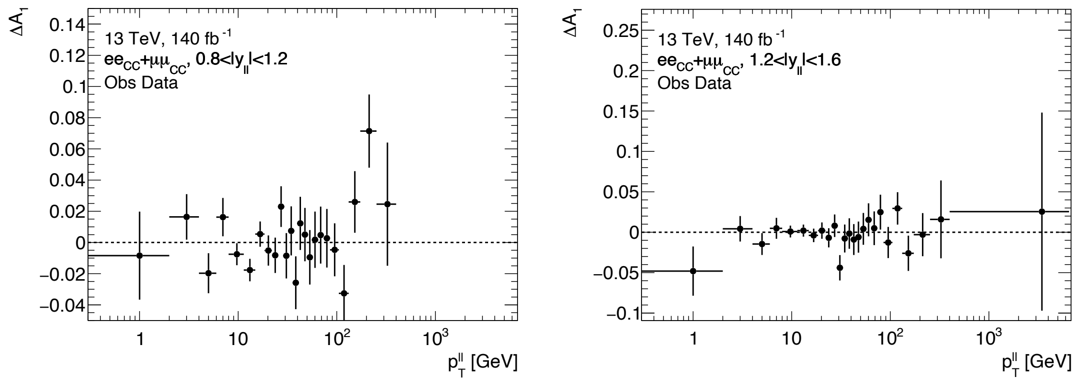

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

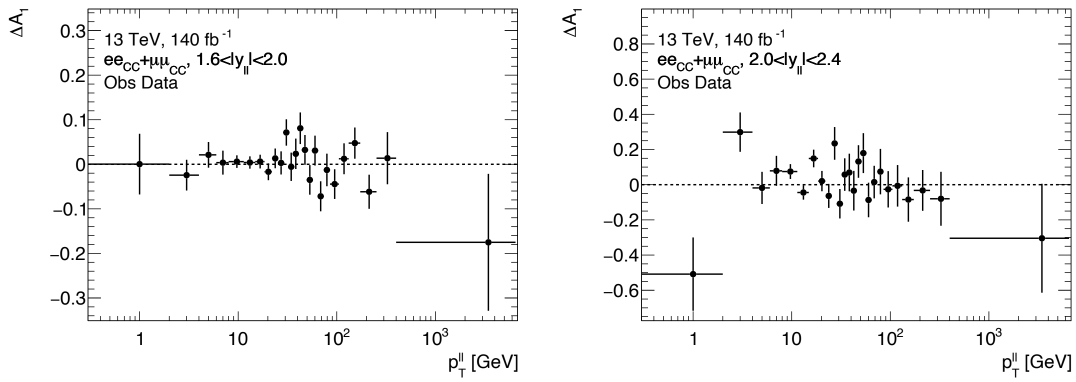

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 23: Fitted difference in A_1 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

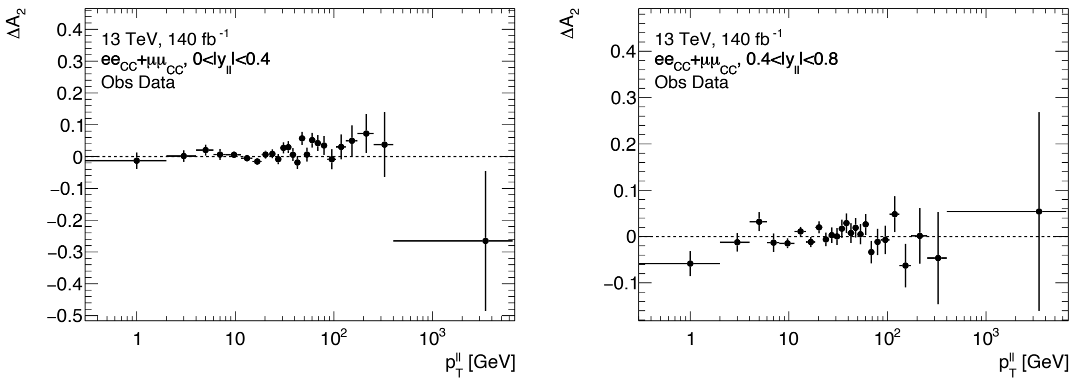

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

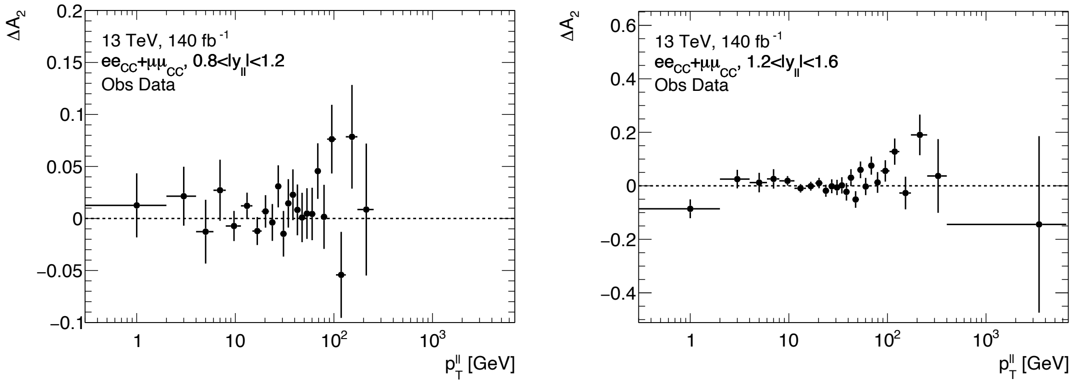

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

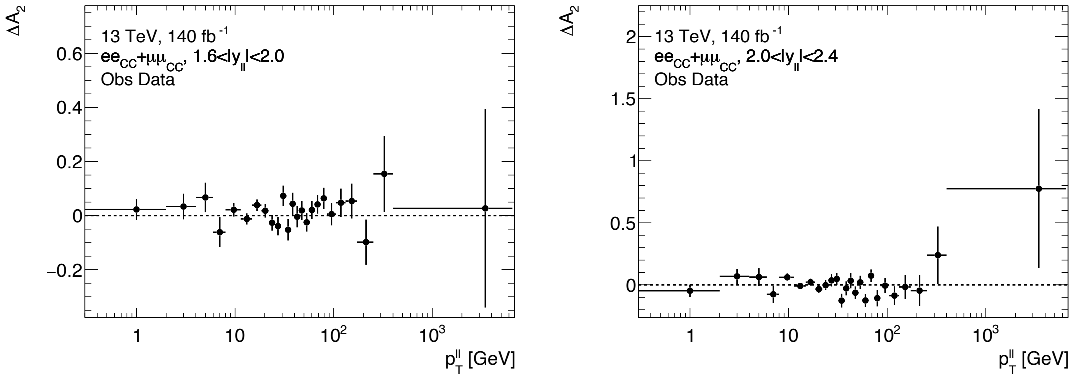

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 24: Fitted difference in A_2 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

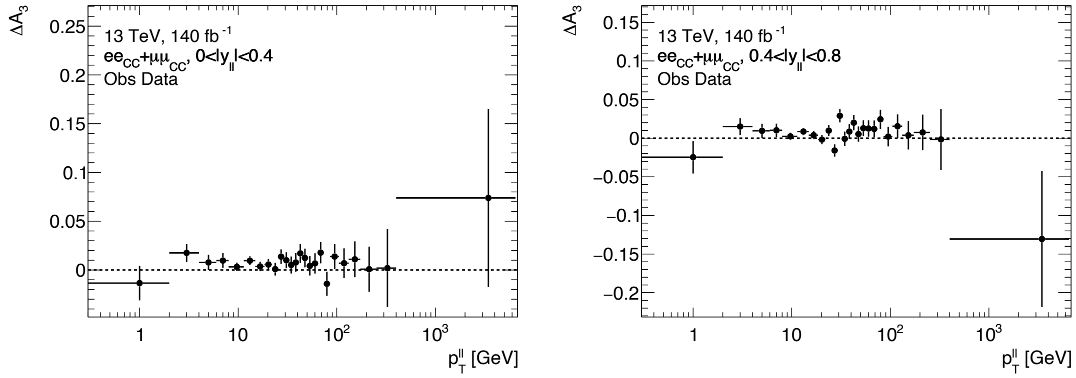

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

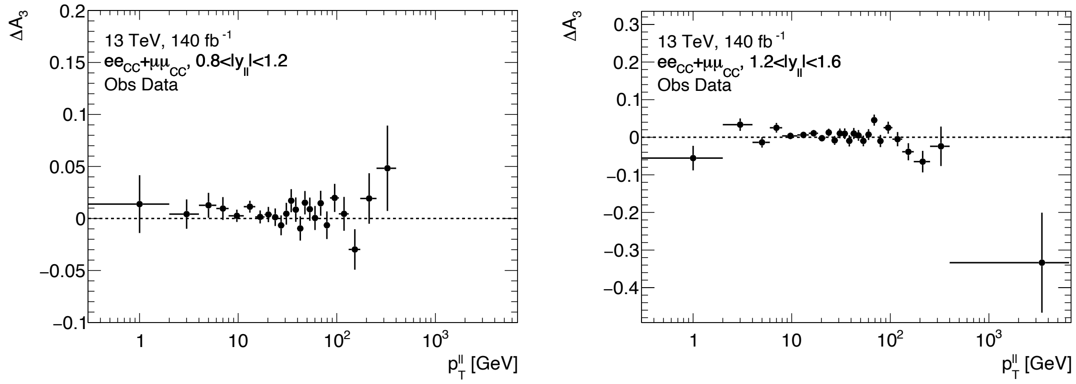

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

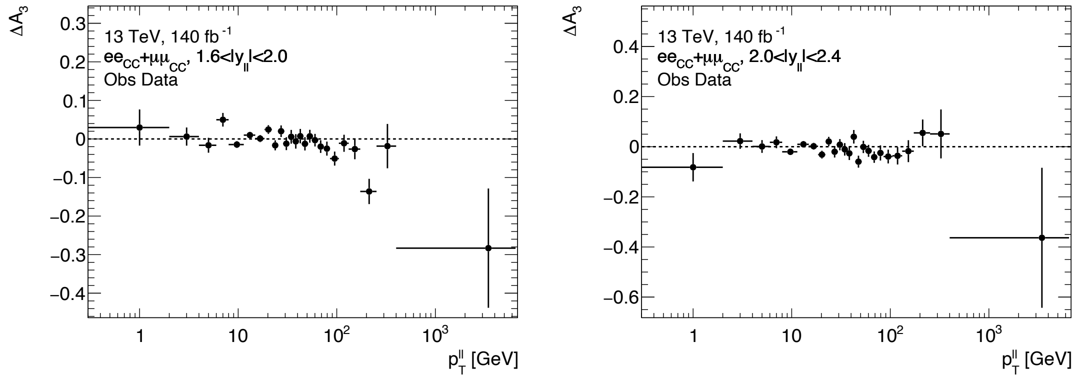

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 25: Fitted difference in A_3 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

