

Measurement of Vector Boson Production Cross Sections and Muon Calibration with the CMS Experiment

Zur Erlangung des akademischen Grades eines
DOKTORS DER NATURWISSENSCHAFTEN (Dr. rer. nat.)

von der KIT-Fakultät für Physik
des Karlsruher Instituts für Technologie (KIT)
genehmigte

DISSERTATION

von M.Sc. Jost von den Driesch
aus Karlsruhe

Bearbeitungszeitraum: 1. Juli 2022–23. April 2025
Tag der mündlichen Prüfung: 30. Mai 2025

Referent: Prof. Dr. Markus Klute Institut für Experimentelle Teilchenphysik, KIT
Korreferent: Prof. Dr. Ulrich Husemann Institut für Experimentelle Teilchenphysik, KIT



This document is licensed under a Creative Commons
Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0):
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

Acknowledgments

This thesis would not have been possible without the support of many.

First and foremost, I would like to thank my supervisor, Markus Klute, for guiding me through these three exciting and productive years. I deeply appreciate the trust you placed in me and the countless opportunities for both scientific and personal development—in particular for supporting my time at CERN, the opportunity to substitute for your lecture, and the unforgettable experience of attending the Oman Science Festival.

I am also grateful to Ulrich Husemann, not only for preparing me so well for this PhD during my bachelor’s and master’s studies, but also for your continued support whenever needed—especially your many helpful comments during the writing of this thesis.

My thanks also go to the “Karlsruhe School of Elementary Particle and Astroparticle Physics: Science and Technology” for the generous financial and intellectual support, and to the Institute of Experimental Particle Physics (ETP) for stepping in to cover the rest of the funding.

I would like to express my appreciation to everyone directly involved in the research presented here, especially Minseok, Nils and Xunwu, as well as the students I had the pleasure to supervise: Dorian, Erik, and Jonas. I am also thankful to everyone who provided feedback on various aspects of this work, in particular Cedric, Dorian, Hannes, Lucas, Nils, and Thorsten for carefully reading through it and offering valuable suggestions. Many thanks to those who contributed during discussions and collaboration, including colleagues from ETP-SMP, CMS-SMP, CMS-MUO, and our analysis review committee chaired by Paolo Azzurri.

The PhD years were also a lot of fun, and that is largely thanks to the great environment at the institute—especially my awesome office mates in 9-23, Lars and Nils!

Last but not least, I want to thank my family and friends for their continuous support—especially my partner, for always having my back and being so understanding and supportive during the final weeks of this journey.

Contents

1	Introduction	1
2	Preface	3
3	The Standard Model of Particle Physics	5
3.1	Particles and Interactions	5
3.1.1	Mathematical Foundations	6
3.1.2	Fundamental Interactions	8
3.1.3	Higgs Mechanism	10
3.2	Theoretical Predictions in the Standard Model	14
3.2.1	Cross Section and Decay Width	14
3.2.2	Perturbation Theory	15
3.2.3	Details of Proton-Proton Collisions	16
3.2.4	Uncertainties	19
3.3	Limitations of the Standard Model	20
4	Experimental Environment	21
4.1	The Large Hadron Collider	22
4.2	The Compact Muon Solenoid Experiment	24
4.2.1	Detector Setup	24
4.2.2	Data Acquisition and Trigger System	28
4.2.3	Event Reconstruction	30
5	Statistical Methods	33
5.1	General Concepts of Probability Theory	33
5.1.1	Probability Interpretation	33
5.1.2	Hypothesis Tests	34
5.1.3	Conditional Probabilities and Likelihood	34
5.1.4	Parameter Estimation with the Likelihood Method	35
5.1.5	Cross Section Measurement with Likelihood Maximization	37
5.2	Monte Carlo Methods	39
5.2.1	Integration	39
5.2.2	Bootstrapping	39
6	Muon Corrections in the CMS Experiment	43
6.1	Kinematic Behavior of Muons	43
6.1.1	Momentum Distribution of Muons from Resonances	43
6.1.2	Optimizing Muon Reconstruction	47
6.2	Tag and Probe Method	49
6.2.1	General Concept	49

6.2.2	Statistical Model	50
6.2.3	Uncertainty Estimation	52
6.2.4	Application of the Efficiency Corrections	52
6.3	Muon Momentum Calibration	54
6.3.1	Scale and Resolution	54
6.3.2	Correction Procedure	55
6.3.3	Uncertainty Estimation	67
6.3.4	Consistency Checks	71
6.4	Discussion	75
7	Measurements of the Cross Sections of W and Z Boson Production	77
7.1	Introduction and Overview	77
7.1.1	Motivation	77
7.1.2	Event Topology	79
7.1.3	Measurement Strategy	80
7.1.4	Ingredients for the Measurement	82
7.2	Data Calibration and Tuning of the Prediction Model	87
7.2.1	Pileup Reweighting	87
7.2.2	Boson Momentum Reweighting	88
7.2.3	Muon Efficiency Correction	91
7.2.4	Muon Inefficiency Correction	92
7.2.5	Muon Momentum Correction	93
7.2.6	Azimuthal Correction of the Missing Transverse Momentum	97
7.2.7	Recoil Calibration	98
7.2.8	Estimation of the QCD Multijet Background	101
7.2.9	Further Sources of Uncertainty	105
7.3	Statistical Fit	106
7.3.1	Fit Performance	106
7.3.2	Cross Section Measurement	109
7.3.3	Cross Section Ratios	113
7.4	Discussion	118
8	Conclusions	121
Bibliography		123
Appendix		133
A	Estimations for Relevant Final-State Particles	133
B	Muon Selection Efficiencies	136
C	QCD Multijet Background Estimation	139
D	Correlation Plots of Maximum Likelihood Fit	142

List of Figures

3.1	Summary chart of fundamental particles in the standard model.	6
3.2	Higgs potential simplified to one dimension.	11
3.3	Parton distribution functions.	17
3.4	Example Feynman Diagram.	18
4.1	Sketch of the CMS experiment.	23
5.1	Poissonian distributions as functions of different dependent variables.	35
5.2	Logarithmic Poissonian likelihood function with uncertainty construction. .	37
5.3	Example for the bootstrapping procedure.	40
5.4	Distributions of the mean values and standard deviations from a set of bootstrapped samples.	41
6.1	Distribution of the p_T values of muons from a Z boson decay in its rest frame. .	44
6.2	Distribution of the muon momenta from the decay of a Z boson and J/ψ meson in momentum space.	45
6.3	Example fits for the extraction of the efficiency from the data.	51
6.4	Simplified example of the correction procedure applied to a hypothetical measurement of a specific true value.	56
6.5	Scheme for the derivation of the muon momentum scale and resolution corrections.	57
6.6	Sketch of the impact of misalignment on the track curvature.	59
6.7	Example fits in the resolution extraction steps in the J/ψ region and in the Z region.	61
6.8	Distributions of the dimuon mass in the J/ψ region and in the Z region. .	62
6.9	Example plots of the evolution of the mean difference between the dimuon mass distribution at the reconstruction-level and generator-level.	64
6.10	Dimuon mass distributions in the J/ψ region and in the Z region, before and after the p_T scale correction.	65
6.11	Distributions before and after the full correction procedure for the J/ψ and the Z region.	66
6.12	Mean and standard deviation values of the fine-tuning correction parameters with respect to statistical and systematic variations.	68
6.13	Exemplary comparison plots for the calculation of the uncertainty in the correction of the scale and resolution.	71
6.14	Distributions of the corrected dimuon mass including variations of the statistical uncertainty in the scale and resolution parametrization.	72
6.15	Example distributions of fine-tuning correction parameters in the J/ψ region. .	73
6.16	Crosscheck of the additive term in the muon scale correction between the medium- p_T and the high- p_T region.	74

7.1	Example Feynman diagrams for Z, W ⁺ , and W ⁻ boson production.	80
7.2	Example Feynman diagrams for muonic Z, W ⁺ , and W ⁻ boson decays.	80
7.3	Schematic concept of the different phase spaces in the analysis.	81
7.4	Example Feynman diagrams of the background processes.	84
7.5	Impact of the pileup reweighting on the distribution of reconstructed vertices.	88
7.6	Impact of the pileup correction on the signal distributions.	89
7.7	Impact of the boson p_T correction on the boson p_T distribution.	90
7.8	Impact of the boson p_T correction on the signal distributions.	90
7.9	Selection efficiency for negative muons in simulation and data.	92
7.10	Example dimuon mass fits before and after the muon p_T correction.	95
7.11	Dependence of the dimuon mass peak width on the additional smearing factor.	96
7.12	Impact of the muon p_T correction on the dimuon mass distribution.	97
7.13	Extraction of the azimuthal p_T^{miss} correction parameters.	98
7.14	Impact of the azimuthal p_T^{miss} correction on the azimuthal p_T^{miss} distribution.	99
7.15	Impact of the recoil corrections on distributions related to the hadronic recoil.	102
7.16	Conceptual design of the QCD extrapolation method.	104
7.17	Pre- and post-fit distributions in the three signal regions.	107
7.18	Distribution of the pull and impact values of the leading nuisance parameters with respect to the Z and W $^{\pm}$ signal strength parameters.	111
7.19	Comparison between the measured and predicted cross sections.	114
7.20	Distribution of the pull and impact values of the leading nuisance parameters with respect to the cross section ratios.	115
7.21	Comparison between the measured and predicted cross section ratios.	117
7.22	Comparison of measured total cross sections at different center-of-mass energies with the corresponding theoretical prediction.	119
A.1	Muon energy loss in copper.	134
B.2	Standalone muon reconstruction efficiency in data and simulation.	136
B.3	Global muon reconstruction efficiency in data and simulation.	136
B.4	Tight identification selection efficiency in data and simulation.	137
B.5	Tight isolation selection efficiency in data and simulation.	137
B.6	Muon trigger efficiency in data and simulation.	138
C.7	Example plots for the first step of the QCD extrapolation procedure.	139
C.8	Example plots for the second step of the QCD extrapolation procedure.	140
C.9	Extrapolated QCD templates.	141
D.10	Correlation matrix for the Z cross section measurement.	142
D.11	Correlation matrix for the W ⁺ cross section measurement.	143
D.12	Correlation matrix for the W ⁻ cross section measurement.	143
D.13	Correlation matrix for the measurement of the W ⁺ /Z cross section ratio.	144
D.14	Correlation matrix for the measurement of the W ⁻ /Z cross section ratio.	144
D.15	Correlation matrix for the measurement of the W ⁺ /W ⁻ cross section ratio.	145
D.16	Correlation matrix for the individual signal strength parameters.	145

List of Tables

6.1	Selection requirements for the muon p_T calibration procedure.	57
7.1	Predictions for the product of the cross sections and branching fractions.	86
7.2	Predicted ratios of the product of cross sections and branching fractions.	86
7.3	Pre-fit event yields in the fiducial region.	108
7.4	Post-fit event yields in the fiducial region.	109
7.5	Measured cross section values in the fiducial and total phase space.	110
7.6	Post-fit uncertainties in percent for the fiducial cross section measurement.	113
7.7	Ratios of the measured cross sections in the fiducial and total phase space.	114
7.8	Post-fit uncertainties for the fiducial cross section ratio measurement.	117

Acronyms

c.d.f. cumulative distribution function
CERN Conseil européen pour la recherche nucléaire
CKM Cabibbo-Kobayashi-Maskawa
CMS Compact Muon Solenoid
CSC cathode strip chamber
DT drift tube
ECAL electromagnetic calorimeter
EWK electroweak
FSR final-state radiation
GEM gas electron multiplier
HCAL hadron calorimeter
HLT high-level trigger
ISR initial-state radiation
LHC Large Hadron Collider
LO leading order
LS long shutdown
MC Monte Carlo
ML maximum likelihood
NLO next-to-leading order
NNLL next-to-next-to-leading logarithmic
NNLO next-to-next-to-leading order
npvGood number of good reconstructed primary vertices
p.d.f. probability density function
PDF parton distribution function
PF particle-flow
pp proton-proton
PV primary vertex
QCD quantum chromodynamics
QED quantum electrodynamics
RPC resistive plate chamber
SM standard model
ST single-top
TnP tag-and-probe
VBF vector boson fusion
VEV vacuum expectation value

1 Introduction

Driven by the desire to understand its surroundings, humankind has developed a great many of theories on various aspects of nature. Over time, these theories have evolved through rigorous experiments and continual refinement. As the theoretical predictions become more and more precise, probing them tends to require increasingly sophisticated experiments.

In the descriptions of processes at smallest scales, this progress has lead to today's [standard model \(SM\)](#) of particle physics, which describes the elementary particles and their interactions with remarkable accuracy. The [Large Hadron Collider \(LHC\)](#) in Geneva, Switzerland provides an experimental environment to probe the [SM](#), thereby helping to advance our understanding of nature at the smallest scales. It does so by accelerating protons in two counter-rotating beams and colliding them at four major experiments, one of which—the [Compact Muon Solenoid \(CMS\)](#) detector—is used throughout this thesis. With each proton reaching ultra-relativistic energies of 6.8 TeV, these collisions can produce high-mass, short-lived particles, such as the W^\pm and Z bosons. These three bosons mediate the weak interaction within the [SM](#); the W^\pm bosons are, for instance, responsible for β decays in atomic nuclei. As the name suggests, the weak interaction is characterized by a relatively low occurrence rate compared to the other fundamental forces. To enable precise studies of such rare processes, the [LHC](#) is optimized not only for high beam energies but also for high collision rates.

A large share of the studies conducted with the [CMS](#) experiment focuses on unstable particles, whose defining characteristic is their decay into lighter particles. To reconstruct the properties of these short-lived particles, it is essential to measure their decay products with high precision. One such product that frequently appears in decays of massive unstable particles is the muon. Although the muon is unstable itself, it typically lives long enough to be detected by the [CMS](#) experiment with high efficiency.

Probing the [SM](#) or theories beyond usually involves statistical hypothesis testing, where theoretical models are translated into expected signatures using detailed simulations of both the physics processes and the detector response. Despite the high accuracy of these simulations, residual differences between the simulation and the real measurement often remain. Given the high precisions reached in modern measurements, even subtle mismodeling in the detector description can affect the outcome of hypothesis tests by introducing systematic effects. To mitigate this, residual differences in the description

of simulated and measured muons need to be addressed. One major contribution of this thesis is the precise calibration of muon transverse momenta to mitigate such residual discrepancies.

The motivation for the muon calibration sparked in the context of the analysis of a dataset corresponding to an integrated luminosity of 5.01 fb^{-1} [1]. This dataset was collected early in Run 3 of the [LHC](#), which began delivering [proton-proton \(pp\)](#) collisions at the unprecedented center-of-mass energy of 13.6 TeV —just four days after the start of this thesis in July 2022. Following the second long shutdown (2018-2022), it was essential to start analyzing the new data early on in order to identify and correct potential issues in the detector response before recording the bulk of the data. To this end, and to probe the [SM](#) at high precision, the inclusive cross sections of W^\pm and Z boson production are measured in this thesis. Furthermore, their ratios are studied, offering high sensitivity towards the description of the proton’s substructure. These measurements focus on the muonic decay channels of the bosons, further motivating the need for a precise calibration of the muon momenta.

This thesis starts by placing the work in the broader context of the [CMS](#) Collaboration and outlining my specific contributions in Chapter 2. Then, Chapter 3 introduces the [SM](#), which is later probed through the cross section (ratio) measurements. The experimental environment provided by the [LHC](#) and the [CMS](#) detector is described in Chapter 4. In addition, Chapter 5 introduces statistical methods that are central to many aspects of this thesis.

Building on this foundation, the main contributions of this thesis are presented. The first major contribution is the precise calibration of muon momenta in [LHC](#) Run 3, discussed in Chapter 6. This chapter provides a robust understanding of muons in the [CMS](#) experiment, setting the stage for the cross section measurements of W^\pm and Z boson production with muons in the final state, as detailed in Chapter 7. Finally, Chapter 8 provides a summary and an outlook on future developments.

2 Preface

Data analyses in high-energy physics, such as those conducted in this doctoral thesis, are a massive collaborative effort. The [LHC](#), which provides the particle collisions essential for studies with the [CMS](#) experiment, was built, and is operated and maintained by the [Conseil européen pour la recherche nucléaire \(CERN\)](#) accelerator departments. The [CMS](#) experiment itself was built, and is operated and maintained by the [CMS](#) Collaboration. Furthermore, the collaboration produces the simulated and recorded datasets, which are processed using the infrastructure provided by the Worldwide [LHC](#) Computing Grid. The software environment necessary for the studies with [CMS](#) data also plays an important role and is created and maintained largely by the collaboration. Thus, numerous individuals are involved in any data analysis, which is reflected by the number of close to 2000 scientific authors in the [CMS](#) Collaboration [2].

This effort relies on the dedication of every member of the [CMS](#) Collaboration. To support the continuous data taking and processing of the [CMS](#) experiment, I contributed as a technical shifter in the [CMS](#) control room for multiple shift-blocks in 2023 and 2024, and as a GridKa shifter in 2022 and 2023.

Moreover, I was involved in the [CMS](#) MUO group, which focuses on understanding the behavior of muons in the [CMS](#) detector and improving their description in simulation. Based on the expertise developed during the W^\pm and Z boson cross section measurement, I led the effort toward muon p_T scale and resolution corrections in Run 3. As the persons responsible for the correction in previous runs did not continue their effort—and as the existing corrections had been derived in a private analysis framework that was not provided to the collaboration—we restarted the calibration procedure mainly based on the information provided in Reference [3]. During this period, the code base was developed from scratch by a master's student, Dorian Guthmann, whom I co-supervised, and myself, and we ultimately made it available to the collaboration. In particular, I constructed the software architecture, participated in code development, and optimized the code with respect to performance. I also proposed a simplification of the uncertainty calculations, developed the relevant code, and co-supervised a bachelor's student to evaluate the performance of the simplified method.

Furthermore, as a muon L3 convener responsible for muon calibration and commissioning, I co-maintained the code of the [tag-and-probe \(TnP\)](#) tool, which enables the analysis-specific estimation of muon efficiency scale factors in the [CMS](#) Collaboration.

When I joined the cross section measurement effort, a significant part of the preparation had already been done by Dr. Minseok Oh, especially the general workflow of event selection, correction, and fit setup. My first major contribution to the analysis was the precise calibration of the muon momenta, which became essential with the new data and sparked the interest in the muon service work. After Minseok left the analysis group, I took over responsibility for the analysis. Together with the analysis contact, Dr. Nils Faltermann, we worked closely to understand every detail of the analysis and quickly pushed an already well-understood part of the measurement (the cross section of Z boson production) to a preliminary publication [4]. Dr. Xunwu Zuo provided the theoretical cross section calculations, which served as both input to the measurement and a benchmark for the results.

The majority of the work, however, was focused on the other part of the measurement, where the undetectable neutrino poses greater challenges to the precision measurement. On the way to the publication of the full work [5], we placed even more emphasis on understanding the details of the measurement. In this phase, I made the main contributions in optimizing the corrections, which included, in particular:

- detailed studies of different pileup correction techniques, ultimately implementing the pileup corrections, which are centrally provided by the collaboration,
- extracting and incorporating the boson momentum corrections,
- calculating the muon efficiency corrections and implementing the muon trigger inefficiency correction, which was provided by the collaboration,
- deriving the azimuthal corrections of the missing transverse momentum and providing both the corrections and the derivation code to the collaboration,
- re-performing and improving the recoil calibration procedure, e.g., by implementing the correlation with the pileup correction,
- improving the [quantum chromodynamics \(QCD\)](#) estimation procedure, including both idea and implementation to derive separate efficiency scale factors in the non-isolated region.

Lastly, while a basic fit implementation was already available, I invested significant effort in optimizing the selection process, implementing additional nuisance parameters, and generally understanding its behavior.

In recent years, artificial intelligence tools have become integrated into various day-to-day activities. This thesis also incorporates the use of such tools, particularly for specific coding tasks and for grammatical or stylistic improvements of the text. More specifically, I used ChatGPT¹ and Gemini² for both text and programming improvement. Additionally, Le Chat³ and Github Copilot⁴ were used to assist with programming.

These tools provide substantial benefits, especially in efficient programming, improved communication, and summarizing specific topics. However, I am aware of their potential risks for research quality, particularly hallucination, bias, and research privacy, as outlined in Reference [6]. I have therefore used these tools responsibly, including cross-checking code suggestions and ensuring that text improvements do not introduce new information beyond the original input—thereby maintaining highest research standards.

¹OpenAI. ChatGPT. <https://chat.openai.com/>

²Google. Gemini. <https://gemini.google.com/>

³Mistral AI. Le Chat. <https://mistral.ai/>

⁴GitHub. GitHub Copilot. <https://github.com/features/copilot/>

3 The Standard Model of Particle Physics

The [standard model \(SM\)](#) of particle physics is a fundamental theory that provides the mathematical framework for describing elementary particles and their interactions. It accounts for all known fundamental particles and describes three of the four fundamental interactions with remarkable accuracy. Section 3.1 provides an overview of the fundamental particles and interactions as well as their mathematical description. In order to validate the [SM](#), it is essential to obtain predictions, such as cross sections, which can be tested with experiments. For this reason, a conceptual introduction to such calculations is given in Section 3.2. Even though the [SM](#) performs exceptionally well in the description of interactions, there are some long-standing observations hinting at physics beyond the [SM](#). Such limitations of the [SM](#) are discussed briefly in Section 3.3. Explanations and motivations provided in this chapter are inspired by References [7–11].

3.1 Particles and Interactions

In contrast to classical theories, the fundamental building blocks of the [SM](#) are quantum fields. Rather than being clearly localized objects, particles are described as quantized excitations of these fields, with a probability distribution determining their presence in a given region of space-time. Interactions, in turn, arise from the coupling of matter fields to gauge fields, which mediate the fundamental forces. These interactions occur probabilistically and are dictated by the symmetries of the underlying gauge theory.

All fields in the [SM](#) transform under certain representations of the Lorentz group, following from specific spin properties of the corresponding particle. Integer-spin particles are called bosons and can be further classified as vector bosons (spin-1) or scalar bosons (spin-0). Vector bosons mediate the fundamental interactions: the gluon mediates the strong interaction, while the electroweak interaction is carried by the photon, as well as the W^\pm and Z bosons.

Fermions, on the other hand, are half-integer spin particles. They can generally be divided into two groups—quarks and leptons—each consisting of three generations. Each generation can then be divided into two subgroups, respectively. Quarks can be of either up- or down-type, while leptons are neutral or charged with respect to the electromagnetic interaction. An overview of all known fundamental particles, as well as their intrinsic properties, is provided in Figure 3.1.

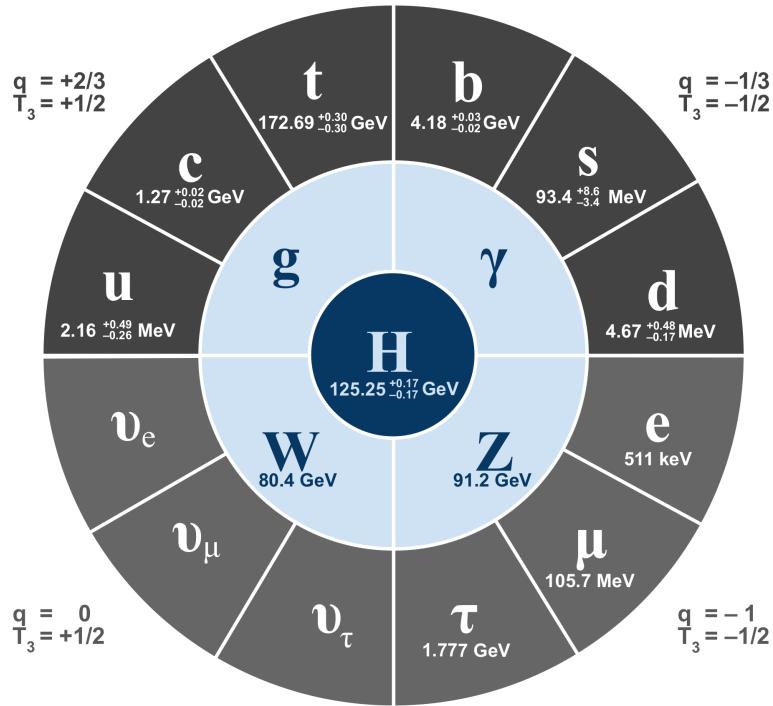


Figure 3.1: Summary chart of the fundamental particles in the SM. Adapted from Reference [12] with masses from Reference [7]. Particles without a given mass are considered massless in the SM. Masses without a given uncertainty value have an uncertainty smaller than the given precision.

3.1.1 Mathematical Foundations

The fundamental framework for the mathematical description of the SM consists of the Lagrangian density as well as assumed continuous symmetries. The Lagrangian density is the generalization of the Lagrangian from classical mechanics to quantized fields, encoding both their dynamical properties and their interactions. Symmetries of the Lagrangian density are enforced by requiring its invariance under certain transformations, ensuring that the physical behavior of the system remains unchanged. According to Noether's theorem [13], every continuous symmetry corresponds to a conserved quantity, and vice versa. A prime example is translational invariance: if the system's dynamics remain unchanged under spatial translation, then momentum is conserved. Symmetries connected to such transformations, which are independent of the space-time position, are called global symmetries. In contrast, local symmetries depend on the space-time coordinates and are therefore generalizations of global symmetries.

Such continuous symmetries can be described by the more abstract concept of Lie groups, which are deeply connected to their corresponding Lie algebras. The elements of the Lie algebra can be expressed as a linear combination of the basis elements T^a , which are called generators and satisfy the commutation relation:

$$[T^a, T^b] = T^a T^b - T^b T^a = i f^{abc} T^c, \quad (3.1)$$

where $[\cdot, \cdot]$ denotes the Lie bracket, and f^{abc} are the structure constants of the algebra. Elements of the corresponding Lie group can be obtained by exponentiating the Lie algebra

elements. For the Lie groups $U(n)$, which the SM relies on, this takes the simple form:

$$A(\theta) = \exp\left(i \sum_a^n \theta^a T^a\right), \quad (3.2)$$

with real valued parameters θ^a . On the other hand, the generators of the Lie algebra can be derived by taking the directional derivatives at the neutral element $A(0)$ of the group:

$$T^a = -i \frac{d}{d\theta^a} A \Big|_{\theta=0}. \quad (3.3)$$

Consequently, the dimensions of the Lie group and the associated Lie algebra are identical and reflect the number of generators. Moreover, following Noether's theorem, the dimension of the symmetry group is also equal to the number of conserved currents. Integrating a conserved current over space-time yields a conserved charge, which defines a quantum number associated with the symmetry group.

The connection between Lie algebras and Lie groups implies, that if and only if the structure constants of the algebra are zero, the Lie group is abelian. This means that all generators commute, and consequently, the order of applying different symmetry operations becomes irrelevant.

For illustration of the impact of these mathematical concepts, consider a Lagrangian density for a massless spin-1/2 field:

$$\mathcal{L} = i\bar{\psi}\gamma^\mu\partial_\mu\psi, \quad (3.4)$$

where $\bar{\psi} = \psi^\dagger\gamma^0$, with \dagger denoting the Hermitian conjugate (the combination of a complex conjugation and transposition), ∂_μ is the space-time derivative, and γ^μ are the gamma matrices. The Lagrangian density in Equation 3.4 is constructed such that its Euler-Lagrange equations reproduce the Dirac equation for a massless field. The application of a global symmetry operation $\psi \rightarrow U\psi$, where U is an element from a unitary Lie group, leaves the Lagrangian density unchanged, as $U^\dagger U = 1$.

In contrast, consider now a local unitary transformation $\psi \rightarrow U(x)\psi$, where the transformation matrix $U(x) = \exp(iu(x))$ is an element of a Lie group. Following the relation in Equation 3.2, $u(x)$ is a linear combination of the basis elements of the corresponding Lie algebra, where the parameters depend on the space-time position x . Such a transformation introduces a term that does not cancel:

$$\mathcal{L} \rightarrow \bar{\psi}(i\gamma^\mu\partial_\mu - \gamma^\mu(\partial_\mu u(x)))\psi. \quad (3.5)$$

To avoid this behavior, the usual space-time derivative is replaced with the covariant derivative $D_\mu = \partial_\mu + iA_\mu$, where A_μ is called gauge field. If the gauge field transforms like $A_\mu \rightarrow A_\mu - \partial_\mu u(x)$, then the additional term in Equation 3.5 cancels. This example shows the connection between the imposed local symmetry and the introduction of new fields via the Lie group and its corresponding Lie algebra.

Assuming the gauge field itself is a linear combination of the generators $A_\mu(x) = A_\mu^a(x)T^a$, a gauge invariant quantity can be constructed by considering the commutator of the covariant derivatives:

$$\begin{aligned} F_{\mu\nu} &= i[D_\mu, D_\nu] \\ &= \partial_\mu A_\nu - \partial_\nu A_\mu - i[A_\mu, A_\nu], \\ F_{\mu\nu}^a &= \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + f^{abc} A_\mu^b A_\nu^c, \end{aligned} \quad (3.6)$$

where in the last line a single component of the algebra is picked to illustrate the connection with the structure constants. A Lagrangian density proportional to $\text{tr}(F^{\mu\nu}F_{\mu\nu})$ is invariant under gauge transformations and yields the equations of motion for a massless spin-1 (vector) field. If the dimension of the Lie group is n , then there are n independent contributions to the gauge field, which are identified with one massless gauge boson each. The vector field needs to be massless because additional mass terms proportional to $m^2 A_\mu A^\mu$ would violate gauge invariance.

To summarize, imposing a local symmetry on the Lagrangian density leads to the introduction of new massless gauge fields. The resulting terms, where different fields appear together, can be interpreted as interaction terms, as they describe the coupling between the fields. Specifically, for non-abelian groups, where the structure constants are non-vanishing, the gauge fields can also interact with themselves.

3.1.2 Fundamental Interactions

In its symmetric state, the **SM** is assumed to be invariant under independent transformations from three Lie groups— $\text{SU}(3)_C$, $\text{SU}(2)_L$, and $\text{U}(1)_Y$ —where the first is used for the description of the strong interaction, while the latter two are the symmetry groups describing the electroweak interaction.

3.1.2.1 Strong Interaction

The strong interaction is described by **quantum chromodynamics (QCD)** [14–16], a theory based on the symmetry group $\text{SU}(3)_C$. This symmetry corresponds to the exchange of three color charges—red, green and blue—between quarks. The group has a dimension of eight, meaning there are eight generators and independent contributions to the total gauge field, which are called gluons. A commonly chosen set of generators are the eight Gell-Mann matrices λ^a . With these matrices, and the corresponding gluon fields G_μ^a , the Lagrangian density of the strong interaction reads:

$$\mathcal{L}_{\text{QCD}} = i\bar{\psi}_i \gamma^\mu \partial_\mu \delta_{ij} \psi_j + \sqrt{4\pi\alpha_s} \bar{\psi}_i \gamma^\mu (\lambda^a)_{ij} \psi_j G_\mu^a - \frac{1}{4} F^{a,\mu\nu} F_{\mu\nu}^a, \quad (3.7)$$

where δ_{ij} is the Kronecker delta with color-indices i and j running from 1 to 3. Furthermore, α_s is the strong coupling constant, and the field strength tensor of the strong interaction is given by:

$$F_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + \sqrt{4\pi\alpha_s} f^{abc} G_\mu^b G_\nu^c, \quad (3.8)$$

with the structure constants f^{abc} of the Lie algebra $\mathfrak{su}(3)$.

Since the symmetry group of **QCD** is not abelian, gluons can interact with each other, altering the effective behavior of the strong interaction at different energy scales. These self-interactions can be absorbed into the coupling constant, which becomes energy scale dependent. At small distances (high energy scales), the coupling constant is small, a behavior known as asymptotic freedom, and allowing the use of perturbation theory [17–19]. At large distances (low energy scales), the coupling constant increases, leading to confinement. For these energy scales, perturbation theory is no longer applicable, and methods like lattice **QCD** are used.

A practical consequence of this behavior is the formation of jets in hadronic events. When a pair of color-charged particles is produced and moves apart, the strong coupling constant increases with distance. As a result, the energy stored in the field eventually becomes large enough to create new particle-antiparticle pairs. In high-energy strong interaction processes, this leads to the production of multiple collimated particles, which cluster into collimated structures known as jets.

3.1.2.2 Electroweak Interaction

The [electroweak \(EWK\)](#) interaction is introduced by postulating local gauge invariance under transformations of the $SU(2)_L \times U(1)_Y$ group [20–22]. The $U(1)_Y$ symmetry is generated by the conserved quantum number hypercharge Y and has dimension one. The associated gauge field is denoted with B_μ , and the coupling constant with g' . Using these definitions, the corresponding terms in the Lagrangian density read:

$$\mathcal{L}_Y = i\bar{\psi}\gamma^\mu\partial_\mu\psi + \frac{g'Y}{2}\bar{\psi}\gamma^\mu\psi B_\mu - \frac{1}{4}F^{\mu\nu}F_{\mu\nu}, \quad (3.9)$$

where the field strength tensor of the hypercharge symmetry is given by $F_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$, reflecting the abelian nature of the symmetry group.

The $SU(2)$ group, on the other hand, encodes the weak isospin symmetry (T) and has dimension three, corresponding to the three gauge fields $W_\mu^{1,2,3}$. A set of generators are the Pauli matrices $\sigma^{1,2,3}$, whose structure constants are described by the totally antisymmetric Levi-Civita symbol $f^{abc} = \epsilon^{abc}$. With the coupling constant g , the Lagrangian density can be written as:

$$\mathcal{L}_L = i\bar{\psi}\gamma^\mu\partial_\mu\psi + \frac{g}{2}\bar{\psi}\gamma^\mu\sigma^a\psi W_\mu^a - \frac{1}{4}F^{\mu\nu}F_{\mu\nu}, \quad (3.10)$$

with field strength tensor

$$F_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g f^{abc} W_\mu^b W_\nu^c. \quad (3.11)$$

Similar to spin, creation and annihilation operators can be constructed, by convention changing the third component of the weak isospin by ± 1 :

$$\sigma^\pm \equiv \frac{1}{\sqrt{2}}(\sigma^1 \mp i\sigma^2). \quad (3.12)$$

The third generator σ^3 is diagonal and therefore does not change the isospin.

The $SU(2)$ group is conventionally labeled with an index L , indicating that only left-handed fermions couple to the weak-isospin gauge fields in the [SM](#). In this context, “handed” refers to chirality, an intrinsic quantum property of particles that is Lorentz-invariant but not conserved for massive Dirac particles.

This parity-violating behavior is implemented by assigning left-handed fermions to weak isospin doublets, while right-handed fermions are treated as isospin singlets. Consequently, terms with isospin-singlets in Equation 3.9 remain present (except for right-handed neutrinos, as they do not carry hypercharge), whereas those in Equation 3.10 vanish. Mixed contributions, corresponding to the coupling of a left-handed and a right-handed particle, vanish in both equations—correctly reflecting that the gauge interactions preserve chirality.

According to the CPT theorem [23–25], the couplings of left-handed fermions and right-handed antifermions (and vice versa) must be identical under the combined symmetry operations of charge conjugation (C), parity transformation (P), and time reversal (T). This symmetry is reflected in the fact that the CPT conjugates of fermions (i.e., antifermions with opposite chirality) appear in the theory with the same weak interaction structure. Thus, right-handed antifermions transform as isospin doublets, while left-handed antifermions transform as isospin singlets.

3.1.3 Higgs Mechanism

The SM as constructed up to this point is fully symmetric with respect to the assumed symmetry groups. However, this symmetry requires all vector bosons to be massless fields, which is in strong contradiction to the observed massive W^\pm and Z bosons. The Higgs mechanism [26–29] can solve this contradiction by introducing additional fields and breaking the electroweak $SU(2)_L \times U(1)_Y$ symmetry down to an effective $U(1)_{\text{em}}$ symmetry.

3.1.3.1 Electroweak Symmetry Breaking

The Higgs field is introduced as a complex scalar doublet of the weak isospin:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (3.13)$$

where ϕ^+ corresponds to an electrically charged complex Higgs field component, whereas ϕ^0 describes an electrically neutral complex Higgs field component.

The Higgs field enters the Lagrangian density via the term:

$$\mathcal{L}_{\text{Higgs}} = \underbrace{(D_\mu \phi)^\dagger (D^\mu \phi)}_{\mathcal{L}_{\text{Higgs,kin}}} + \underbrace{\mu^2 \phi^\dagger \phi - \lambda (\phi^\dagger \phi)^2}_{-V(\phi)}, \quad (3.14)$$

where D_μ describes the covariant derivative of the electroweak interaction, μ and λ are real positive parameters, $\mathcal{L}_{\text{Higgs,kin}}$ is the kinetic term of the Higgs Lagrangian density, and $V(\phi)$ is the Higgs potential.

Due to the specific shape of the Higgs potential, which is illustrated in a simplified form in Figure 3.2, the symmetry point of the potential does not coincide with its ground state. Consequently, it is energetically beneficial to break the symmetry of the system. As the gradient vanishes at the symmetry point, there is no preferred direction, making the symmetry point a meta-stable state rather than a true ground state. However, in a spontaneous process, the symmetry can be broken, lowering the potential energy and leaving behind a stable state. The system's ground state no longer occurs as the symmetry point, but at one of the points along the circular symmetry of the potential, characterized by a non-zero **vacuum expectation value (VEV)**. After the symmetry breaking, the Higgs field can be described in terms of a single radial degree of freedom and is often brought to the following form using $SU(2)$ gauge transformations:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix}, \quad (3.15)$$

where $v = \sqrt{\frac{\mu^2}{\lambda}}$ is the **VEV** of the Higgs field, corresponding to the point, where the potential is minimized, and h reflects the remaining radial degree of freedom that corresponds to the Higgs particle. The choice of an electrically uncharged **VEV**, where v is a real non-zero value, reflects the fact that the universe is electrically neutral in its vacuum state. This condition is consistent with the observation that the electromagnetic interaction remains unbroken, as it is necessary for the mass of the photon to vanish.

In this broken state, the kinetic term of the Lagrangian density reads:

$$\begin{aligned} \mathcal{L}_{\text{VEV, kin}} &= (D_\mu \phi)^\dagger (D^\mu \phi) \\ &\supseteq \left| \left(\frac{i}{2} g \sigma^k W_\mu^k + \frac{i}{2} g' B_\mu \right) \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \right|^2 \\ &= \frac{v^2}{8} \left[g^2 ((W_\mu^1)^2 + (W_\mu^2)^2) + (g W_\mu^3 - g' B_\mu)^2 \right], \end{aligned} \quad (3.16)$$

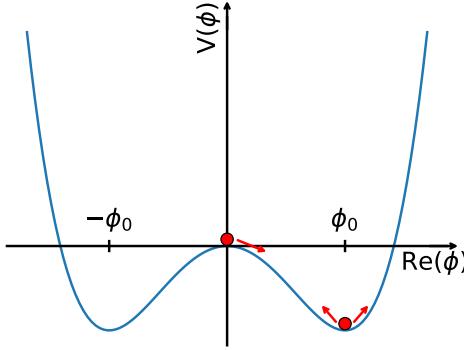


Figure 3.2: Higgs potential simplified to one dimension. The symmetry point is in the center of the image at $\phi = 0$. The symmetry can be broken spontaneously, when the state moves to the energetically favored minimum at $\phi_0 = v/\sqrt{2}$, called **VEV**. From here, the Higgs field can oscillate around the **VEV**, as indicated by the small red arrows.

where terms related to the Higgs field h are not considered as they are not important for the construction of the masses of the vector bosons. As the **VEV** does not depend on the space-time position, the derivatives vanish. The result can be interpreted as mass terms for the fields of the physically observed W^\pm bosons and the Z boson:

$$\begin{aligned} W_\mu^\pm &\equiv \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2), \\ Z_\mu &\equiv \frac{1}{\sqrt{g^2 + g'^2}}(gW_\mu^3 - g'B_\mu). \end{aligned} \tag{3.17}$$

By comparison with the general mass term for a spin-1 field B_μ , which reads $m^2 B_\mu^\dagger B^\mu$, the mass of the W^\pm bosons is $m_W^\pm = \frac{gv}{2}$, while the mass of the Z boson is $m_Z = \frac{v}{2}\sqrt{g^2 + g'^2}$.

The massless photon field is finally given by the linear combination of W^3 and B which is orthogonal to the Z boson field:

$$A_\mu \equiv \frac{1}{\sqrt{g^2 + g'^2}}(g'W_\mu^3 + gB_\mu), \tag{3.18}$$

which is still massless as there is no mass term in the Lagrangian density and thus represents the fully conserved $U(1)_{\text{em}}$ symmetry.

Finally, the remaining terms from the Higgs field in the Lagrangian density are considered. The Higgs potential gives rise to the mass of the Higgs boson; for this only terms of order h^2 are taken into account:

$$\begin{aligned} V(h) &= -\frac{1}{2}\mu^2(v + h)^2 + \frac{1}{4}\lambda(v + h)^4 \\ &= \lambda v^2 h^2 + \mathcal{O}(\neq h^2), \end{aligned} \tag{3.19}$$

corresponding to a Higgs field with mass $m_h = \sqrt{2\lambda}v$. Remaining terms give rise to interactions between the Higgs field, the **VEV** and the gauge fields, which will be further explained in Section 3.2.