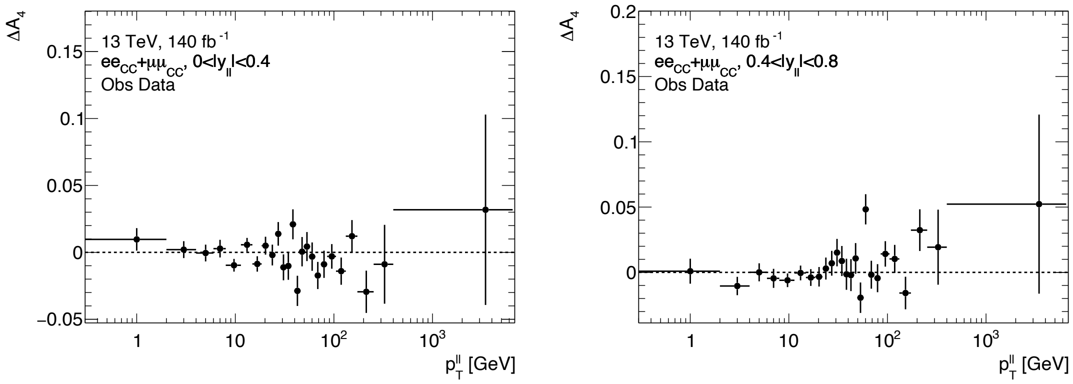

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

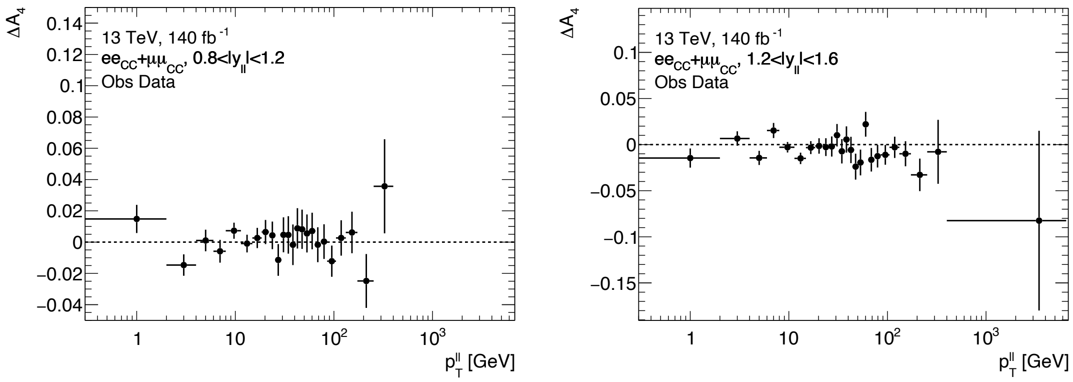

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

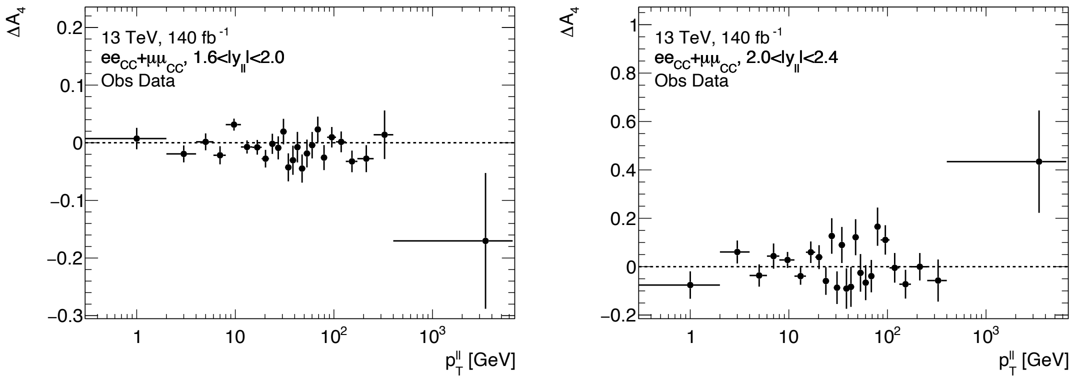

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 26: Fitted difference in A_4 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

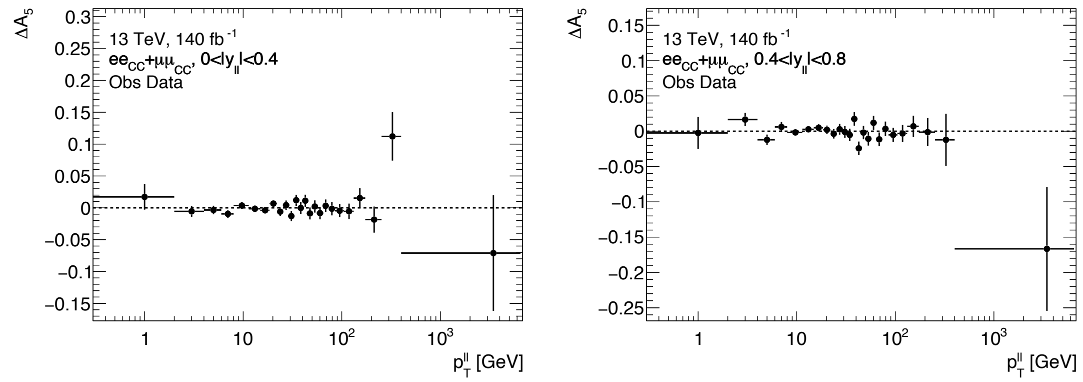

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

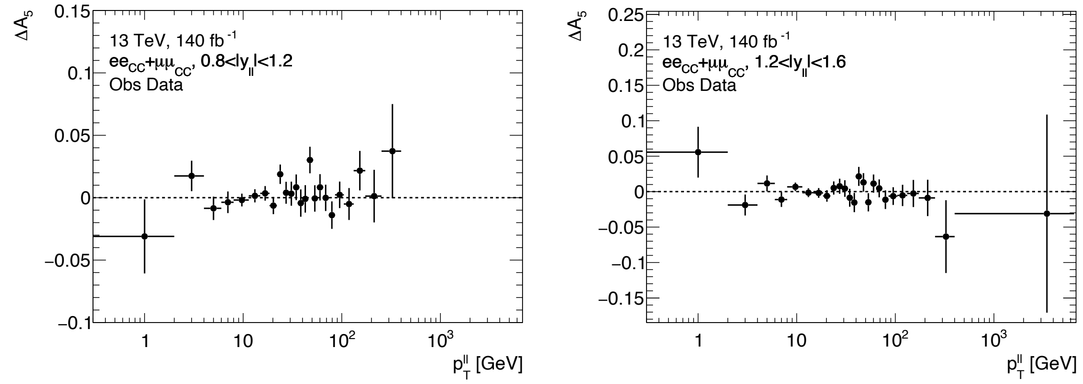

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

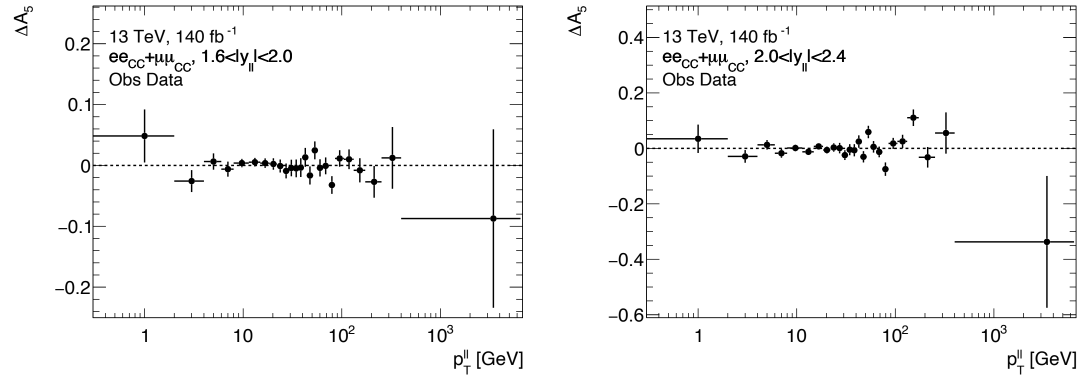

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 27: Fitted difference in A_5 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