3.1.3.2 Yukawa Coupling

As mentioned before, left-handed fermion fields transform as an isospin doublet under $SU(2)_L$, whereas right-handed fermion fields transform as a singlet. In general, the Dirac spinor consists of both left- and right-handed parts:

$$\psi = \psi_L + \psi_R, \quad (3.20)$$

where the individual contributions can be extracted using the projection operator, constructed with the gamma matrix γ^5 :

$$\begin{aligned} \psi_L &= \frac{1}{2}(1 - \gamma^5)\psi, \\ \psi_R &= \frac{1}{2}(1 + \gamma^5)\psi. \end{aligned} \quad (3.21)$$

Consequently, the naive mass term in the Lagrangian density, which would reconstruct the Dirac equation for a massive fermion, can be manipulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{mass}} &= -m\bar{\psi}\psi \\ &= -\frac{m}{4}\bar{\psi}((1 - \gamma^5) + (1 + \gamma^5))((1 - \gamma^5) + (1 + \gamma^5))\psi \\ &= -\frac{m}{4}\bar{\psi}\left[(1 - \gamma^5)^2 + (1 + \gamma^5)^2\right]\psi \\ &= -m[\bar{\psi}_R\psi_L + \bar{\psi}_L\psi_R], \end{aligned} \quad (3.22)$$

since $\gamma^5\gamma^5$ corresponds to the identity matrix. Here, it is important to note that $\bar{\psi}_R$ transforms as an isospin singlet, whereas ψ_L transforms as an isospin doublet. As a result, the corresponding cross terms are not invariant under local $SU(2)_L$ gauge transformation and thus violate gauge symmetry. To restore gauge invariance and allow for consistent fermion mass generation, the naive mass term is typically replaced by the Yukawa interaction term [30]. This mechanism couples the fermion field ψ to the Higgs field ϕ , and the resulting Yukawa mass term takes the following form:

$$\mathcal{L}_{\text{Yukawa}} = -y_f\bar{\psi}_L\phi\psi_R + \text{h.c.}, \quad (3.23)$$

where y_f represents the Yukawa coupling constant associated with the fermion type in question, and ‘h.c.’ denotes the Hermitian conjugate terms. In contrast to the naive mass terms, this Yukawa interaction preserves gauge invariance, because the Higgs field transforms as an isospin doublet and therefore provides the necessary counterpart for the fermion doublet. After symmetry breaking, the Higgs field acquires its non-zero **VEV**. Inserting this into the Yukawa interaction term results in a simplified expression:

$$\mathcal{L}_{\text{Yukawa}} = -\frac{y_f v}{\sqrt{2}}\bar{\psi}_L^d\psi_R + \text{h.c.}, \quad (3.24)$$

where $\bar{\psi}_L^d$ describes the second element of the left-handed fermion doublet doublet $\bar{\psi}_L$ that couples to the right-handed singlet ψ_R . By comparison with the standard mass terms, this leads to the generation of fermion masses of the size:

$$m_f = \frac{y_f v}{\sqrt{2}}. \quad (3.25)$$

3.1.3.3 Quark Generation Mixing

A generalization from one quark generation to the observed number of three generations can be implemented by absorbing the Yukawa coupling and the Higgs **VEV** into a mass matrix M . Then, the Yukawa term for down-type quarks can be written as:

$$\mathcal{L}_{\text{Yukawa}} = -\bar{q}_L M_d q_R, \quad (3.26)$$

where $q_{L/R} = (d_{L/R} \ s_{L/R} \ b_{L/R})^T$. As there is no underlying principle prohibiting interactions with the Higgs field between different generations, the mass matrix is, in general, not diagonal. Introducing two well-chosen unitary matrices V_d and U_d , the Yukawa term can be expanded to

$$\mathcal{L}_{\text{Yukawa}} = -\underbrace{\bar{q}_L}_{\bar{q}_L^m} \underbrace{V_d^{\dagger}}_{M_d^{\text{diag}}} \underbrace{M_d}_{M_d^{\text{diag}}} \underbrace{U_d^{\dagger}}_{q_R^m} q_R, \quad (3.27)$$

where M_d^{diag} is the diagonalized mass matrix, and $q_{L/R}^m$ the corresponding mass eigenstates. An analogous calculation can be performed for the mass terms of up-type quarks. If the unitary matrices are different from identity, then this means that the mass eigenstates—the physically observed quarks—are different from the eigenstates of the Lagrangian before the electroweak symmetry breaking.

To assess the implication of this observation, the interaction terms of the Lagrangian can be investigated. Due to the unitary property of the matrices, they cancel in interactions that contain the same particle twice if the interaction is independent from the quark generation. This is the case for the strong interaction, as well as for the part of the electroweak interaction mediated by Z boson and photon.

However, interactions mediated by W^\pm bosons contain mixed terms of an up-type and a down-type quark. While the W^\pm coupling does not depend on the generation, the mixed terms lead to a new contribution, the **Cabibbo-Kobayashi-Maskawa (CKM)** matrix V_{CKM} [31, 32]:

$$\bar{q}_u X q_d = \bar{q}_u^m V_u^{\dagger} X V_d q_d^m = \bar{q}_u^m X \underbrace{V_u^{\dagger} V_d}_{V_{\text{CKM}}} q_d^m, \quad (3.28)$$

where $q_{u/d}$ denote three-vectors containing the left-handed up-/down-type quark fields, respectively. Note that the **CKM** matrix itself is unitary again and therefore allows for rigorous consistency tests by measuring its parameters.

To summarize, the introduction of the fermion masses leads to generation mixing in the charged currents of the weak interaction. This means that for interactions mediated by W^\pm bosons, there is a certain probability for interactions between quark mass eigenstates from different generations. The size of the **CKM** matrix elements dictates this probability and is generally small for transitions to and from the third generation, and larger for transitions between the first and second generation. However, diagonal elements are overall dominant, indicating that transitions between different quark generations are generally suppressed.

3.2 Theoretical Predictions in the Standard Model

To probe the **SM**, it is crucial to derive experimentally accessible predictions. While some of the parameters in the **SM** can be tested for consistency, e.g., the parameters of the **CKM** matrix, a main connection between theory and experiments is the probability of certain interactions.

3.2.1 Cross Section and Decay Width

Cross sections and decay widths both describe transitions between different states. While the cross section quantifies interactions between two (or more) particles in the initial state, the decay width is a property of unstable states. A common feature of both quantities is their dependence on the probability amplitude \mathcal{M}_{fi} for the transition from an initial state i to a final state f , which will be explained in greater detail in the next section.

A cross section is the rate of a defined set of scattering processes in relation to the effective flux. A differential element of the cross section can be expressed in terms of the flux Φ , a transition amplitude \mathcal{M} , and the differential phase space volume $d\Pi$:

$$d\sigma(i \rightarrow f) = \frac{1}{\Phi} |\mathcal{M}_{fi}|^2 d\Pi. \quad (3.29)$$

Notably, four-momentum conservation leads to a restriction of the phase space which is available to the final-state particles. By performing the integration over the phase space, all allowed combinations of energy and momentum in the final state are considered.

If the scattering process proceeds via an intermediate, unstable particle, the cross section shows a resonant behavior. The width Γ of this resonance is inversely proportional to the particle's lifetime. If this width is much smaller than the mass of the intermediate particle, the resonance can be approximated by a delta distribution. Under this ‘narrow width approximation’ [33] the cross section factorizes into the production of the intermediate particle X and its subsequent decay:

$$\sigma(i \rightarrow X \rightarrow f) \approx \sigma(i \rightarrow X) \mathcal{B}(X \rightarrow f), \quad (3.30)$$

where $\sigma(i \rightarrow X)$ denotes the production cross section, and $\mathcal{B}(X \rightarrow f)$ is the branching fraction describing the relative probability for the decay of X to the final state f .

Branching fractions can be determined by comparing the decay width Γ_f of a specific final state to the total decay width:

$$\mathcal{B} = \frac{\Gamma_f}{\sum_n \Gamma_n}. \quad (3.31)$$

Decay widths are intrinsic properties of particles and quantify the decay probability per time interval. The partial width for a specific decay mode f is defined by:

$$\Gamma_f = \frac{1}{2m} \int |\mathcal{M}_{fX}|^2 d\Pi, \quad (3.32)$$

where m is the mass of the intermediate particle, and \mathcal{M}_{fX} and $d\Pi$ are defined as before.

The transition amplitude \mathcal{M} plays a crucial role in both scattering and decay processes. However, due to the complexity involved in its full calculation, approximations are usually necessary. In high-energy physics, the most widely adopted approximation framework is perturbation theory, which treats interactions as small corrections to a simpler system.

3.2.2 Perturbation Theory

Since the Lagrangian density provides a comprehensive description of the dynamics of states in the SM, the time evolution of any state can, in principle, be derived from it. When examining transitions, the kinetic and mass terms of the Lagrangian density play a relatively minor role, as they describe the propagation of a free particle between interactions. The remaining terms in the Lagrangian density are collectively referred to as interaction Lagrangian \mathcal{L}_{int} . The time evolution of any quantum state $|\psi(t)\rangle$ is given by the time dependent Schrödinger equation:

$$i\partial_t |\psi(t)\rangle = \hat{H} |\psi(t)\rangle, \quad (3.33)$$

where \hat{H} is the Hamilton operator and natural units ($\hbar = 1$) are used. Assuming the Hamilton operator is generally time-dependent and the time dependence of the state can be extracted into a function $U(t, t_0)$, the simplified Schrödinger equation reads:

$$i\partial_t U(t, t_0) = \hat{H}(t)U(t, t_0). \quad (3.34)$$

Integrating over the time interval $[t_0, t]$ gives:

$$U(t, t_0) - U(t_0, t_0) = \int_{t_0}^t \hat{H}(t')U(t', t_0)dt'. \quad (3.35)$$

Since the initial condition is $U(t_0, t_0) = 1$, substituting this equation into itself iteratively reveals a time-ordered exponential series, which can be compactly written as:

$$U(t, t_0) = \hat{\mathcal{T}} \exp\left(-i \int_{t_0}^t \hat{H}(t')dt'\right). \quad (3.36)$$

In this discussion, certain computational details—such as the time ordering of integrals and calculations in the interaction picture—are neglected for clarity. However, since the Hamilton operator can be expressed in terms of the Lagrangian density, the key conceptual point remains: the time evolution of a quantum state is expressed as an infinite series, fundamentally governed by the interactions defined in the Lagrangian density.

Finally, the transition amplitude for the transition from an initial state $|i\rangle$ to a final state $\langle f|$ is given by:

$$\mathcal{M} = \langle f| U(t, t_0) |i\rangle. \quad (3.37)$$

The complexity in the calculation of this transition is determined by how often the interaction terms from the Lagrangian density appear in the series expansion: The leading non-trivial contribution, is referred to as [leading order \(LO\)](#). Subsequent corrections are called [next-to-leading order \(NLO\)](#), [next-to-next-to-leading order \(NNLO\)](#), and so on, based on the order of the term in the expansion. Since higher-order terms refine the result, the precision of a calculation generally improves as more terms are included. However, each additional term introduces more occurrences of the interaction terms from the Lagrangian, making computations increasingly complex and resource-intensive.

Each term in the series can be graphically illustrated as a Feynman diagram by identifying its constituents with a certain symbol. In such a diagram, each straight line corresponds to a fermion, where an arrow indicates whether it is a particle or an antiparticle. Vector bosons of the electroweak interaction are depicted as curvy lines, while gluons are shown as curls, and the Higgs boson as a straight dashed line. Couplings between the different fields are depicted as points.

3.2.3 Details of Proton-Proton Collisions

Protons are compositions of quarks and gluons. Since the description of the proton is not deterministic, the collisions of protons cannot be described deterministically either.

3.2.3.1 Parton Distribution Functions

A [parton distribution function \(PDF\)](#) $f(x, Q^2)$ provides the probability density to find a specific parton within the proton carrying a momentum fraction x of the proton's total momentum, where the squared four-momentum transfer is given by Q^2 . Figure 3.3 shows the different contributions at momentum transfers typical for collisions at the [Large Hadron Collider \(LHC\)](#). At large values of x , the valence quarks dominate. However, due to vacuum fluctuations, pairs of quarks and antiquarks can be created. The smaller the momentum share, the larger the probability to find such quark pairs. Similarly, the probability to find a gluon that is being radiated off rises with decreasing momentum share. At a typical [LHC](#) center of mass energy of 13.6 TeV the creation of heavy particles such as the W^\pm , Z or Higgs boson with roughly 100 GeV mass requires a momentum share of less than 1 %. At such momentum shares, gluons are the dominating partons. Consequently, at the typical energy scales of the [LHC](#), whenever allowed in terms of conserved quantities, the production in the gluon channel usually dominates.

This is, for example, the case in Higgs boson production, where the dominant channel is gluon fusion via a top-quark loop. In contrast, gluon-fusion production of W^\pm and Z bosons requires more complicated diagrams to conserve both angular momentum and the electroweak charges. As a result, this contribution to W^\pm and Z boson production is suppressed. Furthermore, W^\pm and Z boson production is not suppressed by small fermion masses, allowing significant contributions from valence quarks. This leads to a preference for W^+ over W^- production at the [LHC](#), since up quarks are more abundant in the proton than down quarks.

In general, the dependence of the [PDFs](#) on the scale Q^2 is predicted by perturbative [QCD](#), but there is no theory describing the behavior with respect to the momentum share x . Consequently, the currently best way to find this behavior is by measurement. Because of the known dependency on Q^2 , measurements at different energies can be combined to put further constraints on the unpredicted part. The uncertainties connected to these measurements, as well as the uncertainties in the theoretical descriptions have to be considered when calculating the total cross section.

Ultimately, the total cross section for the production of a state f in proton-proton (pp) collisions can be expressed as:

$$\sigma_{pp \rightarrow f} = \sum_i \sum_j \int_0^1 dx_i \int_0^1 dx_j f_i(x_i, \mu_F) f_j(x_j, \mu_F) \int d\sigma_{ij \rightarrow f}(\mu_R, \mu_F). \quad (3.38)$$

Here, μ_F and μ_R denote the factorization and renormalization scales, respectively. The factorization theorem provides a framework to separate short-distance (perturbative) from long-distance (non-perturbative) [QCD](#) effects, where the artificial boundary between both domains is introduced via by μ_F . In the perturbative regime, [QCD](#) loop corrections can lead to divergences, which are handled through renormalization at the scale μ_R .

Both μ_F and μ_R are unphysical scales introduced as part of the corresponding mathematical procedures. Physical observables should, in principle, not depend on their specific values. However, since the calculations are typically performed at finite order in perturbation theory, a residual dependence remains.

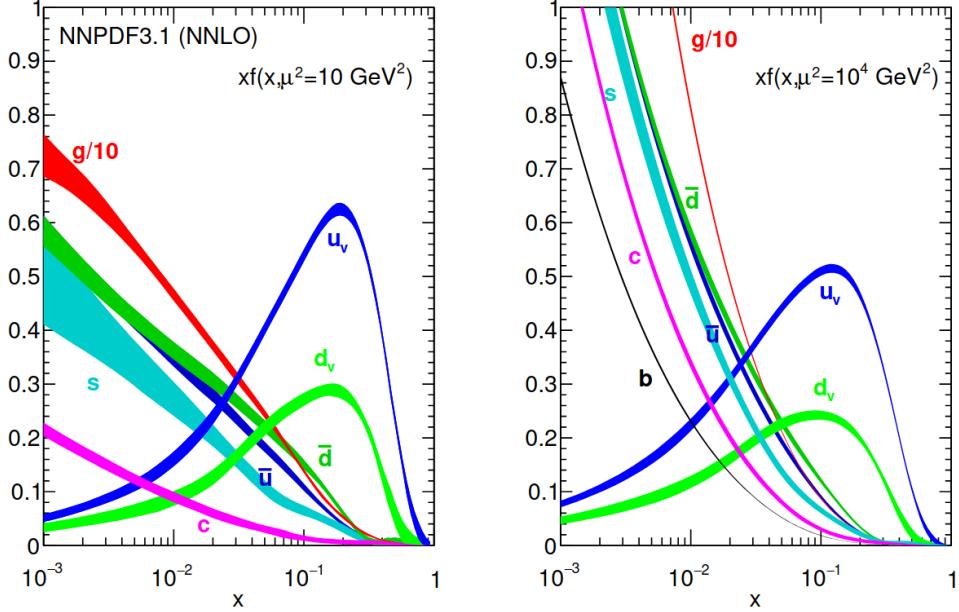


Figure 3.3: Proton PDFs evaluated with the NNPDF3.1 package at NNLO, taken from Reference [34]. On the left, the squared momentum transfer is relatively small with 10 GeV^2 , while it is close to the mass scale of the heavy SM bosons on the right with 10^4 GeV^2 . At smaller momentum fractions x , gluons are the dominant contribution (note that they are scaled down by a factor of ten). At larger momentum fractions x , the valence quark's contribution is dominant and reflects the presence of two up and one down quark.

3.2.3.2 Shower and Hadronization

Similarly to the challenges in the description of the partons before the collisions, describing the behavior of partons after the interaction poses a challenge as well. The reason for that is the running coupling constant of the strong interaction, which does not allow single color charges to exist. Instead, the energy in the field between the color-charged particles increases as their distance grows, ultimately leading to the creation of new color-charged particles.

The theoretical description of this process contains two steps: parton shower modeling and hadronization. In the first step, the parton shower is evolved from the energy scale of the process down to the energy scale of hadronization (infrared cut-off), typically chosen at 1 GeV . This is done by introducing [initial-state radiation \(ISR\)](#) and [final-state radiation \(FSR\)](#), where [ISR](#) describes the radiation of partons before, and [FSR](#) after the hard process. Then the shower is evolved by describing the splitting into two partons, with a certain probability following the Sudakov form factor that represents the probability of a parton not emitting radiation during the evolution between two energy scales.

When the hard process is generated at higher orders in perturbative [QCD](#), partons may be accidentally double-counted. As an example, consider the production of a W^+ boson with one jet in the final state, as illustrated in Figure 3.4. In the simulation, there are three different contributions to this process:

- as [LO](#) term in an event with 0 hard partons and 1 parton shower
- as [NLO](#) term in an event with 0 hard partons and 0 parton showers
- as [LO](#) term in an event with 1 hard parton and 0 parton showers

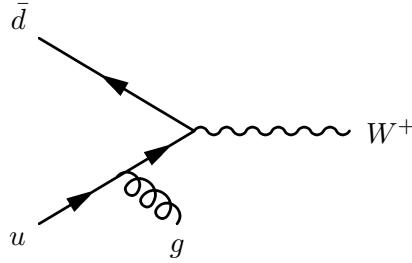


Figure 3.4: Example Feynman diagram for the production of a W^+ boson with emission of a gluon in the initial state.

Overlaps in the descriptions need to be accounted for, which is usually done by applying matching and merging algorithms. While matching algorithms addresses double-counting through parton shower modeling (corresponding to the first two bullet points), merging algorithms reduce the impact of double-counting from different generated processes and is typically applied after the parton matching.

At the energy scale of 1 GeV and below, the parton behavior starts to be dominated by non-perturbative effects, eventually leading to the formation of color-neutral hadrons. Modeling of these non-perturbative effects requires a certain level of simplification. There are different approaches to deal with this analytically (so far) unsolvable problem. Two commonly used approaches are string models, such as the Lund model [35, 36], and clusters models [37–39].

All simulated samples used in this thesis rely on string models, which focus on the linear behavior of the strong coupling constant for large distances. Similar to the lines of an electromagnetic field between two charges, the [QCD](#) field is built between two color-charged particles, and simplified to a one-dimensional representation (string). With increasing distance, the energy in the field is increased and at a certain point, a new pair of color-charges is created, similar to a rubber band, which tears when pulled apart sufficiently. Cluster models, on the other hand start by splitting all gluons at the hadronic energy scale into quarks and then cluster the different quarks to color-neutral objects. The clusters then decay into hadrons, where the available phase space dictates the behavior.

Since the final state particles studied in this thesis are muons and neutrinos, the dominant [FSR](#) effects are [quantum electrodynamics \(QED\)](#) radiation processes connected to the muons. These are analogously to [QCD](#) emissions, but are conceptually simpler as the photons do not self-interact. Consequently, neither cut-off scales nor hadronization models are required for [FSR](#) in this context, and the propagation can be described entirely with perturbative methods. The corresponding radiation behavior is modeled using Sudakov form factors, similar to the treatment of [QCD](#) radiation.

3.2.3.3 Underlying Event and Pileup

In the theoretical description of [pp](#) collisions, proton remnants must also be considered. Their behavior, along with soft interactions at low momentum transfers (such as multi-parton interactions), is collectively referred to as underlying event. Due to the non-perturbative nature of [QCD](#) at low energy scales, these processes cannot be fully described using first-principle calculations. Instead, phenomenological models are used, whose parameters are tuned to match experimental data and theoretical calculations in the perturbative regime at higher momentum scales.

Additional low-momentum pp interactions, which arise in the high-luminosity regime of modern particle accelerators, are called pileup. In simulations, these additional events, which are totally independent from the hard process, are considered by adding so-called minimum bias events.

3.2.4 Uncertainties

To assess possible influences of the theoretical modeling on the result, different approaches are followed. The dependence on the choice of μ_R and μ_F on the cross section is evaluated by varying both values independently by a factor of two, where the extreme case of opposite variations is omitted. Then, the envelope is taken as an estimate of the impact of the choice of scale on the theoretical prediction. A similar approach is followed to estimate the uncertainty in the parton shower scales.

As the [PDF](#) parameters are obtained by fitting models to data, there is an inherent uncertainty from this measurement. Such uncertainties can be considered using different approaches. The most common among these are the [Monte Carlo \(MC\)](#) approach and the Hessian approach. In the [MC](#) approach, a large number (typically at least 100) of [PDF](#) replicas is provided, each corresponding to a random fluctuation. The prediction of an observable then corresponds to the mean value over all results using the different replicas, and the uncertainty is given by the standard deviation. In the Hessian approach, one central value is provided, along with a (typically smaller) number of disentangled members, where correlations between the different variations are removed by orthogonalizing the variation matrix. Consequently, the total uncertainty is evaluated by summing the squared differences between the central value and the evaluation with the different member.

Additionally, the [PDFs](#) are subject to variations depending on the treatment of α_S in the fit to the data. The corresponding uncertainties are incorporated by evaluating the [PDF](#) with a central value of $\alpha_S(m_Z) = 0.118 \pm 0.002$.

3.3 Limitations of the Standard Model

Despite its huge success, the **SM** is not consistent with all theoretical observations. A well-known phenomenon is the oscillation of neutrinos, which was observed for solar and atmospheric neutrinos. The most practical solution to this problem is to assume that the neutrinos have non-negligible masses. Thus, the eigenstates in which the neutrinos are created, which are the eigenstates of the electroweak interaction, could be shifted with respect to the mass eigenstates, in which the neutrinos propagate as long as they do not interact weakly. Since the neutrinos are considered massless in the **SM**, their oscillations showcase one of its limitations. However, measurements of the neutrino mass have found an upper limit of 0.45 eV at 90 % confidence level [40], which is more than five orders of magnitude smaller than the mass of the next heavier elementary particle, the electron. Thus, the impact of this mass in measurements done within the scope of this thesis is negligible.

Furthermore, there is strong evidence for the presence of additional matter in the universe, which is not part of the **SM**. Such evidence can be found on different orders of magnitude, from the rotation velocity of galaxies [41, 42] to gravitational observations in galaxy clusters [43, 44] up to the structure formation of the whole universe [45, 46]. However, there has been no observation of this dark matter yet, indicating cross sections in the order of or smaller than those of the weak interaction. Due to the small interaction probabilities, this potential limitation of the **SM** does not affect the precision of the analyses presented in this thesis.

Finally, there could exist other effects that are incompatible with the **SM**. One option to figure out effects despite their likely small size is to perform highly precise measurements. If there is a deviation from the expected behavior, a new theory might be needed. This is the approach followed in this thesis, which measures processes in the **SM** at very high precision in order to possibly find deviations from the expectations based thereon.

4 Experimental Environment

For millennia, investigating the building blocks of matter has sparked human interest both in a scientific and philosophical sense. But with smaller structures, the creation of additional knowledge became increasingly challenging. The main reason for that is the resolution limit, which is mainly driven by the effective length of the waves used for measuring a structure in comparison to the size of the structure. For objects with a size in the order of micrometers, such as human cells, visible light can be used as the wavelength is below the size of the structure. While it is possible to investigate smaller objects with light by decreasing the wavelength even further, it is less technically challenging to utilize the wave properties of matter. Any particle with momentum p can be assigned a wavelength λ given by the de-Broglie formula [47]:

$$\lambda = \frac{2\pi}{p}. \quad (4.1)$$

For electrons at 50 % of the speed of light, the corresponding wavelength is already in the nanometer domain. In order to study even smaller scales, the momenta of the particles need to be increased. This ultimately led to the rise of particle colliders, which make use of the kinematic benefit in the collision energy, introduced by the moving target. Such colliders allowed for the observation of all the fundamental particles required to experimentally confirm the full particle content of the [SM](#).

The era of accelerators culminated (for now) with the construction of the [LHC](#) at the [Conseil européen pour la recherche nucléaire \(CERN\)](#) in Geneva, and the observation of the Higgs boson in 2012 [48, 49] with the [Compact Muon Solenoid \(CMS\)](#) and [ATLAS](#) experiments. Section 4.1 explains the experimental environment provided by the [LHC](#). The [CMS](#) detector is one of the experiments used to investigate collisions produced by the [LHC](#), and is discussed in greater detail in Section 4.2. This chapter is mainly based on the original technical design reports [50, 51], and the corresponding updates with respect to the high-luminosity era [52–55].

4.1 The Large Hadron Collider

The [LHC](#) at [CERN](#) accelerates protons in two opposing directions up to energies of 6.8 TeV. Equally distributed around the accelerator ring are eight interaction points, four of which are dedicated to collisions. Here, the two beam lines intersect each other, and the products of the $\sqrt{s} = 13.6$ TeV-collisions are measured with detectors. The other four interaction points are dedicated to acceleration, beam cleaning, and beam dump.

These high proton energies pose multiple challenges to the design of the [LHC](#). In order to guide the protons around the circular-shaped accelerator, powerful dipole magnets are required. Furthermore, bremsstrahlung is emitted during the deflection process, leading to significant radiation emission. To limit both the energy loss due to radiation and the power requirements for the dipole magnets, the [LHC](#) was built with a large circumference of 26.7 km.

The effective size of the proton, with a radius of roughly 1 fm, is very small compared to the beam pipe with a diameter of roughly 5 cm. To put this into perspective, the size ratio corresponds to that of a marble in comparison with Earth's trajectory around the Sun. In order to collect a sizable amount of data, not single protons but bunches of about 10^{11} protons each are brought to collision at high rates.

The number of bunches per beam is limited by three aspects: first, the 400 MHz frequency of the radiofrequency cavities, which are used to accelerate the protons. As each oscillation of the electric field can accelerate one bunch, the minimal spacing between two bunches is limited by this criterion to 75 cm.

Second, the detector readout rate is limited since a certain minimal amount of time is required for the collision products to produce a signal and for its subsequent processing. The detectors were built with the requirement to analyze events at a rate of 40 MHz. Consequently, only every tenth oscillation of the radiofrequency cavities is occupied with a bunch and the effective minimum distance between two bunches corresponds to 7.5 m.

Third, for technical reasons, additional empty bunch positions are required, leading to the formation of bunch trains (consecutive lines of bunches). The distance between trains is mostly governed by the rise-time of injection kickers—magnets leading the trains from the pre-accelerator into the [LHC](#)—as well as their limited operation time. Bunches that are positioned in the rise-time of the injection kicker, would not be optimally guided into the [LHC](#) and could potentially destroy the technical equipment. One of the gaps between trains needs to be even larger in order to match the rise-time of the dump kicker, which is aimed at quickly emptying the beam by guiding the protons out of the accelerator. These limitations lead to a reduction of the maximum potential filling number by about a third, depending on the actual filling scheme [56]. Most notably for this thesis, however, the train structure of the bunches in the [LHC](#) can be used for studying dead-time effects of the detectors; for the first bunch in a train there are no immediate preceding bunches, ensuring that the activity in the detector is decreased to a minimum.

Different protons within a bunch repel each other as a consequence of their positive electric charge. The resulting elongation and broadening of bunches needs to be addressed to run the experiment at full efficiency for hours. To reduce the spread transverse to the beam axis, quadrupole magnets re-focus the beam.

Elongation, on the other hand, is avoided by a special setup in the radiofrequency cavities: The [LHC](#) cavities generate a voltage of 8 MeV oscillating at the aforementioned 400 MHz frequency. Each proton bunch experiences acceleration over half of an oscillation cycle, ensuring that it passes through the cavity only when the electric field is oriented in the

correct direction to provide acceleration. By shifting the phase of the oscillating electric field in the cavity so that it aligns with the back of the bunches rather than their center, the slower protons receive greater acceleration than the faster ones. This differential energy gain compresses the bunch in the longitudinal direction, reducing its spread and improving beam stability.

The central quantity describing the rate of collisions in particle colliders is the instantaneous luminosity L . More specifically, it describes how often particles are brought together in a certain area. Large luminosity values can be achieved at the LHC by increasing the number of protons per bunch, the frequency of bunch intersections, or by focusing the beam. However, through the collisions, the number of protons per bunch declines gradually over time, leading to a decreasing instantaneous luminosity. If the instantaneous luminosity falls below a certain threshold, it becomes beneficial to dump the beam and start a new filling procedure.

The integrated luminosity $L_{\text{int}} = \int L dt$ is a measure of the total amount of data collected in a period of time. The number of events of a certain process can be calculated as the product of its cross section σ and the integrated luminosity:

$$N = \sigma L. \quad (4.2)$$

While the accelerator provides the collisions, the experiments record the collision data. If an experiment is for some reason not running or taking faulty data, then the luminosity recorded by that experiment may be smaller than the luminosity provided by the LHC. To maximize the scientific output of the collisions provided, the LHC has to be run in close cooperation with the experiments.

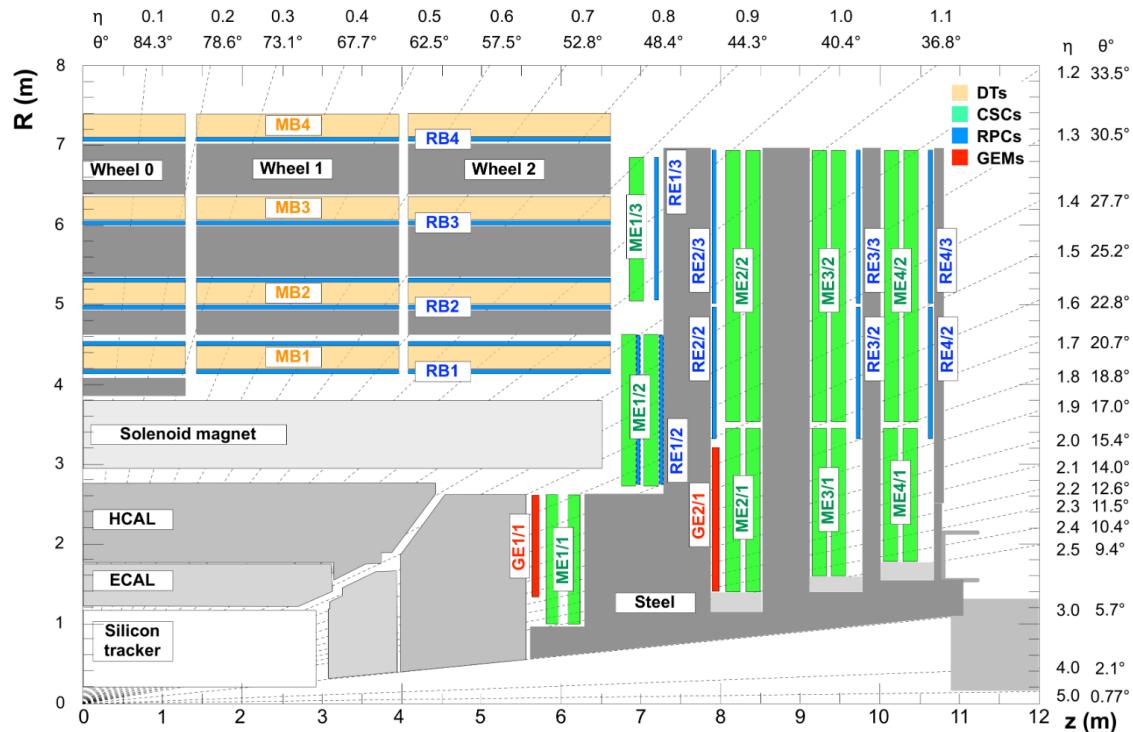


Figure 4.1: Sketch of the CMS experiment with colored muon system, taken from Ref. [57]. The sketch shows a quadrant of the cross section along the beam axis. The red GE2/1 has only been installed partially for demonstration purposes [58, 59].

4.2 The Compact Muon Solenoid Experiment

The [CMS](#) detector is a multi-purpose experiment aimed at an optimized measurement of the collisions provided by the [LHC](#). Its name refers to three remarkable properties. Despite its large cylindrical shape with a length and diameter of roughly 21 m and 15 m, respectively, it is comparably compact with a mass of 14 000 t. The remaining two components of the name refer to muons, which it can measure exceptionally well, and the large solenoid magnet.

4.2.1 Detector Setup

The detector consists of several layers that are dedicated towards different tasks. At the center of the detector, localized directly around the intersection point is the tracker. The next detection layers are [electromagnetic calorimeter \(ECAL\)](#) and [hadron calorimeter \(HCAL\)](#), followed by the solenoid. The outermost detection layer consists of the muon chambers. These chambers are embedded in the iron yoke, which is designed to return the magnetic field lines. An overview over the layers is provided in Figure 4.1.

4.2.1.1 Solenoid

A large superconducting solenoid magnet generates a magnetic field of 3.8 T in the central part of the detector. In the outer region, beyond the solenoid, the field strength is about 2 T, enhanced by iron return yokes that are installed to guide and amplify the field lines. The magnetic field lines are generally aligned parallel to the beam axis, which forces charged particles to bend in the plane perpendicular to the beam direction, also referred to as transverse plane. This results in helix-shaped trajectories for charged particles, where the circular component of their motion is provided by the Lorentz force F_L acting as the centripetal force:

$$F_c = F_L \Leftrightarrow p_T = qBr, \quad (4.3)$$

with the transverse momentum p_T , particle charge q , a constant magnetic field strength B , and the radius of the circular movement r . This relationship allows the transverse momentum to be estimated based on the curvature of the particle's path, as the curvature is inversely proportional to the radius. Furthermore, the sign of the particle's electric charge determines the direction of the curvature, allowing to distinguish positively and negatively charged particles based on the measurement of the particle's trajectory.

The strength of the magnetic field has to be chosen carefully. If the magnetic field is too weak, high-momentum charged particles will have a small curvature, making it more difficult to distinguish between different momenta and reducing the overall measurement precision. On the other hand, a strong magnetic field could trap charged particles in the central region of the detector. More importantly, extremely strong magnetic fields are challenging to reach and expensive in terms of both material and operating costs. The magnetic field strength of 3.8 T generated by the [CMS](#) solenoid provides a reasonable compromise. Interestingly, the minimum transverse momentum for charged particles to escape the 6 m-diameter solenoid is approximately $p_T > 1.7 \text{ GeV}$. As most events that are relevant to analyses usually contain large transverse momenta, this is not limiting the performance significantly. For particles with transverse momenta much larger than 100 GeV, however, the limitations due to the finite magnetic field strength become significant. This will be further discussed in the tracker description.

4.2.1.2 Tracker System

The tracker system consists of silicon detector modules and aims to measure the trajectory of charged particles from hits in individual tracker layers. As discussed in the previous section, both particle charge and transverse momentum can be obtained from the knowledge of a charged particle's path in the magnetic field. Furthermore, knowing the trajectories of several particles from the same bunch crossing, common vertices can be determined. Subsequently, each trajectory can be assigned to one vertex, improving the discrimination between an interesting collision and uninteresting ones (pileup) within the same bunch crossing. In addition, the vertex information is also useful to identify metastable particles, which exist for a measurable distance and then decay still within the tracker.

The tracker system consists of two subsystems: the pixel and the strip tracker. Both have the same fundamental working principle: a p-n junction under high voltage applied such that the depleted zone at the junction is increased to span over the main part of the detector material. In this zone a passing charged particle may interact with the material creating electron-hole pairs, which are—due to the high voltage—pulled towards the electrodes, where they are read out.

The main difference between pixel and strip detectors consists in the shape of the sensing elements. While the pixel detector consists of small pixels, enabling a precise two-dimensional measurement of the transit point, a strip detector only has a decent resolution in one direction. Consequently, the pixel detector generally has a precision advantage over the strip detector. In return, the strip detector is cheaper to build and has fewer read-out channels, which limits the demand for information throughput and computing power. Therefore, pixel detectors are used in the vicinity of the interaction point, where the improved resolution is crucial for the discrimination of different interaction vertices. Outside of the pixel detector, the cheaper strips are located, primarily aiming at a precise measurement of the curvature in the transverse direction.

The relative resolution of the tracker-based transverse momentum measurement is approximated by the Gluckstern formula [60]:

$$\frac{\sigma_{p_T}}{p_T} \approx \frac{\sigma_x[m] \cdot p_T[\text{GeV}]}{0.3 \cdot B[\text{T}] \cdot L^2[\text{m}^2]} \sqrt{\frac{720}{N+4}}. \quad (4.4)$$

Here, $\sigma_x = \frac{d}{\sqrt{12}}$ denotes the spatial resolution of a single hit measurement, governed by the width d of the detector element. The smaller the width of a single pixel or strip, the more precise the position can be determined and therefore also the curvature of the track. For larger p_T , it becomes increasingly difficult to differentiate between small variations of the curvature as the bending of the tracks decreases. On the other hand, an increased magnetic field strength increases the curvature and therefore makes it easier to distinguish different curvatures. A measurement over an extended length of the tracker L increases the precision as well, as a larger share of the circular path can be measured. Finally, the statistical benefit of more hit measurements is reflected by the number of layers in the tracker N .

While larger tracker dimensions would generally improve the momentum resolution (as more layers can fit inside and the length is increased), this comes at a cost. Increased tracker size reduces the available space for other detector components and adds to the material budget. Both may affect the precision of subsequent measurements outside of the tracker and therefore require careful optimization.

Since the tracker upgrade in 2017, the CMS tracker system consists of four pixel layers and ten strip layers, six of which are double-sided to enhance the resolution of the reconstructed

tracks [61]. Depending on the location within the detector, the material budget of the tracker corresponds to between 0.3 and 2 electromagnetic interaction lengths, and between 0.1 and 0.6 hadronic interaction lengths [62]. This contributes an additional uncertainty to the momentum resolution described in Equation 4.4. As an example, the CMS tracker is designed to achieve relative momentum resolutions for muons in the barrel region ($|\eta| < 0.9$) on the order of 1-2 % for p_T below 100 GeV [63, 64].

4.2.1.3 Calorimeter System

The **ECAL** and the **HCAL** are located in subsequent layers outside of the tracker and aim to measure the energy of particles. Conceptually, this measurement is done by converting the energy of the incoming particles to low-energy photons, whose number gives rise to the initial energy. The precision of this measurement therefore highly depends on the conversion efficiency, i.e., if a particle does not convert all its energy to photons, the measurement is not as precise as possible. Some particles, such as muons at typical energies in collisions at the **LHC**, or neutrinos, cannot be stopped within the calorimeter, leading to an imprecise energy measurement (neutrino basically cannot be measured at all, as motivated in Appendix A.2).

The **ECAL** is aimed specifically at a precise measurement of electron and photon energies [65]. High-energy electrons (and positrons) lose energy primarily through bremsstrahlung interactions. The energy loss dE in a material with respect to the mass path length $X = \rho x$, with mass density ρ and path length x is given by:

$$-\frac{dE}{dX} = \frac{E}{X_0}, \quad (4.5)$$

with the particle's current energy E , and the material-specific radiation length X_0 . Solving the differential Equation 4.5 yields the energy of the particle with respect to the mass path length:

$$E = E_0 \exp\left(-\frac{X}{X_0}\right), \quad (4.6)$$

with initial energy E_0 . The **ECAL** has a depth of about $25 X_0$, corresponding to a remaining energy fraction of 1.39×10^{-11} after traversal. In order to minimize the space requirement of the **ECAL**, it is made of the highly dense but transparent lead tungstate (PbWO_4). High-energy photons primarily undergo pair-creation procedures, where in the vicinity of a recoil partner, such as a nucleus, a photon is converted into an electron-positron pair. The kinetic threshold for this process is slightly larger than two times the electron mass, because the recoil partner needs to take a certain share of the four-momentum as well. As opposed to electrons, photons are converted in the interaction. Consequently, the general behavior in material is not described by the energy loss, but by a conversion probability. The latter depends on the mean free path $\langle l \rangle$, which is closely related to the radiation length via $\langle l \rangle \approx \frac{9}{7} X_0$. The probability that a photon has converted after having passed a length x in a material with mean free path $\langle l \rangle$ is given by:

$$P_{\text{conv}}(x) = 1 - \exp\left(-\frac{x}{\langle l \rangle}\right). \quad (4.7)$$

Consequently, both high-energy electrons and photons initiate the production of large numbers of secondary photons and electrons. When the energy of the secondary photons drops below the threshold of electron pair production, the probability for further interactions decreases drastically, effectively limiting the shower evolution. The resulting low-energy

shower particles are finally captured by the material and their energy is converted to photons through scintillation. As the intensity of the resulting light signal is proportional to the energy of the initial electron or photon, the photons need to be collected and measured. However, the photon intensity is typically small, which requires the usage of fast signal amplifiers in the measurement, such as the avalanche photodiodes and vacuum phototriodes used to measure the photon rate in the [ECAL](#).

The [HCAL](#) aims at measuring the energy content of hadrons. Unlike the [ECAL](#), which consists of single-block crystals, the [HCAL](#) is a sampling calorimeter, where dense layers of absorber material (brass) alternate with less dense layers of scintillator material. Hadronic showers are initiated mainly in the absorber material and lead to electromagnetic sub-showers, which behave similar to those in the [ECAL](#) and are ultimately measured with photo sensors.

The thickness relative to the interaction probability is comparably small for the [HCAL](#); in the direction perpendicular to the beam axis about 6 hadronic interaction lengths. While tracker and [ECAL](#) effectively increase the total stopping power by roughly one hadronic interaction length, the combined interaction length corresponds to a final energy of hadronic particles after having passed the [HCAL](#) of about 1%. As a result of the limited stopping power, remnants of the hadronic shower can lead to signatures in the muon chambers, called hadronic punch-through.