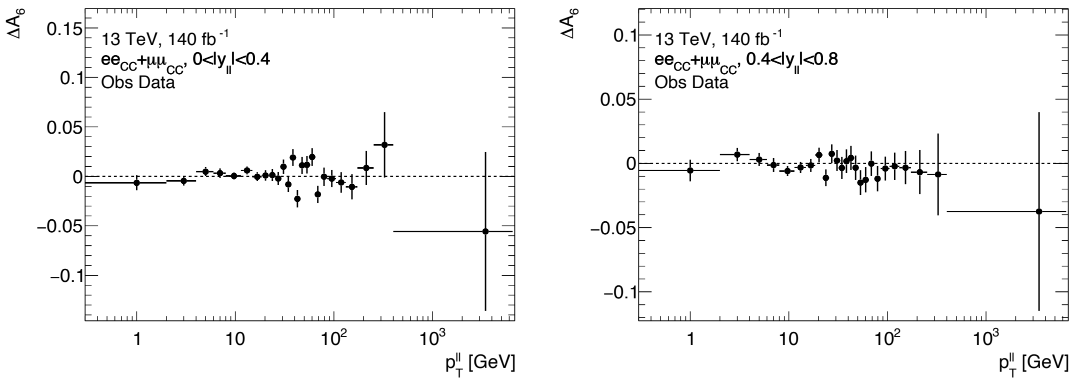

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

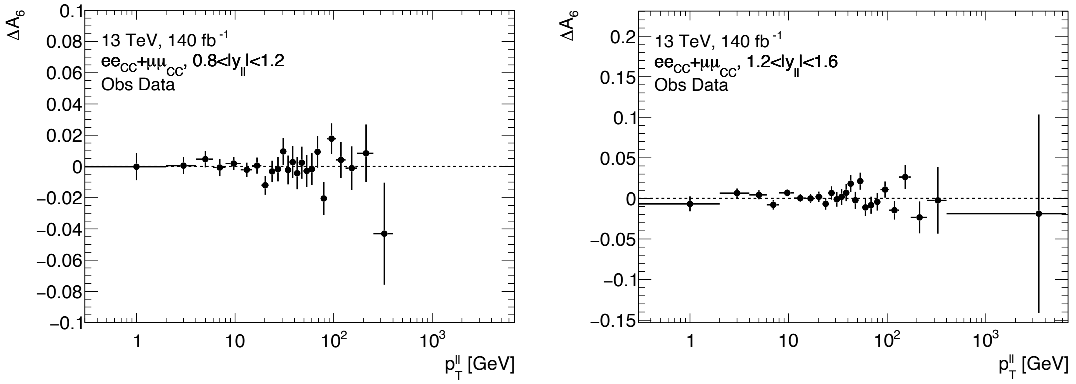

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

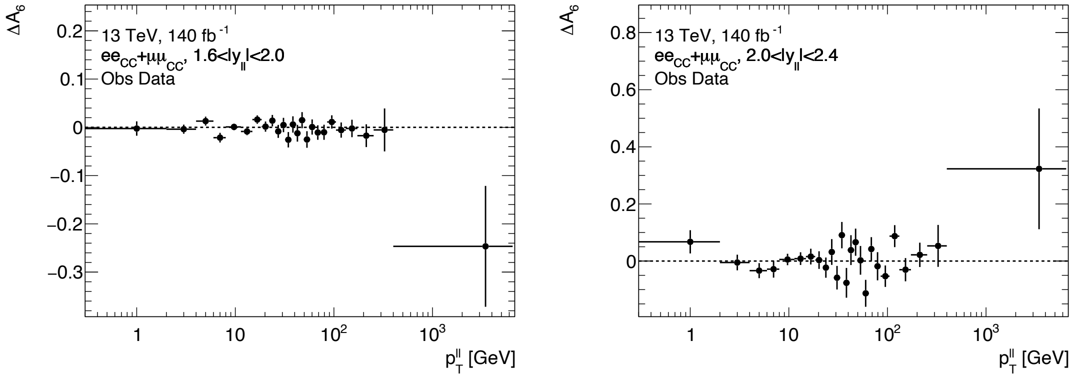

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 28: Fitted difference in A_6 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

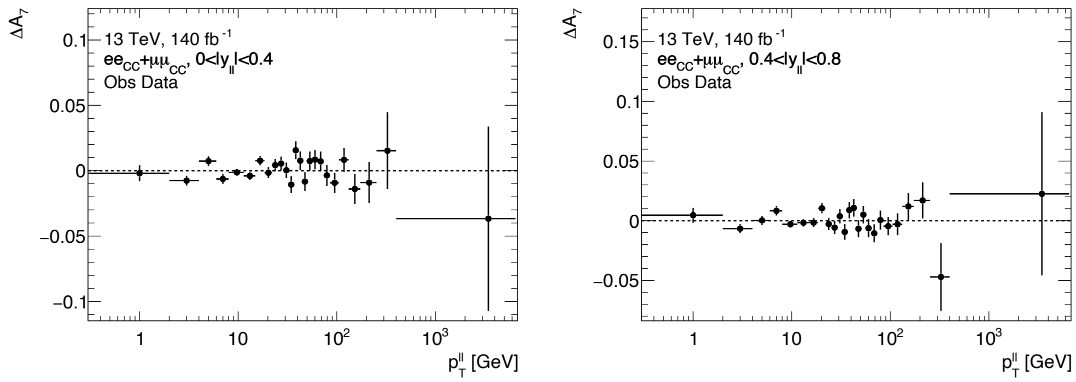

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

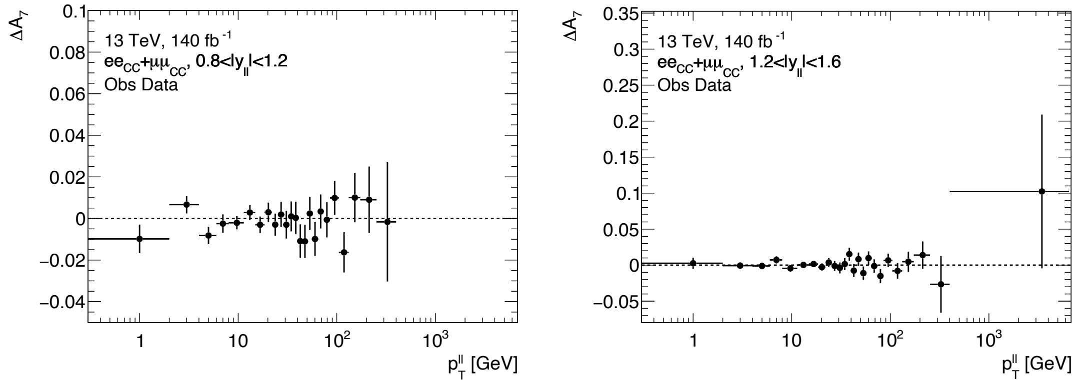

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

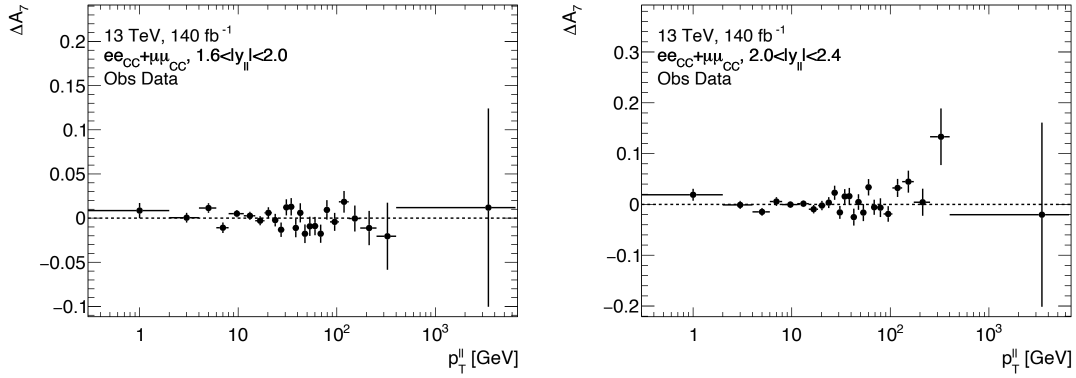

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 29: Fitted difference in A_7 between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

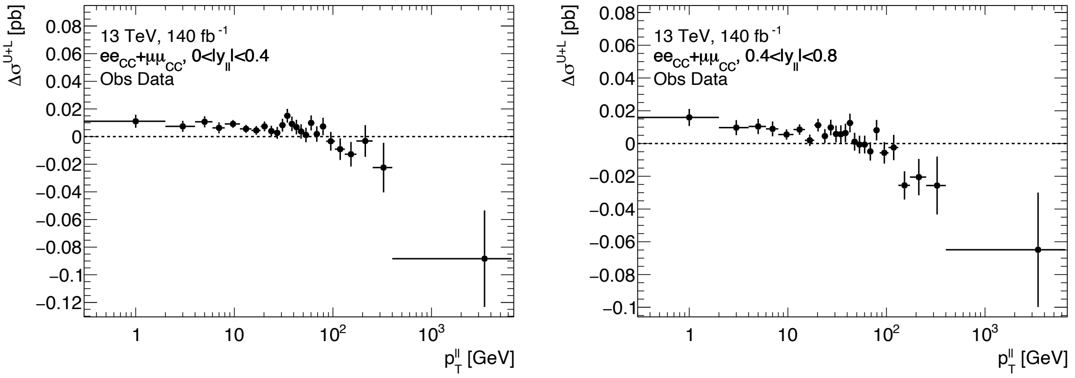

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

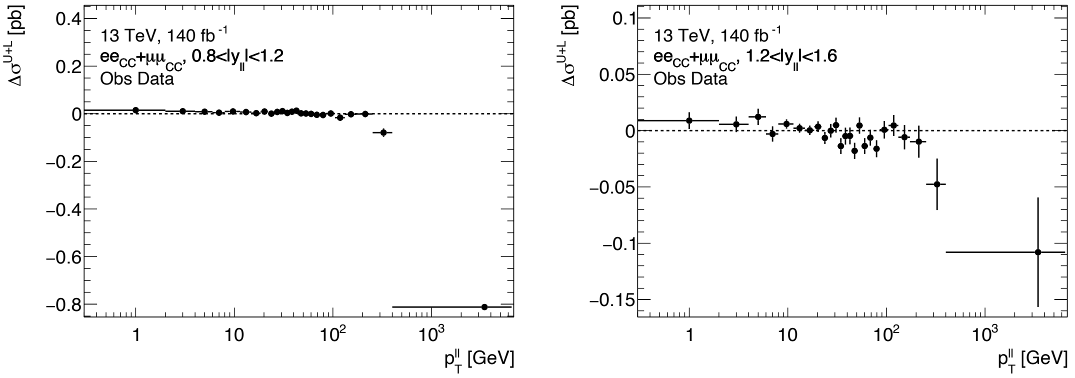

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

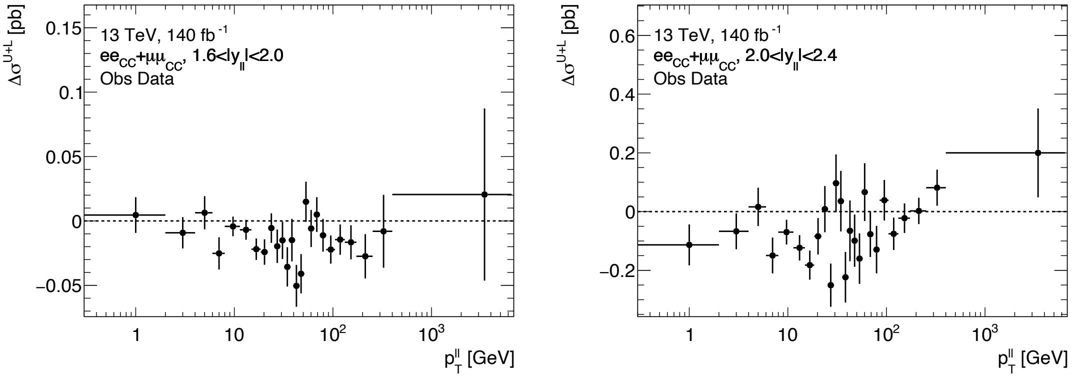

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 30: Fitted difference in σ^{U+L} between the electron and muon channels, as a fraction of the electron channel result and as a function of the dilepton transverse momentum. The error bars correspond to the combination of statistical and systematic effects.