4.2.1.4 Muon Chambers

The outermost layer of the [CMS](#) detector is aimed at measuring muons, which are minimally ionizing particles at typical energies in interesting collisions. This behavior can be explained through their mass, which is roughly 200 times larger than that of the electrons, effectively suppressing the energy loss via bremsstrahlung by several orders of magnitude in comparison. Consequently, muons with moderate p_T in the order of 10 GeV and higher consistently reach the muon chambers (see Appendix A.1).

Since other detectable particles are efficiently stopped in the calorimeter system, the minimally ionizing behavior enables muon identification at high efficiencies. Furthermore, the muon chambers contribute to the track measurement of muons, increasing the effective length of the combined muon tracking system, albeit with substantially smaller spatial resolution than in the tracker. This allows for a more precise measurement, in particular of high- p_T muons, whose curvature is easier to measure over a longer path.

Four different gas-based detectors are used to determine such muon tracks outside of the solenoid: [resistive plate chambers \(RPCs\)](#), [drift tubes \(DTs\)](#), [cathode strip chambers \(CSCs\)](#), and [gas electron multipliers \(GEMs\)](#). All these gas detectors share the common principle of detecting ionization in the gas volume caused by traversing charged particles—primarily, but not exclusively, muons. Through an applied high voltage, the ions and electrons are pulled to the electrodes causing avalanches of charged particles on their way, which are ultimately detected by the readout electronics. Signals from different chambers can then be further processed to construct segments and tracks. In comparison with the tracker, the resolution of the muon system is considerably worse. This is mostly due to the typically larger spacing between the single units. In addition, the energy content of muons is reduced on their path to the muon chambers and through the iron return yokes, limiting the achievable resolution further. However, the gas detector technology is cheaper than the silicon tracking technology, making it a reasonable solution for the muon system.

Differences between the gas detectors are mostly given by the design of the electrodes, and the technical implementation of the readout [51, 57]:

- **RPCs** consist of two resistive plates that form the electrodes and create an electric field. Electrons from the ionized atoms are accelerated in this field, cause avalanches that can pass through the resistive plate and are read by detection strips located outside the gas volume. Although this architecture introduces a delay between the impact and the measurement, the bias can be corrected due to the precise knowledge of the delay. As a result, **RPCs** can reach time resolutions at the order of 1 ns.
- **DTs** consist of single tubes, each 4 cm wide and separated by cathodes. In each tube, a central wire acts as the anode. The electric field between the negative cathodes and the positive anodes causes electrons created from a passing muon to drift towards the cathode. This drift can be evaluated, where especially the information on the time of the electron drift is essential as it hints at the exact position within the tube.
- **GEMs** are used in the forward region of the detector, where radiation hardness is required. Their central feature is a system of foils with tiny holes, in which a strong electric field leads to a multiplication of the electrons from the ionization process. The large electric field in the holes can intuitively be understood by comparison with a flow of air that is forced through a small hole, which leads to a high pressure. The **GEMs** of the **CMS** detector are made of three such layers ensuring a large multiplication of the signal before the measurement.
- **CSCs** have cathodes in the form of strips, which are orthogonally crossed by anode wires. The detection principle being similar to those of the other gas detectors, the orthogonally crossed anode wires of this detector type allow a good 2d-resolution.

While **CSCs** provide the main contribution to the track measurement in the endcaps, this part is done by the **DTs** in the central detector region. The main purpose of **RPCs** and **GEMs** is the fast evaluation of muon information, which is particularly important for triggering.

4.2.2 Data Acquisition and Trigger System

The vast number of individual detecting units in combination with the high rate of collisions poses a great challenge to the data acquisition. During data taking, the **CMS** detector material is generally active and measures signals, which are collected and read out in fixed time intervals. A large share of the collected data is not meaningful, i.e., noise, and can be suppressed in a first processing step without major loss of information. The remaining signals can be further processed and structured, to gain a first level of information, e.g., by combining the size of a signal with the position of the corresponding detector piece in the detector. This first data structure is called raw data, with a corresponding storage space of roughly 1 – 2 MB per detector readout.

However, since the effective collision rate is at about 30 MHz, the corresponding data stream would amount to more than 30 TB/s. This is infeasible both in terms of bandwidth and total storage space available (for comparison, the total **CERN** storage space of roughly 1 EB would be filled in just under ten hours). Consequently, a large share of the data acquired by the **CMS** detector needs to be neglected to fit the computing infrastructure. However, the reason why the data rates are chosen so high is that interesting events are rare, the production of a Higgs boson, e.g., occurs in about one of a billion **pp** collisions. Thus, the vast majority of events are not worth saving, e.g. collisions with low momentum transfers, which have already been studied extensively in the past.

The corresponding filter process is implemented by applying a two-level trigger system. The first trigger (L1) is implemented on specialized hardware processors, which are necessary to

compute a pre-defined set of information and automatically signalize in case the information fulfills certain criteria. For this first trigger level, quickly accessible information is used, which distinguishes interesting from uninteresting events. Interesting events typically contain heavy unstable particles, decaying to lighter particles with high energies. Due to the multiplicity of entries in the tracker, finding tracks requires computationally intensive clustering or track reconstruction algorithms. However, the deployment of such algorithms with the required latency of about 4 μ s is not yet feasible. Instead, information from the calorimeters and muon detectors is used in this first trigger level. While the calorimeter information is already relatively coarse by construction and therefore well-suited for quick evaluation, the muon system is suited mostly due to the lower level of activity. This first trigger level reduces the data rate to about 100 kHz.

If an event is classified as interesting by the L1 trigger, the total detector information is obtained from a buffer, and forwarded to the [high-level trigger \(HLT\)](#) computer farm. Here, a quick analysis of the whole detector information is performed, with algorithms primarily optimized for computation speed and robustness, but at the cost of precision. The resulting stack of information is called [HLT](#) objects and contains a list of high-level variables, mostly reconstructed objects and their properties. Only if the [HLT](#) objects fulfill certain criteria, both raw data and [HLT](#) objects are stored finally.

Some types of events are generally interesting, e.g., for calibration purposes, but occur too often to store them. In that case, corresponding triggers may be pre-scaled, thereby only saving a certain share of the events fulfilling the corresponding trigger criteria.

Overall, this second-level decision reduces the rate to about 5 kHz, corresponding to a raw data stream in the order of 5 – 10 GB/s.

After successful storage, the raw data can be reconstructed with fewer constraints on the time and resources, ultimately obtaining the form that is analyzed in this thesis. The objects in this final data set are reconstructed with higher precision than the [HLT](#) objects. This sometimes leads to mismatches, which needs to be studied and corrected in any analysis, as it would introduce biases otherwise.

The trigger definitions constitute a potential blind fold in the search for new physics beyond the [SM](#). To reduce this self-bias, two further data streams use the possibility of higher data rates under certain conditions: scouting and parking data.

In the scouting data set, the coarse information from the [HLT](#) is stored for future analyses, with smaller biases, namely that of the L1 trigger. However, only the coarse [HLT](#) objects but not the raw data are stored in this data stream, allowing for higher rates. Consequently, this scouting data approach generally benefits the statistical precision of analyses, but comes at the cost of larger systematic uncertainties.

In parking data, the raw event information is written out, but not reconstructed offline immediately. Instead, the reconstruction is performed in times of lower computing demand, in particular when no data is taken and the computing resources are not blocked by the reconstruction of the main data stream. One major purpose of this data stream is to store data from pre-scaled triggers, which are not immediately important but could turn out useful in the long run. The parking data is, for example, used to save events in the domain of well known low-mass resonances, such as the J/ψ resonance. Such processes have a relatively large cross section (compare for instance [66]) and therefore occur too often to be saved and reconstructed all at once. Nevertheless, they are interesting enough to be saved and reconstructed later, as they are useful for calibration purposes, such as those conducted in Chapter 6. In the current Run 3, typically half of the [HLT](#) output rate of roughly 5 kHz is parked, while the other part is reconstructed promptly [67].

4.2.3 Event Reconstruction

The aforementioned full offline reconstruction follows a recipe called **particle-flow (PF)** algorithm. Abstractly, it is a transformation procedure that forms the raw, relatively unstructured data into objects with a clear physical meaning. For this, the **PF** algorithm combines the unstructured data from all detector components in an optimized form, taking into consideration external information about particle's physical behavior.

First, an iterative track reconstruction is performed, combining single hits in the tracker layers with the Kalman filter technique [68–71]. This approach is chosen to obtain both high purity and high efficiency in the large solution space spanned by the vast number of hits in an event. It starts by reconstructing tracks that fulfill tight quality criteria. The hits associated with the reconstructed tracks are then masked from the use in subsequent iterations to reduce the combinatorial multiplicity. Subsequent iterations are then performed with relaxed quality criteria. The resulting objects are called tracker tracks.

The track reconstruction in the muon chambers is conducted independently from that in the tracker and relies on the formation of so-called segments. These segments are formed from individual hits in the corresponding system through the fit of a straight line. Starting from such segments, tracks are formed using a Kalman-filter approach that takes into account the different hits from all muon subsystems. The resulting objects are called standalone-muon tracks.

As muons from collisions usually leave a track in the tracker as well, the information from the two detector systems can be combined. Tracker tracks that can be matched with at least one muon segment are called tracker muons. If a standalone muon can be matched to a tracker track, a new fit is performed combining information from both tracker and muon systems, and the resulting object is called global muon.

The calorimeter systems are required to measure the energy of neutral particles, as these leave no track in the tracker. However, the calorimeter systems do not exclusively measure the energy content of neutral particles because the charged particles typically lose only a part of their energy on their path through the tracker. Consequently, it is crucial to not only measure the energy in the calorimeters but also separate and identify the different energy deposits. The separation is achieved through the application of clustering algorithms. The resulting energy clusters are then matched to tracker tracks for identification. Depending on the resulting combination, different particle types can be identified subsequently.

- An **electron** produces a **track in the tracker** and an **ECAL** cluster, but no **HCAL** cluster.
- A **photon** corresponds to an **ECAL** entry without a matched track and no associated **HCAL** signal.
- **Charged hadrons** and **neutral hadrons** can be distinguished based on the presence of a track. A charged hadron will have a track, an **ECAL** cluster, and an **HCAL** cluster, while a neutral hadron is identified by an **HCAL** cluster without a matched track.

Once the individual **PF** candidates are classified, jet clustering algorithms can be applied, such as the anti- k_t algorithm [72].

The next step consists of the final object selection. Based on the different tracks, common vertices are constructed. The vertex with the largest value of the summed squared transverse momenta of associated tracks is defined to be the primary vertex. Particles connected to

this primary vertex are kept in a special collection, constituting the basic objects used in most analyses. For neutral particles, such an assignment is not possible, as they lack tracker information. These neutral contributions can have substantial effects on jet reconstruction, as all kinds of information within a cone around the jet are in principle added to the jet to reflect its collimated structure. However, contributions from uninteresting collisions can be mitigated through dedicated techniques, e.g., using profiles from neutral event content based on the distribution of the charged particles.

After the event reconstruction, an additional quantity can be constructed, the missing transverse momentum \vec{p}_T^{miss} . This quantity makes use of the fact that the collision partners have almost no momentum in the direction transverse to the beam axis before the interaction. Due to the conservation of momentum, the vectorial sum of the transverse momenta of all particles from a common vertex must be zero after the collision as well. If a particle cannot be measured, such as a neutrino or possible dark matter candidates, then the detectable vectorial transverse momentum sum can only vanish by chance, e.g., if two undetected particles leave in exact opposite directions with the same momentum. Thus, the measurement of the negative vectorial sum of all transverse momenta

$$\vec{p}_T^{\text{miss}} = - \sum_i \vec{p}_{T,i} \quad (4.8)$$

gives rise to undetectable particles. However, mismeasurements of transverse momenta, additional interactions in the event, as well as their modeling are possible sources of bias limiting the resolution of this quantity.

5 Statistical Methods

In the SM, all fundamental interactions are inherently stochastic. As a result, the study of such processes relies on the statistical analysis of their occurrence in nature. As introduced in Chapter 3, each process is characterized by a probability amplitude. Measurements of such amplitudes can be performed implicitly by investigating cross sections and decay widths.

These physical observables have well-defined values in nature—limited only by the quantum mechanical uncertainty principle [73]. On the theoretical side, uncertainties arise from the limited precision in perturbation theory, or from the limited knowledge about the SM parameters. From a statistical point of view, a theoretical prediction is a hypothesis that can be tested with experiments.

Section 5.1 introduces fundamental concepts of probability theory and outlines the statistical framework of likelihoods used in particular for the cross section measurement in this thesis. The following Section 5.2 covers different application cases of MC methods.

5.1 General Concepts of Probability Theory

In probability theory, two widely used interpretations exist: the Bayesian and the frequentist approach. Agreeing on a specific interpretation is crucial when discussing probabilities, as it influences both the conceptual understanding and the mathematical modeling of a problem.

5.1.1 Probability Interpretation

In the Bayesian interpretation, the probability associated with a hypothesis A is denoted as $P(A)$ and represents the degree of belief in that hypothesis. A common example is weather forecasting, where a certain hypothesis—such as “it will rain tomorrow”—is assigned a probability that reflects the confidence in that prediction, given the current meteorological data and models.

The probability is not fixed but can depend on various parameters and prior information, as is the case for the weather prediction. Importantly, the exact same set of conditions may never occur again. When new experimental data becomes available, the probability can be updated, refining the degree of belief from a prior to a posterior probability.

An alternative approach is the frequentist interpretation, where probability is defined differently. In this framework, the probability of an event is determined by the relative frequency of its occurrence in an infinite number of identical, independent trials. For instance, the probability of a fair coin landing on heads is interpreted as the relative frequency of heads in an infinite sequence of coin flips. For the frequentist interpretation to be meaningful, all trials must be conducted under identical or sufficiently similar conditions; otherwise, this interpretation is not well-defined. The environment provided by the **LHC** satisfies this requirement, as high rates of **pp** collisions occur under repeatable and controlled conditions. Therefore, the frequentist approach offers a robust and meaningful statistical interpretation for the analyses presented in this thesis and is adopted throughout.

5.1.2 Hypothesis Tests

In the field of natural science, any theory is intrinsically either right or wrong, but this information is generally unknown. Hypothesis tests provide the mathematical framework to deal with this challenge. Their aim is to reject a certain (null) hypothesis H_0 based on pre-defined criteria in the form of error probabilities. By construction, there are four possible cases when comparing the result of a hypothesis test with the (unknowable) true situation. In case the theory is actually true, it can be correctly accepted or falsely rejected. If the theory is wrong, the test can correctly reject it or falsely accept it.

The decision to reject or accept the hypothesis is usually done based on a pre-defined significance level α , which correspond to the probability of falsely rejecting the null hypothesis despite its truth. In particle physics, processes are usually excluded at a 5% significance level, while new physical processes are accepted (corresponding to a rejection of the model without the corresponding process) at the 5σ level. The probability for an $n\sigma$ deviation is generally given by the area under a normalized Gaussian function \mathcal{N} further than n standard deviations σ away from the mean μ :

$$\alpha = 1 - \int_{\mu-n\sigma}^{\mu+n\sigma} \mathcal{N}(x; \mu, \sigma) dx. \quad (5.1)$$

For 5σ , the corresponding probability of falsely claiming the existence of a new process is about 2.87×10^{-7} .

However, this thesis aims neither at an exclusion of a process nor at an observation, but at precision measurements of already known processes. In terms of hypothesis tests, the measurements conducted in this thesis look for deviations from a known, non-zero value, whereas observations search for deviations from zero, and searches exclude variations from zero. While this conceptual difference could be implemented into hypothesis tests, it is more straightforward to instead conduct the measurement and compare the resulting values and associated uncertainties to different theoretical predictions. The hypothesis test is then implicitly shifted towards the direct comparison between the measurement and the theoretical predictions.

5.1.3 Conditional Probabilities and Likelihood

In many cases, probabilities are evaluated conditionally, meaning they depend on the occurrence of another event. Given two events A and B , the probability of event A occurring, given that event B has occurred, is defined as:

$$P(A|B) \equiv \frac{P(A \cap B)}{P(B)}. \quad (5.2)$$

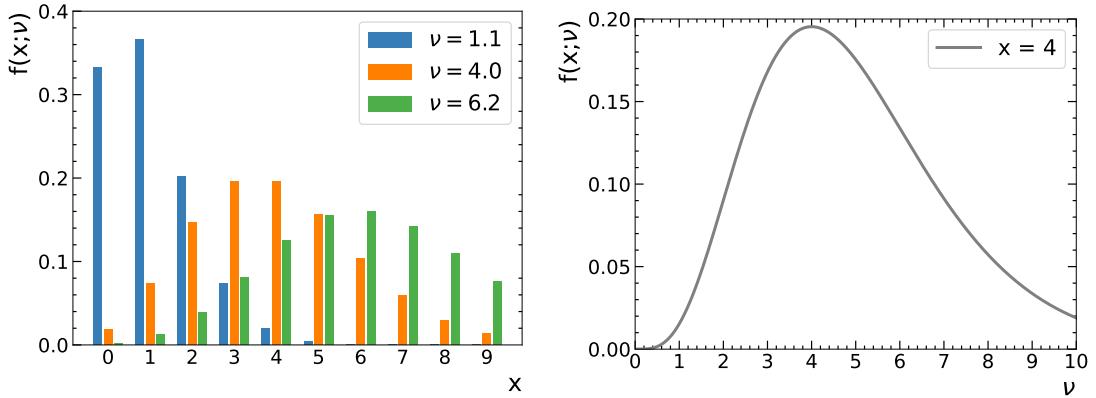


Figure 5.1: Poissonian distributions as functions of different dependent variables. On the left, the Poissonian is shown as a function of the observation x for three different hypotheses ν , while the right shows the dependence on the hypothesis ν for a specific observation of x .

In the frequentist interpretation, the conditional probability is estimated by counting the number of times both A and B occur together and dividing this count by the number of occurrences of event B .

An important application of conditional probability is Bayes' Theorem:

$$P(B|A) = \frac{P(B \cap A)}{P(A)} \frac{P(A)}{P(B)} = \frac{P(A \cap B)}{P(B)} \frac{P(B)}{P(A)} = \frac{P(A|B)P(B)}{P(A)}. \quad (5.3)$$

Despite its name, this theorem not only holds in Bayesian statistics but is a general mathematical result that is also true in the frequentist interpretation.

In high-energy physics, conditional probabilities are often used to compare a measured dataset x with a model, represented by its parameter vector θ . The probability of obtaining a result like the observed data, given the current model $P(x|\theta)$ is called likelihood. In the frequentist interpretation, it corresponds to the expected share of experiments reproducing the measurement x under the assumption that the model is true. The likelihood is a valuable tool as it allows to compare different hypotheses, or to optimize a single hypothesis by finding the parameters that maximize the likelihood.

5.1.4 Parameter Estimation with the Likelihood Method

Due to the independence of different collisions at a known rate, high-energy processes at the [LHC](#) generally follow Poissonian statistics. For this reason, the parameter estimation with the likelihood method is started with a simplified example of a Poissonian distribution in a one-bin counting experiment:

$$f(x; \nu) = \frac{\nu^x e^{-\nu}}{x!}. \quad (5.4)$$

This function provides the probability of measuring a certain process x times, when the expectation of the process is actually to measure it ν times. While the expected value ν can be of any positive real value, x needs to be a positive integer as the process can be either measured or not. The Poissonian distribution for a few different values of ν is provided in Figure 5.1.

The Poissonian distribution can be interpreted as a likelihood to measure a value $x > 0$, where the model expectation is given by $\nu > 0$. To identify the value of ν , which maximizes the likelihood to measure x , the first derivative is set to zero:

$$\frac{df}{d\nu} = \frac{x\nu^{x-1}e^{-\nu} - \nu^x e^{-\nu}}{x!} = 0 \quad \Leftrightarrow \quad \nu = x, \quad (5.5)$$

where evaluation of the second derivative shows that this solution is a maximum of the likelihood.

Generalizing from such a one-bin counting experiment to a histogram with n bins yields the likelihood:

$$f(x; \nu) = \prod_{i=1}^n \frac{\nu_i^{x_i} e^{-\nu_i}}{x_i!}, \quad (5.6)$$

where now x and ν are vectors with n entries x_i and ν_i , where $i \in \{1, \dots, n\}$, respectively. In this example, each bin is completely independent from each other, therefore maximization of the likelihood yields $x_i = \nu_i$

However, in high-energy physics, the model prediction is usually not independent between different bins. Assuming full correlation between the different bins—expressed by a common multiplicative factor μ , with otherwise fixed values of ν_i —the likelihood function reads:

$$f(x; \nu; \mu) = \prod_{i=1}^n \frac{(\mu \nu_i)^{x_i} e^{-\mu \nu_i}}{x_i!}. \quad (5.7)$$

Maximization of this likelihood function becomes computationally more straightforward by considering the logarithmic likelihood. Since the logarithm is increasing strictly monotonous with the argument, and is well-defined as the likelihood is always positive, the position of the maximum is identical. The log likelihood reads:

$$\log f(x; \nu; \mu) = \sum_{i=1}^n x_i \log(\mu \nu_i) - \mu \nu_i - \log(x_i!). \quad (5.8)$$

and the maximization with respect to the single model parameter μ yields:

$$\frac{d \log f}{d \mu} = 0 \quad \Leftrightarrow \quad \mu = \frac{\sum_{i=1}^n x_i}{\sum_{i=1}^n \nu_i}. \quad (5.9)$$

Thus, the ideal multiplication factor turns out to scale the total number of events in the model to that in the observed dataset.

The likelihood model can be further refined by incorporating additional information, such as an additive contribution b_i , which could be scaled with a separate parameter μ_b . Furthermore, additional information can be introduced to the likelihood by adding constraining terms with previous knowledge to the log likelihood. There are two common options to introduce such constraints. Parameters that only affect the rate can be constrained with a log-normal prior, effectively introducing a penalty increasing in size as the rate parameter moves away from an expected value.

Shape deviations, on the other hand, can be introduced as two additional templates, where one corresponds to an up-variation of a quantity, and the other corresponds to a down-variation of the quantity. Intermediate values are then calculated by interpolating between the nominal template and the variations, or by extrapolating to values out of the given range.

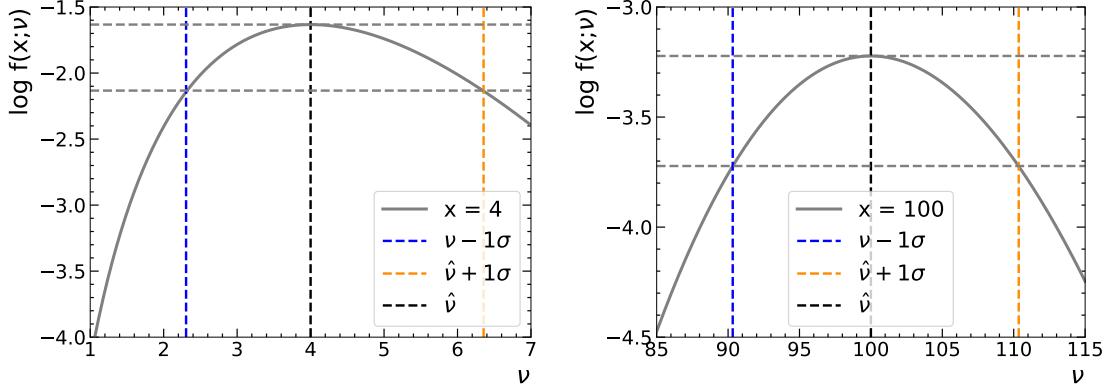


Figure 5.2: Logarithmic Poissonian likelihood function with uncertainty construction. The distribution on the left corresponds to the right plot in Figure 5.1 and shows a strong deviation from a parabolic shape. The right plot shows a similar construction for an observation with a larger statistical sample size, showing almost symmetric uncertainties.

In both cases, it is necessary to provide a range in which the parameter is supposed to be with a certain probability. For the shape deviation, the given additional templates correspond to a variation of $\pm 1\sigma$. For rate parameters, the effective relative effect of a deviation on the rate is provided by $e^{\pm 1\sigma}$, ensuring that the rate must not become negative.

For large statistical sample sizes, the Poissonian distribution approaches a Gaussian distribution. This behavior enables uncertainty estimation based on the logarithm of the likelihood function, which in this limit takes on a parabolic form:

$$\log \mathcal{N}(\nu; x, \sigma) \propto -\frac{1}{2} \left(\frac{\nu - x}{\sigma} \right)^2, \quad (5.10)$$

where x and σ represent the position of the optimum and the standard deviation with respect to this parameter. As a result, the standard deviation can be extracted by varying the parameter ν away from its optimum x until the log-likelihood decreases by 0.5 from its maximum value. This behavior is illustrated in Figure 5.2 for a one-bin example at small and large samples sizes. Importantly, this approach is not limited to one-bin examples but also generalizes to multiple bins sharing a common parameter μ , including cases with additional parameters, if the corresponding term in the likelihood is approximately Gaussian.

5.1.5 Cross Section Measurement with Likelihood Maximization

Cross sections are a measure to describe the probability of a certain interaction. If a certain process 1 has a larger cross section than another process 2, the expected number of processes 1 is consequently larger by the same factor than that of 2. The proportionality constant is called instantaneous luminosity and is a measure for the size of a dataset. As a result, theoretical cross section calculations for a certain process can be probed by detecting a certain process and counting the number of its occurrences in a given dataset of luminosity L_{int} :

$$N = L_{\text{int}} \sigma. \quad (5.11)$$

However, several experimental challenges require adaptations of this measurement strategy for the application in the CMS experiment.

First and foremost, the detector coverage is limited by construction through the beam line and therefore cannot cover the full space. To deal with this, the theoretical cross section can be defined and calculated in the subset of the total phase space, that matches the detector coverage, called fiducial volume. The ratio between the resulting fiducial cross section σ_{fid} , and the total cross section σ_{tot} is called acceptance:

$$A = \frac{\sigma_{\text{fid}}}{\sigma_{\text{tot}}} = \frac{N_{\text{fid}}}{N_{\text{tot}}}, \quad (5.12)$$

where the second step follows from Equation 5.11 and the fact that the integrated luminosity is an accelerator-related quantity and does not depend on the detector.

Furthermore, any detector is limited in terms of efficiency ϵ . This quantity is defined as the ratio between the number of events that is reconstructed in the detector N_{reco} and the true number of events N_{fid} in the detector:

$$\epsilon = \frac{N_{\text{reco}}}{N_{\text{fid}}}. \quad (5.13)$$

Combining these definitions, the initial cross section equation can be modified to:

$$N_{\text{reco}} = L_{\text{int}} \underbrace{\sigma_{\text{tot}} A}_{\sigma_{\text{fid}}} \epsilon. \quad (5.14)$$

Another challenge in cross section measurements comes from the presence of additional (background) processes that mimic the (signal) process under investigation. Considering this connection, a total model of the expected number of events N_{exp} in a measurement can be constructed as:

$$\begin{aligned} N_{\text{exp}} &= \mu N_{\text{sig}} + N_{\text{bkg}} \\ &= L_{\text{int}} (\mu \sigma_{\text{fid, sig}} \epsilon + \sigma_{\text{fid, bkg}} \epsilon), \end{aligned} \quad (5.15)$$

where μ ensures the flexibility of the model with respect to the signal cross section $\sigma_{\text{fid, sig}}$. Given a measurement N_{obs} , the signal strength modifier μ can be optimized by maximizing the likelihood with respect to that parameter as described in the previous section. The cross section of the background process $\sigma_{\text{fid, bkg}}$ is not equipped with a signal strength modifier since the fit would have no clear minimum otherwise. However, additional flexibility can be provided in the fit through the introduction of nuisance parameters, which are allowed to float within a certain range but introduce penalty terms to the likelihood function.

Often, using a histogram of certain physical quantities instead of a single bin improves the overall precision in the fit. The reason for this lies in the potential discriminating power of the different underlying kinematic behavior of the processes.

The result for the cross section measurement is then given by the product of the post-fit signal strength parameter and the theoretical signal cross section in the fiducial region used as input to the fit:

$$\sigma_{\text{fid, meas}} = \mu \sigma_{\text{fid, sig}}. \quad (5.16)$$

It is interesting to note that this result is completely independent from the theoretical cross section that is used as input for the signal process normalization. If, for instance, only half the value is used in the fit, then the optimal value of μ will be larger by a factor of two (assuming the fit finds the optimal minimum in both cases). Consequently, any normalization uncertainty in the cross section will be reflected by the signal strength modifier. For this reason, theoretical uncertainties in the normalization of the signal model do not need to be considered in the fit.

5.2 Monte Carlo Methods

Some problems in high-energy physics are not analytically solvable and therefore require numerical methods. One common approach consists in the generation of pseudo-random numbers and the clever application of functionalities to approximate the behavior of a system. Such MC methods are frequently used for integration tasks or estimation of key figures of random distributions.

5.2.1 Integration

The general principle behind MC integration is illustrated in the following using a one-dimensional function $f(x)$. If this function is continuous on an interval $[a, b]$, then its mean value over this interval is defined as:

$$\langle f(x) \rangle = \frac{1}{b-a} \underbrace{\int_a^b f(x) dx}_I. \quad (5.17)$$

Applying this definition in reverse, the integral I can be approximated by estimating the mean value $\langle f(x) \rangle$. This estimate can be obtained by drawing N values x_i randomly from a uniform distribution over the interval $[a, b]$. The corresponding approximation for the integral is then given by:

$$I \approx (b-a) \frac{1}{N} \sum_{i=1}^N f(x_i) \equiv I_N. \quad (5.18)$$

By construction, the expectation value of I_N is equal to the true value I in the limit $N \rightarrow \infty$. However, the statistical uncertainty of this estimate scales as $1/\sqrt{N}$, meaning that the method converges slowly and at high computational costs. To achieve a precision of 1 %, for example, approximately 10 000 random numbers have to be generated. Improving the precision by an order of magnitude requires increasing the number of samples by a factor of 100, highlighting the computational costs of this method.

An important application case is the calculation of cross sections. According to Equation 3.29, this requires the integration of the transition amplitude \mathcal{M} over the phase space Π of the final state particles:

$$\sigma = \frac{1}{\Phi} \int |\mathcal{M}|^2 d\Pi, \quad (5.19)$$

where Φ is the flux of incoming particles. Due to the complexity of the transition amplitude, this integration step is usually not analytically solvable. Instead, random points in the allowed phase space are sampled. Then, the matrix element is evaluated at those points and the average value is taken to approximate the cross section.

5.2.2 Bootstrapping

In order to estimate key figures of a given random distribution, resampling methods can be applied. An approach used later in this thesis is the bootstrapping method. The goal of this approach is the approximation of the distribution of a quantity which is not directly accessible in an analytic procedure.

For this, from a given sample, varied samples are generated (resampling) by drawing with replacement. This procedure can be simulated by drawing a weight w_i from a Poissonian

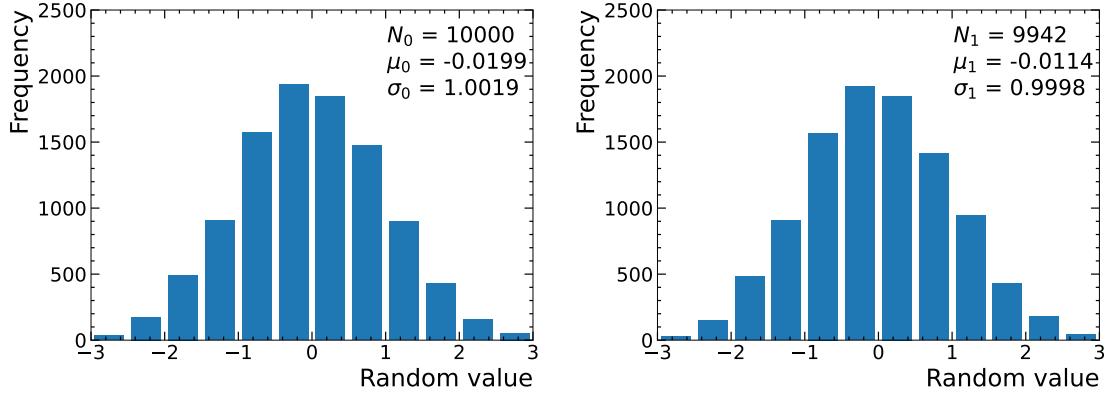


Figure 5.3: Example for the bootstrapping procedure. On the left, a random distribution sampled from a Gaussian distribution with $\mu = 0$ and $\sigma = 1$ is shown, with a sample size of $N = 10\,000$. The distribution mean μ_0 and standard deviation σ_0 slightly vary from the true values due to the limited size of the statistical sample. On the right, a bootstrapped sample is shown. The Poissonian reweighting leads to a fluctuation in the sample size in the order of \sqrt{N} . The expected deviation of its mean μ_1 from the original value μ_0 corresponds to the uncertainty on the mean $\sigma_0/\sqrt{N} \approx 1\%$, while the expected fluctuation of the standard deviation is $\sigma_0/\sqrt{2N - 2} \approx 0.7\%$.

distribution with mean one for each element i of the sample with size N . Then, the new sample consists of $N' = \sum_{i=1}^N w_i$ elements, with expectation value $E(N') = N$ and variance

$$\sigma^2(N') = \sum_{i=1}^N \sigma^2(w_i) = N, \quad (5.20)$$

corresponding to a standard deviation of \sqrt{N} . Thus, for large sample size N , the relative difference in size between the original and the bootstrapped sample becomes negligible. In this procedure, the random selection of the weights along with their Poissonian construction with mean one is of great importance. These properties lead to the conservation of the probability moments, which is shown in the following for a general k -th moment:

$$\langle X^k \rangle = \frac{1}{N} \sum_i x_i^k \quad \rightarrow \quad \langle X'^k \rangle = \frac{1}{N'} \sum_i w_i x_i^k. \quad (5.21)$$

Since the weights are uncorrelated and their expectation value is one, the expectation value of the k -th moment in a bootstrapped sample is consistent with that of the original sample. Figure 5.3 shows a random distribution and a corresponding bootstrapped distribution.

The value of the bootstrapping method comes into play when considering a larger number n of bootstrapped samples. Since the expectation value of any estimator is conserved in the bootstrapping procedure, it can be approximated by calculating the average estimator value of all bootstrapped samples. The standard deviation of the estimator can be approximated correspondingly. However, the precision of these values is subject to the number of bootstrapped samples: for the mean the uncertainty is σ/\sqrt{n} , for the standard deviation $\sigma/\sqrt{2n - 2}$. For illustration, Figure 5.4 shows the distribution of means and standard deviations from bootstrapped samples imitating the original sample from Figure 5.3.

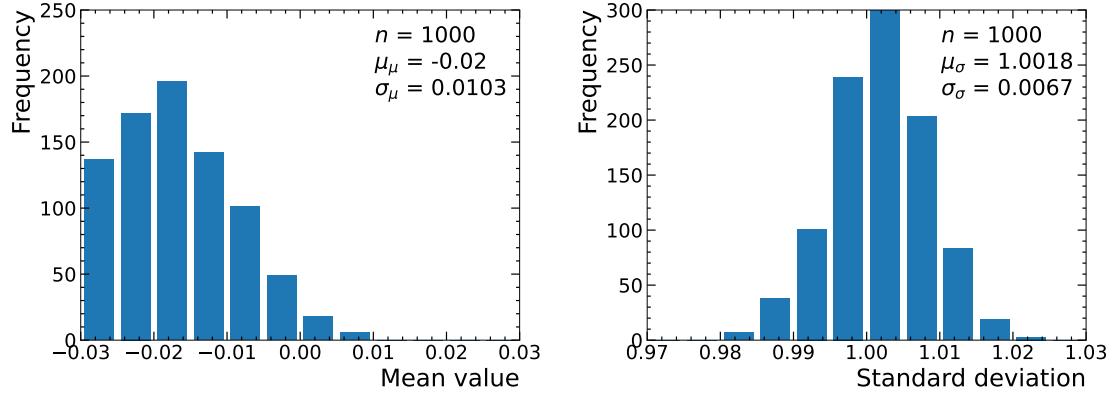


Figure 5.4: Distribution of the mean values (left) and standard deviations (right) from a set of $n = 1000$ bootstrapped samples from the original distribution in Figure 5.3 (left). The central values of the two distributions align well with those of the original distribution. The standard deviations correspond to the expected deviations from the mean and standard deviation in Figure 5.3 (right), i.e., 1% for the mean (left) and 0.7% for the standard deviation (right), showing good agreement with the expectation. The precision of these uncertainty estimations could be improved by increasing the number of bootstrapped samples.

To illustrate the technical implementation with the weights, consider a sample X with a smaller size $N = 4$:

$$X = [2, 4, 7, 3]. \quad (5.22)$$

Then a bootstrapping weight vector with Poissonian weights could be:

$$w = [0, 1, 0, 2], \quad (5.23)$$

which by drawing each entry of X as many times as the corresponding weight entry dictates, results in the following bootstrapped sample:

$$X_b = [4, 3, 3]. \quad (5.24)$$

The new sample size in this case would be $N' = 3$, as indicated by the sum of weights, and the number of entries in X_b .

In this thesis, the bootstrapping method is applied to estimate statistical uncertainties in the correction procedure of the muon scale and resolution corrections in Section 6.3.

6 Muon Corrections in the CMS Experiment

Muons are important final state particles in many high-energy analyses conducted by the CMS Collaboration. Understanding their behavior in the detector is therefore essential. Beyond their quantum numbers, the defining properties of muons are their mass of approximately 106 MeV and the mean lifetime of about 2.2 μ s [7]. This unique combination plays a key role in the reconstruction process.

To enhance the understanding of muons, Section 6.1 explores the implications of their kinematic behavior in the detector, as well as methods to improve reconstruction based on these effects. To achieve the most accurate description of the data using simulated events in analyses involving muons, two fundamental types of residual corrections are typically applied: efficiency corrections and momentum corrections. The former is addressed by applying scale factors to the simulation, which are derived using the [tag-and-probe \(TnP\)](#) technique described in Section 6.2. Finally, a framework for deriving muon momentum corrections is introduced in Section 6.3.

6.1 Kinematic Behavior of Muons

The kinematic behavior of muons in the detector is strongly influenced by their momentum, which depends on the production process under study. In many cases, muons originate from the decay of short-lived resonances.

6.1.1 Momentum Distribution of Muons from Resonances

In the rest frame of the resonance, the daughter particles are emitted isotropically and back-to-back. The magnitude of their momentum $p \equiv |\vec{p}|$ is determined by the mass of the resonance m_{res} and the masses of the daughter particles $m_{1/2}$. Applying four-momentum conservation yields:

$$m_{\text{res}}^2 = m_1^2 + m_2^2 + 2(E_1 E_2 + \vec{p}^2) \quad \Rightarrow \quad |\vec{p}| \approx \frac{m_{\text{res}}}{2}, \quad (6.1)$$

where the approximation holds when the mass of the resonance is significantly larger than the masses of the daughter particles.

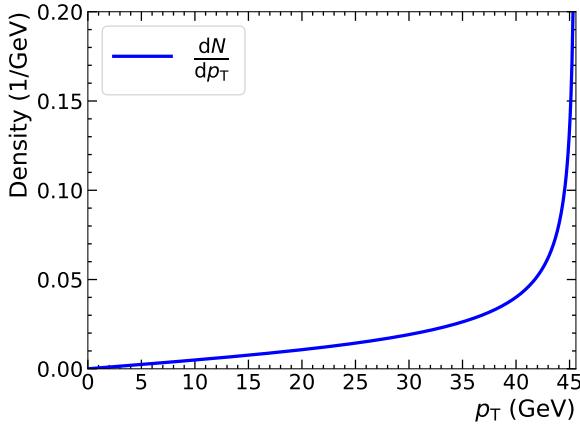


Figure 6.1: Distribution of the p_T values of muons from a Z boson decay in its rest frame. The distribution shows a divergence at $p_T \approx m(Z)/2$.

The transverse momentum p_T is defined in terms of p as:

$$p_T \equiv p \sin \theta, \quad (6.2)$$

where θ is the polar angle between the direction of the muon momentum and the z axis. Thus, the transverse momentum can range from 0 to approximately $m_{\text{res}}/2$.

Naively, one might expect that muons are produced uniformly between these two values. However, higher p_T values are preferred due to the isotropic decay of the resonance. Since each solid angle element $d\Omega = \sin \theta d\theta d\phi$ has an equal probability of containing the muon, the p_T spectrum follows from the chain rule:

$$\frac{dN}{dp_T} = \underbrace{\frac{dN}{d\Omega}}_{\text{const.}} \frac{d\Omega}{dp_T} \propto \frac{\sin \theta}{p \cos \theta} = \frac{p_T}{\sqrt{p^2 - p_T^2}}, \quad (6.3)$$

where the second step makes use of the ϕ -independence of p_T . This result reveals a divergence as $p_T \rightarrow p$, which can be observed in Fig 6.1. While divergences can sometimes indicate unphysical behavior, this particular divergence is not problematic. In experimental measurements, muons are always counted within finite p_T intervals, and integration over such intervals eliminates the divergence naturally.

Given the kinematic muon distribution in the resonance's rest frame, the corresponding distribution in the detector frame is obtained via a Lorentz boost. In pp collisions, this boost can be decomposed into two components: parton momentum imbalance and radiation effects.

- The **parton momentum imbalance** arises because the colliding partons may carry unequal momentum fractions x_1 and x_2 of the proton's energy. This results in a Lorentz boost along the beam axis.
- The **radiation effects** primarily involve **ISR** and **FSR**. **ISR** occurs when additional gluons (higher-order **QCD**) or photons (higher-order **QED**) are radiated off before the resonance is produced. As a result, the resonance recoils against the emitted radiation, introducing an additional boost with no preferred direction. The magnitude of this boost depends on both the resonance mass and the momentum of the recoiling particles. In contrast, **FSR** can be neglected at first order for muons as their relatively large mass—compared to electrons—suppresses bremsstrahlung.

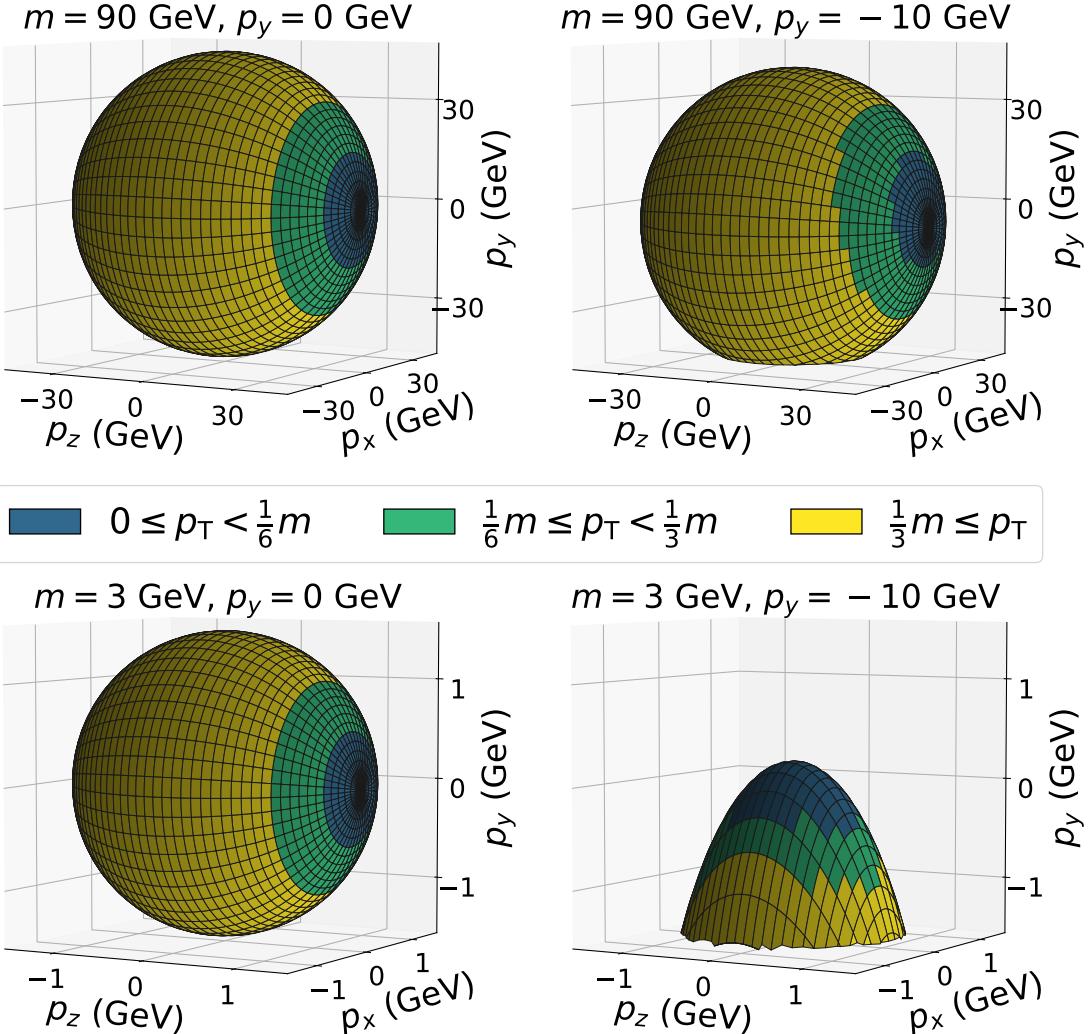


Figure 6.2: Distribution of the muon momenta from the decay of a Z boson (top), and J/ψ meson (bottom) in momentum space. In the rest frames, which are depicted on the left, the distributions are consistent by construction (except for the axis range). Here, the tendency towards larger p_T values is demonstrated once more by the yellow area being considerably larger than the other two despite covering the same p_T range. On the right, the corresponding distributions are shown for the case of a 10 GeV momentum of the mother particle in the detector frame. Here, the distribution of the Z boson decay particles is only slightly distorted, while that of the J/ψ meson is boosted significantly stronger.

The boost factor β is given by the ratio of momentum and energy

$$\beta = \frac{p}{E} = \frac{p}{\sqrt{m^2 + p^2}} = \frac{1}{\sqrt{\frac{m^2}{p^2} + 1}}. \quad (6.4)$$

From this expression, it follows that heavy resonances require a large recoil or parton momentum imbalance to achieve a significant boost. Consequently, high-mass resonances are generally less boosted than low-mass ones.

Figure 6.2 illustrates the muon momentum distribution in momentum space for a J/ψ meson and a Z boson, both in their rest and a boosted frame. The boost distorts the originally spherical p_T distribution, increasing the p_T values in the direction of the boost and decreasing them in the opposite direction. Effects such as the limited detector resolution and the intrinsic width of the resonance additionally distort the p_T spectrum of the daughter muons.

To summarize, although muons are emitted with a well-defined momentum in the resonance's rest frame, various effects contribute to the observed p_T spectrum. While p_T values are not completely governed by the resonance mass, the mass still has a decisive effect on the spectrum as they are still strongly correlated.

Since muons from different p_T regions may behave differently, it is useful to categorize them based on their transverse momentum:

- low- p_T : $\lesssim 10$ GeV
- medium- p_T : $10 \text{ GeV} \lesssim 100 \text{ GeV}$
- high- p_T : $\gtrsim 100$ GeV

This separation is not strict but rather serves as a rule of thumb. Similarly, resonances are often grouped into different categories:

- The J/ψ meson can be used in particular for the investigation of low- p_T muons, due to its low mass of 3.097 GeV [7]. Since its decay to D mesons is kinematically not allowed, it has a relatively long lifetime (on the scale of hadronic physics) and consequently a narrow decay width of about 93 keV [7], which makes it highly suitable for calibration tasks. Due to the low mass, the J/ψ tends to be significantly boosted even for small momenta. As a result, the decay products exhibit strong angular correlations, a feature that can be exploited by trigger selections.
- The Z boson can be used in particular for medium- p_T muons, due to its mass of 91.2 GeV [7]. The decay width of 2.50 GeV [7] is more than four orders of magnitude larger than that of the J/ψ . In comparison with the J/ψ , the Z boson is usually less boosted in the CMS experiment. Consequently, the daughter muons are usually emitted back-to-back, leading to large angular separations. While there is no single resonance that can be exploited at high- p_T , the spectrum of Z boson production reaches to the high- p_T range as well and can therefore—with statistical limitations—be used to calibrate high- p_T muons.

An important feature of both resonances are their large production cross sections, which allow for high statistical precision when using these resonances for calibration purposes. Furthermore, by carefully selecting a phase space that reflects the properties of a resonance, other contributions that do not resonate in this phase space can be relatively reduced. As the corresponding systematic uncertainties from the background processes are suppressed as well, this property is particularly beneficial for calibration.

6.1.2 Optimizing Muon Reconstruction

As muons traverse the detector, they leave a trail of energy deposits in the different subdetectors. However, to efficiently identify a true muon among thousands of tracks in a typical event, it is crucial to combine information from both the muon chambers and the tracker. As summarized in Section 4.2.3, there are three different types of muon objects on reconstruction level:

- **tracker muons:** a tracker track that has been matched with a muon segment,
- **standalone muon:** a track reconstructed only in the muon system,
- **global muon:** a track in the muon system matched to a track in the tracker; a global fit is performed using all corresponding hits.

The definitions have in common that they rely on the muon leaving some information in the muon chambers. Due to the large relevance of the muons reaching the chambers, the corresponding probability is examined briefly in Appendix A.1. Optionally, an additional constraint on the muon can be imposed by the beam spot position to further refine the muon reconstruction.

For prompt muons produced in the hard scattering process, the tracker muons are the most efficiently reconstructed muon objects as they require only a single matched segment in the muon system. This is particularly beneficial for low- p_T muons, which may not reach the outer chambers. However, tracker muons also have the highest misidentification rate due to the abundance of tracks in an event that may accidentally match a segment.

Standalone muons are reconstructed by only using the muon system. As a result, their efficiency is generally lower than that of tracker muons, since some true muons may not reach the muon chambers. However, they also have a lower misidentification rate, as the probability of reconstructing an entire track in the muon chambers by chance is relatively small. Above the low- p_T range, where muons are prone to not fully traversing the muon chambers, the standalone reconstruction efficiency is relatively stable up to high- p_T [74]. At values $p_T \gtrsim 200$ GeV, the probability for radiative losses rises, frequently inducing electromagnetic showers in the return yoke that complicate the standalone reconstruction. Therefore, dedicated algorithms are applied for such muons [75].

The precision of the momentum measurement depends on both the tracker and muon chambers, but their relative performance varies with p_T . For low- and intermediate- p_T muons, the tracker provides much better momentum resolution than the muon system. This is because muons lose energy before and while passing through the muon chambers due to interactions with the detector material, including the iron return yoke. In addition, the tracker has a better spatial resolution compared to the muon system. For very high- p_T muons, the muon system may provide a more accurate measurement, particularly if the tracker cannot resolve differences in the curvature anymore.

Global muons combine the advantages of both reconstruction methods. They inherit the standalone muon's small misidentification rate, while also profiting from the high-precision momentum measurement of the tracker. This makes global muons the preferred choice for many physics analyses. There are few exceptions where other methods may be more suitable, such as displaced muons, where the track may not be recovered by the tracker system and the standalone muons could be the preferred choice. Furthermore, at very low p_T , where the muons may not have sufficient energy to be reliably reconstructed as a standalone muon, the usage of tracker muons can be of advantage due to the higher efficiency.

6.1.2.1 Muon Identification

To optimize the balance between efficiency and misidentification rate, additional identification criteria can be applied. These can be either based on simple selections or include multi-variate optimization [76]. In this thesis, only cut-based identification criteria are used, in particular the soft, loose and the tight identification criteria.

The soft identification aims to select low- p_{T} muons and corresponds to a tracker muon, where the track fulfills a set of quality criteria. In contrast, the tight identification processes is aimed at providing a very pure collection of prompt energetic muons, which comes at the cost of lower efficiency, in particular for low- p_{T} muons. This high purity is obtained by applying very strict criteria, e.g., requiring global muons with a good fit quality ($\chi^2/\mathrm{dof} < 10$). Muons that are selected by the **PF** algorithms and are a tracker or a global muon satisfy the loose identification criteria. This identification is more relaxed than the tight identification and therefore has a higher efficiency but a lower purity. A complete description of the selection criteria is provided in Reference [77].

6.1.2.2 Muon Isolation

To suppress muons originating from jets, additional selection criteria based on muon isolation can be applied. Isolation is defined as the sum of transverse momenta from particles within a cone of radius $R < 0.4$ in the η - ϕ plane around the muon:

$$I_{\mathrm{rel}} = \frac{1}{p_{\mathrm{T}}(\mu)} \left[\sum_{\mathrm{PV-had.}} p_{\mathrm{T}} + \max \left(0, \sum_{\mathrm{neutr. had.}} E_{\mathrm{T}} + \sum_{\gamma} p_{\mathrm{T}} - 0.5 \sum_{\mathrm{non-PV-had.}} p_{\mathrm{T}} \right) \right]. \quad (6.5)$$

In this equation, all transverse momentum contributions from charged hadrons associated with the **primary vertex (PV)** are included directly, as the tracker provides efficient vertex assignment for them. The neutral contributions—such as neutral hadrons and photons—are summed separately, while a conservative estimate of pileup contamination is accounted for by subtracting half of the summed transverse momentum of charged hadrons not associated with the **PV**. To prevent negative isolation values due to over-subtraction, the maximum function ensures that the total neutral contribution remains non-negative. The final isolation value is then normalized by the transverse momentum of the muon itself.

A lower relative isolation value indicates a more isolated muon, meaning there are fewer nearby particles with significant transverse momentum. Consequently, contributions from non-prompt muons can generally be suppressed by rejecting muons that exceed a certain isolation threshold. In general, the tighter this threshold, the fewer non-prompt muons remain in the sample—at the cost of reduced efficiency. Commonly used thresholds are 0.25, 0.2, and 0.15, corresponding to loose, medium, and tight isolation criteria, respectively. For a fixed threshold, the efficiency of the isolation selection typically depends on the muon p_{T} , as the relative isolation tends to decrease with increasing p_{T} (except at very high p_{T} , where bremsstrahlung effects become significant). An example plot showing the dependence of the isolation efficiency on the muon p_{T} value can be found in Figure B.5 in the appendix.

6.2 Tag and Probe Method

Muon-based event selection involves various sources of limitation in the efficiency. In this context, the efficiency ϵ is defined as the probability P that a real muon is correctly identified and included in the analysis. This efficiency can be calculated using the frequentist approach, expressed as a conditional probability:

$$\begin{aligned}\epsilon &= P(\text{selected}|\text{real}) = \frac{P(\text{selected and real})}{P(\text{real})}, \\ \Rightarrow \hat{\epsilon} &= \frac{N_{\text{selected,real}}}{N_{\text{real}}},\end{aligned}\tag{6.6}$$

where $\hat{\epsilon}$ denotes the estimator for ϵ and is calculated from the number of selected real muons $N_{\text{selected,real}}$, and the total number of real muons N_{real} (both selected and unselected).

Since analyses typically compare theoretical simulations with experimental results, it is essential to ensure the consistency between the selection efficiency in the simulations and the real data. To improve the consistency, the efficiency must be estimated in both cases. In simulations, the presence of truth-level (generator) information simplifies this process, allowing the efficiency to be directly computed using Equation 6.6. However, in real data, the truth information is not available, indicating the need for alternative approaches to determine the efficiency. The **TnP** method [78] provides a possible approach to overcome this limitation and to estimate the efficiencies directly from the data.

6.2.1 General Concept

The central idea of this approach is to use a set of reconstructed objects that are highly likely to be real muons—a process called tagging. By assuming these objects are indeed real muons, Equation 6.6 can be formally applied to determine the corresponding efficiency. Since this assumption is critical for the validity of the result, it must be carefully evaluated, especially by assessing the impact of potential limitations.

Phase space regions with two oppositely charged muons with a dimuon mass close to a resonance, such as the Z or J/ψ peak, are intrinsically well-suited for this approach, due to the high abundance of the resonance events in the corresponding region. One of the two muons is additionally required to satisfy a set of tight identification criteria to further suppress non-prompt muon contributions. Then, the probability for the other (probe) muon, which is only required to satisfy basic criteria, to not be a real muon is very small:

$$P(\text{real}|\text{probe}) \approx 1.\tag{6.7}$$

If this relation holds, then any subset of probe muons should consist mostly of real muons as well, in particular the subset satisfying the selection criterion under study:

$$P(\text{real}|\text{selected probe}) \approx 1.\tag{6.8}$$

Furthermore, it is assumed that the selection efficiency of real muons is well described by the subset of real muons that is also a probe:

$$\epsilon = P(\text{selected}|\text{real}) \approx P(\text{selected}|\text{real probe}).\tag{6.9}$$

Then, applying Bayes' theorem, along with the assumptions from above, the efficiency can be further approximated:

$$\begin{aligned}
\epsilon &\stackrel{6.9}{\approx} \frac{P(\text{real selected probe})}{P(\text{real|probe})P(\text{probe})} \\
&\stackrel{6.7}{\approx} \frac{P(\text{real|selected probe})P(\text{selected probe})}{P(\text{probe})} \\
&\stackrel{6.8}{\approx} P(\text{selected|probe}), \\
\Rightarrow \hat{\epsilon} &\approx \frac{N_{\text{selected,probe}}}{N_{\text{probe}}},
\end{aligned} \tag{6.10}$$

where $N_{\text{selected,probe}}$ is the number of selected probe muons, and N_{probe} is the total number of probe muons.

To determine the number of probe muons before and after applying the selection criterion, several strategies can be followed. The simplest approach is the cut-and-count method, in which all tagged events within the defined resonance phase space are used to determine the efficiency. This approach is most useful in case of small statistical sample sizes, where more sophisticated methods may experience stronger limitations. In case of decent statistical coverage, the **TnP** method can be further improved by reducing the impact of background events in the data. Such events could otherwise be a limitation to the validity of the assumptions, as they do not share the same properties as the resonance events. A common approach is to extract the number of signal processes from a fit of a model—containing both signal and background contributions—to the dimuon mass distribution in the data.

6.2.2 Statistical Model

Statistical models to extract the number of signal processes can be broadly divided into two groups: simulation-based and analytical models. Simulation-based models are constructed from **MC** simulated events and converted into histogram templates that represent the expected shape of the corresponding contribution. Since the simulation might not perfectly model all aspects of the detector behavior—such as the muon p_T scale and resolution—these models typically incorporate additional degrees of freedom. For example, the simulated templates can be convolved with a Gaussian function to account for such discrepancies in scale and resolution.

In contrast, analytical models are based on fundamental theoretical considerations on the underlying physics process. A commonly used function to model resonances is the Breit-Wigner distribution [79]. The Breit-Wigner function naturally arises from calculations involving an unstable mediator particle, whose mass M and width Γ determine the shape of the resonance. The probability density function as a function of the energy E of the process is given by:

$$f(E; M, \Gamma) \propto \frac{1}{(E^2 - M^2)^2 + M^2 \Gamma^2}. \tag{6.11}$$

Analytical models are simple and efficient, relying on a few parameters and minimal computing resources, unlike simulation-based models that are more resource intensive. However, they may lack accuracy, as they do not fully account for detector effects and higher-order corrections. For example, the detector's finite resolution changes the observed resonance shape, meaning that a pure Breit-Wigner model does not accurately describe the experimental data. To address this, the Breit-Wigner distribution is typically convolved with a Gaussian function, resulting in a Voigtian distribution that better represents the

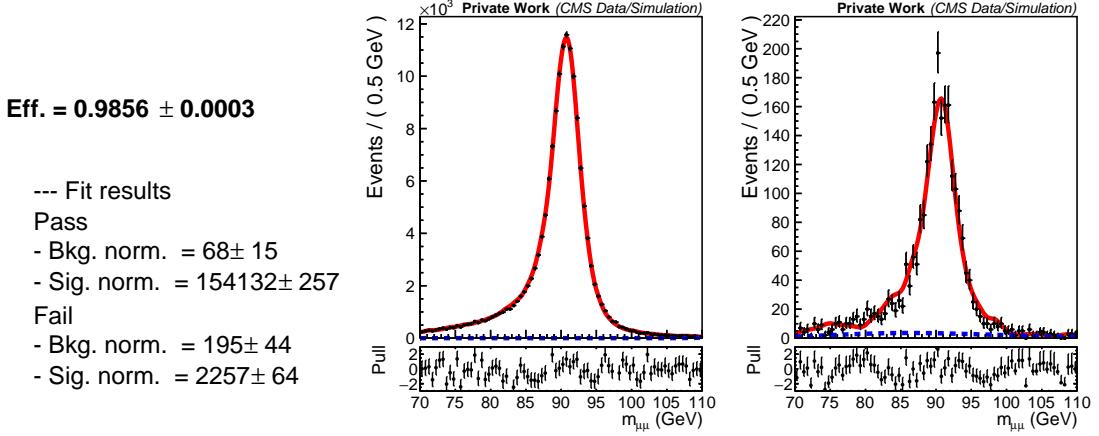


Figure 6.3: Example fits to extract the efficiency from data. Events include global muons as probes with $-1.6 < \eta < -1.2$ and $40 < p_T < 50$ GeV, and passing (failing) the tight identification criteria in the plot on the left (right). Upper panels show the dimuon mass distributions in the data (black points), with the combined post-fit (red line) and the background-only (blue dashed line) models. Lower panels show the pull distribution, defined here as the difference between yield in data and post-fit model, scaled by the statistical uncertainty in the data, generally indicated by black bars. The post-fit normalizations and resulting efficiencies are shown on the left for signal (Sig) and background (Bkg).

measured spectrum. When applying such an analytical approach to estimate the number of muons in the data, it is essential to apply the same method to the simulated events as well. Directly comparing events at the generator and reconstruction level could otherwise introduce biases, as detector-related systematic effects are present in the data but may be missing from the signal model in the simulation.

The standardized approach in the CMS MUO group uses a simulation-based signal model, and an analytical function to model background contributions. The dataset, consisting of tagged events, is then split into two parts: one where the probe muons pass the selection criteria, and one where the probe muons fail the selection criteria. For each category, a fit is performed to extract the signal contribution and mitigate effects from background contamination. The result of these fits are two quantities:

- N_{pass} : the number of probes satisfying both baseline and selection criteria.
- N_{fail} : the number of probes satisfying baseline but failing selection criteria.

The efficiency estimate is then given as:

$$\hat{\epsilon} = \frac{N_{\text{pass}}}{N_{\text{pass}} + N_{\text{fail}}}. \quad (6.12)$$

By construction, this approach ensures that the efficiency is always constrained within the valid range $[0, 1]$. Figure 6.3 presents an example fit for the efficiency extraction from the data.

Since the behavior of muons in the detector strongly depends on their p_T , the efficiencies generally are expected to depend on p_T as well. Similarly, the detector composition and geometry depend primarily on the pseudorapidity η . Variations in the material budget and number of detector hits in different regions may lead to non-uniform efficiency distributions. Therefore, efficiency measurements are typically parametrized as a function of both η and p_T to capture these effects accurately.

6.2.3 Uncertainty Estimation

The standardized approach used in the CMS Collaboration evaluates several different sources of uncertainty which fall into three categories:

- Tag selection uncertainties: to ensure the purity of the probe events, the tag muon must satisfy a combination of tight selection criteria. In general, each of these criteria can be correlated with the efficiency of the probe muon. For instance, events with lower overall activity may lead to both the tag and probe muons more easily fulfilling the isolation requirement. In this case, the tag isolation criterion implicitly influences the probe isolation. Such effects are studied by varying the properties of the tag muon and determining the impact on the efficiency.
- Modeling uncertainties: the choice of the model can have an influence on the result, for instance if background contributions are systematically overestimated. To estimate the impact, the models are varied and the difference between the variations is considered as the associated uncertainty.
- Binning uncertainties: the impact of the choice of the bin size and bin range in the fit of the dimuon mass distribution is estimated by varying both.
- Statistical uncertainties: the finite number of events limits the precision of the measurement and is evaluated directly through the fit.

The different uncertainties are generally considered as uncorrelated. However, the assumption of no correlation bears the risk of double-counting statistical uncertainties from the different evaluations of the systematic uncertainty. In particular if the systematic uncertainty is significantly smaller than the statistical uncertainty, the evaluation of the corresponding systematic variation can on average not be smaller than the statistical uncertainty. To avoid this double-counting, systematic contributions are only included in the total uncertainty if their size exceeds that of the statistical uncertainty.

6.2.4 Application of the Efficiency Corrections

To ensure that simulated events accurately reflect the real data, efficiency corrections are applied in the form of scale factors. These scale factors are used as event weights in the simulation, adjusting for discrepancies between data and simulation. The scale factor SF_x for a given criterion x is defined as the ratio of the efficiency in the data $\varepsilon_{x,\text{data}}$ to the efficiency in the simulation $\varepsilon_{x,\text{sim}}$:

$$SF_x = \frac{\varepsilon_{x,\text{data}}}{\varepsilon_{x,\text{sim}}}. \quad (6.13)$$

For events with multiple selection criteria, the application of scale factors depends on the logical relationship between these criteria. If the selection criteria x and y must both be satisfied, then the total efficiency is given by the product of the individual efficiencies. As a consequence, the total scale factor is also the product of the single scale factors:

$$SF_{x \wedge y} = \frac{\varepsilon_{x \wedge y,\text{data}}}{\varepsilon_{x \wedge y,\text{sim}}} = \frac{\varepsilon_{x,\text{data}} \varepsilon_{y,\text{data}}}{\varepsilon_{x,\text{sim}} \varepsilon_{y,\text{sim}}} = SF_x SF_y. \quad (6.14)$$

If an event is selected when at least one of the criteria x and y is satisfied, then the total efficiency must reflect that an event is only not selected if neither x nor y are fulfilled. Thus, the total scale factor is given by:

$$SF_{x \vee y} = \frac{\varepsilon_{x \vee y,\text{data}}}{\varepsilon_{x \vee y,\text{sim}}} = \frac{1 - (1 - \varepsilon_{x,\text{data}})(1 - \varepsilon_{y,\text{data}})}{1 - (1 - \varepsilon_{x,\text{sim}})(1 - \varepsilon_{y,\text{sim}})}. \quad (6.15)$$

This is often the case in trigger efficiency corrections, where an event is selected if at least one object satisfies a trigger condition.

When estimating the impact of uncertainties on the final result of an analysis, different sources of uncertainty are usually combined for each scale factor and treated as correlated across different η - p_T bins. Different scale factors themselves are assumed to be uncorrelated. This approach is motivated by the simplifications it provides in the fitting procedure when determining the optimal scale factors. Since the impact of muon efficiency scale factors on the overall results is rarely dominant, additional refinements in the treatment of their uncertainties are usually not necessary. It is important to note that a correct treatment of the uncertainties would mean to treat

- statistical uncertainties as independent between different bins in η - p_T , and
- both statistical and individual systematic uncertainties as correlated between different scale factors.

This distinction ensures that uncertainties are not artificially inflated due to incorrect assumptions about their correlation structure.

6.3 Muon Momentum Calibration

The precise measurement of muon momenta is a complex procedure that depends on multiple factors. As a muon propagates through the detector, its trajectory is affected by both the magnetic field and interactions with the detector material. These effects must be properly accounted for in the reconstruction of the muon track. In addition, the spatial resolution of detector hits is a limiting factor to the precision of momentum reconstruction. For an accurate representation of muon momenta in simulation, all these aspects need to be modeled with high precision. Even though the description is highly optimized, residual discrepancies between simulation and real data naturally arise due to the finite accuracy of these models. This section provides an approach to mitigate such discrepancies.

As opposed to efficiency corrections, which are applied via event reweighting, the muon momentum calibration is performed by modifying the p_T of reconstructed muons directly. The overall calibration procedure can be categorized into two fundamentally different types of corrections—scale and resolution—depending on the nature of the underlying effect that is addressed by the correction.

6.3.1 Scale and Resolution

In this thesis, scale effects refer to the deterministic modifications of the measured momentum, whereas resolution effects introduce random, non-deterministic variations. Therefore, the main difference between these two lies in their reproducibility: for exactly the same original muon, scale effects are reproducible, whereas resolution effects are not.

To formalize this, consider a hypothetical true value x_{true} of an observable, which is measured by an idealized detector. In this context, idealized refers to the detector response being completely described by a Gaussian distribution with a mean μ and standard deviation σ . Applying the definition introduced above, the scale effect is described by the quantity $\mu - x_{\text{true}}$, whereas σ summarizes the resolution effect. In a perfect measurement both scale and resolution effects would be negligible, but this is usually not the case in real-world experiments. Instead, in high-energy physics, where most experiments rely on comparisons between simulations and the data, a more practical goal is to eliminate obvious differences between both.

To correct for scale differences between simulation and data, one can either transform both data and simulation such that their mean values match x_{true} , or shift one distribution to align with the other. The choice between these approaches is mostly philosophical, as both achieve the goal of improved data-simulation agreement.

Resolution corrections, however, are more challenging due to their random nature. Unlike scale effects, resolution cannot be improved through ad-hoc corrections (only by improving the detector or reconstruction methods), but only decreased by introducing additional smearing. This can be done by adding a random Gaussian-distributed term $\mathcal{N}(0, \sigma_{\text{add}})$ with mean 0 and standard deviation σ_{add} to each value x . Alternatively x can be multiplied with a Gaussian distributed term with mean 1 and standard deviation σ_{add}/x because:

$$x + \mathcal{N}(0, \sigma_{\text{add}}) = x \left[1 + \frac{1}{x} \mathcal{N}(0, \sigma_{\text{add}}) \right] = x \left[1 + \mathcal{N}\left(0, \frac{\sigma_{\text{add}}}{x}\right) \right] = x \mathcal{N}\left(1, \frac{\sigma_{\text{add}}}{x}\right). \quad (6.16)$$

When modifying all individual values in a distribution according to this procedure, the effect on the total distribution of values can be described through a convolution with a Gaussian distribution of mean 0 and resolution σ_{add} . Thus, the new distribution has

a mean consistent with that of the old distribution, but the new standard deviation is increased as follows:

$$\sigma_{\text{new}} = \sqrt{\sigma_{\text{old}}^2 + \sigma_{\text{add}}^2}, \quad (6.17)$$

where σ_{old} is the standard deviation of the old distribution. Thus, to perfectly match the resolution between data (σ) and simulation (σ'), the required additional smearing is given by:

$$\sigma_{\text{add}} = \sqrt{\sigma^2 - \sigma'^2}. \quad (6.18)$$

Here, it was assumed that the standard deviation in the data is larger than in the simulation. This assumption is usually true because the simulation by construction only considers effects that are understood (as they are fed into the simulation), while the data may also contain unmodeled effects. Figure 6.4 illustrates the general concepts of the scale and resolution correction with a hypothetical example.

Unlike indicated by the hypothetical example, the true value for the real data is usually unknown, making a direct calibration challenging. A more practical approach is to iteratively adjust correction parameters and assess the impact on agreement between simulation and data. Dimuon resonances are particularly useful for such an approach, as they provide a well-defined reference point: the peak position of the resonance reflects the momentum scale, while the peak width is related to the momentum resolution. This makes them an ideal object for validating and refining calibration procedures.

6.3.2 Correction Procedure

The muon momentum correction procedure presented in this work is based on the method described in Reference [3], and was refined in the scope of this thesis in collaboration with a master's thesis [80]. The procedure consists of two main steps:

1. Extraction of information on the scale and resolution based on distributions of single-muon quantities.
2. Refining of information from the first step by gradually improving the agreement between data and simulation in the dimuon mass spectrum close to a resonance.

This process is also illustrated in Figure 6.5, reflecting that in each of the main steps the scale is targeted before the resolution.

The corrections are derived separately in two phase spaces, where the properties of resonances can be fully exploited:

- Medium- p_{T} phase space: this is the main target of this part of the work. It can be investigated at the Z boson resonance, using standard triggers in the main data stream.
- Low- p_{T} phase space: this serves mainly as a cross check of the method validity of the results, but may become of greater value in the future. Here, the J/ψ resonance is used from a parking data set, as the corresponding low- p_{T} triggers are prescaled in the main data stream.

To focus on these processes while reducing background contributions, both phase spaces require the presence of two opposite-sign muons with a dimuon mass close to the corresponding resonance peak position. Table 6.1 summarizes the selection requirements for the two regions.

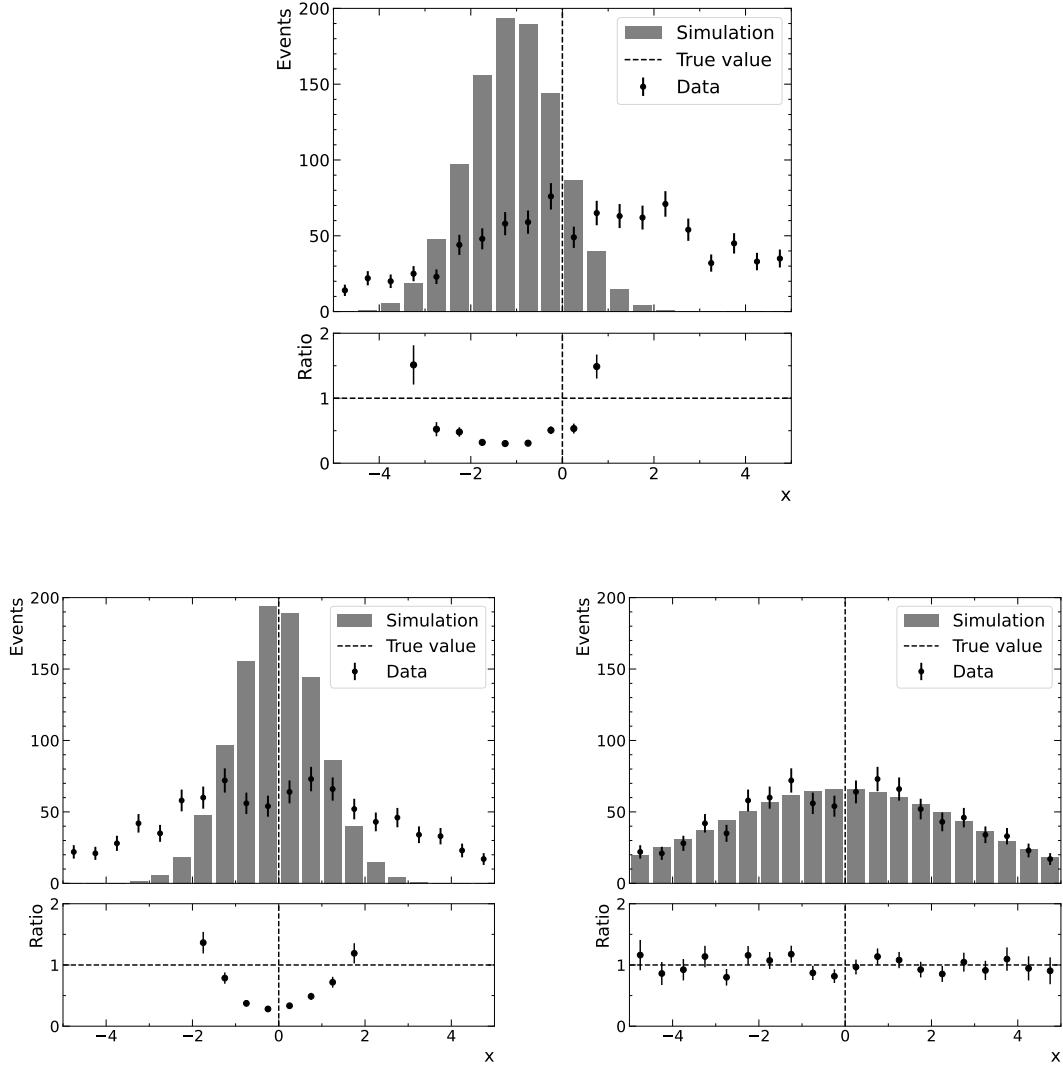


Figure 6.4: Simplified example of the correction procedure applied to a hypothetical measurement of a specific true value (here 0). At the top, the original distributions of the measurement of a certain true value in data and simulation are shown. The scale of the distributions is corrected by applying a shift by the size of $x_{\text{true}} - \mu$ to each value, where μ is the mean value of the corresponding distribution. The result of this correction is shown at the bottom on the left and yields good agreement of the central value. The resolution is then optimized by multiplying each value in the simulated sample with a random value from a normal distribution, with mean one and standard deviation σ_{add} .

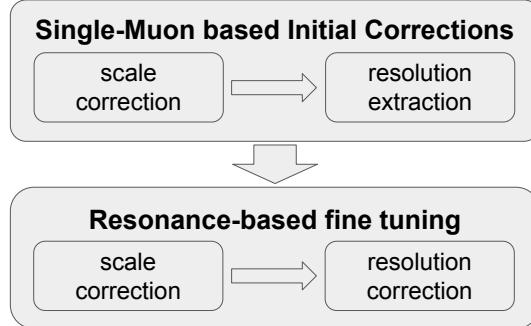


Figure 6.5: Scheme for the derivation of the muon momentum scale and resolution corrections in the CMS experiment.

Table 6.1: Baseline selection requirements for the derivation of the muon scale and resolution corrections at low and medium p_T . The selections are refined in the individual correction steps, e.g., using a smaller range around the dimuon mass.

Requirement	Low- p_T region	Medium- p_T region
Trigger	HLT_DoubleMu4_3_LowMass	HLT_IsoMu24
Offline p_T (GeV)	> 4	> 26
Muon $ \eta $	< 2.4	< 2.4
# muons	2	2
Sum of muon charges	0	0
Relative muon isolation	—	< 0.15
Identification	soft	tight
Dimuon mass (GeV)	$\in [2, 4]$	$\in [50, 130]$

To minimize systematic differences between data and simulation that are not a result of deficiencies in the muon p_T description, several corrections are applied: **NLO** samples are used, wherever possible, to better approximate nature’s infinite-order behavior. Muon efficiency scale factors are applied to the simulation to align the efficiency description between data and simulation. Furthermore, the pileup distribution is reweighted to match the expectation from the measured luminosity and the **pp** cross section. Additionally, the transverse momentum of the dimuon system is reweighted to match the data, thereby covering differences in the recoil descriptions arising from missing higher-order corrections. In the Z region, background contributions are generally subtracted from the data. As the J/ψ region serves as a proof of concept rather than being as optimized as the medium- p_T region, the background contribution is not subtracted there. The missing background subtraction in the low- p_T region presents a systematic effect but is not considered a general show-stopper for future developments.

6.3.2.1 Single-Muon-Based Scale Correction

After event selection and simulation-tuning, the first step in the correction procedure aims to correct the muon p_T scale. As the muon p_T is reconstructed from the track curvature using Equation 4.3, two primary sources can lead to momentum scale effects:

- Magnetic field mismodeling: affects the radius of the trajectories independent of the charge sign.
- Detector misalignment: leads to deviations in the track reconstruction, which appear as charge-dependent effects.

Additionally, approximations in the simulation—while computationally efficient—may introduce minor biases. However, the corrections introduced in this thesis focus on first-order effects and therefore neglect such simulation approximations.

The effect of alignment issues is assumed to be additive in the curvature. As illustrated by Figure 6.6, the effect has a different sign for differently charged muons. Since the curvature is related inversely to the transverse momentum, the alignment effect is additive in $1/p_T$. Given Equation 4.3, differences in the magnetic field can be propagated to the momentum as a multiplicative term. The size of this term is equal for both positive and negative muons, as an increased magnetic field would lead to a larger muon track curvature independently of the charge. In summary, the naive correction scheme follows the relation:

$$\frac{1}{p_T} \rightarrow \frac{M}{p_T} + qA, \quad (6.19)$$

with a multiplicative correction term M and an additive term A . Since both effects generally vary with the detector position, the **CMS** detector is divided into bins of η and ϕ to derive separate correction parameters for each region.

As explained in Section 6.1, the p_T of muons from resonances follows a specific shape, governed by the interplay of the resonance mass and its boost. Assuming the recoil and parton momentum imbalance are described correctly by the simulation, a set of initial corrections can be extracted directly from Eq. 6.19. For this, it is assumed that the parameters A and M remain constant within each η - ϕ bin, independent of the charge. Then, taking the mean of the $1/p_T$ distribution for both muons, the correction parameters can be estimated by solving the system of linear equations:

$$C_{\text{new},\pm} \stackrel{!}{=} M \cdot C_{\text{old},\pm} \pm A, \quad (6.20)$$

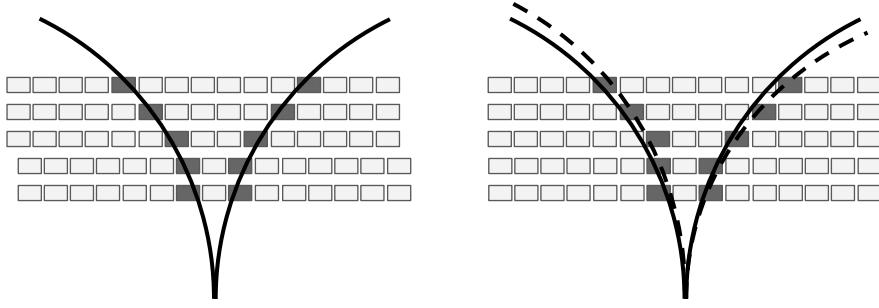


Figure 6.6: Sketch of the impact of misalignment on the track curvature. The left part of the sketch displays the true but unknown alignment of a simplified tracker, showing a small shift in the positioning of several layers. True tracks for a positive and negative muon at a certain momentum are represented by the black uninterrupted lines. The darker tracker elements correspond to the hits from the true muons. The right part of the sketch displays the alignment in reconstruction, where the shift is not accounted for. As a result of the misalignment unaccounted for in reconstruction, the reconstructed tracks (represented by dashed lines) are not aligned with the true tracks anymore. In particular, the impact on the curvature has a different sign for the two differently charged muons.

where $C \equiv \langle 1/p_T \rangle$ for better overview, and the labels 'new' and 'old' depend on the chosen reference. In this analysis, both data and reconstructed simulation are corrected to match the generator-level distribution. This is the most general approach, as other transformations (e.g., correcting only data to simulation) can be obtained from these parameters. The solution of this system of equations is given by:

$$M = \frac{C_{\text{new},-} + C_{\text{new},+}}{C_{\text{old},-} + C_{\text{old},+}} \quad (6.21)$$

$$A = C_{\text{new},+} - C_{\text{old},+} \cdot \frac{C_{\text{new},-} + C_{\text{new},+}}{C_{\text{old},-} + C_{\text{old},+}}.$$

Generally, the accuracy of this step is limited by the assumption that the modeling of the parton momentum imbalance and the recoil is sufficient. Consequently, it is only used as an initial correction and fine-tuned in a further step with the help of a resonance.

6.3.2.2 Single-Muon-Based Resolution Extraction

The approach for obtaining resolution corrections differs from that used for scale corrections due to their different cause. While scale corrections address systematic shifts in measured momenta, resolution corrections focus on describing the spread of reconstructed values around their true momenta. In the resolution correction procedure, the first step consists of the parametrization of the muon p_T resolution in simulation. The resolutions are then tuned to achieve optimal agreement between the dimuon resonance in the simulation and in the data.

The resolution of simulated muons is obtained by comparing the true (generator-level) p_T to the reconstructed value. This comparison is mathematically expressed by calculating the ratio R for each muon:

$$R = \frac{p_{T\text{gen}}}{p_{T\text{reco}}}. \quad (6.22)$$

If the resolution is high, then the reconstructed momenta closely match the generated values, resulting in a distribution of R with a small variance.

Following the discussions in Section 4.2 leading to Equation 4.4, the muon p_T resolution depends on multiple factors. Apart from its intrinsic p_T dependence, additional dependencies can be effectively formulated as a function of $|\eta|$. Consequently, the variance of R values is computed in bins of $|\eta|$, its dependence on p_T is extracted using a second-order polynomial fit, which includes the linear expectation from the Gluckstern Equation 4.4 as a special case but was observed to describe the behavior better in some bins. However, binning solely in $|\eta|$ may not fully capture effects from a difference in the number of hits in the tracker layers. If a muon leaves fewer hits in the tracker, then the corresponding p_T measurement is less precise than that of a similar muon with more hits. To account for this, the resolution is further binned in the number of tracker layers n_L .