D | Electron and Muon Channel Combination Results

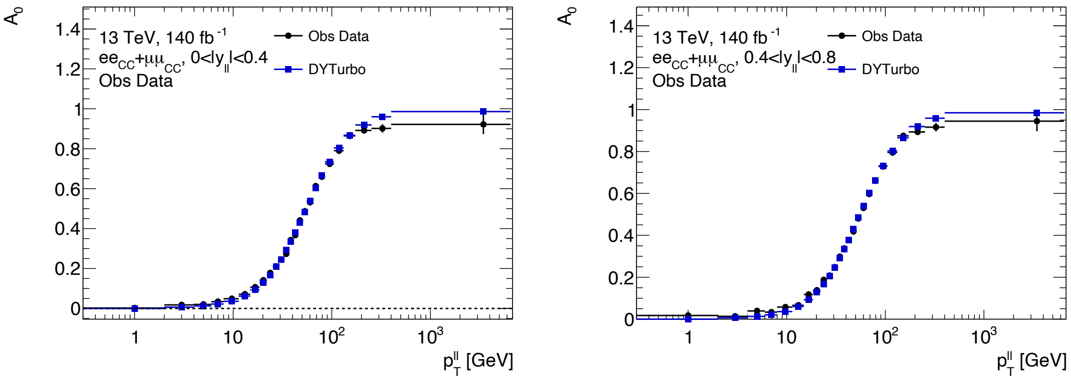

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

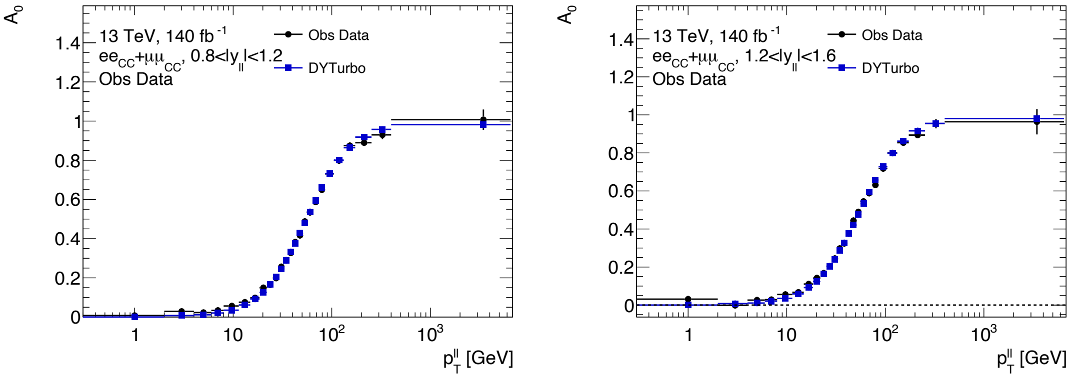

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

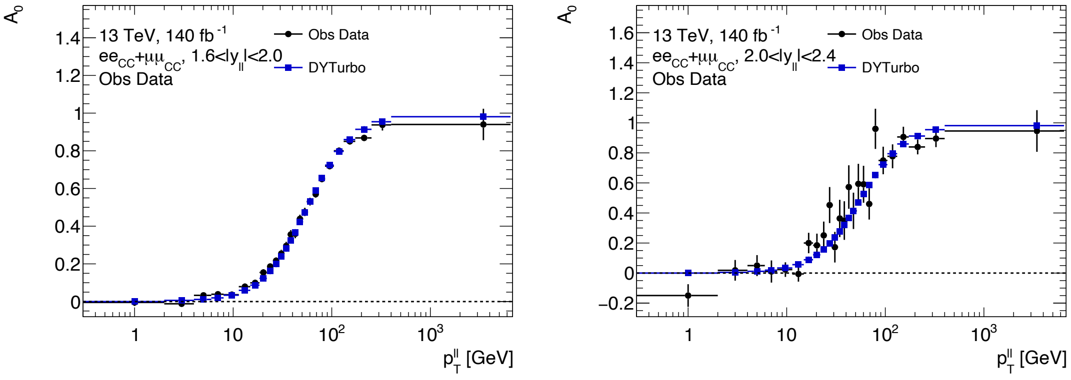

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 31: The observed A_0 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

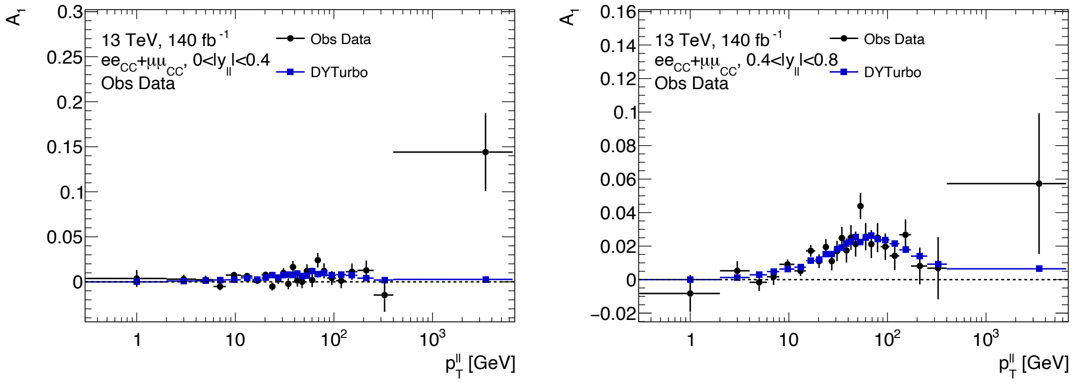

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

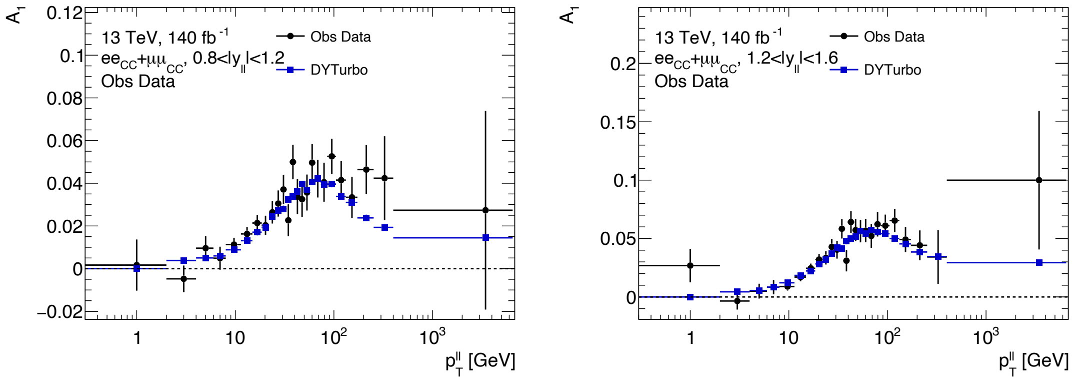

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

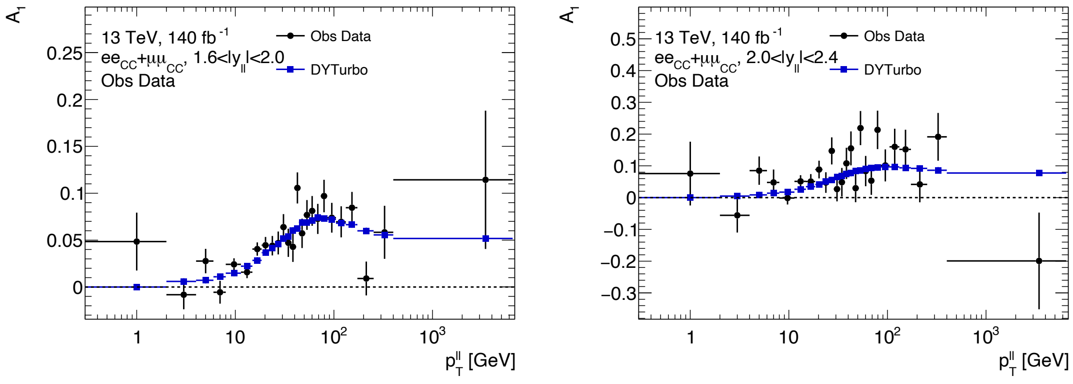

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 32: The observed A_1 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

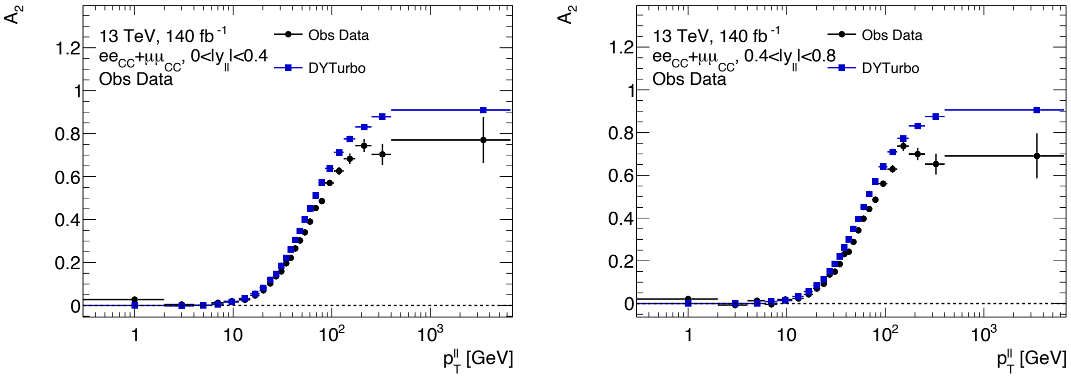

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

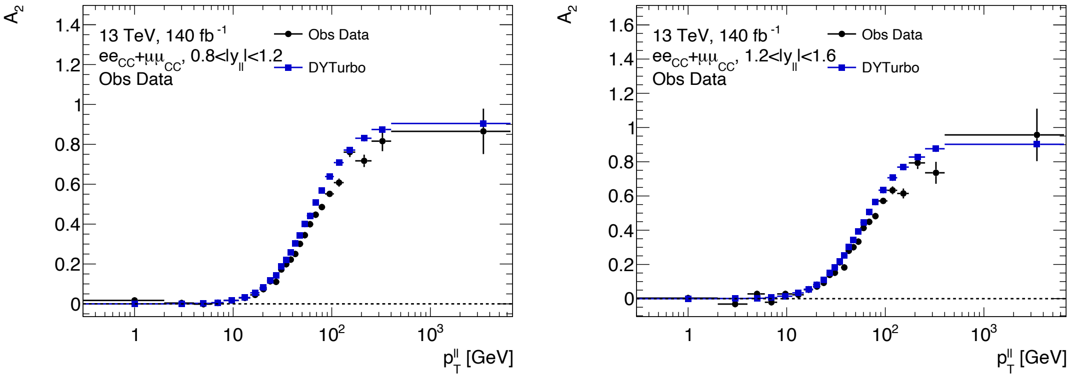

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

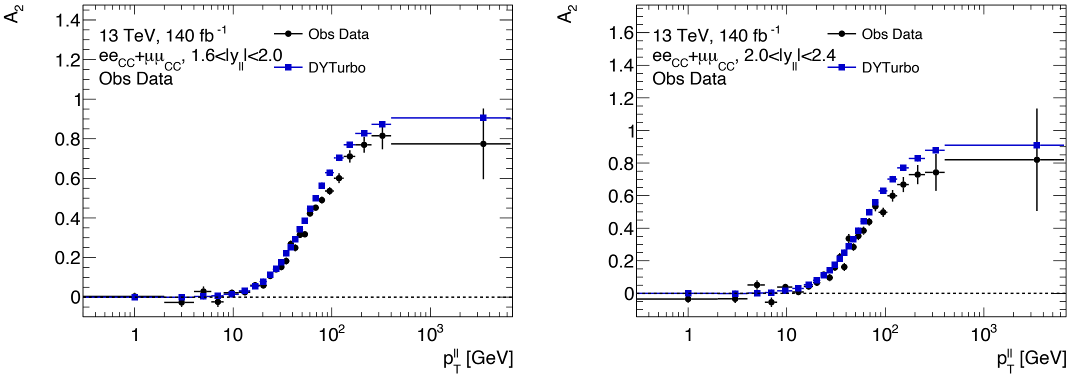

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 33: The observed A_2 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

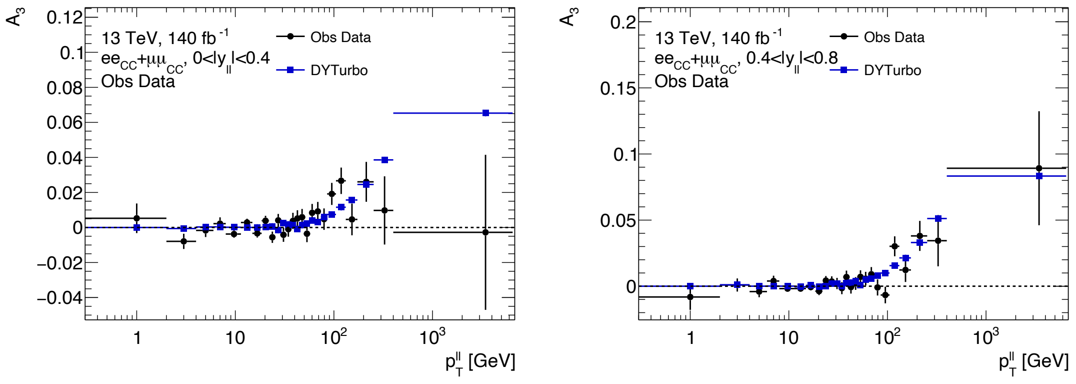

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

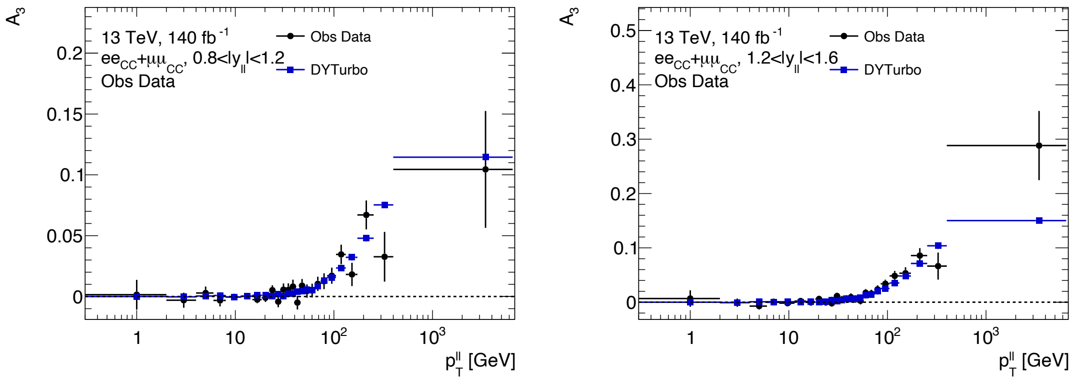

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

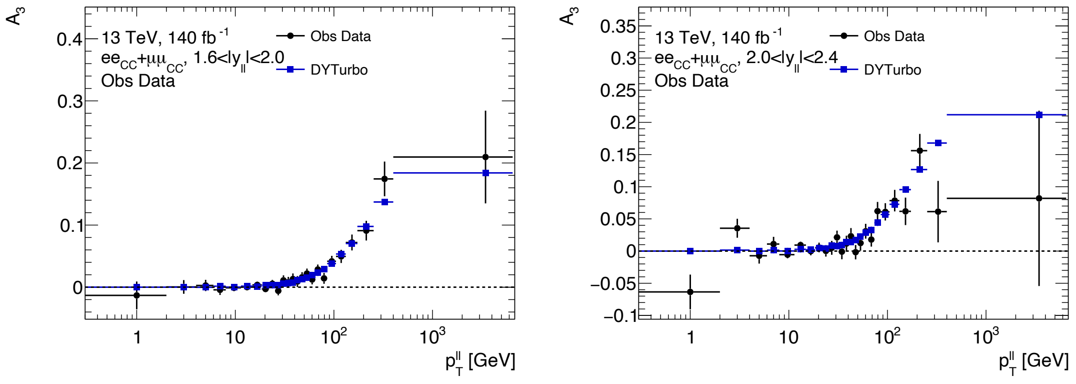

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 34: The observed A_3 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

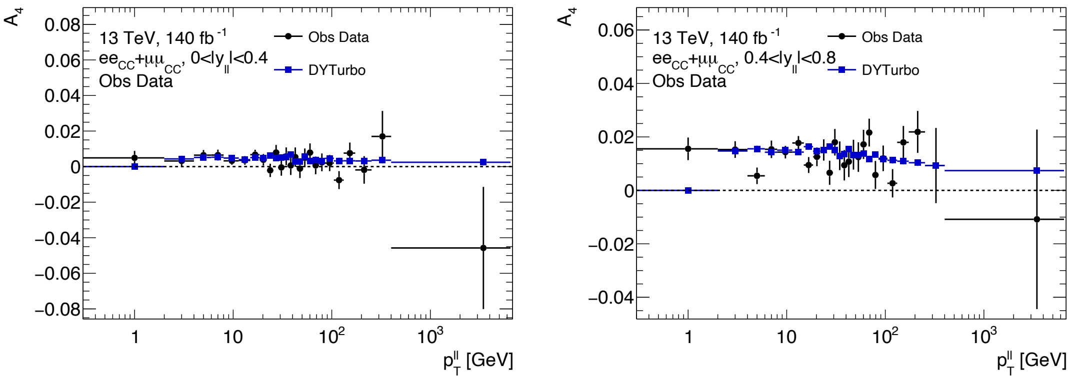

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

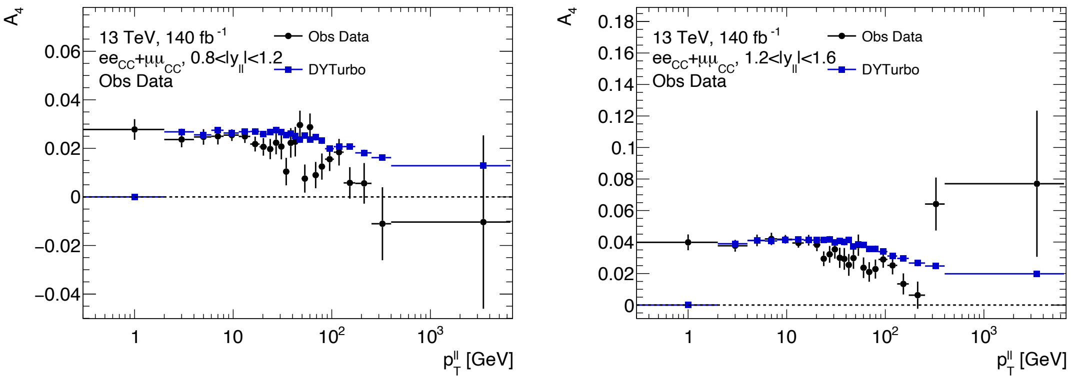

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

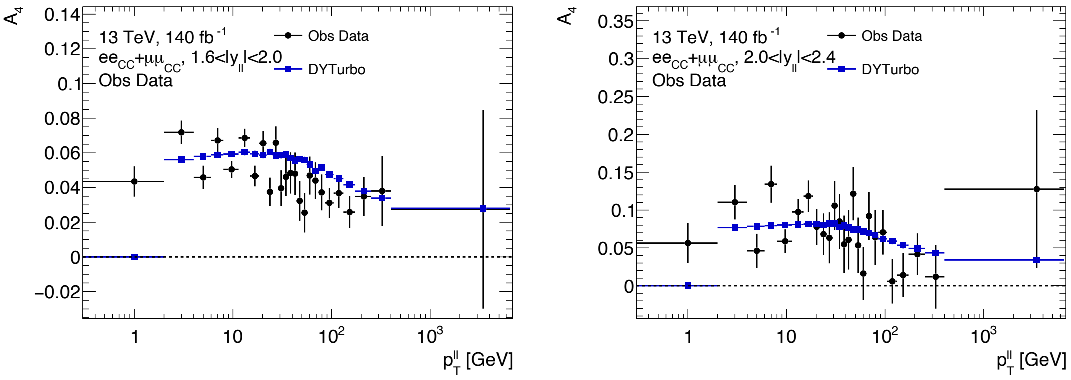

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 35: The observed A_4 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