The extracted standard deviations alone are not sufficient, as the distribution of R values is not purely Gaussian. To account for residual effects, a double-sided Crystal Ball function [81, 82] is fit to the pull distribution

$$\text{Pull} = \frac{R - \langle R \rangle}{\sigma(R)}, \quad (6.23)$$

where $\langle R \rangle$ and $\sigma(R)$ denote the mean and standard deviation of R . The double-sided Crystal Ball function contains a Gaussian core with power-law tails on both sides, making it well suited for modeling non-Gaussian effects such as radiative energy losses.

Figure 6.7 illustrates example fits for both the quadratic parametrization of the resolutions and the double-sided Crystal Ball fit applied to the pull distribution. While the resolution fits in the shown bins exhibit a linear behavior, this is not representative for all bins. In particular, less-populated bins often show non-linear effects, that are better captured by the quadratic fit model compared to the idealized linear one. Nevertheless, the quadratic fit provides a good model of the underlying linear behavior in the presented bins due to the sufficient number of events in these bins. Closure plots for the resolution parametrization are presented in Figure 6.8 and compare the original and smeared generator-level dimuon mass distributions to the reconstruction-level distribution. In both regions, the dominant share of the resolution difference is mitigated after the corrections, while a small discrepancy remains, especially in the scale. These remaining differences are targeted in the following fine-tuning steps.

6.3.2.3 Resonance-Based Scale Fine-Tuning

Unlike the p_T spectrum of individual muons, the invariant mass of a resonance does not explicitly depend on the modeling of its momentum. However, the reconstructed mass of the dimuon system is derived from the single muon momenta using the formula:

$$m \approx \sqrt{2p_{T1}p_{T2}[\cosh \Delta\eta - \cos \Delta\phi]}, \quad (6.24)$$

which follows from the definition of the dimuon mass and relativistic approximations. In this equation, p_{T1} and p_{T2} are the relativistic momenta of the two muons, while $\Delta\eta$ and $\Delta\phi$ represent their pseudorapidity and azimuthal separation, respectively. Thus, in contrast to the true invariant mass of the resonance, the reconstructed mass implicitly depends on the modeling of the momentum. However, this dependence arises solely due to the position of the muons in the detector. Since the scale corrections are generally derived and applied in bins of p_T and η , this residual correlation should be eliminated to a large extent, allowing for highly precise fine-tuning of the scale corrections at the resonance level.

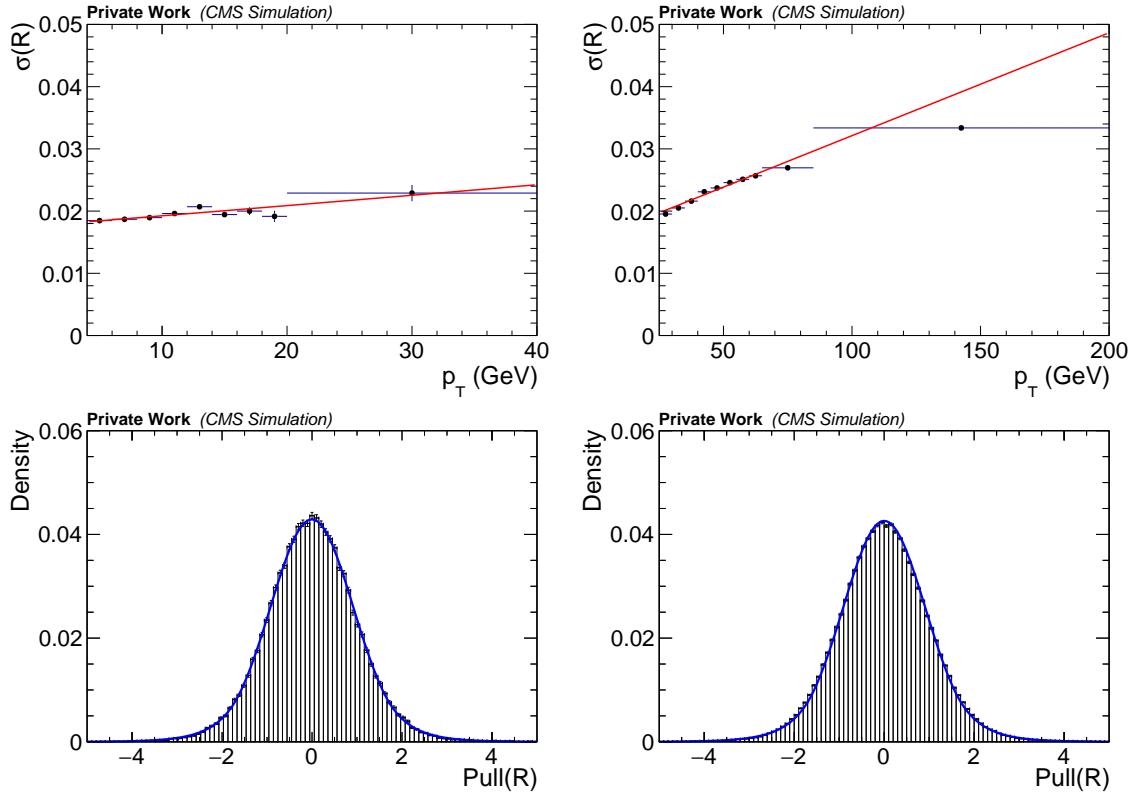


Figure 6.7: Example fits in the resolution extraction steps in the J/ψ region (left) and in the Z region (right). Muons in these fits are characterized by hits in 13 tracker layers, and $\eta \in [1.6, 1.8]$. At the top, the standard deviation values of the distributions R are shown as data points. The dependence on the momentum is extracted with a quadratic fit, represented by the red line. The bottom plots show the modeling of the non-Gaussian behavior of the R distribution. Here, the histograms represent the normalized pull distribution of R , which is modeled with a double-sided Crystal Ball function (blue line). The modeling between the two regions is generally consistent in this bin.

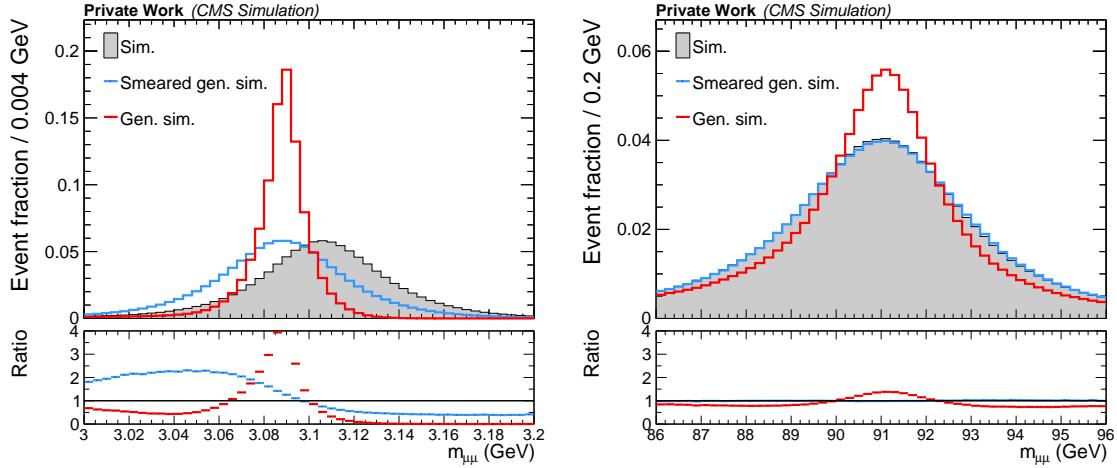


Figure 6.8: Distributions of the dimuon mass in the J/ψ region (left) and in the Z region (right). The upper panels show the normalized distributions of the reconstructed, the smeared generated, and the original generated dimuon masses. The lower panels display the ratio of the generator-level distributions to the reconstruction-level distribution.

Despite the advantages of using a dimuon resonance for calibration, the correlations of scale adaptations for both muons must be addressed simultaneously. A naive approach that optimizes the muon p_T scale for only one muon at a time is inherently flawed, as the second muon would remain uncorrected. Therefore, it is crucial to de-correlate the two muons in order to enable an accurate scale optimization.

Building on the assumption from Eq. 6.19 that the p_T scale correction consists of both an additive and a multiplicative component, the impact of this correction on the dimuon mass can be expressed as:

$$m \rightarrow m \underbrace{\frac{1}{\sqrt{M_+ + A_+ p_{T+}}} \frac{1}{\sqrt{M_- - A_- p_{T-}}}}_{\approx 1 + \Delta m/m}, \quad (6.25)$$

where the index $+$ refers to the positively charged muon and the index $-$ to the negatively charged muon from the resonance. The correction parameters depend only on the position of the muon. For instance, the notation A_+ is shorthand for $A(\eta_+, \phi_+)$, where η_+ and ϕ_+ indicate the muon's position in the detector, with analogous notation for the other parameters. Expanding the square root via a Taylor series yields the following expression for the mass shift:

$$\Delta m \approx m \left[1 - \frac{M_+ + A_+ p_{T+} - 1}{2} - \frac{M_- - A_- p_{T-} - 1}{2} - 1 \right] \quad (6.26)$$

$$\Rightarrow -2\Delta m \approx m [M_+ - 1 + A_+ p_{T+} + M_- - 1 - A_- p_{T-}]. \quad (6.27)$$

Thus, the effect of small changes in the alignment or magnetic field on the dimuon mass is parametrized. However, the issue of correlation between the two muons still persists.

To eliminate this correlation, an iterative approach is used to solve Eq. 6.27 under specific assumptions. First, all positive muons within a certain η - ϕ bin are considered, and the average value across these muons is computed:

$$-2\langle \Delta m \rangle = \langle m \rangle (M_+ - 1) + \langle m p_{T+} \rangle A_+ + \langle m (M_- - 1) \rangle - \langle m p_{T-} A_- \rangle. \quad (6.28)$$

Here, all positively charged muons share the same correction parameters M_+, A_+ , while the negatively charged muons may have different parameters. Since the contribution of the negative muons is relatively small when averaging over all positive muons in the same bin, their effect could, in principle, be neglected entirely. Alternatively, the second muon can also be assumed to fulfill similar properties as the muons in the current bin. Under this assumption, the multiplicative terms double in size, while the additive correction parameters A_{\pm} cancel out. Previous versions of this method used a hybrid approach, assuming that the multiplicative correction for the second muon is the same as for the first, while the additive terms of the second muon cancel on average. Since this approach has been found to perform well, it is also used in this work.

Applying these assumptions and performing the same procedure for the negative muon yields a solvable system of equations:

$$\begin{aligned} -2\langle\Delta m\rangle_+ &= 2\langle m\rangle_+(M-1) + \langle mp_T\rangle_+ A, \\ -2\langle\Delta m\rangle_- &= 2\langle m\rangle_-(M-1) - \langle mp_T\rangle_- A. \end{aligned} \quad (6.29)$$

Here, Δm represents the difference between the target mass, which follows the smeared generator-level distribution, and the reconstructed mass at the current iteration. Since the target values are independent from the iteration, the mean of the differences is equivalent to the difference of the means. To ensure unbiased starting values, the initial corrections are used as an input for the first iteration.

The updated parameters are determined by solving the equation system:

$$\begin{aligned} M &= 1 - \left(\frac{\langle\Delta m\rangle_+}{\langle mp_T\rangle_+} + \frac{\langle\Delta m\rangle_-}{\langle mp_T\rangle_-} \right) / \left(\frac{\langle m\rangle_+}{\langle mp_T\rangle_+} + \frac{\langle m\rangle_-}{\langle mp_T\rangle_-} \right), \\ A &= \left(\frac{\langle\Delta m\rangle_+}{\langle m\rangle_+} - \frac{\langle\Delta m\rangle_-}{\langle m\rangle_-} \right) / \left(\frac{\langle mp_T\rangle_+}{2\langle m\rangle_+} - \frac{\langle mp_T\rangle_-}{2\langle m\rangle_-} \right). \end{aligned} \quad (6.30)$$

Due to the assumptions regarding the second muon, multiple iterations are required to reach the target values. With each iteration, the impact of these assumptions on the dimuon mass decreases, as the description of the second muon is also updated and moves closer to the target.

While the mean is commonly used and a robust estimator for central values, it can become biased when applied to asymmetric distributions, such as those encountered in dimuon resonance spectra. For example, the peak of the Z boson is systematically shifted towards lower mass due to **FSR** effects, where photons carry away a fraction of the energy that is not recovered in the dimuon mass reconstruction. Furthermore, the spectrum includes off-shell photon contributions that create a smoothly falling continuum beneath the resonance. Since these low-mass contributions are indistinguishable on an event-by-event basis, they further distort the distribution and shift the mean away from the true peak position.

To reduce this bias, the dimuon mass window is restricted to a narrow window around the resonance. For more precise measurements, the peak position could instead be determined by fitting the mass spectrum with a suitable model—such as a Breit-Wigner function convolved with a Gaussian or Crystal Ball function.

Figure 6.9 shows the evolving value of the mean of the reconstructed dimuon mass distribution with respect to the mean of the generator-level distribution over 20 iterations. The impact of the correction on the reconstructed distributions in data and simulation is shown by comparing the dimuon resonance mass distribution before and after applying the correction in Figure 6.10.

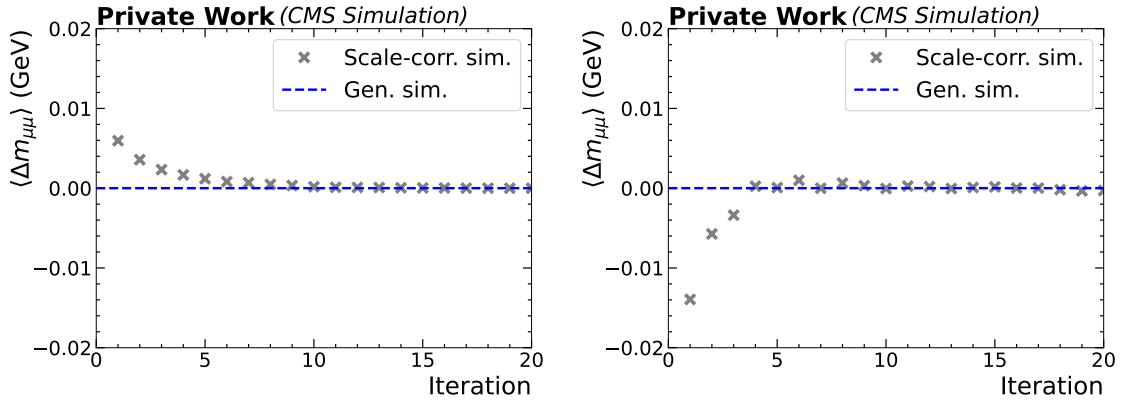


Figure 6.9: Example plots of the evolution of the mean difference between the dimuon mass distribution at the reconstruction-level and generator-level. Both distributions are from the same correction bin, defined by $-2.2 < \eta < -1.85$ and $-2.8 < \phi < -2.4$. The left plot is from the J/ψ region, while the right plot is from the Z region. Both converge towards zero after a few iterations.

6.3.2.4 Resonance-Based Resolution Fine-Tuning

As a final step, the resolution parametrization from the simulation is tuned using the dimuon spectrum. A scaling factor k is applied in bins of $|\eta|$ to adjust the width of the smeared generator-level distribution to match either data or reconstructed simulation. The corresponding scaling factors, k_{data} and k_{sim} are determined by minimizing the χ^2 between the additionally smeared generator-level and the reconstructed distribution (either data or simulation). Since no analytical solution is available, the optimal k values are found via an iterative grid refinement:

1. the χ^2 value is calculated for five different values of k , evenly spaced between 0.9 (slightly decreased smearing) to 1.4 (largely increased smearing).
2. the two values of k with the highest χ^2 values are discarded, and two new intermediate values are introduced between the remaining three.
3. This procedure is repeated five times, refining k to a precision of $\Delta k \approx 0.2\%$, halving the step size in each iteration.

For applying these corrections, two strategies can be used. One applies the tuned resolution to generator-level muons, yielding a momentum distribution aligned with data but potentially missing detector-specific effects. Alternatively, additional smearing is applied to reconstructed muons in simulation, accounting more realistically for detector response. In this case, the required extra smearing is smaller than for the generator-level muons, since the reconstructed muons are already affected by detector resolution. The combined effect is handled via Gaussian convolution properties:

$$k_{\text{add}} = \sqrt{k_{\text{data}}^2 - k_{\text{sim}}^2}. \quad (6.31)$$

In case $k_{\text{data}} < k_{\text{sim}}$, no additional smearing is applied.

The distributions of the dimuon mass in the resonance region before and after the application of all corrections are provided in Figure 6.11. In general, the discrepancies between the data and the simulated distributions are mitigated to a large extent.

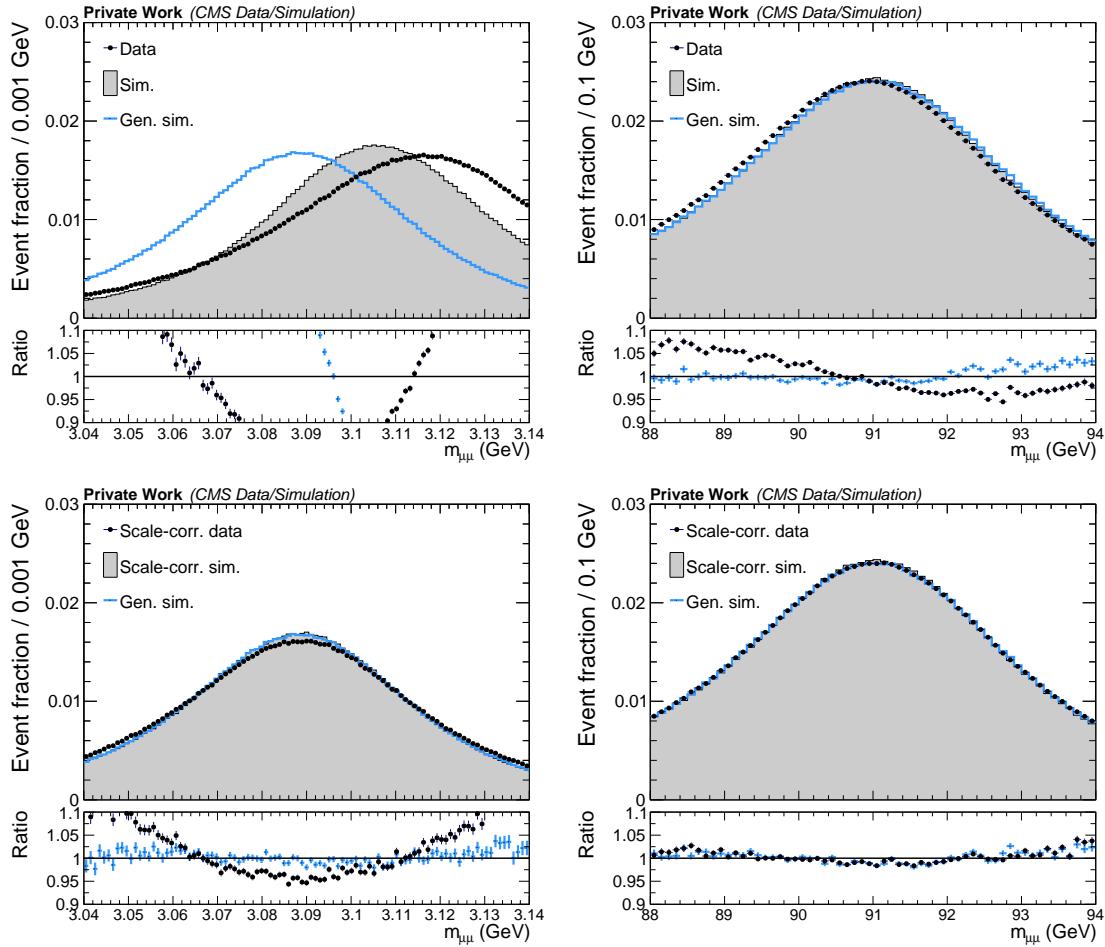


Figure 6.10: Dimuon mass distributions in the J/ψ region (left) and in the Z region (right). The plots show the normalized distributions of the data, and simulation, where the latter is present at both generator- and reconstructed level. While the top plots show the distributions before the iterative scale correction, the bottom plots display the same distributions after the correction. In addition to the normalized distribution, their ratios are provided in the panels below. Uncertainty bars include only the statistical uncertainty due to the limited size of the dataset.

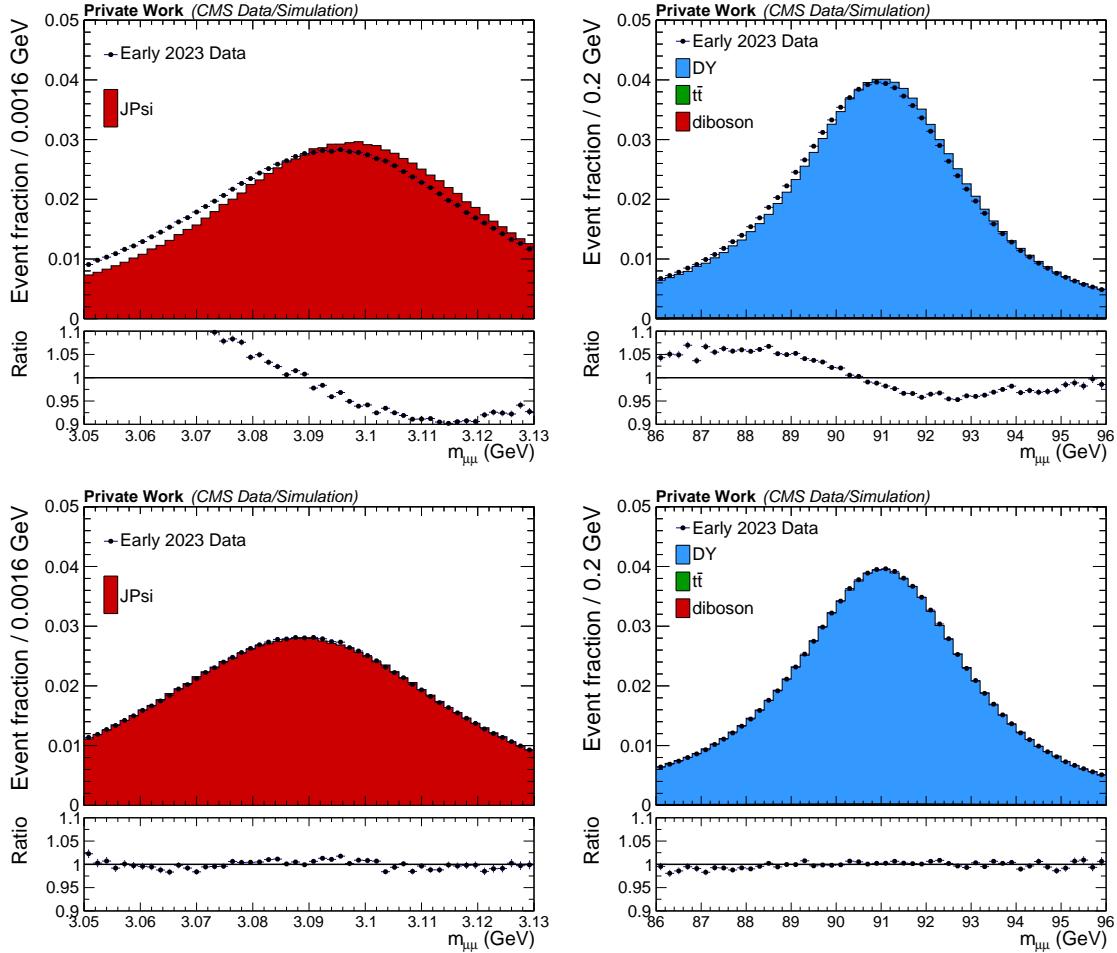


Figure 6.11: Distributions before (top) and after (bottom) the full correction procedure for the low-mass (left) and the high-mass region (right). The upper panels display the normalized distributions of the data and the normalized distribution of the simulated samples, while the lower panels show the ratio of the normalized data distributions and the combined simulated distributions.

6.3.3 Uncertainty Estimation

The scale and resolution corrections are encoded in multiple parameters, making it essential to carefully account for possible correlations when assessing their impact on results. However, if these correlations are small, they can be neglected to simplify the overall model.

Several factors contribute to uncertainties in the correction parameters: on the one hand there are statistical uncertainties that are primarily driven by the sample size. On the other hand, the result is influenced systematically, which can have various reasons, such as: the choice of binning used to obtain the parameters, the propagation of further systematic uncertainties (for instance from the selection efficiency scale factors), and uncertainties related to limitations of the fit models. Since there are no strong reasons to believe in large correlations between these sources, the total uncertainty can be obtained by summing them in quadrature. Hence, the largest uncertainties dominate the total uncertainty and smaller sources of uncertainty can be neglected without major underestimations of the total uncertainty. Therefore, second order effects such as the uncertainty on the efficiency scale factors (primarily affecting normalization with a limited impact on shapes) are not evaluated in this thesis.

In summary, the uncertainty estimation can be simplified by neglecting correlations between different parameters, and by neglecting small effects in general. Both simplifications, however, need to be applied with care, as they introduce certain biases. Whether these biases are acceptable depends on the magnitude of the effect and the required precision of the results.

In the following, first statistical and systematic uncertainties on the correction parameters are examined. Then, a simplified approach, which drastically reduces the complexity and computational demands, is presented and compared to the full propagation of the uncertainties.

6.3.3.1 Uncertainty Derivation

To estimate the statistical uncertainties of the correction parameters, a bootstrapping technique is applied, following the principles introduced in Section 5.2.2. This procedure consists of $n = 100$ independent repetitions of the correction procedure with different samples bootstrapped from the original distributions. The resulting parameter distributions can then be used to estimate the statistical uncertainty in the correction procedure. The choice of n balances the computing demands and the requirement for precision in the uncertainty of the results; for this choice of n , the uncertainty on the standard deviation of the parameter distributions obtained with the bootstrapping technique corresponds to roughly 7%.

The systematic uncertainties from the choice of histogram range are evaluated by varying the range. Second-order effects from uncertainties on the corrections applied during the derivation of the correction parameters are evaluated and found to be smaller than those from the range variation.

The statistical and systematic uncertainties on the parameters of the fine-tuning correction steps are illustrated in Figure 6.12. One key observation is that the statistical uncertainty dominates over the systematic uncertainties across different bins. Following the line of arguments introduced in Section 6.2.3, the systematic uncertainties can therefore be neglected in the further discussion.

In addition to the aforementioned uncertainties, a potential bias in the simulation model can be evaluated by comparing the correction parameters derived from different simulation

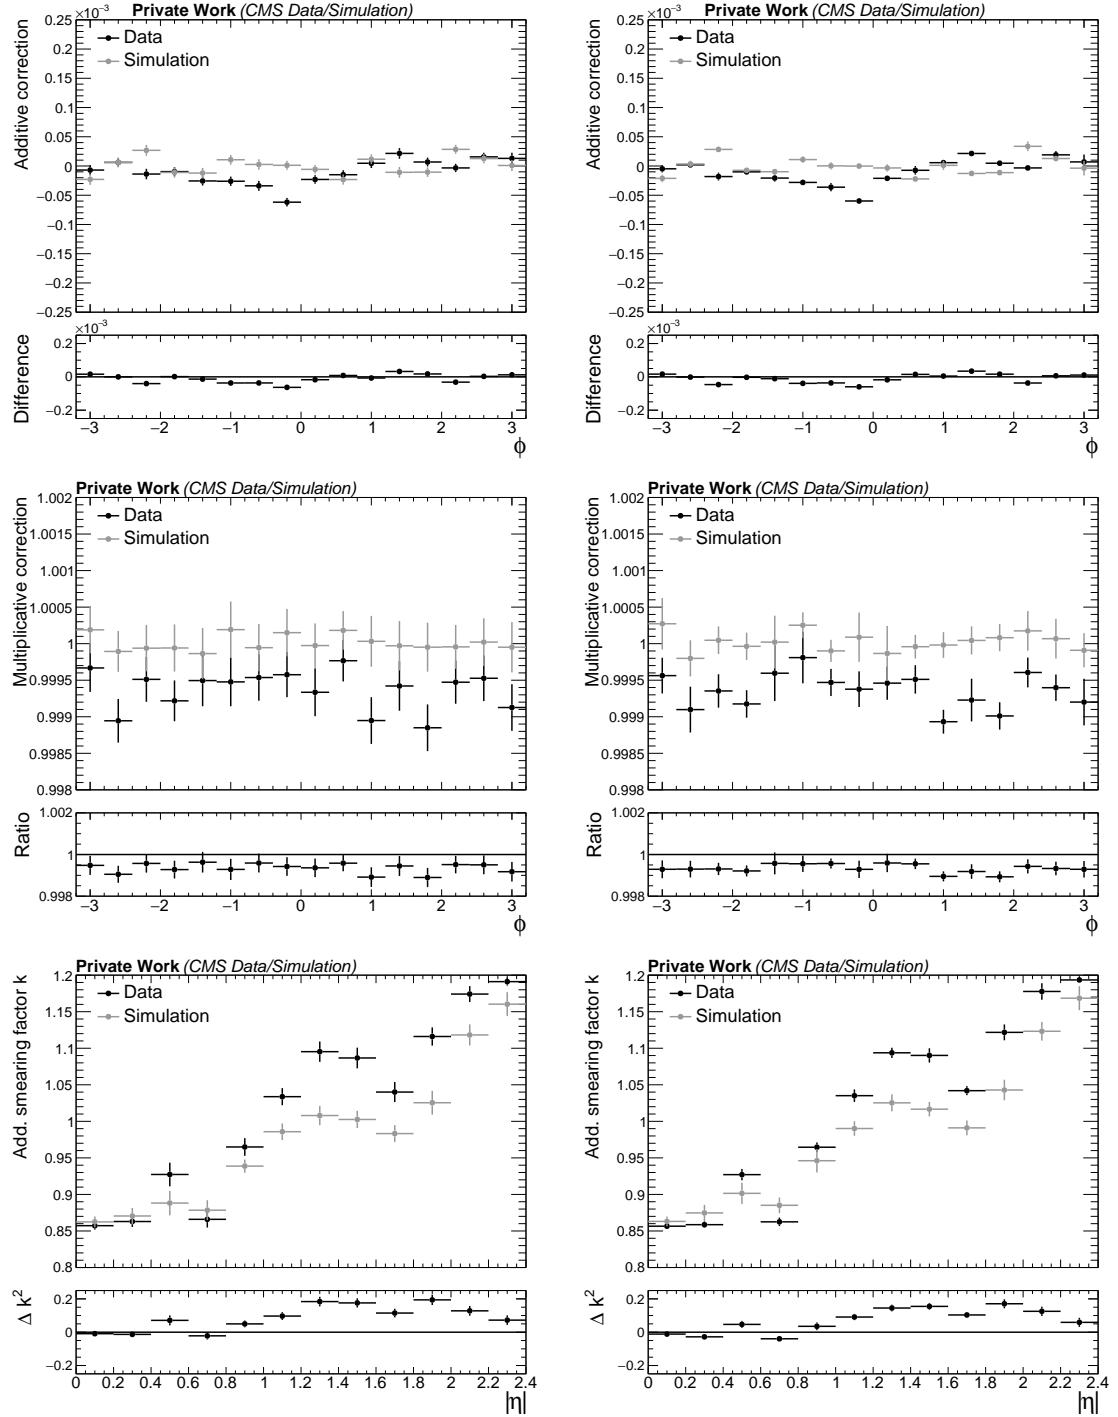


Figure 6.12: Mean and standard deviation values of the parameter distributions from bootstrapping (left) and dimuon mass range variation in the fine-tuning corrections (right) in the Z resonance region. Each row represents a parameter, with distributions in the upper panel for data and simulation. The lower panels show the correction needed for simulation if the data is left uncorrected. The top and middle rows display the additive and multiplicative parameters, for muons with $-0.4 < \eta < 0$ as a function of ϕ . The bottom row shows the additional smearing factor k that needs to be applied to pre-smeared generator muons to match the Z boson peak width in data and simulation. No additional smearing is applied to the simulation if $\Delta k^2 < 0$.

samples. Since the scale and resolution corrections are specifically intended to address effects related only to the muon object’s scale and resolution, it is essential to use the most precise simulation samples available. In this context, samples generated at **NLO** precision in **QCD** are considered suitable. At the time of the correction derivation, however, only one centrally available **NLO** precision sample existed within the collaboration, along with additional samples with **LO** precision in **QCD**. Using **LO** samples to estimate the impact of simulation modeling would lead to an artificial overestimation of the associated uncertainty, as the discrepancy between **LO** and **NLO** is typically much larger than that between **NLO** and **NNLO**. For this reason, such uncertainties are not considered in this thesis. In general, the evaluation of this uncertainty is straightforward to implement and can be studied once sufficiently large samples with at least **NLO** precision in **QCD** become available.

6.3.3.2 Application of Uncertainties

Since statistical uncertainties dominate for the muon scale and resolution corrections, the focus in the uncertainty application is placed on them. Other sources of uncertainty can be incorporated in a similar or even more straightforward manner.

The bootstrapping method is applied to produce $n = 100$ self-consistent sets of correction parameters. In principle, all of these variations could be provided to the collaboration and applied directly in analyses. To quantify the impact on the corresponding final observable, each analysis would need to be repeated n times; once per correction parameter set. The effect of the correction could then be estimated from the distribution of results obtained through this procedure. While this approach has been used in the past, it requires substantial computing resources—an often unnecessary cost due to the typically negligible impact of the associated uncertainties in most analyses.

To address this, a simplified approach is introduced in this thesis, reducing the computational demand of the uncertainty evaluation for analyzers by a factor of 25 (for the target uncertainty precision of 7%; even more for higher precision in the uncertainty). Instead of performing $\mathcal{O}(100)$ computations for each bootstrapped sample, only four calculations are needed: one upward and one downward variation each for both scale and resolution corrections.

The simplified approach is based on two main assumptions, which ultimately allow simple propagation of uncertainties:

1. **Independence of scale and resolution corrections:** this assumption is well-motivated since scale corrections primarily affect the means of distributions, whereas resolution corrections impact their standard deviations.
2. **Neglecting uncertainties from initial steps:** this assumption is well justified for scale corrections, as the fine-tuning step directly updates the initial scale parameter sets, largely reducing the impact of the original single-muon based corrections. For resolution corrections, however, this assumption is less robust, since the final correction is defined by both the initial resolution parametrization and the subsequent fine-tuning. Therefore, it is particularly important to validate the results built on this assumption.

Both assumptions largely reduce the complexity of the uncertainty model. In particular, ignoring the uncertainty on the resolution parametrization is crucial for the simplified approach. This is due to the non-analytical step of sampling a random number that follows a double-sided Crystal Ball function. Consequently, the related uncertainty cannot be propagated.

With these simplifications, the number of correction parameters that are subject to uncertainties is reduced to three: two for the scale correction, and one for the resolution correction. Since the scale correction is parametrized using two parameters, their correlation must be taken into account. The uncertainty propagation formula for the scale correction uncertainty is given by:

$$\sigma_{p_T} = \frac{1}{(\frac{M}{p_T} + qA)^2} \sqrt{\frac{1}{p_T^2} \sigma_A^2 + \sigma_M^2 + \frac{2}{p_T} \rho_{A,M} \sigma_A \sigma_M} \quad (6.32)$$

$$\approx p_T^2 \sqrt{\frac{1}{p_T^2} \sigma_A^2 + \sigma_M^2 + \frac{2}{p_T} \rho_{A,M} \sigma_A \sigma_M}, \quad (6.33)$$

where σ_A and σ_M are the standard deviations of the corresponding parameter samples, and $\rho_{A,M}$ is the correlation coefficient between A and M .

While the scale correction can be implemented in different ways—correcting simulation to data, data to simulation, or both to the generator-level distribution—the associated uncertainties are typically assigned only to the simulation. Consequently, even if the data are corrected to match the simulation, the uncertainty from this procedure must still be applied to the simulation. This is implemented by evaluating the uncertainty based on the parameter values that would correct the simulation to the data, regardless of the specific correction strategy. By construction, the dominant contributions to this uncertainty arises from the (statistical) discrepancies between the data and signal distributions, with the generator-level distribution playing a less significant role. This behavior aligns well with the primary correction goal of correcting residual differences in the muon p_T between data and simulation.

For the resolution correction, this inconsistency in the treatment of corrections and the application of their uncertainties does not arise, as the resolution corrections are applied exclusively to muons in simulation. Nevertheless, the uncertainty in the resolution correction consists of two components as well: one originating from the fit of the pre-smeared generator-level dimuon distribution to data, and one from the fit to the reconstruction-level simulation. Similarly to the scale correction, the uncertainty can be evaluated without separately considering the individual contributions k_{data} and k_{sim} , but instead by evaluating the uncertainty in the effective correction parameter k_{add} . This approach accounts for statistical correlations between the two individual components, which may arise from the common parametrization used to pre-smear the generator-level distributions. It is also consistent with the assumption (still to be validated) that the final correction is independent of the specific parametrization.

To validate the simplified uncertainty estimation, its results are compared to those obtained from the bootstrapped approach for individual muons. Following the distinction between scale and resolution corrections, the two components are evaluated separately.

The bootstrapped approach to evaluate the scale correction involves of two steps: first, a distribution of scale corrected muons is obtained by applying each individual correction set to the uncorrected muon. Then, the standard deviation of this distribution is taken as the uncertainty estimate. In the simplified approach, the uncertainty is estimated by applying the scale correction twice—once using the upward variation of the parameters and once using the downward variation. The difference between these two corrected values serves as an estimate of the uncertainty in the scale correction.

For the resolution uncertainty, both approaches require an additional step: estimating the resolution corresponding to each specific parameter set. This is done by applying the

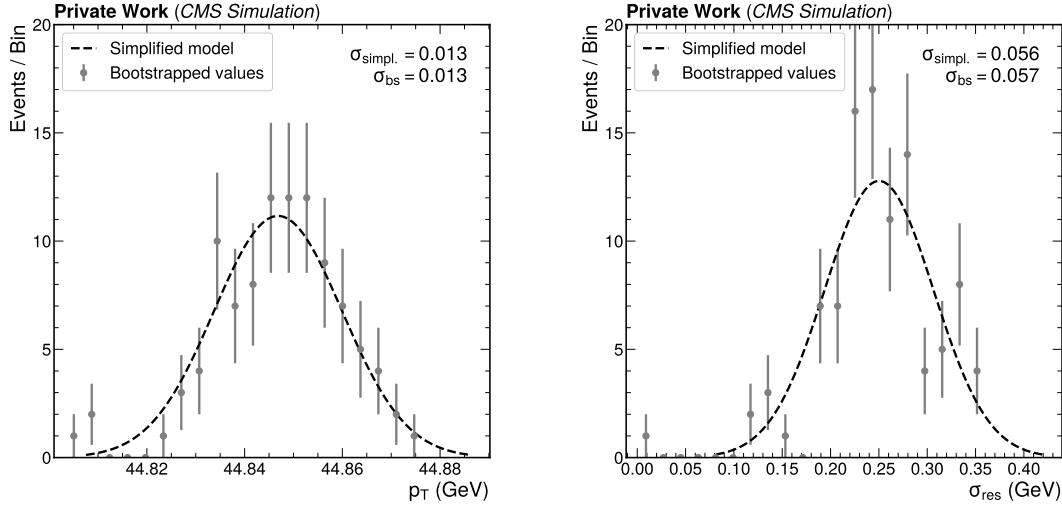


Figure 6.13: Exemplary comparison plots for the calculation of the uncertainty in the correction of the scale (left) and resolution (right). Each point corresponds to the evaluation of the scale or resolution with one bootstrapped set. The simplified model is shown as a Gaussian distribution, where the standard deviation is calculated from the difference between the upward and the downward variation of the simplified method. To improve the visual comparison of the distribution widths, the mean value of the Gaussian does not correspond exactly to the nominal correction value, but to the mean of the bootstrapped distribution.

resolution correction 100 times per muon and calculating the standard deviation of the resulting distribution. The remaining steps are conducted similarly to the scale correction: the distribution of resolution estimates from the bootstrapped samples is compared to the up and down variations derived from the simplified method.

Examples of this comparison are shown in Figure 6.13 and show good agreement between the p_T distribution obtained by correcting the original p_T value with the different bootstrapped correction parameter sets, and the simplified method. This good level of agreement is representative for muons in all detector regions for the scale correction. For the resolution correction, on the other hand, the good agreement is only representative for bins in which the additional smearing factor in the data is significantly larger than that in the simulation. In the other case, which for example affects muons in the central detector region of $|\eta| < 0.4$, the resolution is not corrected at all and therefore does not provide meaningful uncertainty variations. Figure 6.14 shows the impact of the individual uncertainty variations of the scale and resolution corrections on the dimuon mass distribution.

6.3.4 Consistency Checks

In addition to the self-consistency check of the correction, cross-validations can be performed using different momentum regions. However, these checks must be interpreted carefully, as differences in muon properties across p_T regions may dominate over discrepancies between data and simulation.