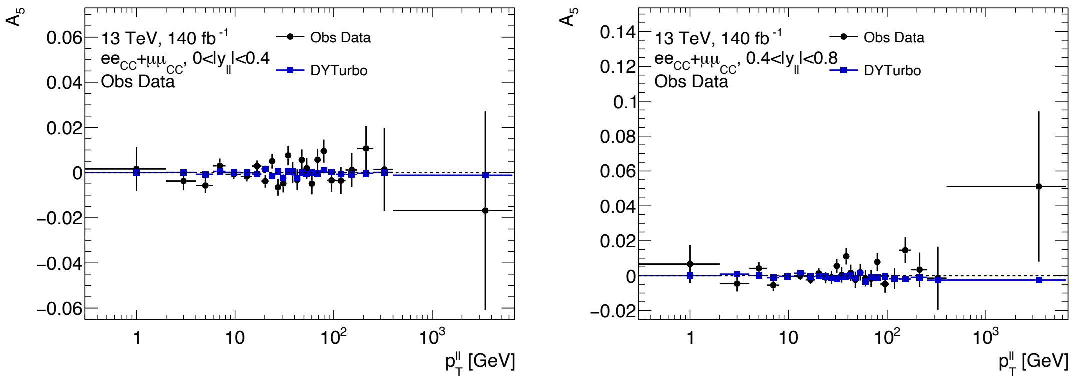

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

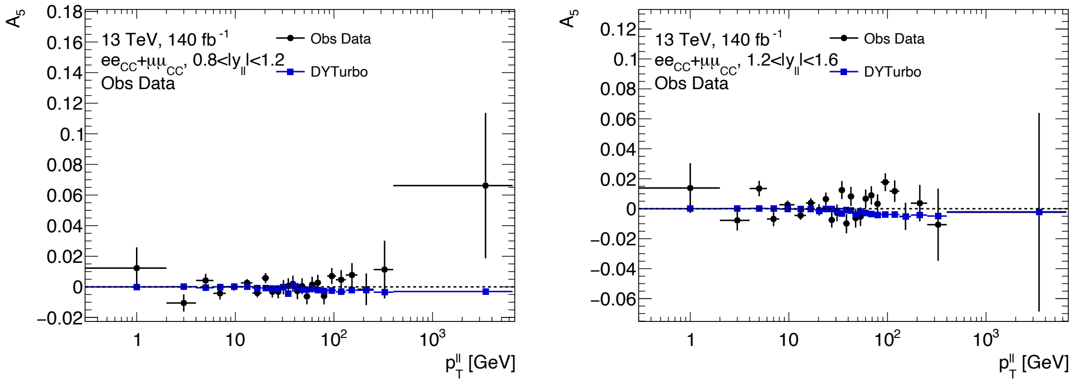

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

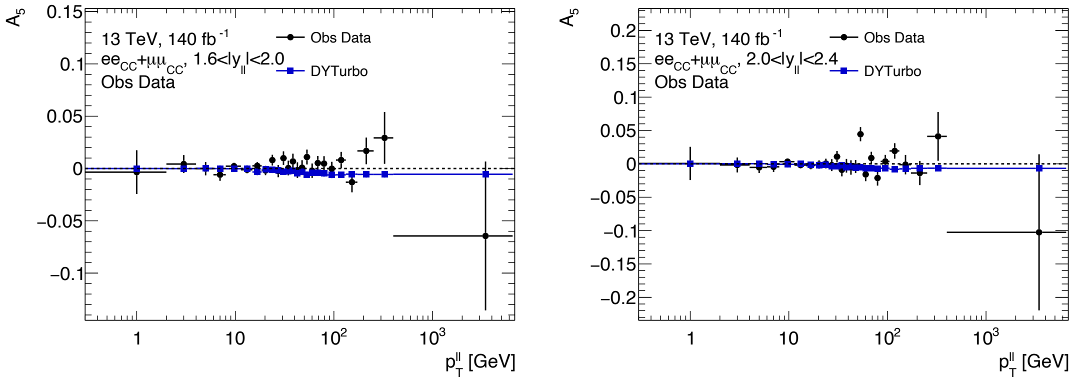

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 36: The observed A_5 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

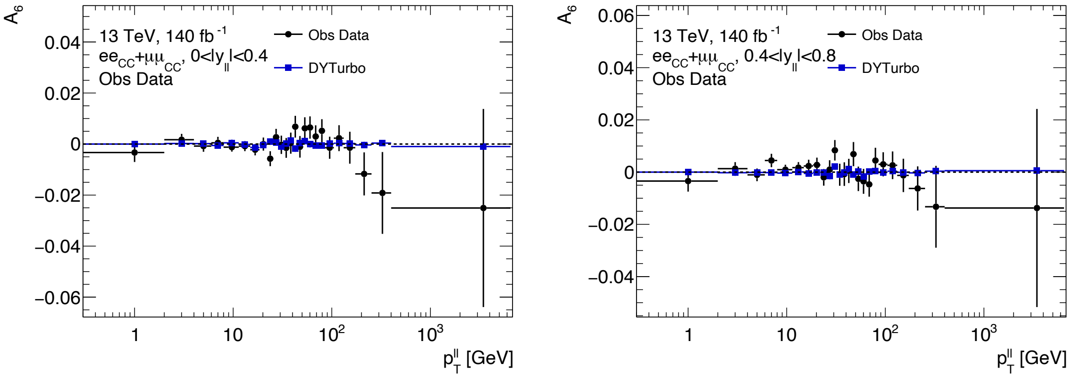

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

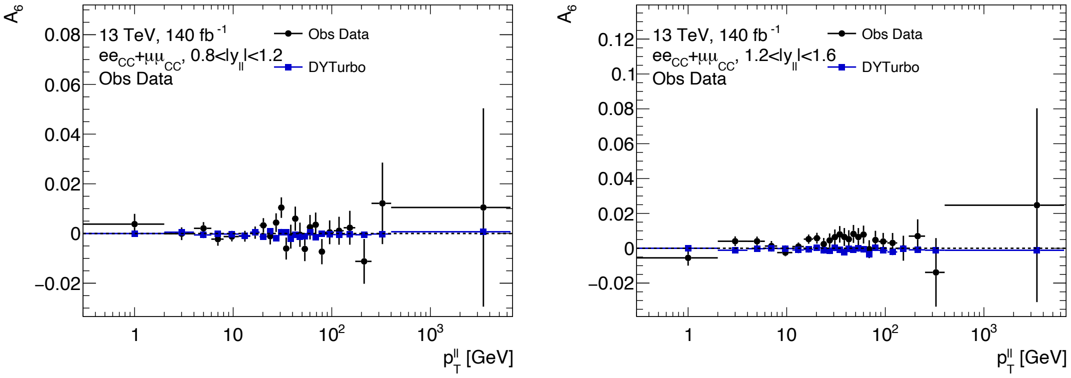

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

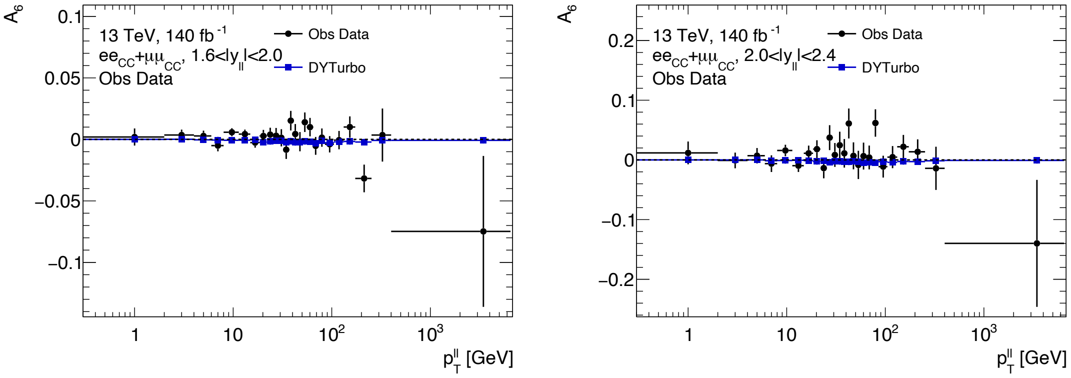

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 37: The observed A_6 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

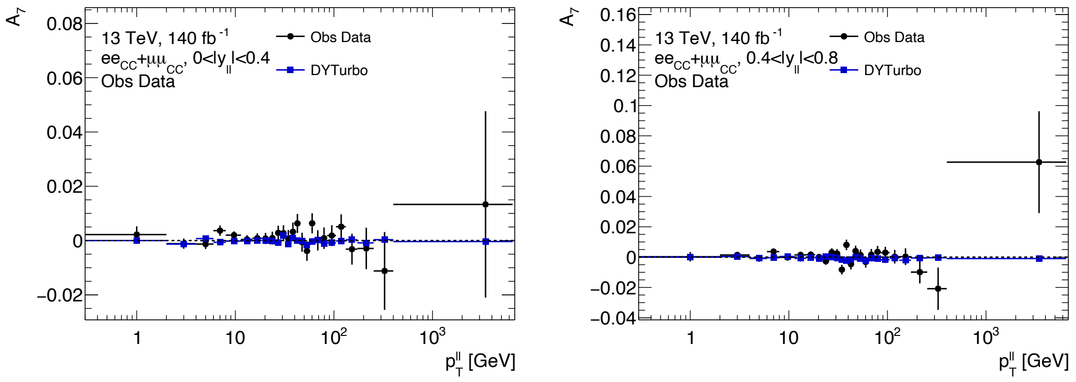

(a) $0 < Y_{ll} < 0.4$

(b) $0.4 < Y_{ll} < 0.8$

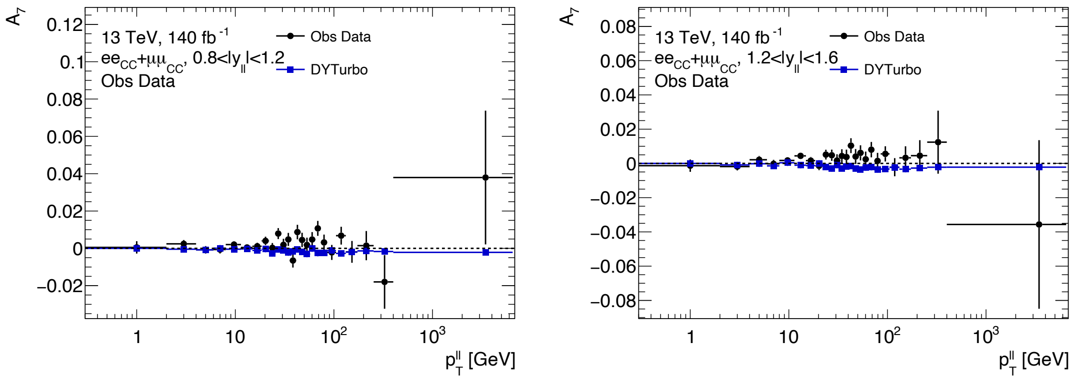

(c) $0.8 < Y_{ll} < 1.2$

(d) $1.2 < Y_{ll} < 1.6$

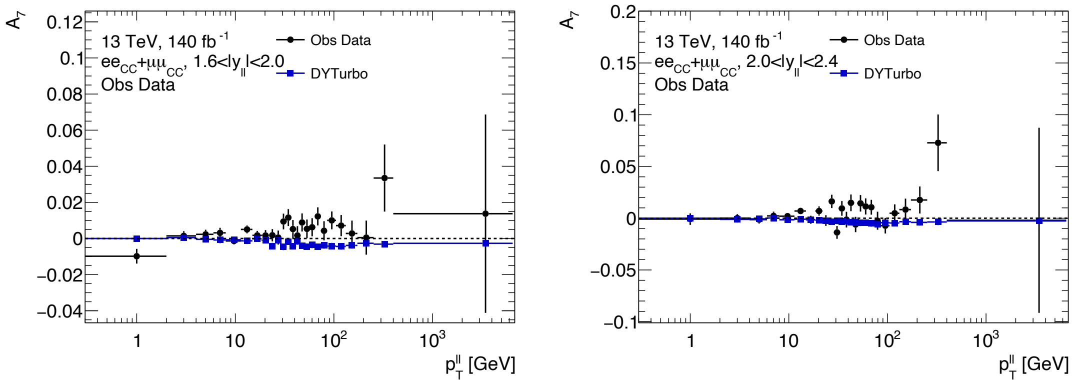

(e) $1.6 < Y_{ll} < 2.0$

(f) $2.0 < Y_{ll} < 2.4$

Figure 38: The observed A_7 values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

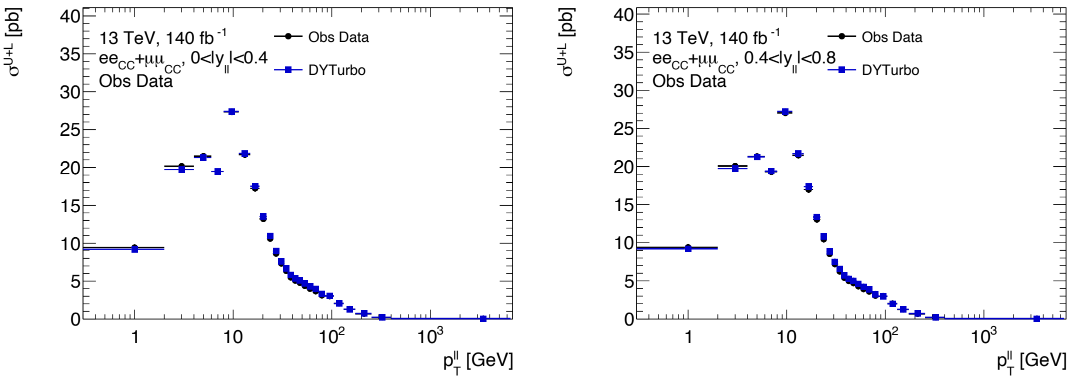

(a) $0 < Y_{ll} < 0.4$ (b) $0.4 < Y_{ll} < 0.8$

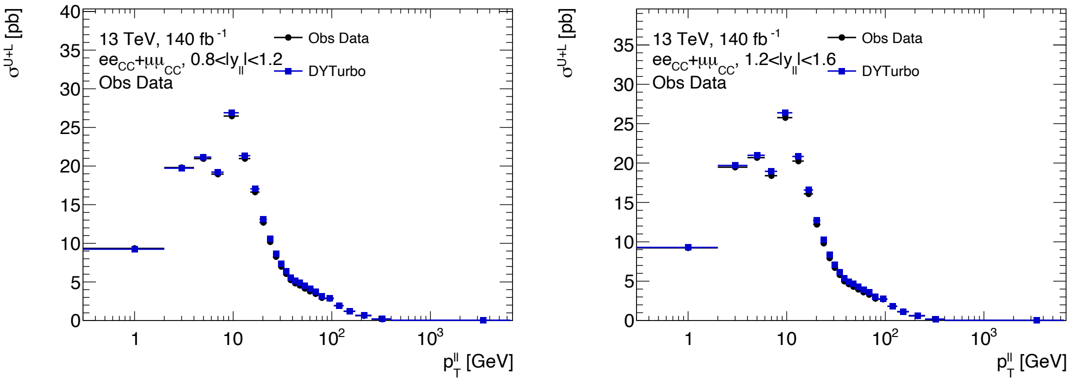

(c) $0.8 < Y_{ll} < 1.2$ (d) $1.2 < Y_{ll} < 1.6$

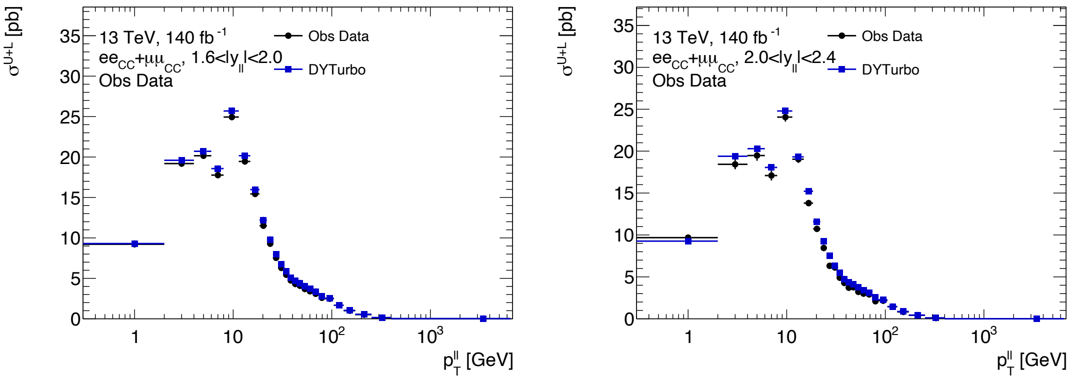

(e) $1.6 < Y_{ll} < 2.0$ (f) $2.0 < Y_{ll} < 2.4$

Figure 39: The observed σ^{U+L} values for the combined electron and muon channels and analytical DYTurbo predictions as a function of the dilepton transverse momentum.

Bibliography

[1] Morad Aaboud et al. Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data. *JINST*, 14(03):P03017, 2019.

[2] Morad Aaboud et al. Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C*, 79(8):639, 2019.

[3] G. Aad et al. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.

[4] G. Aad et al. The ATLAS Simulation Infrastructure. *Eur. Phys. J. C*, 70:823–874, 2010.

[5] G. Aad et al. Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C*, 83(10):982, 2023.

[6] Georges Aad et al. Measurement of the angular coefficients in Z-boson events using electron and muon pairs from data taken at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 08:159, 2016.

[7] Georges Aad et al. Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data. *JINST*, 14(12):P12006, 2019.

[8] Georges Aad et al. Alignment of the ATLAS Inner Detector in Run-2. *Eur. Phys. J. C*, 80(12):1194, 2020.

[9] Georges Aad et al. Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C*, 81(7):578, 2021.

[10] Georges Aad et al. Operation and performance of the ATLAS semiconductor tracker in LHC Run 2. *JINST*, 17(01):P01013, 2022.

[11] Georges Aad et al. Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C*, 83(8):686, 2023.

[12] Georges Aad et al. Tools for estimating fake/non-prompt lepton backgrounds with the ATLAS detector at the LHC. *JINST*, 18(11):T11004, 2023.

[13] Georges Aad et al. Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment. *JHEP*, 05:162, 2024.

[14] Georges Aad et al. Electron and photon energy calibration with the ATLAS detector using LHC Run 2 data. *JINST*, 19(02):P02009, 2024.

[15] R. Aaij et al. First Measurement of the $Z \to \mu^+\mu^-$ Angular Coefficients in the Forward Region of pp Collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. Lett.*, 129(9):091801, 2022.

[16] H. Abdolmaleki et al. xFitter: An Open Source QCD Analysis Framework. A resource and reference document for the Snowmass study. Technical report, 2022. 20 pages, 12 figures. Contribution to the DPF Snowmass Community Planning Exercise, `https://cds.cern.ch/record/2847462`.

[17] ALEPH, DELPHI, L3, OPAL, SLD Collaborations, LEP Electroweak Working Group, SLD Electroweak Group, and SLD Heavy Flavour Group. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006.

[18] Simone Alioli, Paolo Nason, Carlo Oleari, and Emanuele Re. A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX. *JHEP*, 06:043, 2010.

[19] G Altarelli and G Parisi. Asymptotic freedom in parton language. *Nucl. Phys. B.*, 126:298–318, 1977.

[20] Andrej B. Arbuzov. Quantum Field Theory and the Electroweak Standard Model, 2018. arXiv:1801.05670.

[21] Aaron Armbruster. Faster Than Light (FTL) Fitting Framework, 2020. Non-public coding framework.

[22] Aaron Armbruster. Forward electron performance for 2012 ATLAS data. Technical report, CERN, 2021. `https://cds.cern.ch/record/2771945`.

[23] Richard D. Ball et al. Parton distributions with LHC data. *Nucl. Phys. B*, 867(2):244–289, 2013.

[24] Richard D. Ball et al. Parton distributions with QED corrections. *Nucl. Phys. B*, 877:290–320, 2013.

[25] Richard D. Ball et al. The path to proton structure at 1% accuracy. *Eur. Phys. J. C*, 82(5):428, 2022.

[26] S. P. Baranov, O. Duenger, H. Shooshtari, and J. A. M. Vermaseren. LPAIR: A generator for lepton pair production. In *Workshop on Physics at HERA*, 1991.

[27] Johannes Bellm et al. Herwig 7.0/Herwig++ 3.0 release note. *Eur. Phys. J. C*, 76(4):196, 2016.

[28] A. Bodek, J. Han, A. Khukhunaishvili, and W. Sakumoto. Using Drell–Yan forward–backward asymmetry to reduce PDF uncertainties in the measurement of electroweak parameters. *Eur. Phys. J. C*, 76(3):115, 2016.

[29] D. Boussard and T. Linnecar. The LHC Superconducting RF System. *Adv. Cryog. Eng. A*, 45:835–844, 2000.

[30] Andy Buckley, James Ferrando, Stephen Lloyd, Karl Nordström, Ben Page, Martin Rüfenacht, Marek Schönherr, and Graeme Watt. LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J. C*, 75:132, 2015.

[31] Manuel Bähr and others. Herwig++ physics and manual. *Eur. Phys. J. C*, 58(4):639–707, November 2008.

[32] Stefano Camarda, Leandro Cieri, and Giancarlo Ferrera. Drell–Yan lepton-pair production: qT resummation at N4LL accuracy. *Phys. Lett. B*, 845:138125, 2023.

[33] J.M. Campbell et al. Event Generators for High-Energy Physics Experiments. 2024. 164 pages, 10 figures, contribution to Snowmass 2021, `https://cds.cern.ch/record/2807002`.

[34] M Capeans et al. ATLAS Insertable B-Layer Technical Design Report. Technical report, CERN, 2010. `https://cds.cern.ch/record/1291633`.

[35] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π Decay of the K_2^0 Meson. *Phys. Rev. Lett.*, 13:138–140, Juyl 1964.

[36] ATLAS collaboration. Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung. Technical report, CERN, 2012. All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2012-047`.