In particular, low- p_T muon paths have a considerably larger curvature, making them less sensitive to misalignment effects, which are more prominent at high p_T , where curvature

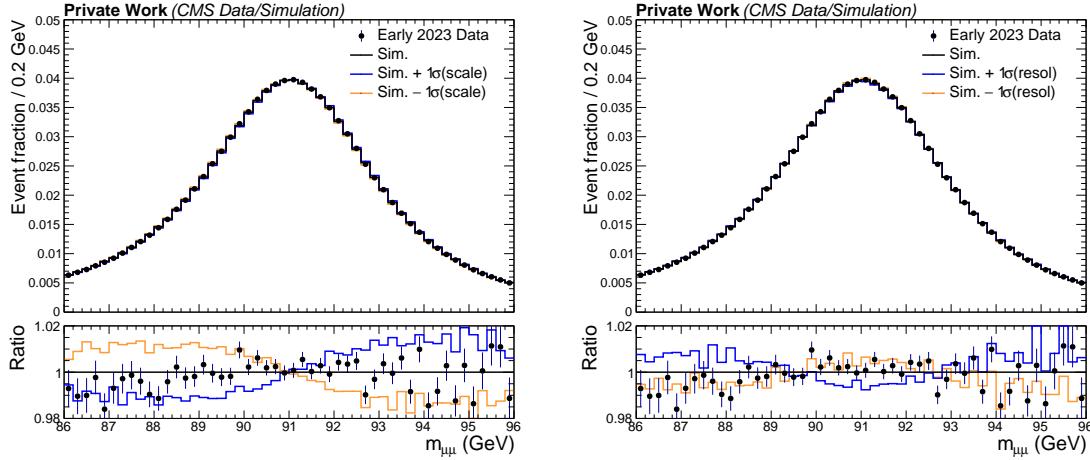


Figure 6.14: Distributions of the corrected dimuon mass including variations of the statistical uncertainty in the scale (left) and resolution (right) parametrization. The lower panels display the ratio of the normalized data distributions or variations in the simulation to the normalized and nominally corrected simulation. The difference between the nominal simulation and the data are covered by each of the corrections individually.

is small. Effects such as mismodeling of the magnetic field, on the other hand, affect all muon p_T ranges similarly, as the magnetic field governs the relationship between curvature and p_T . Furthermore, any unaccounted effects in the model may be absorbed differently across p_T regions, which carries the potential for inconsistencies.

Figure 6.15 displays example distributions of the scale and resolution fine-tuning parameters in the low- p_T region. In comparison with Figure 6.12, which shows the same distributions in the medium- p_T region, the multiplicative scale correction parameters appear consistent, while the additive parameters show a relative modulation between the two regions. This modulation is partly covered by the significantly larger statistical uncertainties in the low- p_T domain, which reflect the smaller impact of the misalignment on the stronger curved low- p_T muons. The uncertainties in the multiplicative parameters, on the other hand, are generally smaller in the low- p_T region. This behavior is a consequence of the smaller impact from misalignment, allowing to additionally constrain the magnetic field correction. The interplay between the two correction terms is encoded in the correction parametrization from equation 6.19, where the multiplicative term becomes relatively more important for smaller p_T values.

The differences observed in the additional smearing factor between low- and medium- p_T corrections are less intuitive than those for the scale parameters. This is partly due to the fact that resolution fine-tuning is less well motivated and parametrized compared to the scale corrections. Furthermore, it is highly sensitive to the underlying resolution parametrization, since the required additional smearing to match the reconstructed dimuon mass distribution can vary. If, for instance, the parametrization underestimated the true resolution, the extracted smearing factors will be larger. A full consideration of the differences in the resolution correction thus requires the inclusion of the resolution model itself. However, this aspect is beyond the scope of this thesis and is left for future work.

Apart from the comparison in the low- p_T region, a partial crosscheck can also be performed with the high- p_T regime, based on preliminary results from the CMS Collaboration in Run 3 [83]. This independent analysis employs the general endpoint method [77] to extract

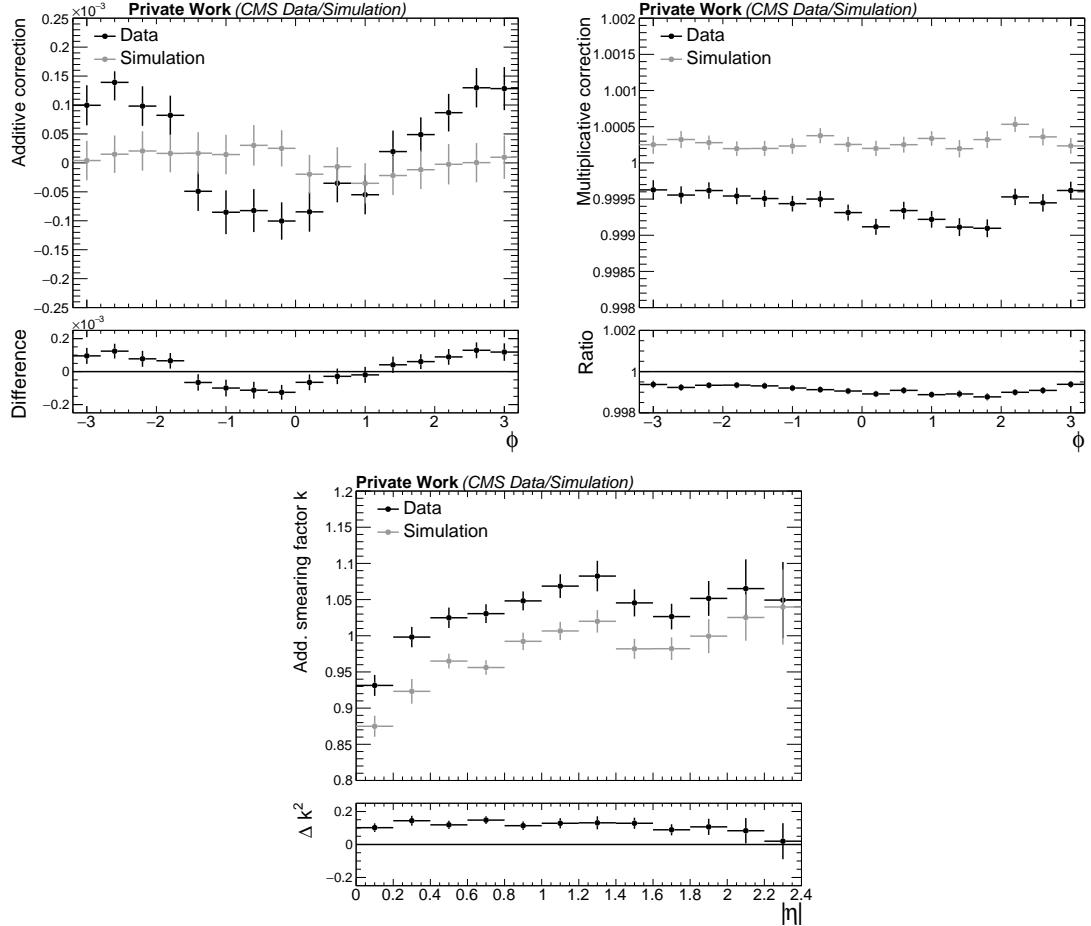


Figure 6.15: Example distributions of fine-tuning correction parameters in the J/ψ region. The top row shows the additive (left) and multiplicative (right) scale correction parameters for muons with $-0.4 < \eta < 0$. The bottom plot displays the additional smearing factors applied to generator-level muons, pre-smeared according to the extracted resolution parametrization. Points in the upper panels represent the correction parameters; uncertainty bars indicate their standard deviation obtained via bootstrapped samples. The lower panels illustrate the behavior of the parameters when applying the correction to simulated muons at reconstruction-level, aiming to match the description of muons in data.

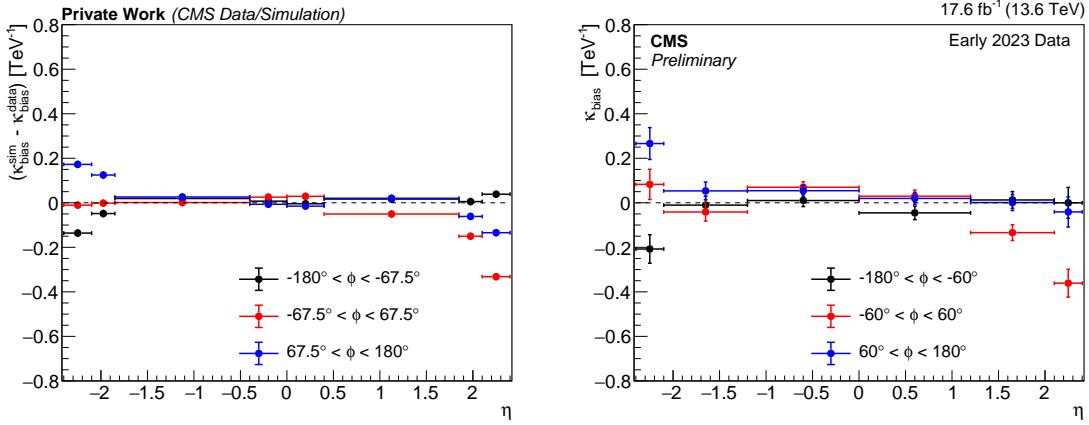


Figure 6.16: Crosscheck of the additive term in the muon scale correction between the medium- p_T region analyzed in this chapter (left) and the high- p_T region (right), taken from Reference [83]. The points indicate the additive curvature correction applied to simulation to improve agreement with data. Unlike the high- p_T analysis, where corrections are extracted by directly comparing reconstructed simulation and data, the medium- p_T corrections are derived by aligning both reconstructed simulation and data to the generator-level simulation. This difference is addressed by considering the difference of the additive curvature terms between simulation and data in the medium- p_T region. Furthermore, the parameters in the medium- p_T region have been re-binned to approximate the binning scheme used in the high- p_T analysis for visual comparison. Statistical uncertainties are indicated by vertical bars in both regions.

scale corrections at high p_T . Conceptually, the method injects an additive term to the muon curvature in the simulation on reconstruction-level and optimizes it to achieve the best agreement between the curvature distributions in data and simulation for dimuon events. Due to the fundamentally different approach of the high- p_T correction, a direct one-to-one comparison is not straightforward. Instead, the additive correction parameters obtained in this work are re-binned to match the binning used in the high- p_T analysis as closely as possible. Figure 6.16 compares the parameters for one data-taking period, observing a reasonable level of agreement. The smaller uncertainties in the medium- p_T values indicate the benefit of higher statistical precision in the Z boson region, which outweighs the larger systematic sensitivity of the high- p_T muons.

6.4 Discussion

The work presented in this chapter represents a substantial contribution to the understanding of muon objects in the CMS experiment in Run 3. It was conducted within the expert group of the CMS Collaboration and regularly presented at internal meetings. Preliminary results of the momentum scale and resolution corrections in the medium- p_{T} region for 2022 and 2023 have been published in a detector performance summary [83] and are in use of physics analyses of that data.

While the efficiency correction using the TnP method is a well-established procedure—maintained and adjusted but not originally developed in the scope of this thesis—the momentum scale and resolution corrections were developed in this work based on loosely defined recipes from earlier data-taking periods. The framework for deriving these corrections was built and validated in collaboration with a master’s thesis [80], reproducing key features of previous corrections. Additionally, a new method to simplify the propagation of uncertainties for analyzers was proposed and tested in collaboration with a bachelor’s thesis [84].

To conclude the chapter, the impact of these corrections on different analyses is briefly discussed: obviously, the corrections are only relevant for analyses that rely on muon properties. When muon-based event selection is used, the efficiency correction via scale factors becomes essential. If an analysis makes use of muon p_{T} , scale and resolution corrections should generally be applied to mitigate residual biases between data and simulation.

For analyses focused on process rates, scale factors often have the larger impact, as they directly affect normalization. In contrast, shape-dependent analyses are usually more sensitive to the specific use case. In mass measurements, for example, scale corrections dominate due to their direct influence on the reconstructed mass. For these it can make sense to improve the description of the scale beyond the level provided by the method in this thesis, as done in the W^{\pm} boson mass measurement [85]. In searches for new processes, resolution corrections can be more impactful, as improved resolution allows to pronounce certain features in distributions, such as resonances, enhancing the signal-to-background discrimination.

Beyond direct shape changes, scale and resolution corrections can also affect normalization indirectly, particularly if they cause events to shift across selection thresholds. The size of this effect depends on the magnitude of the correction (uncertainty) and the shape of the p_{T} distribution. For the scale correction, a large share of the events near the threshold increases the effect. For the resolution correction, the impact on the normalization is strongest when the p_{T} distribution is steep near the threshold, causing asymmetric migrations.

Cross checks of the correction values across different p_{T} regions show reasonable agreement in the scale correction, while the differences in the resolution corrections require further investigation. These comparisons also underline the expected varying level of precision in different p_{T} regimes. Future studies could therefore benefit from a combined approach, that uses the strengths of each region for the determination of common underlying parameters.

7 Measurements of the Cross Sections of W and Z Boson Production

This chapter presents a measurement of the fiducial and total cross sections of W^\pm and Z boson production, and their corresponding ratios, which was developed within the CMS Collaboration in the scope of this thesis [5]. The analysis is conducted with a dataset collected early in the recent LHC Run 3 and corresponding to an integrated luminosity of 5.01fb^{-1} . The measurement primarily focuses on muons and is performed inclusive with respect to other particle types.,

Section 7.1 introduces concepts that are relevant to this measurement. To enhance the accuracy of the theoretical model, several optimizations are applied, as detailed in Section 7.2. The fitting procedure and results are summarized in Section 7.3, and discussed in Section 7.4.

7.1 Introduction and Overview

The discovery of the massive W^\pm and Z bosons at the CERN Proton-Antiproton Collider SppS by the UA1 and UA2 collaborations in 1983 [86–89] marked a crucial milestone in the development of the SM of particle physics. These direct mass measurements provided the experimental confirmation for the long-standing theory of electroweak symmetry breaking. Ever since, the properties of the W^\pm and Z bosons have been measured with increasing precision. The most precise determination of the Z boson mass was achieved at the Large Electron-Positron Collider [90]. In contrast, the W^\pm boson mass remains less precisely known due its decay structure—containing either a neutrino or jets in the final state—which inherently limits the precision compared to the Z boson decay into two charged leptons. Recently, the W^\pm boson mass was measured by the CMS Collaboration with high precision [85]. When combined with other precise measurements, these results allow for stringent consistency tests of the electroweak sector in the SM.

7.1.1 Motivation

The production cross sections and decay widths of W^\pm and Z bosons provide fundamental insights into their coupling strengths to fermions. However, in pp collisions, cross sections depend not only on these couplings but also on additional factors, as outlined in Section 3.2.

In particular, they are sensitive to the proton’s [PDFs](#), making them dependent on the center-of-mass energy of the collisions. As a result, these cross sections have been measured at various energies ($\sqrt{s} = 2.76, 5.02, 7, 8, 13$, and 13.6 TeV) by the [ATLAS](#), [CMS](#), and [LHCb](#) collaborations [91–106]. To date, all measured cross sections have been consistent with theoretical predictions within uncertainties.

Despite the strong agreement between theoretical predictions and experimental results so far, precise measurements of these cross sections remain essential. Most notably, any deviation from the theoretical expectations could indicate theoretical mismodeling or unaccounted experimental effects. Given the large production rates of Z and W^\pm bosons in [LHC](#) collisions, statistical uncertainties in these measurements are generally small—even in partial datasets such as that investigated in this analysis—making it easier to identify systematic effects. However, agreement between theoretical predictions and the measurements does not necessarily confirm the correctness of either. It remains possible that both the theoretical prediction and the measurement are subject to unmodeled effects, which could result in a biased but seemingly consistent result.

Assuming sound theoretical modeling, any discrepancy between measurement and theoretical prediction would instead indicate issues in the experimental analysis. Finding and studying such effects is crucial for future analyses, particularly in cases where the level of trust in the theoretical model is not as great, for instance in exclusion studies. Consequently, this analysis was conducted in close collaboration with experts from different [CMS](#) groups, providing valuable internal feedback on the status of triggers and physics objects at the start of the new Run.

Beyond the interest in the production of W^\pm and Z bosons itself, these processes are important background contributions in many analyses, such as in the measurement of the Higgs boson decay into two muons. The accuracy of such results heavily depends on precise modeling of these backgrounds, making this measurement an important ingredient for increasing sensitivity.

Additional systematic precision is gained by considering cross section ratios rather than their absolute values. While their direct impact on other analyses as background contributions is limited, ratios provide insight into the parton modeling, as the dominating systematic uncertainties cancel. For illustration, consider the uncertainty in the integrated luminosity measurement σ_L , which is the dominant uncertainty in this analysis. If two cross sections σ_1 and σ_2 are measured with the same dataset of integrated luminosity L , their ratio simplifies to:

$$R = \frac{\sigma_1}{\sigma_2} \stackrel{5.14}{=} \frac{N_1}{N_2} \frac{A_2 \epsilon_2}{A_1 \epsilon_1}, \quad (7.1)$$

where $N_{1,2}$ are the observed event counts for each process, whereas A and ϵ denote the acceptance and efficiency of the selection, respectively. Additional positively correlated uncertainties, such as those in the efficiency, also (partially) cancel, enabling highly precise measurements that are particularly sensitive to inconsistencies in the parton modeling.

This analysis focuses on muonic decays of the Z and W^\pm bosons. While this choice introduces a statistical limitation compared to fully-inclusive measurements, the large number of produced bosons ensures that statistical uncertainties remain negligible. Generally, the impact of this restriction to muonic final states is small since the overall precision of muons is exceptional in the [CMS](#) experiment, both in terms of momentum precision and reconstruction efficiency. Adding further decay channels to the analysis would only significantly improve the precision, if those channels offered either smaller intrinsic uncertainties or introduced strong correlations that allow to further constrain the statistical model. Neither

hadronic decay channels of the Z or W^\pm bosons, nor their decays to electrons are expected to have a smaller intrinsic uncertainty than the muonic channel. While some additional constraints could be introduced, they are not expected to significantly benefit the analysis. Given that focusing on a single channel greatly simplifies the measurement, the muonic final state was chosen for this study.

7.1.2 Event Topology

To optimally design the analysis, it is crucial to understand the event topology, which is influenced by two main factors: the kinematic constraints from the initial state, and the characteristics of the final-state particles.

7.1.2.1 W^\pm and Z Boson Production

In pp collisions, W^\pm and Z bosons are produced exclusively through quark-quark interactions, as gluons couple only to color-charged particles and therefore do not directly contribute to electroweak boson production. The production of these bosons typically involves quarks carrying momentum fractions in the order of a few permille of the proton's total momentum.

The most striking difference between the three bosons is their electric charge, which dictates both their production and decay modes. Specifically, the initial state compositions for the production through quark-antiquark annihilation at **LO** are:

- W^+ : an up-type quark and a down-type antiquark
- W^- : a down-type quark and an up-type antiquark
- Z : a quark-antiquark pair of up- or down-type, respectively

Due to the proton's valence quark composition, there is a slight preference for up quarks over down quarks, leading to a higher production rate for W^+ bosons compared to W^- bosons. Contributions from heavier quark generations are suppressed because their significantly larger masses shorten the lifetime of their vacuum fluctuations. The **LO** Feynman diagrams for the production of W^\pm and Z bosons are shown in Figure 7.1.

Since the production mechanisms are similar for W^\pm and Z bosons, higher order corrections are expected to affect their kinematic properties—such as boson recoil—in a similar manner. Furthermore, due to the similar size of the mass, the parton momentum imbalance is expected to be of comparable magnitude as well. As a consequence, their overall momentum profiles are expected to be similar, implying that variations in their boost are not expected to be significant.

7.1.2.2 Boson Decay

This difference in charge between the W^\pm and Z bosons is reflected in their muonic decay final states: the Z boson decays into a pair of oppositely charged muons, while W^\pm bosons decay into a muon and the corresponding muon-neutrino. The **LO** decay processes for each boson are illustrated in Figure 7.2.

For the W^\pm bosons, the branching ratio of muonic decays is approximately 1/9. This follows from the fact that the W^\pm bosons couple exclusively to the weak isospin, which is equal in magnitude for all left-handed quarks and leptons. In contrast, the Z boson couples to a linear combination of the weak isospin and electric charge, resulting in a smaller branching ratio of about 3.4 % for decays into muons.

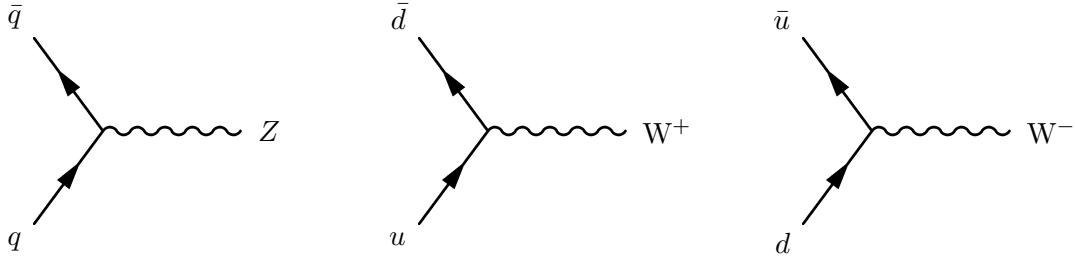


Figure 7.1: Feynman diagrams illustrating the LO mechanisms for Z , W^+ , and W^- boson production at parton level (from left to right).

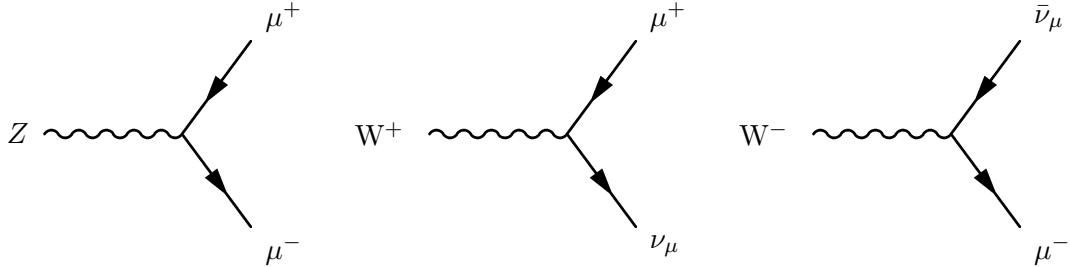


Figure 7.2: Feynman diagrams illustrating the LO mechanisms for muonic Z , W^+ , and W^- boson decays (from left to right).

Since the boson boosts are similar in magnitude, the kinematic distributions of their decay products should also be comparable. The distributions for Z boson decays are discussed in greater detail in Section 6.1. However, for the W^\pm boson decays, only one of the daughter particles is directly observable. The neutrino, on the other hand, leads to an imbalance in the transverse momentum profile, which is measured by the p_T^{miss} .

7.1.3 Measurement Strategy

The general strategy in this analysis is to compare the theoretical prediction with the data by using a [maximum likelihood \(ML\)](#) fit, as described in Section 5.1. Here, the theoretical prediction is represented by [MC](#) simulated samples that typically span a large part of the phase space. However, to ensure a meaningful comparison, the theoretical prediction has to be transformed into a representation that is comparable to the data. That means, three major steps are required:

1. Reduction from the total to the fiducial phase space, which is covered by the detector.
2. Simulation of the detector response.
3. Reconstruction of the event.

In the data, only the reconstructed event content is present, which induces the requirement of performing the comparison with the prediction in this format. Since the data contains also events that (on reconstruction level) are indistinguishable from the signal events, these must be accounted for in the prediction. Figure 7.3 provides a simplified overview of the different samples and their logic connection for a single signal measurement. The [ML](#) fit is then performed by comparing the statistical model, consisting of the reconstructed fiducial signal contribution, as well as reconstructed background processes with the data. However, the result of this fit cannot be interpreted without an in-depth understanding of the [CMS](#) experiment because the result depends on both selection efficiency and fiducial acceptance.

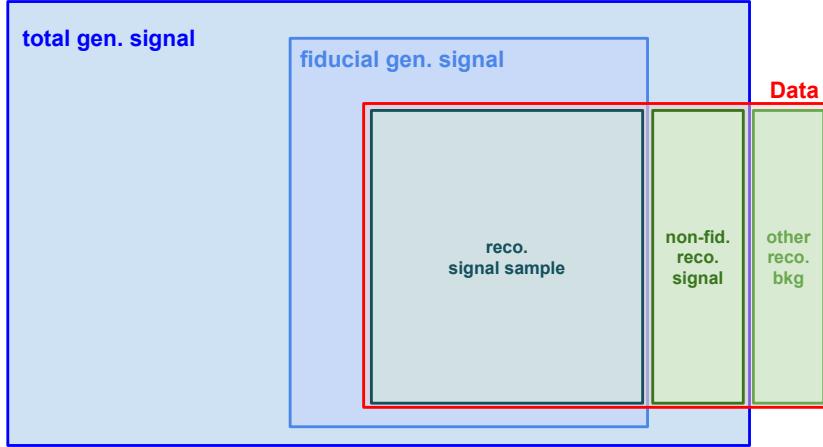


Figure 7.3: Schematic concept of the different phase spaces in the analysis. The measurement is done by comparing the data, with the reconstructed contributions from signal, non-fiducial signal, and backgrounds in the selected phase space.

Consequently, the results on reconstructed level are extrapolated to the generator level in both fiducial and total phase space.

To optimize the information content in the measurement, the dataset is divided into three distinct categories, each corresponding to one boson type. This classification is based on the number and charge of well-reconstructed muons in the event:

- Z region: events with two oppositely charged muons
- W^+ region: events with exactly one positively charged muon
- W^- region: events with exactly one negatively charged muon

This primary separation is highly efficient because the probability of misidentifying a muon charge is generally low for medium p_T . In particular, for muons with $p_T < 100$ GeV in the central region of the detector the charge misassignment rate has been found to be well below 0.1 % [107]. Since the charge misassignment effect is present in both data and simulation, any impact on the analysis would arise only from mismodeling of this effect, which—given the small size of the total misassignment effect—is expected to be negligible and is therefore not considered further. As a result, contamination from W^\pm events in the opposite-sign regions is largely suppressed. However, efficiency and acceptance effects may lead to non-negligible spill-over from the Z region into the W^\pm regions when one muon is not detected. This correlation is accounted for by performing a combined maximum likelihood fit, allowing the simultaneous measurement of all three processes.

In this fit, each category includes templates representing the relevant physics processes. These templates are then adjusted within their uncertainty ranges to achieve the best agreement between the prediction and the observed data. To enhance the precision of the measurement by suppressing background contributions, additional selection criteria are applied, as detailed in the next section. The precision can be further increased by using histogram templates of variables that provide strong separation between signal and background processes. In the Z region, the reconstructed dimuon mass $m_{\mu\mu}$ is a powerful discriminator since the probability of obtaining a dimuon pair from background processes

with an invariant mass near that of the Z boson is small. It can be calculated from the four momenta p_1 and p_2 of the two muons and approximated using their p_T values, as well as their angular differences in η and ϕ :

$$m_{\mu\mu} = \sqrt{(p_1 + p_2)^2} \approx \sqrt{2p_{T1}p_{T2}[\cosh \Delta\eta - \cos \Delta\phi]}. \quad (7.2)$$

While the exact definition is used to construct the fit templates, the approximation formula will be used to explain the muon p_T correction procedure. For the W^\pm regions, the final-state neutrino cannot be directly measured, prohibiting the reconstruction of the mass of the muon-neutrino system. However, p_T^{miss} provides indirect information on the neutrino. By combining it with the muon p_T and including the $\Delta\phi$ information between these two objects, the transverse mass m_T can be constructed:

$$m_T = \sqrt{2p_T p_T^{\text{miss}}[1 - \cos \Delta\phi]}. \quad (7.3)$$

Similar to the dimuon mass in the Z region, the probability to obtain a large m_T value is significantly larger for W^\pm processes than for most background processes, providing effective discrimination. Despite this, misreconstructed events from the Z region in which one muon is not properly reconstructed can mimic W^\pm events, as the misreconstructed muon behaves similar to a neutrino. However, as a result of the combined fit approach, this background contribution is effectively constrained through the parallel fit in the Z region, thereby limiting its impact on the precision of the W^\pm measurement.

Since all relevant processes have already been discovered, a fully blinded analysis is not required. In most analyses, the signal contribution is significantly smaller than the background, marking the need for independent validation studies of the background model in phase spaces orthogonal to the signal region. However, in this analysis, the signal contributions are dominating over the background contributions, allowing for a different blinding strategy. Here, blinding primarily focuses on the normalization of the data, while the shapes of the data distributions remain unblinded.

7.1.4 Ingredients for the Measurement

To conduct the measurement, both experimental input (the data) and theoretical input (the fit model) are required. The results of the measurement can then be compared to theoretically predicted cross sections. Since the signal strength parameters in the fit are allowed to float freely, the signal templates are not strictly required to be normalized to the predicted cross sections. However, incorporating the best available predictions and using them as initial values in the fitting procedure usually improves the fit stability.

7.1.4.1 Dataset and Event Selection

The dataset analyzed in this thesis was collected during the early phase of [LHC](#) Run 3 and corresponds to an integrated luminosity of $(5.01 \pm 0.07) \text{ fb}^{-1}$ [1]. Events are selected based on the presence of at least one isolated muon with $p_T \geq 24 \text{ GeV}$ at the trigger-level. This trigger strategy effectively reduces contributions from background processes while maintaining a high signal efficiency.

Once the full offline reconstruction is applied, all objects are reconstructed with higher precision than at the [HLT](#). Consequently, all subsequent steps, most notably event selection and categorization, are based on the offline reconstruction. To ensure full trigger efficiency for the selected events, the offline muon selection is stricter than at the trigger level,

reducing systematic uncertainties at the trigger threshold. Specifically, muons are required to have $p_T \geq 25 \text{ GeV}$, $|\eta| < 2.4$, and satisfy the tight working points in terms of both isolation and identification.

To improve the signal-to-background ratios, the initial rough categorization into the three signal regions is further refined. In the W^\pm regions, a veto is required on additional loosely identified muons to decrease contributions from only partially reconstructed Z boson decays. In the Z region, the mass of the dimuon system is required to be in the range of $60 \text{ GeV} < m_{\mu\mu} < 120 \text{ GeV}$. This criterion ensures a focus on the Z resonance, while suppressing mass regions that are dominated by photon contributions to dimuon final states.

7.1.4.2 Monte Carlo Simulation

[MC](#) simulated templates represent the theoretical model in the maximum likelihood estimation. To ensure accurate results, the simulated samples must include all significant contributions in the relevant phase space, covering both signal and background processes. The most important background contributions in both regions are the following:

- $t\bar{t}$: events in which at least one top quark from a top-quark-antiquark system decays leptonically, leading to final-state muons; due to the large top quark mass, these daughter muons often enter the selected phase space.
- VV : any pairwise production of W^\pm and Z bosons.
- [vector boson fusion \(VBF\)](#): the production of a W^\pm or Z boson through the annihilation of two vector bosons (W^\pm or Z). While technically a higher-order electroweak correction to the signal process, it is treated as background for consistency with the [LO](#) electroweak theoretical calculations, which do not include it. This distinction is largely of technical relevance, as the [VBF](#) contribution to the signal region is at the permille level, making the difference between its treatment as signal or background negligible.
- [single-top \(ST\)](#): the production of a single top quark, which may decay into a final state that includes muons.
- [QCD multijet](#): [SM](#) events composed exclusively of jets produced through the strong interaction, where a jet is misidentified as a muon.
- $Z \rightarrow \tau\tau, W^\pm \rightarrow \tau^\pm \nu^{(-)}$: the production of W^\pm and Z bosons that decay to tau leptons, which subsequently decay into a muon.
- V (nonfiducial): the production of W^\pm and Z bosons that decay to muons, which are created outside of the fiducial phase space and accidentally enter it due to the finite precision of the reconstruction process.

Example Feynman diagrams for the first four processes are provided in Figure 7.4. Except for the [QCD](#) multijet background, all of these processes are simulated with [MC](#) methods. While the hard processes are generated with different [MC](#) event generator programs, the description of the [PDFs](#), as well as the simulation of the colored residues from the collisions, and detector response are modeled consistently across all samples: proton [PDFs](#) are described with the NNPDF 3.1 package [108] at [NNLO](#). The colored residues, i.e., parton showering, hadronization, and the underlying event are simulated with the PYTHIA 8.306 package [109] and the tune CP5 [110]. Then, the detector response is simulated using

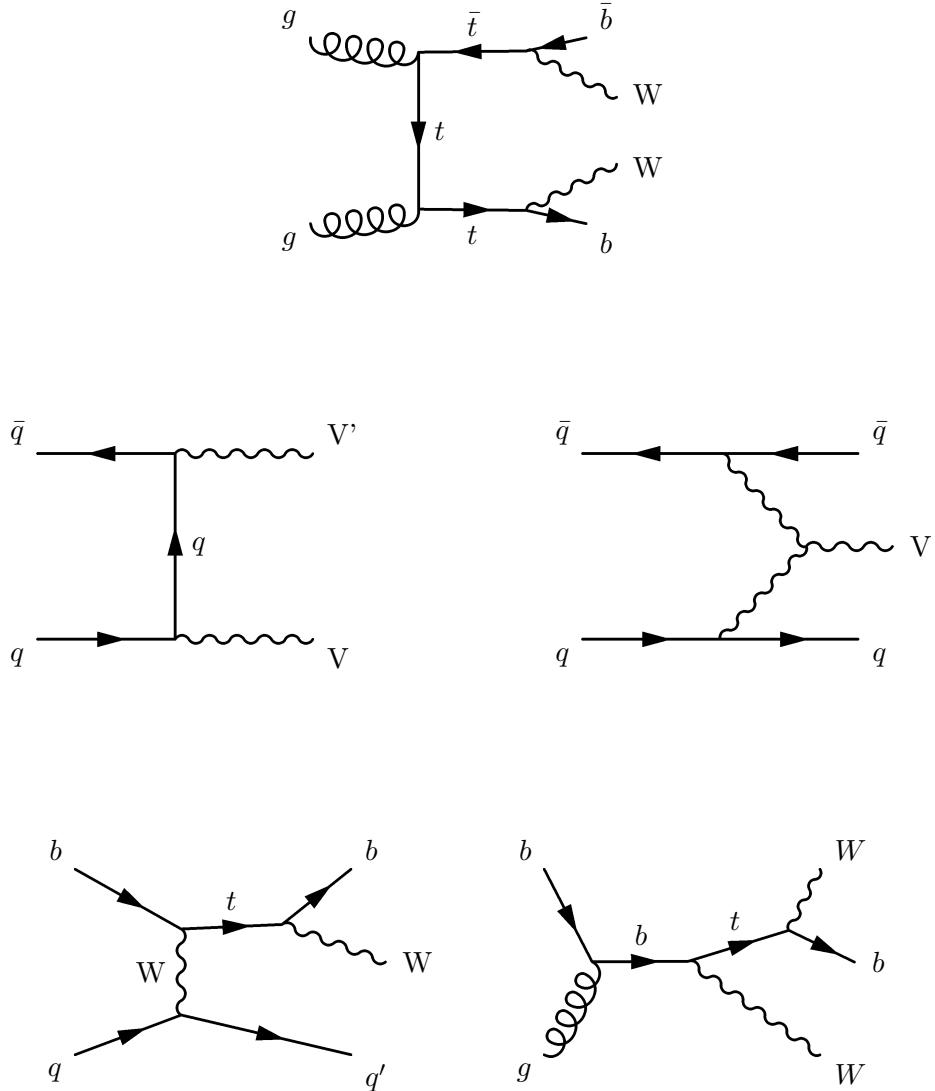


Figure 7.4: Example Feynman diagrams in LO of the background processes. At the top, the production of a top-quark-antiquark system and the subsequent decay is shown. The middle row shows diboson production on the left, and the electroweak production of W^\pm and Z bosons on the right. The latter is considered a background contribution because the corresponding order in electroweak correction is not part of the generator calculation. The lower row shows the production of a single top quark in the t channel (left), and the associated production of a top quark and W^\pm boson (right). All these background Feynman diagrams have in common that one or more W^\pm and Z bosons are created, which—for muonic decays—mimic the selection criteria.

the GEANT4 software package [111]. Lastly, additional minimum bias pp interactions are simulated with PYTHIA and added to the event to reflect the presence of pileup.

The matrix elements for the signal processes are generated with MADGRAPH5_aMC@NLO version 2.9.9 [112] at **NLO** precision in perturbative **QCD**. To avoid double counting of parton emissions from the matrix element and from the parton shower simulation, the FxFx merging scheme is applied [113].

Background processes including the electroweak production of a Z or W^\pm boson are simulated with MADGRAPH5_aMC@NLO version 2.9.13 at **LO** in **QCD**. All the remaining background samples (except for the data-driven **QCD** multijet background) are simulated with POWHEG version 2 [114–119] at **NLO** precision in perturbation theory. The top quark decay is described with MADSPIN [120] for events with a single top quark produced in the t channel.

7.1.4.3 Theoretical Cross Sections

Cross sections of the signal processes are calculated with the generator DYTURBO 1.3.2 [121–124] and the NNPDF 3.1 sets. The precision of these calculations is **LO** in electroweak perturbation theory, **NNLO** in perturbative **QCD**, and **next-to-next-to-leading logarithmic (NNLL)** in q_T resummation. At small values of the boson transverse momentum q_T , fixed-order **QCD** predictions contain large logarithmic terms of the form $\log(Q^2/q_T^2)$, which can lead to divergent or unphysical behavior. The resummation at **NNLL** accuracy accounts for such terms to provide a reliable description, in particular for small values of q_T . Different **PDF** sets are investigated, where the nominal set is the NNPDF 3.1 set, and the alternative sets are the MSHT20 [125] and the CT18 [126] sets. Uncertainties on the cross sections are evaluated with respect to **PDFs**, the scale at which the cross sections are evaluated, and the uncertainty due to the limited sample size in the numerical approximation of the integral.

The predictions for the cross sections and their ratios—evaluated with the NNPDF 3.1 sets—are presented in Tables 7.1 and 7.2, respectively. A notable feature of these predictions is the presence of asymmetric scale uncertainties, indicating that variations in the **QCD** scale induce larger deviations in one direction than the other. The large downward fluctuations in the cross section consistently arise from the downward variation of the factorization scale with fixed renormalization scale. As none of the other scale variations lead to a comparable upward deviation of the cross section, the resulting scale uncertainty—calculated from the envelope of all deviations—shows large asymmetries. These asymmetries were confirmed in discussion with the DYTURBO authors and were cross-validated independently.

In the ratios, where the scale variations are treated as fully correlated between numerator and denominator, the impact of these large outliers partially cancels. For instance, the relative size of the downward fluctuation in the W^-/Z ratio is at roughly 0.5 %, compared to 1 % and 0.7 % for the W^- and Z cross sections individually.

Table 7.1: Predictions for the product of the cross sections and branching fractions, evaluated in the fiducial and total phase space. The provided variations correspond to uncertainties in the [PDF](#), the [QCD](#) scale, and the integration.

Process	$\sigma_{\text{fid}} \mathcal{B} \pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$ in nb	$\sigma_{\text{tot}} \mathcal{B} \pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$ in nb
W^+	$5.490 \pm 0.040 \begin{array}{l} +0.025 \\ -0.048 \end{array} \pm 0.007$	$12.06 \pm 0.10 \begin{array}{l} +0.06 \\ -0.13 \end{array} \pm 0.01$
W^-	$4.222 \pm 0.032 \begin{array}{l} +0.014 \\ -0.045 \end{array} \pm 0.005$	$8.97 \pm 0.07 \begin{array}{l} +0.05 \\ -0.10 \end{array} \pm 0.01$
W^\pm	$9.710 \pm 0.071 \begin{array}{l} +0.033 \\ -0.093 \end{array} \pm 0.009$	$21.03 \pm 0.16 \begin{array}{l} +0.11 \\ -0.23 \end{array} \pm 0.01$
Z	$0.7663 \pm 0.0067 \begin{array}{l} +0.0031 \\ -0.0050 \end{array} \pm 0.0007$	$2.027 \pm 0.015 \begin{array}{l} +0.008 \\ -0.019 \end{array} \pm 0.001$

Table 7.2: Predicted ratios of the product of cross sections and branching fractions, evaluated in the fiducial and total phase space. The provided variations correspond to uncertainties in the [PDF](#), the [QCD](#) scale, and the integration.

Process	$R_{\text{fid}} \pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$	$R_{\text{tot}} \pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$
W^+/Z	$7.158 \pm 0.041 \begin{array}{l} +0.031 \\ -0.016 \end{array} \pm 0.012$	$5.951 \pm 0.027 \begin{array}{l} +0.008 \\ -0.010 \end{array} \pm 0.004$
W^-/Z	$5.508 \pm 0.030 \begin{array}{l} +0.003 \\ -0.025 \end{array} \pm 0.009$	$4.424 \pm 0.016 \begin{array}{l} +0.005 \\ -0.006 \end{array} \pm 0.004$
W^\pm/Z	$12.666 \pm 0.069 \begin{array}{l} +0.034 \\ -0.038 \end{array} \pm 0.016$	$10.375 \pm 0.040 \begin{array}{l} +0.010 \\ -0.016 \end{array} \pm 0.007$
W^+/W^-	$1.300 \pm 0.003 \begin{array}{l} +0.005 \\ -0.000 \end{array} \pm 0.002$	$1.345 \pm 0.004 \begin{array}{l} +0.003 \\ -0.001 \end{array} \pm 0.001$

7.2 Data Calibration and Tuning of the Prediction Model

The [MC](#) simulations can reproduce many different aspects of the real data with high but limited precision. To further optimize this precision, multiple correction methods are applied to the simulated samples. In addition, the data are calibrated such that they are consistent with fundamental physical principles, such as the peak position of the Z resonance and the isotropic nature of particle decays.

As this analysis is among the first to investigate the newly collected data, some essential calibration inputs were initially unavailable. Consequently, certain corrections—usually determined and provided by expert groups within the [CMS](#) Collaboration—had to be developed specifically within the scope of this analysis.

7.2.1 Pileup Reweighting

In the simulation, minimum bias events are added to the hard process to emulate pileup. The number of added minimum bias events follows a certain pileup profile that is chosen to be close to the expected pileup distribution in the data. However, the pileup in the recorded data is subject to variations both within a single fill and between different fills. To correct for these variations, the pileup distribution in simulation is reweighted to match that in the data, significantly mitigating residual differences.

While the true pileup value is an input to the simulation and therefore known explicitly, it must be estimated in data using indirect methods. These methods rely on quantities that are strongly correlated with pileup and are therefore suited to adjust the underlying difference in the pileup distribution.

One option is to use detector-level quantities, such as the number of reconstructed vertices or the energy density in the calorimeters. However, this method is generally limited by detector resolution and reconstruction algorithms. Therefore, it may reduce discrepancies in the detector response rather than the true pileup profile.