[37] ATLAS collaboration. Example ATLAS tunes of Pythia8, Pythia6 and Powheg to an observable sensitive to Z boson transverse momentum. Technical report, CERN, 2013. ”All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2013-017`”.

[38] ATLAS collaboration. Study of the material of the ATLAS inner detector for Run 2 of the LHC. *JINST*, 12(12):P12009–P12009, 2017.

[39] ATLAS collaboration. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *Eur. Phys. J. C*, 77(7):490, 2017.

[40] ATLAS collaboration. Measurement of the effective leptonic weak mixing angle using electron and muon pairs from Z-boson decay in the ATLAS experiment at $\sqrt{s} = 8$ TeV. Technical report, CERN, 2018. All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2018-037`.

[41] ATLAS collaboration. Measurement of the effective leptonic weak mixing angle using electron and muon pairs from Z-boson decay in the ATLAS experiment at $\sqrt{s} = 8$ TeV. Technical report, CERN, 2018. All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2018-037`.

[42] ATLAS collaboration. Operation of the ATLAS trigger system in Run 2. *JINST*, 15(10):P10004, 2020.

[43] ATLAS collaboration. Performance of the muon spectrometer alignment in 2017 and 2018 data. Technical report, CERN, 2021. All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-MUON-PUB-2021-002`.

[44] ATLAS collaboration. Standard Model Summary Plots June 2021. Technical report, CERN, 2021. All figures including auxiliary figures are available at `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-032`.

[45] ATLAS collaboration. ATLAS Luminosity Public Results Run 1, 2024. `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun1` Accessed: 2024-10-15.

[46] ATLAS collaboration. ATLAS Luminosity Public Results Run 2, 2024. `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2` Accessed: 2024-10-15.

[47] John C. Collins and Davison E. Soper. Angular distribution of dileptons in high-energy hadron collisions. *Phys. Rev. D*, 16:2219–2225, Oct 1977.

[48] John C. Collins, Davison E. Soper, and George F. Sterman. Factorization of Hard Processes in QCD. *Adv. Ser. Direct. High Energy Phys.*, 5:1–91, 1989.

[49] Thijs G. Cornelissen, M. Elsing, I. Gavrilenko, J. F. Laporte, W. Liebig, M. Limper, K. Nikolopoulos, A. Poppleton, and A. Salzburger. The global χ^2 track fitter in ATLAS. *J. Phys. Conf. Ser.*, 119:032013, 2008.

[50] G. Cowan. *Statistical data analysis.* Oxford University Press, USA, 1998.

[51] N. Davidson, T. Przedzinski, and Z. Was. PHOTOS interface in C++: Technical and Physics Documentation. *Comput. Phys. Commun.*, 199:86–101, 2016.

[52] Abdelhak Djouadi. The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the standard model. *Physics Reports*, 457:1–216, 2005.

[53] Y Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977.

[54] Michelle J. Dolinski, Alan W. P. Poon, and Werner Rodejohann. Neutrinoless Double-Beta Decay: Status and Prospects. *Ann. Rev. Nucl. Part. Sci.*, 69:219–251, 2019.

[55] Markus A. Ebert, Johannes K. L. Michel, Iain W. Stewart, and Frank J. Tackmann. Drell-Yan q_T resummation of fiducial power corrections at N^{3}LL. *JHEP*, 04:102, 2021.

[56] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–323, 1964.

[57] R. P. Feynman. Space - time approach to quantum electrodynamics. *Phys. Rev.*, 76:769–789, 1949.

[58] R. P. Feynman. Mathematical formulation of the quantum theory of electromagnetic interaction. *Phys. Rev.*, 80:440–457, 1950.

[59] Stefano Frixione, Paolo Nason, and Carlo Oleari. Matching NLO QCD computations with Parton Shower simulations: the POWHEG method. *JHEP*, 11:070, 2007.

[60] Sheldon L. Glashow. The renormalizability of vector meson interactions. *Nuclear Physics*, 10:107–117, 1959.

[61] T Gleisberg, S Höche, F Krauss, M Schönherr, S Schumann, F Siegert, and J Winter. Event generation with SHERPA 1.1. *Journal of High Energy Physics*, 2009(02):007–007, feb 2009.

[62] V. N. Gribov and L. N. Lipatov. Deep inelastic e p scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972.

[63] H1 and Zeus Collaborations. Combination of measurements of inclusive deep inelastic ep scattering cross sections and QCD analysis of HERA data. *Eur. Phys. J. C*, 75, 2015.

[64] Julie Haffner. The CERN accelerator complex. Complexe des accélérateurs du CERN, 2013. `https://cds.cern.ch/record/1621894`.

[65] A. Hayrapetyan et al. Measurement of the Drell–Yan forward-backward asymmetry and of the effective leptonic weak mixing angle in proton-proton collisions at $\sqrt{s} =$ 13 TeV. *Physics Letters B*, 866:139526, 2025.

[66] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[67] Manuel Hohmann. Overview of 13 TeV Z Polarization Analysis Activities, 2021. `https://indico.cern.ch/event/1003893/contributions/4215706/attachments/2183604/3689114/210203_Forward_Electrons_Run2.pdf` Accessed: February, 2024.

[68] Hung-Liang Lai, Marco Guzzi, Joey Huston, Zhao Li, Pavel M. Nadolsky, Jon Pumplin, and C.-P. Yuan. New parton distributions for collider physics. *Phys.Rev.D*, 82:074024, 2010.

[69] C. S. Lam and Wu-Ki Tung. A Systematic Approach to Inclusive Lepton Pair Production in Hadronic Collisions. *Phys. Rev. D*, 18:2447, 1978.

[70] Sascha Mehlhase. ATLAS detector slice (and particle visualisations). 2021. `https://cds.cern.ch/record/2770815`.

[71] Paolo Nason. A New method for combining NLO QCD with shower Monte Carlo algorithms. *JHEP*, 11:040, 2004.

[72] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024.

[73] Jerzy Neyman and Egon Sharpe Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Phil. Trans. Roy. Soc. Lond. A*, 231(694-706):289–337, 1933.

[74] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, 2e edition, 2006.

[75] C. Oleari. The POWHEG BOX. *Nuclear Physics B - Proceedings Supplements*, 205-206:36–41, 2010.

[76] Karl Pearson. X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175, 1900.

[77] Heinz Pernegger. The Pixel Detector of the ATLAS Experiment for LHC Run-2. Technical Report 06, CERN, 2015. `http://cds.cern.ch/record/1985432`.

[78] Andrés Pinto, Zhibo Wu, Fabrice Balli, Nicolas Berger, Maarten Boonekamp, Émilien Chapon, Tatsuo Kawamoto, and Bogdan Malaescu. Uncertainty components in profile likelihood fits. *Eur. Phys. J. C*, 84(6):593, 2024.

[79] J. Pumplin, D. R. Stump, J. Huston, H. L. Lai, Pavel M. Nadolsky, and W. K. Tung. New generation of parton distributions with uncertainties from global QCD analysis. *JHEP*, 07:012, 2002.

[80] Anders Ryd, David Lange, Natalia Kuznetsova, Sophie Versille, Marcello Rotondo, David P. Kirkby, Frank K. Wuerthwein, and Akimasa Ishikawa. EvtGen: A Monte Carlo Generator for B-Physics. 2005. `https://evtgen.hepforge.org/doc/EvtGenGuide.pdf`.

[81] Andrei D Sakharov. Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi*, 34(5):392, may 1991.

[82] Abdus Salam and John Clive Ward. Weak and electromagnetic interactions. *Nuovo Cim.*, 11:568–577, 1959.

[83] Albert M. Sirunyan et al. Measurement of the weak mixing angle using the forward-backward asymmetry of Drell-Yan events in pp collisions at 8 TeV. *Eur. Phys. J. C*, 78(9):701, 2018.

[84] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An introduction to PYTHIA 8.2. *Computer Physics Communications*, 191:159–177, 2015.

[85] S. van der Meer. Calibration of the effective beam height in the ISR. Technical report, CERN, Geneva, 1968. `https://cds.cern.ch/record/296752`.

[86] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[87] Craig John Wells. Introduction of MVA Calibration for Forward Electrons. Technical report, CERN, 2021. `https://cds.cern.ch/record/2789786`.

[88] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Annals Math. Statist.*, 9(1):60–62, 1938.

[89] Kenneth G. Wilson. Confinement of Quarks. *Phys. Rev. D*, 10:2445–2459, 1974.

[90] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental Test of Parity Conservation in Beta Decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.

[91] H. Yukawa. Quantum Theory of Nonlocal Fields. 1. Free Fields. *Phys. Rev.*, 77:219–226, 1950.

Acknowledgements

And so we reach the end of this thesis and with it my work as a PhD student. It's certainly a strange feeling to be writing this! Whilst this work is entirely my own I certainly wouldn't have reached this point without the help and support of many others over the years who more than deserve to be mentioned.

To my supervisor Ludovica, I will be forever grateful for this opportunity you gave me. Your knowledge on both the Standard Model and electron performance in ATLAS is truly unparalleled and I've learnt so much from you because of that. Along every step of this PhD you've supported, encouraged and helped me. As the first of no doubt many PhD students you'll successfully supervise, I hope you've enjoyed becoming my supervisor as much as I have being your PhD student.

To Filip, I was certainly lucky to have the help of a postdoc like you. You have the patience of a saint, as even when you had a million and one other tasks to do you always took the time to help me with debugging or answering some trivial question I had. Thank you as well for reviewing (and improving!) this thesis even if you weren't obligated to.

A big thanks to everyone else past and present in the Z-A_i analysis group is also in order. I always leave each analysis meeting feeling like I've learnt something new or understand something better than before which is a credit to the depth of talent in the group. I'm sure it won't be too much longer until we're ready to have the A_i paper published!

Being able to do my PhD work within the DESY ATLAS group has also been a pleasure and a privilege. A big thank you has to go to the group secretaries Carolin and Sabine who keep everything running like clockwork and make navigating all bureaucracy a breeze. I've made so many good friends amongst the other PhD students and postdocs here, without whom the whole PhD process wouldn't have been anywhere near as fun. A special shout out has to go to my officemate Tom, I've never met a man who can chat so much rubbish yet at the same time make me laugh so much before.

Away from the work side of things I also want to thank everyone back home as well. Mrs Sinclair, it's all your fault that I had to write this thesis in the first place! Thank you for nurturing my love of physics back in school that started off this journey. To all my friends from school and uni, I always look forward to seeing you all when I'm back home even if that isn't as often as I'd like. And finally to Mum and Dad, you have always loved me and supported me in every choice I've made. You might not understand the physics in this thesis but I know you're both proud of me no matter what.