Alternatively, the expected number of [pp](#) interactions per bunch crossing can be computed using the measured bunch luminosity, i.e., the average instantaneous luminosity in fixed time intervals, divided by the number of bunch crossings. The average rate of inelastic [pp](#) collision per bunch is then given by the product of the bunch luminosity and the cross section value. To obtain the average pileup in the corresponding time interval, the average rate is divided by the collision frequency of single bunches (about 11 kHz). By repeating this procedure for all time intervals in the investigated dataset, a pileup profile is obtained. This profile depends on the precision in the measurement of the instantaneous luminosity and that of the inelastic [pp](#) cross section value.

In this analysis, this second, luminosity-based approach is applied. Furthermore, a cross check is performed with the number of reconstructed vertices. The uncertainty in this procedure is estimated—according to [CMS](#) internal recommendations—by varying the value of the inelastic [pp](#) cross section up and down by 4.6 % around the nominal value of 69.2 mb. Figure 7.5 shows the distributions of the [number of good reconstructed primary vertices \(npvGood\)](#) before and after the reweighting procedure. In this context, 'good' refers to the satisfaction of a basic set of quality requirements, such as a small shift with respect to the beamspot. Before the correction, the simulation tends to overestimate the number of primary vertices, suggesting an overestimation of pileup in data. After the correction, the agreement between the simulation and data distributions is improved; differences being generally covered by the assigned uncertainties in the vast majority of the bins.

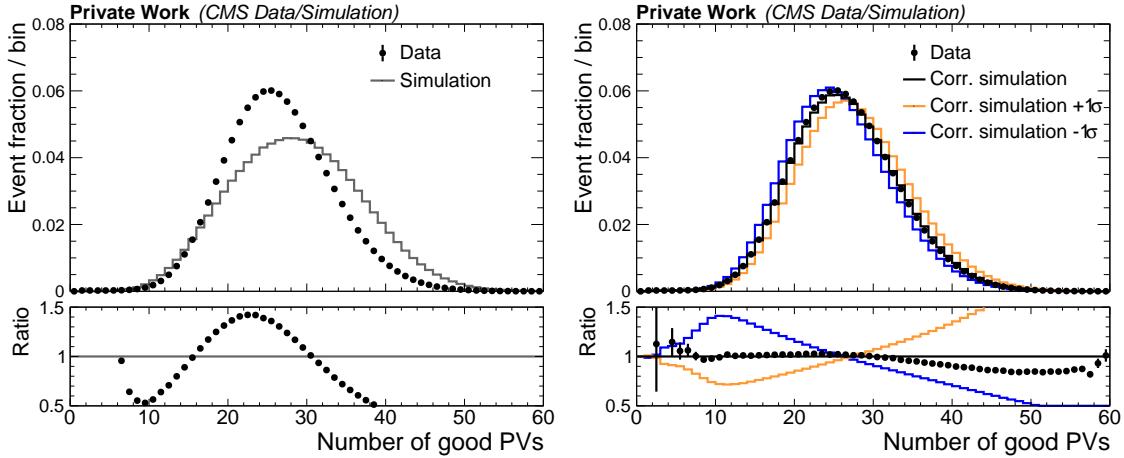


Figure 7.5: Distributions of the number of good reconstructed primary vertices before (left) and after (right) the pileup reweighting. The distributions of both data and simulations in the upper panels are normalized. The lower panels show the ratio of the normalized data to the normalized simulation. On the right, the 1σ variations of the corrections are shown as well.

Figure 7.6 illustrates the impact of the pileup reweighting procedure on the simulated signal samples, showing a tendency towards improved resolution for both. This effect is consistent with the correction lowering the mean of the pileup distribution in the simulation. Lower pileup generally results in fewer tracks and reduced energy deposition in the calorimeters, both leading to improved measurement precision.

The effect on the dimuon mass is generally small, because the muon momentum is only weakly dependent on pileup due to strict muon selection criteria. However, the impact is more pronounced in the m_T distribution of the W^\pm boson samples, where p_T^{miss} is highly sensitive to pileup and can be reconstructed with much higher precision for low pileup. Since it will be further refined through a recoil calibration procedure, the overall impact of this correction is still expected to be small. Nevertheless, to correctly propagate the associated uncertainty (up and down variations of the reweighting procedure), the correlations between the variations and the recoil calibration need to be taken into account.

7.2.2 Boson Momentum Reweighting

The modeling of the W^\pm and Z boson p_T in the simulation is limited by the order of the perturbative calculation applied. Discrepancies between the true p_T spectrum and the simulation are mitigated by using [NLO](#) samples for the main processes, generated with up to two additional jets at [LO](#). This means, the samples contain up to three emitted partons at the matrix-element level, providing a first-order correction to the boson p_T spectrum. Higher-order emissions become increasingly suppressed, but their absence in simulation can lead to distortions, particularly in the low- p_T region, where soft radiation is more frequent.

This mismatch is addressed by reweighting events in the Z boson signal sample such that the shape of the dimuon transverse momentum distribution aligns with that observed in the data. To ensure that this correction accurately reflects the boson kinematics, contributions from non- Z background processes are subtracted from the data before deriving the weights. While the dimuon system can be measured with excellent precision in the data, the p_T of W^\pm bosons presents a challenge due to the presence of an undetectable neutrino. The

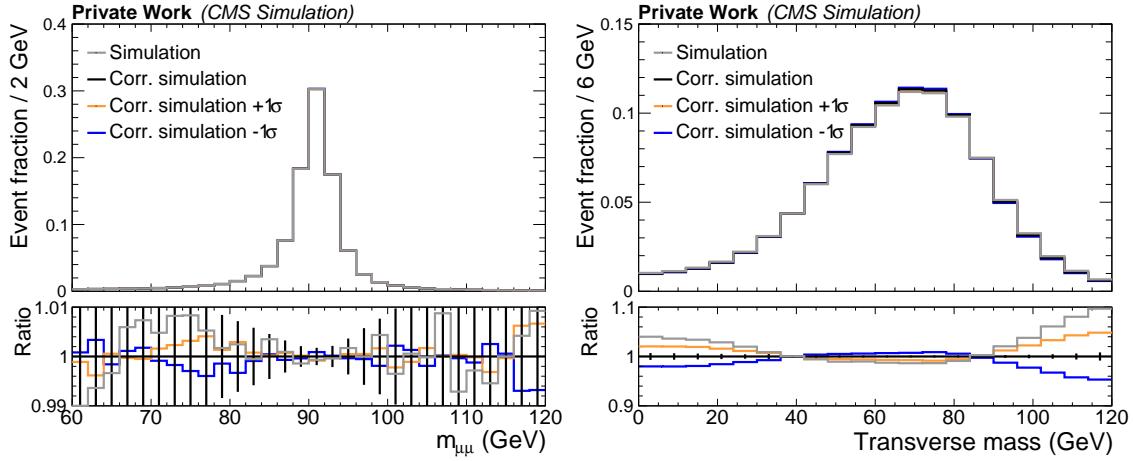


Figure 7.6: Distributions of the reconstructed mass of the dimuon system in the Z signal region (left) and the transverse mass in the W^\pm signal regions (right). The upper plots show the distributions before and after the pileup reweighting, as well as the distributions obtained by varying the correction procedure—all normalized to one. The lower panels show the ratio of the normalized simulation before the correction or after the correction with varied parameters and the normalized simulation obtained with the nominal correction set. To visualize the variations in both ratio panels, the scale is chosen a factor of ten larger for the m_T distributions. In both plots, a u-shaped behavior of the ratio of the uncorrected to the corrected distribution can be observed in the peak region.

resolution of p_T^{miss} is significantly worse due to cumulative uncertainties from multiple detector measurements. Therefore, instead of correcting the reconstructed p_T of the muon-neutrino system in the W^\pm boson signal samples to that in the data, the same weights are applied as for the Z boson signal sample. To align both application methods, the generated W^\pm boson p_T is treated as the reconstructed dimuon momentum in the Z boson sample. This approach is justified by the similarity of the production mechanisms of W^\pm and Z bosons, which differ only in the quark flavors involved. Since QCD interactions are flavor-blind, higher-order effects on the p_T spectra of W^\pm and Z bosons are expected to be similar.

Figure 7.7 shows the background-subtracted dimuon p_T distribution in data compared to the signal distribution in the Z signal region. Event weights for the Z (W^\pm) boson signal sample are obtained from the ratio of these distributions in the corresponding bin of the measured dimuon (generated W^\pm boson) p_T value in the event. The uncertainty on this procedure is estimated conservatively with the full deviation of the weight from unity, applied symmetrically.

Figure 7.8 illustrates the impact of this correction on the signal shapes. Theoretically, the invariant mass of the dimuon system should be independent from its momentum. However, the limited precision in the reconstruction of the dimuon system can introduce systematic effects. This can be observed as a shift in the lower half of the dimuon mass distribution, which is of the order of the statistical uncertainty in the simulated samples. A similar behavior is observed in the low m_T area of the corresponding spectrum in the W^\pm boson signal samples.

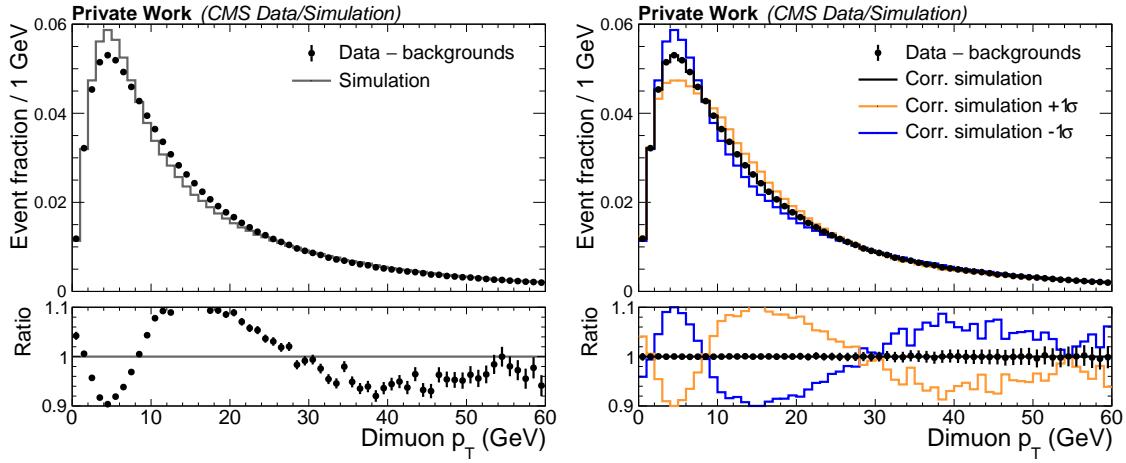


Figure 7.7: Distributions of the reconstructed p_T of the dimuon system in the Z region before (left) and after (right) the boson momentum reweighting. The distributions of both data and simulations in the upper panels are normalized. The lower panels show the ratio of the normalized data to the normalized simulation. On the right, the 1σ variations of the corrections are shown as well, one of which is by construction the uncorrected distribution.

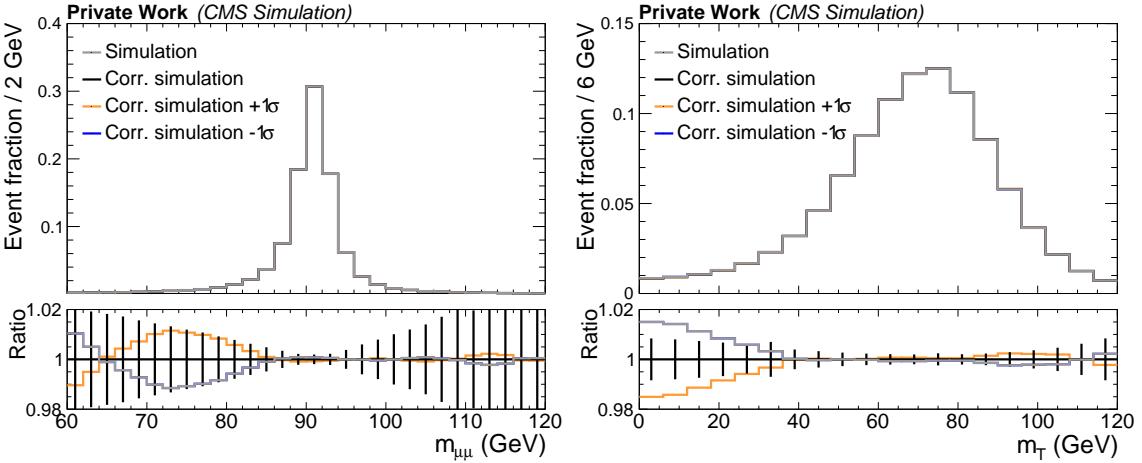


Figure 7.8: Distributions of the reconstructed mass of the dimuon system in the Z signal region (left) and the transverse mass in the W^\pm signal regions (right). The upper plots show the distributions before and after the boson momentum reweighting, as well as the distributions obtained by varying the correction procedure—all normalized to one. The lower panels show the ratio of the normalized simulation before the correction or after the correction with varied parameters and the normalized simulation obtained with the nominal correction set. By construction, the down-variation coincides with the uncorrected distribution in both distributions. Moreover, both plots indicate a residual effect for small values of the corresponding mass distribution. One variation is covered by the uncorrected distribution by construction.

7.2.3 Muon Efficiency Correction

The selection of muon objects in this analysis is subject to efficiencies. Residual differences in the efficiency modeling between data and simulation are targeted by applying muon efficiency scale factors as event weights in the simulation. Unlike the previously discussed reweighting procedures, the scale factors primarily affect the normalization, which is consistent with the principle formula of the cross section measurement in Equation 5.14.

Dedicated scale factors are derived for this analysis, according to the procedures outlined in Section 6.2. The total efficiency is decomposed into individual contributions:

- ε_{sta} : the efficiency of the standalone muon reconstruction for a muon tracker track
- $\varepsilon_{\text{global}}$: the efficiency of the global muon reconstruction for a standalone muon
- ε_{id} : the efficiency of the tight identification process for a global muon
- ε_{iso} : the efficiency of the tight isolation selection for a tightly identified muon
- $\varepsilon_{\text{trigger}}$: the trigger efficiency for a muon fulfilling the criteria above

This approach allows to identify potential inefficiencies or systematic effects in the muon reconstruction, which is particularly relevant since this analysis provides one of the first detailed examinations of muons in Run 3. Since each efficiency builds up on the previous selection criteria, the total efficiency for one selected muon is given by:

$$\varepsilon_{\text{tot}} = \varepsilon_{\text{sta}} \cdot \varepsilon_{\text{global}} \cdot \varepsilon_{\text{id}} \cdot \varepsilon_{\text{iso}} \cdot \varepsilon_{\text{trig}}. \quad (7.4)$$

For events with two muons, the total efficiency is the product of the individual muon efficiencies except for the trigger efficiency. Unlike the other selection criteria, which require both muons to pass simultaneously, the event is saved if at least one muon satisfies the trigger conditions. Consequently, the event-level trigger efficiency in this case is given by:

$$\varepsilon_{\text{trig}} = 1 - (1 - \varepsilon_{\text{trig},1}) \cdot (1 - \varepsilon_{\text{trig},2}), \quad (7.5)$$

which corresponds to the counter-probability of no muon fulfilling the trigger criteria.

Figure 7.9 presents the total efficiency for single muons in data and simulation. A more detailed breakdown is provided in Appendix B. The uncertainties in the muon efficiency scale factors are estimated by varying the bin sizes, signal and background models, and incorporating statistical uncertainties from the [TnP](#) fitting procedure. Systematic uncertainties are considered if they exceed statistical uncertainties. All uncertainties are treated as fully correlated between different p_{T} and η bins, potentially affecting both shape and normalization. Consequently, they are incorporated as shape uncertainties in the fit.

In general, the impact of muon efficiency corrections varies between the Z and W^{\pm} regions, because of the different number of muons in the respective region. For the Z region, where two muons are required, the overall effect of scale factors is more pronounced, except for the trigger efficiency, which benefits from the presence of a second muon. Hence, the uncertainty associated with efficiency scale factors is expected to be larger in the Z region for all components except for the trigger efficiency. In the W^{\pm} regions, the impact is expected to be similar in size but may differ slightly due to variations in the trigger efficiencies for different charge states.

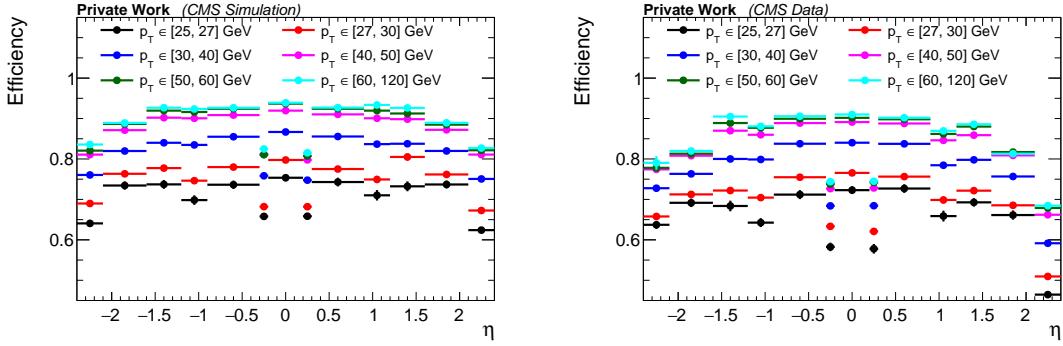


Figure 7.9: Selection efficiency for a negative muon in the simulation (left) and the data (right). The efficiencies are binned in the pseudorapidity η , and provided for different ranges of p_T . The choice of η binning reflects the detector properties. For example, the small bin from 0.2 to 0.3 has limited coverage of the muon chambers and therefore a smaller standalone muon efficiency. In general, the efficiency is higher for larger values of p_T , where the muon isolation criterion is fulfilled more frequently by construction, and the muons are further away from the trigger threshold. The uncertainty on the total efficiency in this plot is obtained by considering the contributions from the individual efficiencies as uncorrelated.

7.2.4 Muon Inefficiency Correction

The high event rate at the L1 trigger requires rapid decision-making and excellent time resolution in the relevant detector components. However, due to limitations of the time resolution, L1 trigger signals are sometimes assigned to the wrong bunch crossing. When this occurs, the event content from the incorrect bunch crossing is read out and passed to the [HLT](#). Since this event is unlikely to contain interesting physics, it is typically rejected by the [HLT](#).

To maintain a manageable data load in the L1 buffer, a trigger rule prohibits more than one accepted event within three consecutive bunch crossings [127]. Consequently, if an L1 trigger signal from an interesting event is mistakenly assigned to the previous bunch crossing—a phenomenon commonly referred to as prefire—the event is lost. This introduces an inefficiency, meaning that any interesting event that fulfills the selection criteria has a nonzero probability of being rejected due to L1 trigger limitations. Since this effect is not accounted for in the simulation, it introduces a bias in cross section measurements.

Unlike standard muon efficiency corrections, this inefficiency cannot be measured using the [TnP](#) technique because the affected events are never recorded. Instead, the trigger prefire rate is estimated using a set of unbiased events. To understand how these events are selected, it is essential to consider the L1 trigger rules. When an event is accepted by the L1 trigger, the L1 objects from the current bunch crossing, as well as those from the two preceding and following bunch crossing, are stored. Additionally, the L1 trigger rules generally veto events in these surrounding bunch crossings. As a result, the first event that can be read out after a triggered event is the one occurring three bunch crossings later. This event is guaranteed to be free of prefire effects, as the L1 trigger rules prevent it. Similarly, trigger prefire is ignored for the first event in a bunch train, since these events have no immediately preceding event.

Because L1 objects from the previous two bunch crossings are also recorded, these unbiased events allow the estimation of the prefire probability. This is done by counting the relative frequency of cases in which the L1 information of the preceding bunch crossing indicates an L1 trigger signal in the set of unbiased events. Furthermore, this probability has to be corrected for [HLT](#) effects, since unbiased events must pass the [HLT](#), while the L1 objects from the preceding bunch crossing do not have this requirement.

However, the probability of observing two triggered events separated by exactly two bunch crossings is extremely low, occurring in fewer than 0.25 % of the L1-triggered events. Thus, prefire corrections are typically statistically limited. The associated uncertainty is binned in η , meaning it may introduce a shape effect, but primarily impacts the event normalization. Following recommendations from trigger experts, the statistical uncertainty is conservatively increased by 20 % to account for potential systematic effects. Although this correction strictly does not address the selection efficiency, it is grouped together with the muon efficiency scale factors. Similar to the non-trigger efficiency scale factors, the impact of the prefire uncertainties increases with the number of muons in the final state.

7.2.5 Muon Momentum Correction

When this analysis was started, the scale and resolution corrections described in Section 6.3 were not available yet. For this reason, a preliminary set of corrections was developed early in this analysis, to account for residual discrepancies in the muon p_T description between data and simulation.

This preliminary method uses the same fundamental principles as the fine-tuning steps of the dedicated scale and resolution corrections, namely aligning the peak position and width of the dimuon mass distribution. Yet, in contrast to the scale and resolution calibration described before, the momentum corrections are derived purely based on the reconstructed muon momenta. Therefore, different approximations are made in order to correct the residual differences.

The general idea is to scale the muon p_T in the data such that the peak position is aligned with that in the simulation. Then, the p_T in the simulation is smeared such that the width of the dimuon mass distribution matches that in the data.

Both corrections use a common binning in p_T and η , assuming the dependence on ϕ is negligible. In each bin, the dimuon mass distribution is fitted with a convolution of a Breit-Wigner distribution and a double-sided Crystal Ball function. Here, the Breit-Wigner distribution models the underlying physical resonance and is therefore fixed to the currently best values. The double-sided Crystal Ball function, on the other hand, accounts for detector resolution effects: in particular, its Gaussian core quantifies detector effects, where the mean models scale effects, and the width models resolution effects. The exponential tails in the Crystal Ball function increase the stability of the fits in comparison with a convolution of the Breit-Wigner and a pure Gaussian function. This behavior can be attributed to effects, such as radiative corrections, which alter the pure Breit-Wigner shape and cannot be absorbed efficiently by a Gaussian model due to their asymmetric nature.

The derivation of the correction is based on the dimuon mass equation:

$$m_{\mu\mu} \approx \sqrt{2p_{T1}p_{T2}[\cosh \Delta\eta - \cos \Delta\phi]}, \quad (7.6)$$

where p_{T_i} is the transverse momentum of muon i , and $\Delta\eta$ ($\Delta\phi$) is the difference in pseudorapidity (azimuthal angle) between the two muons. This approximation is valid for relativistic muons, which is the case for muons in this analysis, where the lower threshold of $p_T > 25$ GeV corresponds to more than 200 times the muon mass.

7.2.5.1 Scale Correction

If the p_T values of both muon are scaled by a common factor α , the dimuon mass is also scaled by α (neglecting any correlations). The optimal α is determined by aligning the extracted dimuon mass peaks in data and simulation:

$$\alpha = \frac{m_Z + \mu_{\text{sim}}}{m_Z + \mu_{\text{data}}}, \quad (7.7)$$

where $\mu_{\text{data/sim}}$ are the mean values of the double-sided Crystal Ball function fitted to the dimuon mass distribution in data and simulation, respectively. Since muons in a given event are not necessarily in the same η - p_T bin, the correction is applied assuming that the integrated effect on the second muon is of similar magnitude to that on the muon under consideration.

7.2.5.2 Resolution Correction

To estimate the impact of the muon p_T smearing $\frac{\Delta p_T}{p_T}$ on the dimuon mass, the mass equation can be log-linearized:

$$\ln m_{\mu\mu} = \frac{1}{2}(\ln p_{T1} + \ln p_{T2} + \text{terms unrelated to } p_T). \quad (7.8)$$

Differentiating both sides yields:

$$2\frac{\Delta m_{\mu\mu}}{m_{\mu\mu}} = \frac{\Delta p_{T1}}{p_{T1}} + \frac{\Delta p_{T2}}{p_{T2}}. \quad (7.9)$$

Using this equation, an implicit connection between the muon p_T resolution in the data and simulation can be extracted by comparing the width of the corresponding dimuon mass peaks.

The total fitted width of the dimuon mass peak consists of the intrinsic Z resonance width Γ_Z , as well as the detector induced resolution $\sigma_{\text{data,sim}}$, in the data or simulation. Usually, the resolution in data is larger than in simulation. Their difference can be extracted from the detector induced resolutions:

$$\Gamma_Z^2 + \sigma_{\text{data}}^2 \stackrel{!}{=} \Gamma_Z^2 + \sigma_{\text{sim}}^2 + \sigma_{\text{extra}}^2, \quad (7.10)$$

$$\Rightarrow \sigma_{\text{extra}} = \sqrt{\sigma_{\text{data}}^2 - \sigma_{\text{sim}}^2}. \quad (7.11)$$

To improve the agreement between the resolution in data and simulation, this difference needs to be reduced. Following Equation 7.9 and assuming no correlation, this can be implemented by applying an additional smearing to the p_T of both muons that is proportional to $p_T/m_{\mu\mu}$:

$$p_T \rightarrow p_T \mathcal{N}\left(1, \frac{\sigma_{\text{extra}}}{m_{\mu\mu}}\right), \quad (7.12)$$

where $\mathcal{N}(\cdot, \cdot)$ denotes the normal distribution which depends on mean and standard deviation, respectively.

This smearing is observed to not describe the data adequately. To account for effects from the simplifying assumptions—such as that of no correlation between the muons—a further degree of freedom is introduced in the form of an additional smearing factor k . This factor is then optimized to yield the best agreement between data and simulation. The technical

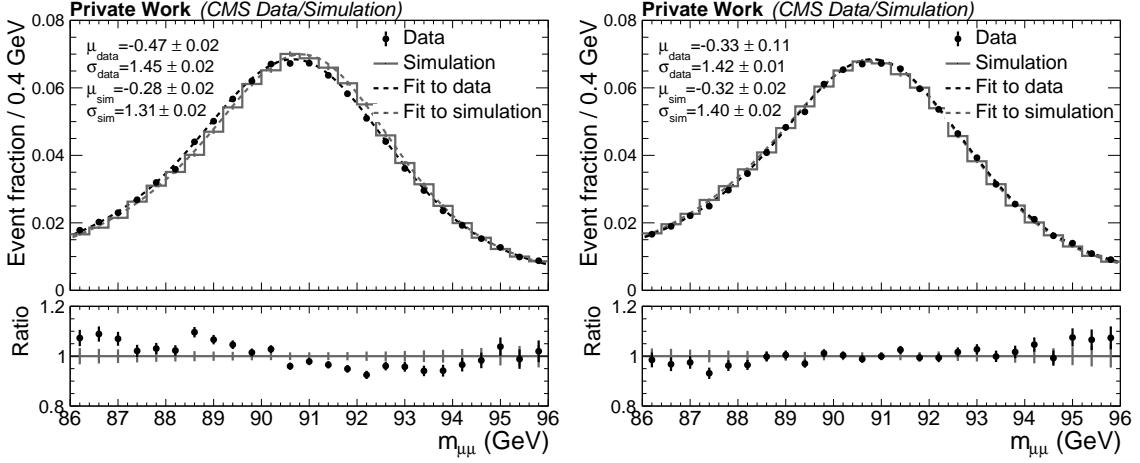


Figure 7.10: Distributions of the reconstructed mass of the dimuon system in the Z signal region before (left) and after (right) application of the muon p_T scale and resolution corrections. Events are included in the distributions if they contain at least one selected muon with $1.2 < \eta < 2.4$ and $39 < p_T < 42$ GeV. The upper plots show the normalized distributions in data, from which background events are subtracted, and the simulated signal sample, as well as fits to both. The fit model is a convolution of a Breit-Wigner distribution and a double-sided Crystal Ball function, where the parameters of the Breit-Wigner distribution are fixed to the mass and width of the Z boson, and the remaining parameters are left freely floating. Fit results for the mean and standard deviation of the double-sided Crystal Ball function are provided as well. The lower panels show the ratio of the normalized data and the normalized simulation.

procedure to derive the optimal agreement consists of testing different values for k and then fitting the resolution with a second order polynomial.

Twenty values between 0 and 4 are injected for k and the resolution is extracted through a fit with the aforementioned model. The added impact of the additional resolution scaling on the actual resolution is assumed in quadrature, due to the Gaussian convolution properties. Hence, using the twenty resolution values at different values of k , the quadratic function can be fitted:

$$\Delta m_{\mu\mu, \text{sim}} = \sqrt{\Gamma_Z^2 + \sigma_{\text{sim}}^2 + (ak)^2}, \quad (7.13)$$

where a and σ_{sim} are parameters extracted in the fit. Finally, the extra smearing factor k is obtained from the intersection of the fitted function and the resolution value σ_{data} in the data.

As this method contains 20 individual fits of the dimuon mass, it is more robust than a single fit. Consequently, instead of extracting σ_{data} from one fit to the dimuon mass, the approach described above is applied to the data as well to obtain a more robust value of the resolution in the data. Figure 7.10 provides two example fits of the dimuon mass, while Figure 7.11 illustrates the procedure to extract the additional smearing factor k .

The uncertainty on the momentum corrections is conservatively estimated as 50 % of the correction value. While the absolute size of the correction is small, it can impact the cross section by shifting events across acceptance thresholds. For instance, a muon with a reconstructed p_T of 25.01 GeV could be corrected down to 24.99 GeV and thereby leave the acceptance region. Since the scale correction is applied to the data but uncertainties are generally applied to the simulation, the uncertainty on the scale correction is applied

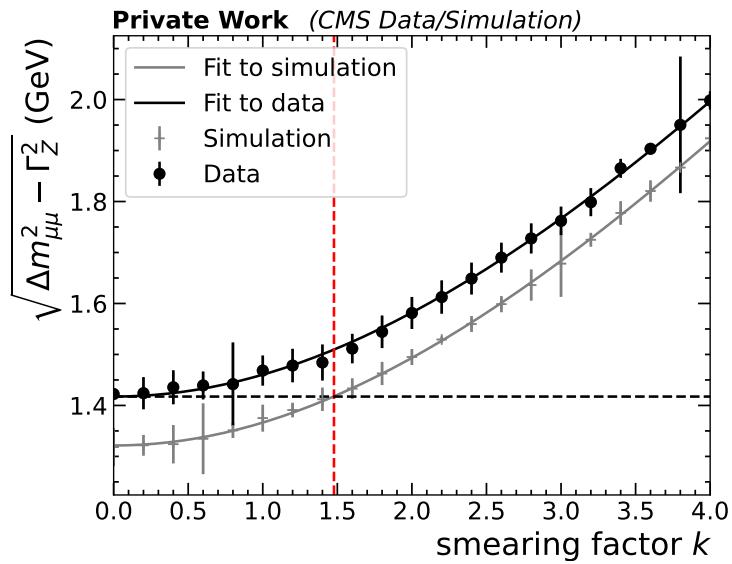


Figure 7.11: Dependence of the dimuon mass peak width on the additional smearing factor, where the intrinsic width of the Z peak is subtracted. The muons used in these fits are characterized by $1.2 < \eta < 2.4$ and $39 < p_T < 42$ GeV. The points are extracted by injecting an additional smearing to the individual muons and evaluating the impact through fits such as those in Figure 7.10. Then, a second-order polynomial is fitted to the individual points for data and simulation to find the optimally injected k in the simulation to match the resolution in the data. The fit is used as a robust method to estimate the resolution in the data (indicated by the grey dashed line), and the optimal smearing factor k in the simulation (indicated by the red dashed line).

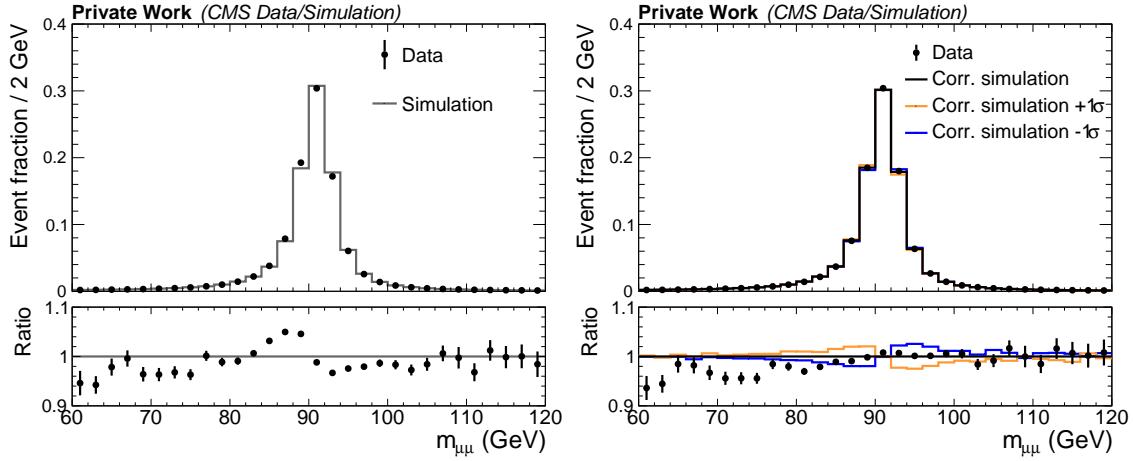


Figure 7.12: Distributions of the reconstructed mass of the dimuon system in the Z signal region before (left) and after (right) the muon p_T corrections. The upper plots show the normalized distributions in the data, from which the background contributions are subtracted, and in the Z boson signal samples. The lower panels show the ratio of the normalized distributions in data and simulation. Moreover, the systematic impact from the scale correction is illustrated in the right plot. After the correction, a slight residual slope is persistent, which is not addressed by the muon p_T corrections.

to the simulation with opposite sign. That means that an up-shift of the muon p_T in the data is modeled by a down-shift in the corresponding bin of the simulation. Thus, the relative increase in the acceptance of the data with respect to that in the simulation is reproduced correctly. The improvement in the agreement between the dimuon mass peak in data and simulation is shown in Figure 7.12, which also illustrates the aforementioned inverse impact of the scale correction uncertainty.

7.2.6 Azimuthal Correction of the Missing Transverse Momentum

In both data and simulation, the azimuthal p_T^{miss} distribution follows a sinusoidal shape. This effect is known from previous data-taking periods and many reasons have been hypothesized [128], including the following:

- anisotropic detector response: e.g. due to misalignment, inactive calorimeter cells or tracking regions,
- displacement of the beam spot: e.g. due to refocusing or proper motion of the detector.

Each of these factors can introduce a preferred direction for measured particle momenta. While the effect may be negligible for single measurements, it becomes significant when considering a large number of particles; the larger the set of particles, the larger the trend. Moreover, since p_T^{miss} is a global event variable—computed from the vector sum of all reconstructed p_T —any systematic distortion in the measurement of individual particles will be reflected by the final p_T^{miss} value. This effect is expected to increase with pileup, as the number of reconstructed particles per event grows, amplifying the effect on p_T^{miss} .

To study this behavior, the x and y components of the p_T^{miss} — p_x^{miss} and p_y^{miss} , respectively—are examined in the Z region as a function of `npvGood`, as depicted in Figure 7.13. This region is dominated by the contribution from Z boson events, which do not contain a neutrino in the main interaction at LO precision in electroweak perturbation theory. The

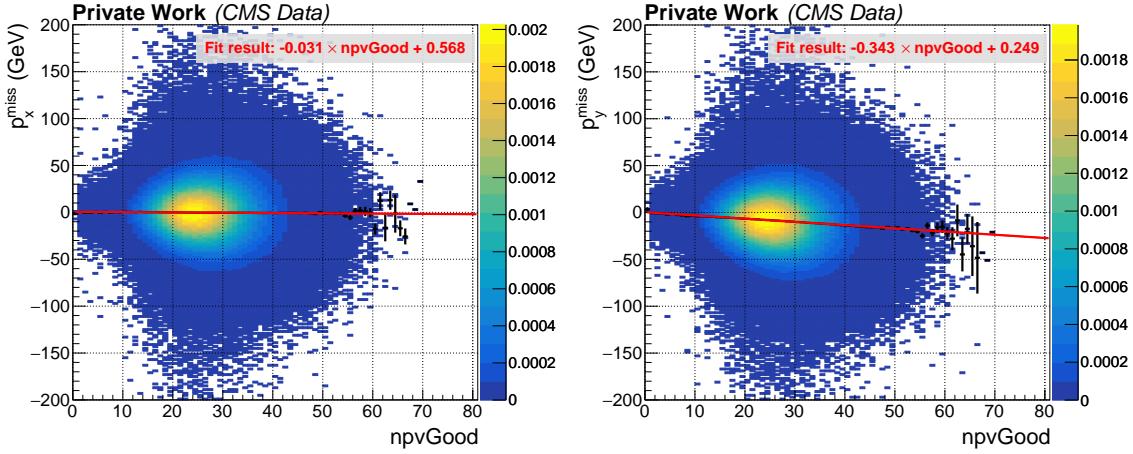


Figure 7.13: Normalized distributions of x (left) and y (right) component of the p_T^{miss} as a function of `npvGood` in the data. The red lines represent linear fits to the profile of the distributions, and the fit result is provided as well.

dominant true contributions to p_T^{miss} come from weak decays inside of jets, and from the background contributions. Since both these effects are generally small, both p_x^{miss} and p_y^{miss} are expected to be symmetrically distributed around zero in the Z region. Instead, both distributions show a deviation from zero, in particular for large values of `npvGood` in p_y^{miss} . To reduce the effect of the observed linear trend, an ad-hoc correction is derived by fitting the profiles of p_x^{miss} and p_y^{miss} as functions of `npvGood` with a linear model f . The correction formula is then given by

$$p_{x/y}^{\text{miss, data/sim}} \rightarrow p_{x/y}^{\text{miss, data/sim}} - f_{x,y}^{\text{data/sim}}(\text{npvGood}), \quad (7.14)$$

and applied to the x and y components for each event in both data and simulation, effectively restoring the mean p_T^{miss} value of both to zero. Events with significant true contributions to p_T^{miss} , such as W^\pm events with a neutrino in the final state follow the same pileup distributions as those in the derivation procedure, only with an overlayed true p_T^{miss} contribution. Consequently, the pileup-dependent bias can be reduced with this correction, thereby improving the estimation of the true p_T^{miss} contributions as well. The impact of this correction is illustrated in Figure 7.14, indicating that the strong sinusoidal pattern in the angular distribution is significantly reduced.

The uncertainty associated with this correction is dominated by the statistical uncertainty in the determination of $f_{x,y}^{\text{data,sim}}(\text{npvGood})$. This uncertainty is negligible in the central region of the `npvGood` distribution, where the high number of events effectively constrains the fit. For small and large numbers of reconstructed primary vertices, the absolute size of the uncertainties increase. However, the relative contribution of these regions to the total event yield is small. Furthermore, this correction is only the first step in the overall p_T^{miss} calibration procedure, which is dominated by the recoil calibration. Given that the uncertainties from the recoil calibration are significantly larger than those from the azimuthal p_T^{miss} corrections, the latter is not explicitly propagated in further systematic uncertainty evaluations.

7.2.7 Recoil Calibration

An additional physics-based correction is applied to the p_T^{miss} by aligning the boson recoil profiles in the simulation with those in the data. The hadronic recoil of the boson is defined

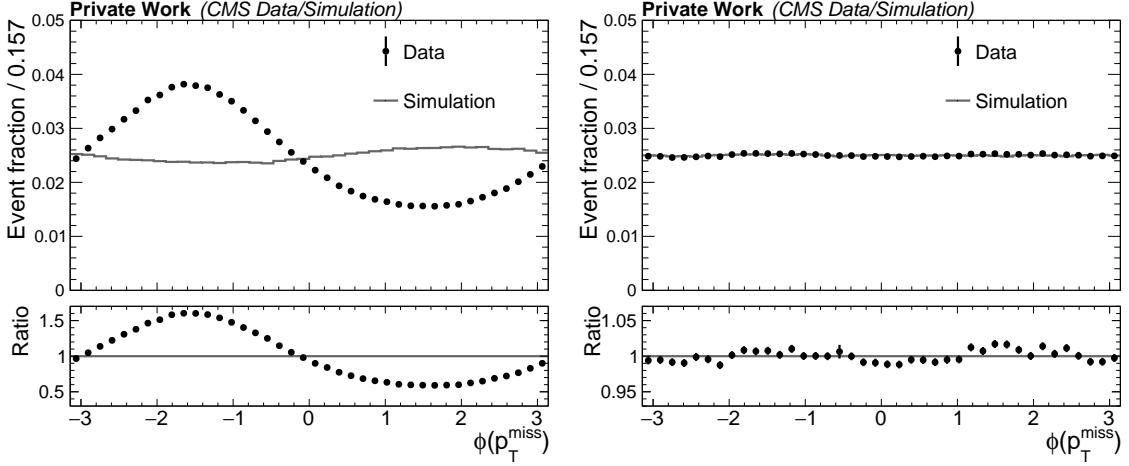


Figure 7.14: Distributions of the p_T^{miss} azimuthal angle in the Z signal region before (left) and after (right) the azimuthal p_T^{miss} correction. The upper panels represent the normalized distributions in the data and simulation, respectively. The lower panels show the ratios of the normalized data distribution and that in the simulation. Before the corrections, a strong modulation in ϕ can be observed, which is mitigated through the corrections to a large extent. Note that the scale of the ratio plot differs by an order of magnitude.

as the vector sum of transverse momenta of all the particles in an event, except for the boson's daughter particles. Ordering the total number n of particles such that the boson daughters are labeled as particles 1 and 2, the recoil can be expressed as:

$$\vec{u} = \sum_{i=3}^n \vec{p}_{T,i} = -(\vec{p}_T^{\text{miss}} + \vec{p}_{T,1} + \vec{p}_{T,2}), \quad (7.15)$$

where $\vec{p}_{T,i}$ is the reconstructed transverse momentum vector of particle i , and therefore vanishes for neutrinos. In an ideal detector with perfect calibration and infinite resolution, the hadronic recoil would exactly match the negative of the boson's transverse momentum.

In the Z region, where no p_T^{miss} is expected, the recoil can be directly mapped to the reconstructed p_T of the dimuon system. Given that the muons' kinematic behavior is already precisely calibrated, the reconstructed p_T of the dimuon system can be estimated with high accuracy and consistency between data and simulation. The recoil, on the other hand, is reconstructed from a large number of individual measurements within one event, which makes it more prone to detector effects and discrepancies between data and simulation.

To improve agreement of the recoil distributions in simulation and data, the relationship between the recoil and dimuon p_T is exploited by decomposing the recoil into its components parallel and perpendicular to the dimuon p_T :

$$\begin{aligned} u_{\parallel} &= \vec{u} \cdot \frac{\vec{p}_{T,\mu\mu}}{|\vec{p}_{T,\mu\mu}|}, \\ u_{\perp} &= \left(\vec{u} \times \frac{\vec{p}_{T,\mu\mu}}{|\vec{p}_{T,\mu\mu}|} \right) \cdot \hat{z}, \end{aligned} \quad (7.16)$$

where $\vec{p}_{T,\mu\mu}$ is the p_T vector of the dimuon system, and \hat{z} the unit vector along the beam axis. The parallel component u_{\parallel} is highly correlated with the dimuon p_T , while the

perpendicular component u_{\perp} is expected to be centered around zero. The resolutions of these components depend on multiple factors, leading to systematic differences in the recoil distributions between data and simulation.

To mitigate these discrepancies, a [cumulative distribution function \(c.d.f.\)](#) matching method is applied to align the [probability density functions \(p.d.f.s\)](#) of the recoil components between data and simulation, based on the Z region:

1. distributions of the recoil components are created in bins of the dimuon p_T for both simulation and data, from which the background contributions are subtracted,
2. the [p.d.f.s](#) of the recoil distributions are determined by fitting them with a sum of three Gaussian functions,
3. for a given simulated signal event, the [c.d.f.](#) values of its recoil components are computed from the simulation [p.d.f.](#),
4. the corresponding value in the data [p.d.f.](#) is found by inverting the [c.d.f.](#)

The recoil component is then adjusted according to:

$$u \rightarrow u + u_{\text{data,ref}} - u_{\text{sim,ref}}, \quad (7.17)$$

where $u_{\text{sim}/\text{data,ref}}$ are the values that solve the equation

$$\underbrace{\int_0^u \mathcal{P}_Z(u') du'}_{\text{c.d.f. value at } u} = \underbrace{\int_0^{u_{\text{sim}/\text{data,ref}}} \mathcal{P}_{Z,\text{sim}/\text{data}}(u') du'}_{\text{c.d.f. value at } u_{\text{sim}/\text{data,ref}}} \quad (7.18)$$

Here, \mathcal{P}_Z represents the [p.d.f.](#) of the current sample in the corresponding boson p_T bin, while $\mathcal{P}_{Z,\text{sim}/\text{data}}$ is the [p.d.f.](#) extracted from the Z boson reference sample in the simulation or data, respectively. Example distributions of the recoil in a specific bin of the dimuon p_T before and after applying the [c.d.f.](#) matching, along with a closure test with the m_T spectrum in the Z region, are shown in Figure 7.15. Discrepancies between data and simulation are mitigated largely through the corrections in all shown distributions.

The recoil corrections are then extended to events in the W^{\pm} region, where a direct mapping between the recoil and the boson p_T is not possible at reconstruction level due to the presence of an undetectable neutrino in the final state. Instead, it is assumed that the differences in the recoil [p.d.f.s](#) between data and simulation for W^{\pm} processes are similar to those observed in the Z region. This allows the application of the reference recoil [p.d.f.s](#) derived from Z events to correct the recoil in W^{\pm} events as well. To determine the appropriate p_T bin of the reference distributions for a given W^{\pm} event, the generator-level p_T of the W^{\pm} boson is used.

Uncertainties associated with the hadronic recoil modeling are evaluated by varying different aspects of the correction. First, statistical uncertainties of the fit are evaluated using the diagonalized covariance matrices from the triple-Gaussian fits in the different p_T bins. Since the nominal model consists of three Gaussians, there are six parameters (three times μ and σ), each associated with an eigenvalue. Each eigenvalue corresponds to a nuisance parameter, whose impact is assessed by shifting the model by $\pm 1\sigma$ and recalculating the template.

The uncertainty in the modeling is evaluated by replacing the triple-Gaussian with a double-Gaussian fit, and taking the shape difference as systematic uncertainty.

Background contamination in the data is an additional source of uncertainty, which is evaluated by calculating the correction without the background subtraction in the Z region. Finally, a potential systematic uncertainty due to varying detector performance in different boson rapidity sections is evaluated by testing alternative binnings:

$$|y| < 0.5, \quad 0.5 < |y| < 1.0, \quad |y| > 1.0, \quad (7.19)$$

while the nominal correction is applied inclusive in boson rapidity. The shape differences induced by these variations are incorporated as systematic uncertainties in the maximum likelihood fit.

7.2.8 Estimation of the QCD Multijet Background

Processes of the strong interaction are highly abundant in [pp](#) collisions at the [LHC](#), spanning a wide energy spectrum. Events in which strong interactions dominate the momentum transfer while further interactions only play a secondary role (e.g. during the particle propagation through the detector), are referred to as [QCD](#) multijet background. Such events can contain a reconstructed muon, which comes from weak decays inside of the jet, or jet remnants that were not fully contained in the [HCAL](#). While most of these events are removed by the tight muon identification and isolation criteria, a residual contamination from [QCD](#) multijet events remains. This background is particularly relevant in the W^\pm region, where only a single muon is required and no further transverse mass criteria are imposed. In contrast, the Z region is largely unaffected, as the requirement of two opposite-sign muons with a combined mass of at least 60 GeV makes accidental [QCD](#) contributions extremely unlikely.

The number of [QCD](#) multijet events generally decreases with increasing muon p_T and p_T^{miss} , as such events require highly energetic jets in combination with reconstruction failures. Consequently, an accurate description of the [QCD](#) multijet background in the W^\pm region is crucial, in particular at low m_T values. However, modeling this background using [MC](#) simulations is challenging due to the limited order in perturbation theory in combination with important non-perturbative contributions requiring tuned models. Furthermore, the high abundance and low selection efficiency (due to the tight muon selection criteria) demands large samples of simulated events, or selective phase-space sampling, which would be subject to additional uncertainties. To avoid these issues, a data-driven approach is adopted in this analysis. This method consists of two main steps:

1. the shape extraction in a [QCD](#)-enriched control region, and
2. the extrapolation to the W^\pm signal region, where the normalization of this background is determined in the final [ML](#) fit.

The [QCD](#) enriched region is constructed by inverting the muon isolation requirement from the W^\pm selection. However, the number of events satisfying both isolation and anti-isolation criteria is limited because events are initially selected based on an isolated muon at the trigger level. To account for potential biases in the selection efficiency, separate efficiency scale factors are determined with the [TnP](#) technique for three muon isolation ranges:

$$[0.2, 0.3, 0.5, 1.0]. \quad (7.20)$$

This binning reflects the decreasing number of events at higher isolation values. A consistency test is performed in the Z region with inverted isolation criterion, confirming improved agreement between data and simulation.

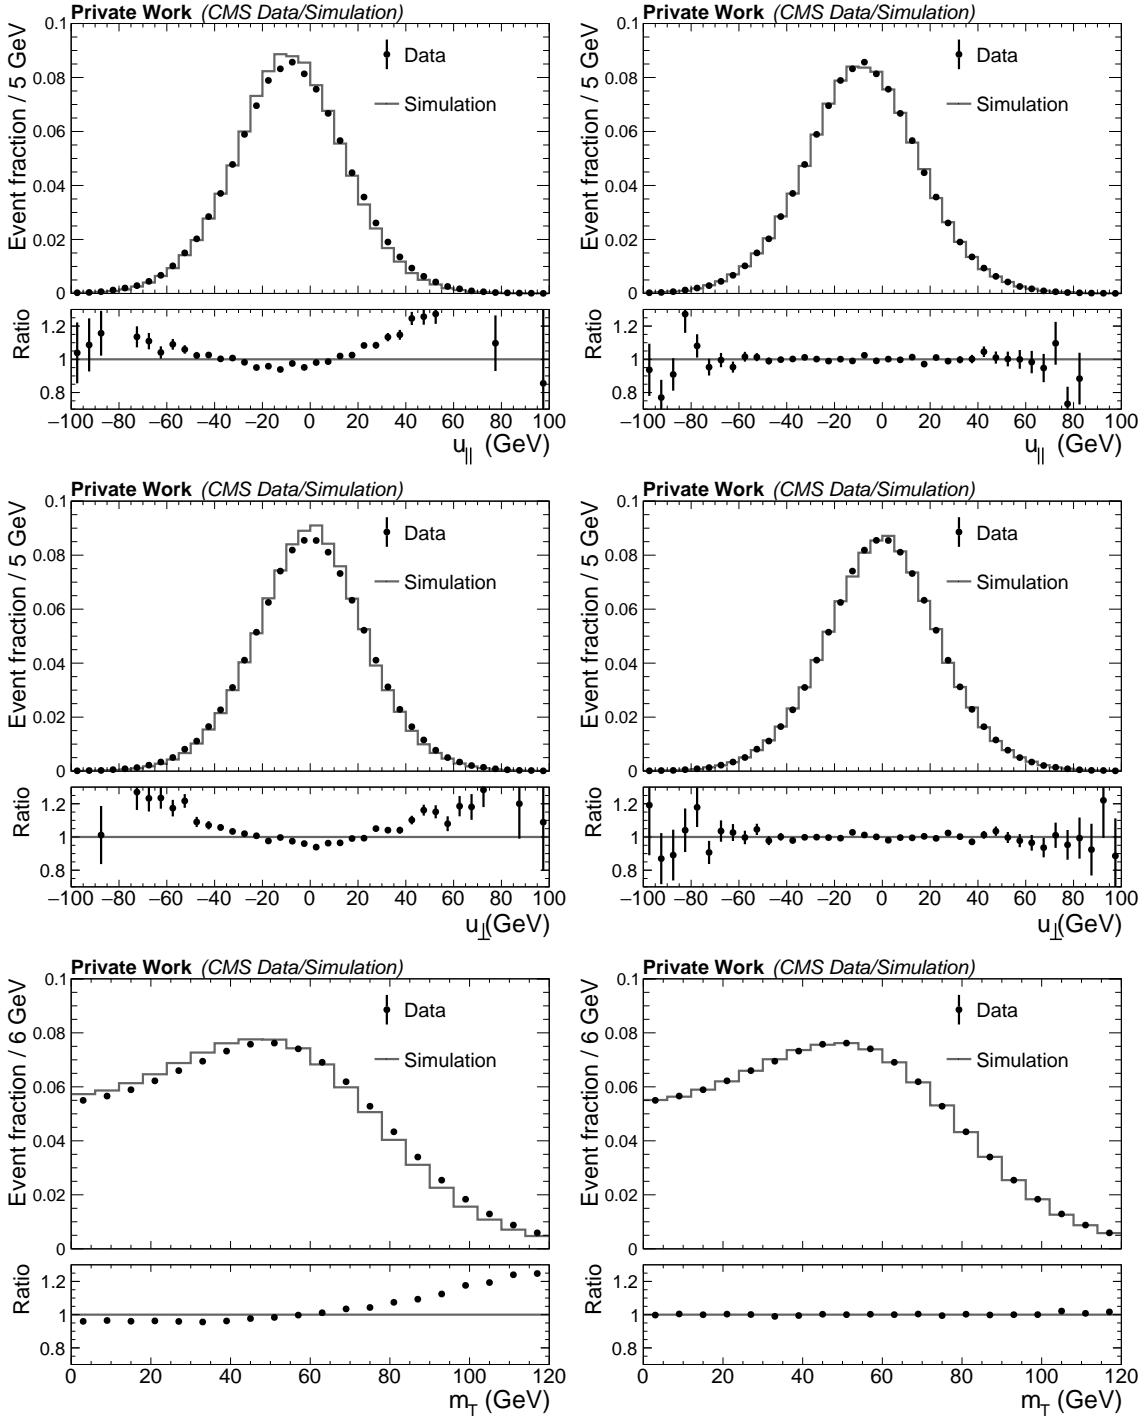


Figure 7.15: Distributions related to the hadronic recoil before (left) and after (right) the recoil corrections in the Z signal region. The upper two rows show the recoil component parallel and perpendicular to the momentum vector of the dimuon system for dimuon p_T values between 10 and 12.5 GeV, respectively. The lower row displays the m_T distributions, integrated over all dimuon p_T values. In each plot, the upper panel indicates the normalized distributions of the data, from which background contributions are subtracted, and the normalized distribution of the Z boson signal sample, while the lower panel shows the ratio between the two.

In the anti-isolated control region, **QCD** templates for the m_T distribution are extracted by producing m_T distributions of the data and all simulated samples, and then subtracting the non-**QCD MC** simulated contributions (including the W^\pm contribution) from the data. To mitigate the impact of the decreasing data normalization for higher relative isolation values, the **QCD** shapes in the non-isolated regions are all normalized to the same value. These extracted **QCD** shapes are then extrapolated bin-wise to the isolated W^\pm signal regions using a linear fit. As the individual shapes are normalized to the expected **QCD** yield in the corresponding W^\pm signal region rather than being normalized to one, the extrapolated result is expected to have a normalization in the correct order of magnitude. Further optimizing the pre-fit normalization of the **QCD** contribution is not expected to improve the results as their normalization is associated with a freely floating signal strength parameter in the fit anyway. Figure 7.16 illustrates the two-dimensional concept of the extrapolation procedure.

A possible bias in the extrapolation model is taken into account by considering the difference to the extrapolation with a quadratic fit. The total uncertainty is then composed of this fit model uncertainty and the statistical uncertainty of the linear extrapolation. While fit model uncertainties could, in principle, be correlated across bins (if the true **QCD** shape deviates systematically from a linear function), the observed variation across bins suggests an uncorrelated approach. This choice allows for a greater degree of flexibility in the **QCD** shape, which is expected to be constrained by the fit but may ultimately limit the precision of the measurement.

In the anti-isolated region, the contribution from the W^\pm signal process is non-negligible. Since the signal cross section both serves as input to the **QCD** background estimation and is subject to the measurement, there is potential for a bias in the extracted **QCD** shape. The optimal solution would be to correlate the signal yield in the anti-isolated region with the signal strength parameter in the fit. However, implementing this correlation is technically complex, as the mapping between the signal cross section in the anti-isolated region and the resulting **QCD** distribution in the isolated region is non-trivial. Furthermore, the signal model in the anti-isolated region is subject to significantly larger uncertainties, both statistically and in the modeling.

An alternative implementation could be to iteratively determine the signal strength parameter by performing the fit and then repeating the fit with the output values from the previous iteration as input value for the **QCD** estimation. Given that the measured cross section is expected to be close to the predicted value, which is used as fit input, the impact of a single iteration is assumed to be small. To verify this assumption, the effect of a 10% bias in the signal cross section in the anti-isolated region is evaluated and found to result in a 0.2 – 0.3% shift of the result. If repeated infinity-many times, the total effect would correspond to

$$\sum_i \left(\frac{0.3}{10} \right)^i - 1 \approx 3.1\% \quad (7.21)$$

of the deviation from unity, which is negligibly small. Nonetheless, a 10% uncertainty on the signal cross section in the anti-isolated region is incorporated into the fit to account for possible modeling imperfections.

Naively, the uncertainties on the **QCD** multijet background should only affect the W^\pm region, as **QCD** contribution in the Z region is assumed to be negligible. However, during the fit, variations in the **QCD** background can indirectly influence the Z boson signal normalization. Specifically, a shift in the **QCD** background alters the balance of the remaining uncertainties, causing the fit to adjust other parameters to achieve convergence.

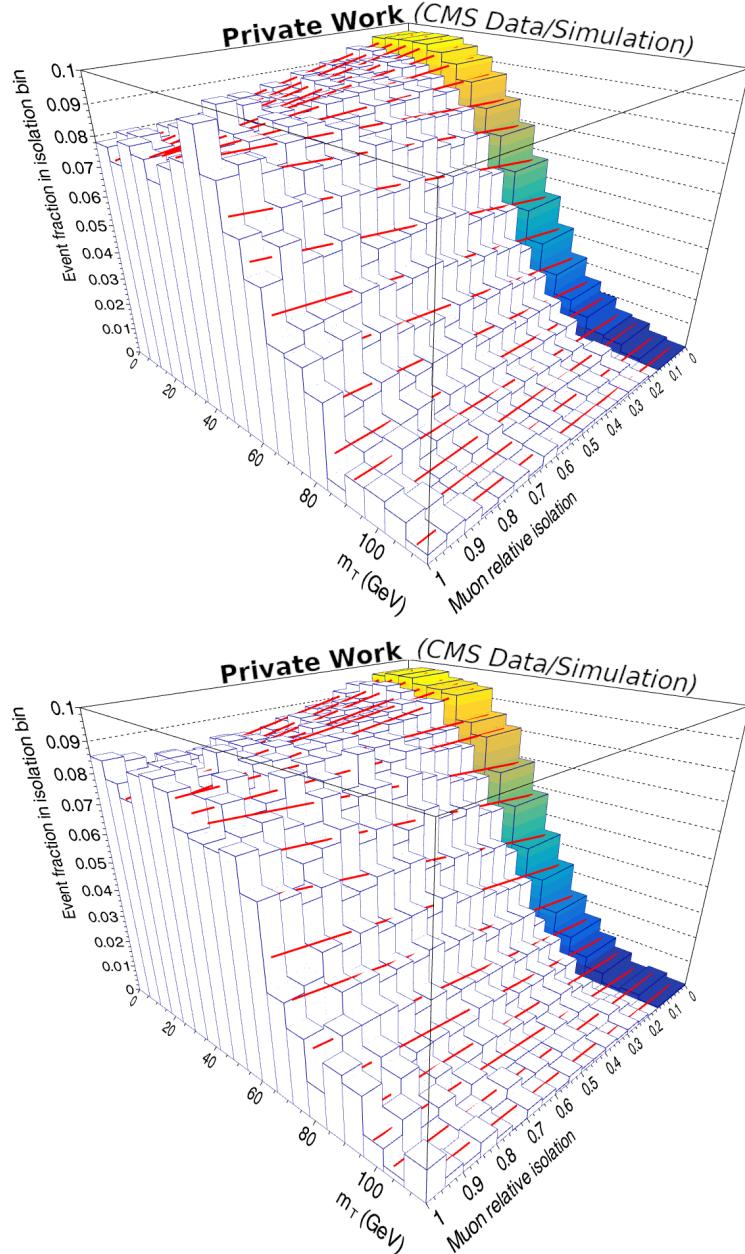


Figure 7.16: Conceptual design of the QCD extrapolation method from the m_T distributions in the non-isolated regions to the isolated region with a linear fit. The plot at the top shows the extrapolation in the W^- region, while the bottom plots shows that in the W^+ region. Each slice in muon isolation shows the normalized m_T distribution of the QCD multijet background. For relative isolation values > 0.2 , the templates (blue-edged white bars) are constructed by subtracting the non-QCD contributions from the data. For the signal region with relative isolation values < 0.15 (colored in this plot), each m_T bin is obtained via linear extrapolation from the non-isolated region. The linear fit is displayed as a (partly hidden) red line. While statistical uncertainties are considered in the fit, they are not displayed here for better overview. However, as a general rule of thumb: the larger the isolation value, the larger the statistical uncertainty and the smaller the effective constraint of the corresponding bin in the linear fit. This plot contains only final states with a negatively charged muon. Exemplary plots for a slice in both m_T and relative isolation of the muon are provided in Figures C.7 and C.8 in the appendix.

This indirect effect is taken into account when assessing the overall uncertainty on the Z boson cross section measurement.

7.2.9 Further Sources of Uncertainty

The uncertainties associated with the various corrections have already been discussed in the previous subsections. Additional systematic uncertainties are briefly introduced in the following.

The luminosity is calibrated with a method based on van-der-Meer scans, limiting the uncertainty in the luminosity measurement to 1.4% [1]. To avoid unphysical constraints or shifts in the central luminosity value during the [ML](#) fit, its nominal value is fixed at 5.01 fb^{-1} . Instead, the luminosity uncertainty is applied to the result post-fit, treating it as completely uncorrelated with other sources of uncertainty. Since the luminosity uncertainty is fully correlated across different cross section measurements, it cancels out when considering cross section ratios.

The predictive power of the [MC](#) simulated samples is limited by statistical uncertainties due to finite sample sizes. This uncertainty is incorporated into the model by application of the Barlow-Beeston lite approach [129].

Uncertainties on the [PDFs](#) are evaluated by applying 100 weights that are provided in the simulated samples and correspond to a Hessian set of uncertainties. Renormalization and factorization scale uncertainties are assessed by varying the corresponding input scales by a factor of two up and down, where the extreme case of variations in opposite directions is omitted. Similarly, parton shower uncertainties are evaluated using event weights provided by the [PYTHIA](#) shower generator. These weights reflect variations in key parameters affecting the parton shower evolution. All these uncertainties are treated as shape uncertainties that affect event kinematics but are constructed to not change the overall normalization. The uncertainties are combined into the uncertainty group called ‘PDF, scales, and parton shower’.

Background processes in the simulation are normalized to the product of their cross section and luminosity, taking into account acceptance requirements and all relevant corrections. Another source of uncertainty in these backgrounds is the cross section uncertainty, which is conservatively considered by varying the normalization up and down by 10%. At first order, an increase in the background cross sections reduces the measured signal cross section, with the magnitude of the effect depending on the fraction of background effects.

7.3 Statistical Fit

Following the strategy outlined in Section 7.1, a combined template fit is performed in the three signal regions (W^\pm and Z) using a framework based on the `COMBINE` tool [130]. Common parameters of the fit model—such as the sample normalizations—are treated as correlated between the different regions. The principle fit setup in this cross section measurement and the methodology are described in Section 5.1.

The fit model includes five signal strength parameters: one for each signal process, and one for the `QCD` background contributions in each of the W^\pm regions. Furthermore, uncertainties are incorporated as nuisance parameters and grouped into two categories: pure normalization uncertainties on the background models, and shape uncertainties (in some cases also affecting the normalization).

7.3.1 Fit Performance

During the unblinding procedure, the model was validated by applying the fit to Asimov data (pseudo-data following the model prediction). This self-consistency check revealed sub-percent expected uncertainties on the signal strength parameters (excluding the uncertainty in the luminosity measurement). For the cross section of Z boson production, the largest uncertainties are expected to stem from the muon (in-)efficiency measurements. In contrast, for the W^\pm cross sections, the dominant uncertainties arise from the `QCD` background modeling.

As part of the unblinding procedure, pre- and post-fit distributions are studied (see Figure 7.17). Two dominating traits are observed that require further explanation.

First, a discrepancy is visible in the low m_T range in the W^\pm regions before the fit. This effect is expected, as the pre-fit normalization of the `QCD` contribution is taken directly from the extrapolation and is not further optimized before the fit. Since the normalization is left freely floating, an improved pre-fit estimate would only simplify the fit convergence but would not alter the final result significantly.

The second effect is particularly apparent in the post-fit W^\pm regions and suggests signs of overfitting. The close agreement between the post-fit model and the data, with negligible differences in comparison with the uncertainties, arises from the large and uncorrelated uncertainties in the `QCD` pre-fit model. These uncertainties are significantly larger than the statistical uncertainties in the fit, which allows to greatly reduce differences between the data and the model in the different m_T bins. However, this high degree of flexibility can counteract variations in the signal normalization as well, and consequently limits the fit precision with respect to the signal strength parameters. Since variations of the signal normalization affect all 20 bins of the fit in the corresponding W^\pm region at once, the different `QCD` bin uncertainties are implicitly correlated, diminishing the impact on the signal strength parameters to some extent.

Tables 7.3 and 7.4 provide the pre- and post-fit event yields associated with the different contributions in Figure 7.17. In comparison, the post-fit yield of the W^+ process is decreased by about 1%, while the W^- and Z yields are increased by 5% and 6%, respectively. While this would naively indicate that the cross section results behave similarly with respect to the theoretical prediction, several different effects need to be considered additionally. First, variations in the nuisance parameters may have affected the normalization as well, e.g., in the efficiency scale factors. Second, the pre-fit yields do not exactly correspond to the best theoretical prediction in Table 7.1, since these precise theoretical predictions were calculated at a later stage. As explained before, the cross section results do not depend on

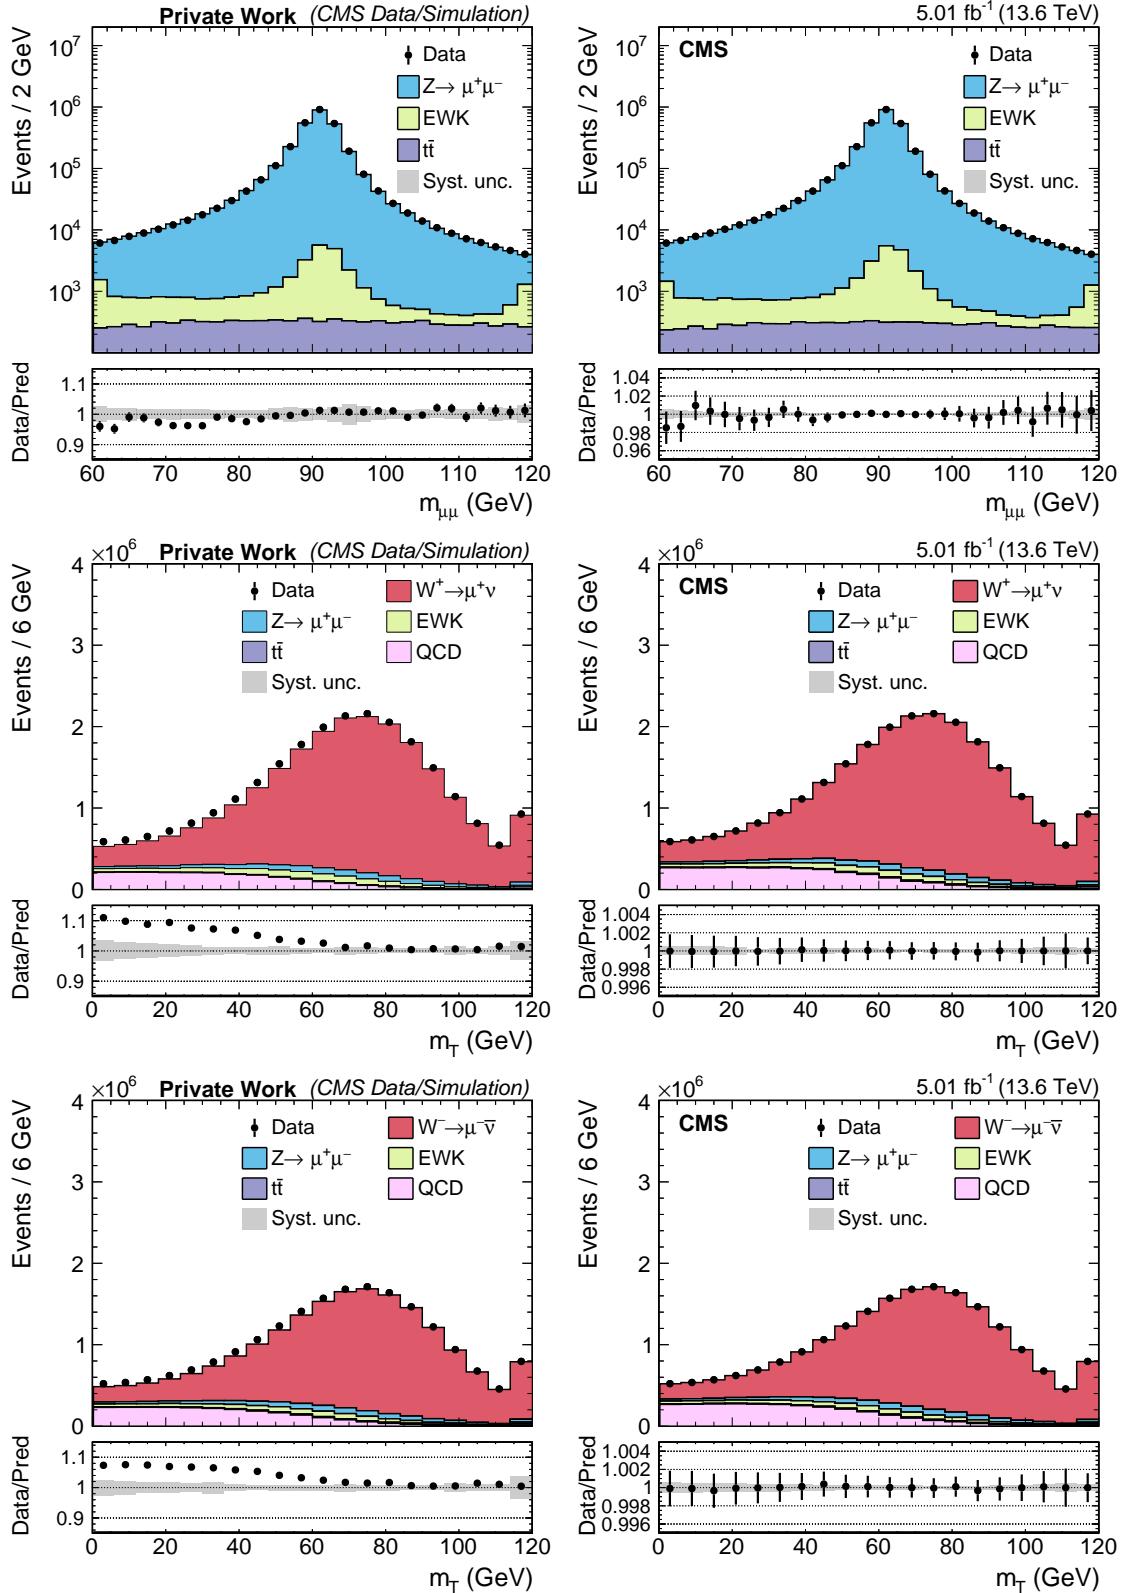


Figure 7.17: Pre-fit (left) and post-fit (right) distributions in the three signal regions. At the top, the dimuon mass distributions is shown in the Z region, in the middle and at the bottom, the m_T distributions are displayed for the W^+ and W^- regions, respectively. The lower panel in each plot shows the ratio of the number of events observed in data to that of the signal and background predictions. Overflow entries are included in the bottom two plots. The vertical uncertainty bars on the data represent the statistical uncertainty.

Table 7.3: Pre-fit event yields in the fiducial region. A dash indicates that the corresponding contribution is found to be negligible in this signal region.

Process	W^+ region	W^- region	Z region
W^\pm	20 035 607	15 591 566	—
Z	958 372	896 734	2 944 611
EWK:	(989 729)	(801 381)	(27 526)
$W^\pm (\rightarrow \tau \nu_\tau)$	608 711	498 487	—
$Z (\rightarrow \tau \tau)$	119 263	114 194	2 177
W^\pm (nonfiducial)	121 304	72 035	—
Z (nonfiducial)	9 000	8 789	16 165
VV	45 062	42 333	6 500
VBF V	21 284	16 522	1 680
Single top quark	65 106	49 021	1 004
$t\bar{t}$	215 409	216 695	9 278

the input fiducial cross sections of the signal processes. Therefore, this approach does not harm the validity of the results.

However, the cross section results do depend on the input background cross sections because these are not left freely floating in the fit. Based on the post-fit reduction of the yield of both **EWK** and $t\bar{t}$ processes by a few percent across all regions, their pre-fit fiducial cross sections seem to have been slightly overestimated. Since the simulated background processes only contribute in the single-digit percent range to the different signal regions, the effective impact from these normalization uncertainties is expected to not exceed the low permille range.

Apart from these observations in the difference of the pre- and post-fit yields, there are a few structural differences between the provided tables:

- **EWK contributions:** before the fit, the individual contributions of different sub-processes are known. However, the sub-processes are treated as a common template in the fit to simplify the model. Consequently, the individual contributions cannot be broken down anymore after the fit.
- **Uncertainties:** the total post-fit uncertainties on the yield consist of the correlated individual post-fit uncertainties on the corresponding event yields. This number cannot be provided for the pre-fit yields as the correlations are not yet known at this stage.
- **QCD multijet contribution:** as discussed before, the normalization of the **QCD** contributions is extracted in the **ML** fit. Since there is no expected normalization, the contribution is not provided in the pre-fit table.

Overall, the systematic uncertainties on the fit results are very small, much smaller than the uncertainties on the normalization of the corresponding simulation samples. While this may seem inconsistent, it is explained by the interplay between the nuisance parameters and signal strengths. If one parameter is varied by a certain amount, the other parameters are tuned to cover up for the effect on the shape and normalization. This behavior can be observed in the correlation matrices provided for the leading systematic uncertainty sources in Appendix D. Consequently, the effect on the process event yield is comparably small, in particular for the total event yield, where also correlations between the different process event yields are considered.

Table 7.4: Post-fit event yields in the fiducial region. The post-fit uncertainties include only statistical and systematic uncertainties, but not the uncertainty in the luminosity measurement. The individual uncertainties in the event yields for a given process are derived by taking the full covariance matrix into account. A dash indicates that the corresponding contribution is found to be negligible in this signal region.

Process	W ⁺ region			W ⁻ region			$\mu^+\mu^-$ region	
W [±]	20 013 000	±	31 000	15 665 000	±	22 000	—	—
Z	966 700	±	2 300	905 000	±	2 300	2 962 980	± 970
EWK	966 000	±	14 000	782 000	±	11 000	26 560	± 460
tt	200 800	±	3 400	202 100	±	3 400	8 630	± 130
QCD multijet	2 973 000	±	33 000	2 920 000	±	24 000	—	—
Total	25 119 600	±	1 500	20 475 200	±	1 300	2 998 160	± 850
Data	25 119 622			20 475 186			2 998 125	

7.3.2 Cross Section Measurement

The fiducial cross section values are obtained by multiplying the best-fit signal strength parameters with the corresponding input fiducial cross sections. Assuming that the efficiency is accurately described after applying the correction via scale factors, this approach effectively removes any dependence on the nominal cross section used in the simulation. However, it is essential to ensure consistency in the definition of the fiducial phase space between the result and the theoretical prediction. In this analysis, this primarily concerns the treatment of generator-level muons, which can be defined either before or after the emission of [FSR](#). Since the theoretical predictions from [DYTURBO](#) are computed at [LO](#) precision in [QED](#), the acceptance region in the result is defined using muons before the emission of [FSR](#). The total cross sections are then extrapolated from the fiducial results using acceptance values derived from [DYTURBO](#).

Table 7.5 summarizes the products of cross sections and branching fractions for both the fiducial and total phase space, along with the acceptance values used for the extrapolation. In both cases, the total uncertainty is dominated by the uncertainty in the luminosity measurement, which is more than twice as large as the combined systematic uncertainties considered in the fit. In each extrapolation from the fiducial to the total cross section, the uncertainties on the theoretical acceptance values are also taken into account. These uncertainties are generally of similar magnitude as the combined systematic uncertainties from the fit. Since the acceptance enters the extrapolation procedure in the denominator, the larger up-fluctuations in the asymmetric scale uncertainties translate to larger down-fluctuations in the total cross section results.

7.3.2.1 Breakdown of the Leading Sources of Uncertainty

Investigating the post-fit values of the nuisance parameters is crucial as significant variations from their pre-fit values may indicate deficiencies in the statistical model. The significance can be interpreted in two ways: with respect to the prior or posterior uncertainties of the nuisance parameter. While the prior uncertainties are an input to the statistical model, post-fit uncertainties can be estimated from the curvature of the likelihood, specifically by evaluating the second derivative with respect to each nuisance parameter at the fit minimum. A small (large) curvature implies that deviations in the parameter value have weak (strong) effect on the likelihood. Reductions in the uncertainty after the fit suggest that the data provide more information than what was initially encoded in the prior. The

Table 7.5: Results for the fiducial and total inclusive cross sections times branching fractions measurements. For the measured values, the quoted uncertainty represents the systematic uncertainty, while the statistical uncertainty is negligible. For the acceptance predictions, the first uncertainty is the [PDF](#) uncertainty, the second is the scale uncertainty, and the third is the integration uncertainty of the calculation.

Process	$\sigma_{\text{fid}} \mathcal{B} \pm (\text{syst})$ $\pm (\text{lumi})$ in nb	Acceptance $\pm (\text{PDF})$ $\pm (\text{scale}) \pm (\text{int})$ in %	$\sigma_{\text{tot}} \mathcal{B} \pm (\text{syst}) \pm (\text{lumi})$ $\pm (\text{acceptance})$ in nb
W^+	$5.428 \pm 0.037 \pm 0.076$	45.48 ± 0.24 $^{+0.11}_{-0.07} \pm 0.07$	$11.93 \pm 0.08 \pm 0.17$ $^{+0.07}_{-0.07}$
W^-	$4.167 \pm 0.028 \pm 0.058$	47.06 ± 0.22 $^{+0.19}_{-0.10} \pm 0.07$	$8.86 \pm 0.06 \pm 0.12$ $^{+0.05}_{-0.06}$
W^\pm	$9.60 \pm 0.06 \pm 0.13$	46.15 ± 0.22 $^{+0.14}_{-0.08} \pm 0.05$	$20.79 \pm 0.14 \pm 0.29$ $^{+0.11}_{-0.12}$
Z	$0.764 \pm 0.003 \pm 0.011$	37.80 ± 0.19 $^{+0.14}_{-0.04} \pm 0.04$	$2.021 \pm 0.009 \pm 0.028$ $^{+0.011}_{-0.013}$

impact of a given nuisance parameter on the signal strength can be quantified by fixing the parameter to its nominal value $\pm 1\sigma$ and re-optimizing the likelihood. The relative impacts of the five leading sources of uncertainty on the Z and W^\pm signal strength parameters are provided in Figure 7.18.

In the Z region, the dominating systematic contributions stem from the muon efficiency corrections. These affect the normalization directly through the scale factors applied to simulated events. Since the Z region contains two muons, the effect is approximately doubled for all efficiencies except the trigger efficiency. The trigger requires only a single muon, so its relative impact is lower when two muons are present. The leading four systematic uncertainties in this region show only small shifts from their initial values and are only mildly constrained by the data, indicating that the fit provides limited additional information. All four have a negative effect on the signal strength, which is expected since a higher value of the scale factor effectively raises the normalization of simulated events, requiring a smaller signal strength to match the data.

The fifth-largest uncertainty in this region arises from the renormalization scale. It is pulled up by about one standard deviation and constrained to roughly 50 % of its pre-fit uncertainty. This behavior is likely due to the scale’s influence on the dimuon mass distribution shape. Raising the renormalization scale lowers the strong coupling constant, reducing the impact of higher-order [QCD](#) corrections. As a result, contributions from [ISR](#) decrease, potentially increasing the energy available for Z boson production, which modifies the shape of the mass distributions by introducing a tendency towards larger dimuon mass values. This residual slope can compensate for the trend observed in the pre-fit distributions (Figure 7.17). However, this adjustment is not physically meaningful, as the renormalization scale is a theoretical artifact rather than an observable. Moreover, the direction of the effect on the Z boson signal strength parameter is ambiguous and arises from second-order effects, as the templates are normalized by construction and any apparent normalization shift must be compensated by other systematic variations.

In the W^\pm regions, the most dominant uncertainty arises from the signal contamination in the [QCD](#) control regions. The fit parameter corresponding to this nuisance is pulled down by approximately 1.2 pre-fit standard deviations and constrained to about 60 % of its pre-fit uncertainty. Generally, a reduced contamination of the W^\pm boson signal processes in the anti-isolated control regions leads to a relatively increased [QCD](#) background at high m_T values. The reason is that the W^\pm boson process—along with the other simulated contributions—is subtracted from the data in the anti-isolated regions. Since the W^\pm boson process is most prominent at high m_T values, this region is relatively affected the most. The strong pull and constraint of the nuisance parameter indicates that the data

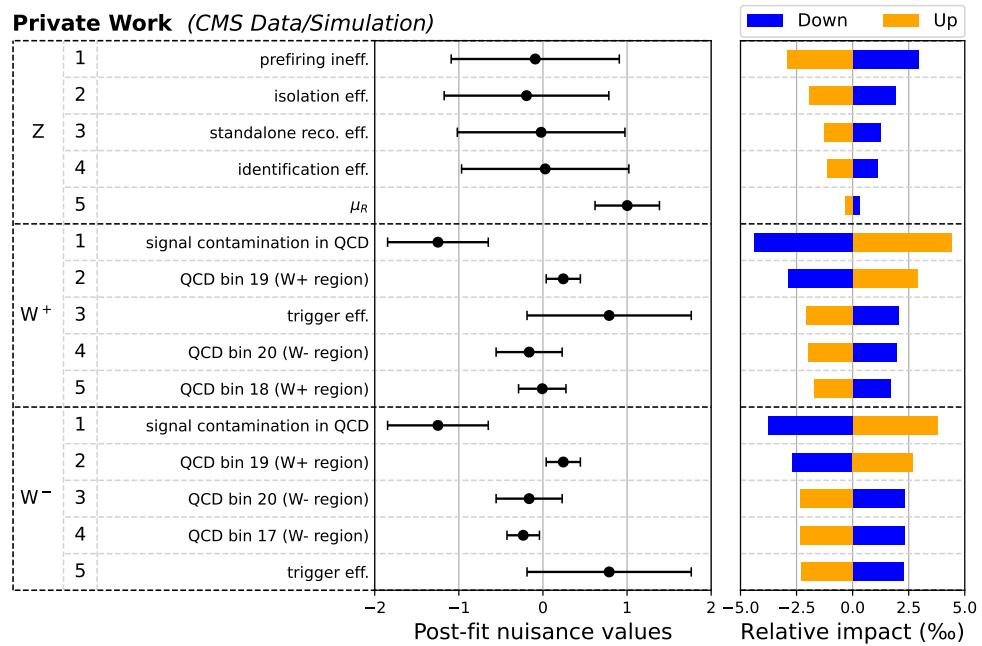


Figure 7.18: Distribution of the pull and impact values of the leading nuisance parameters with respect to the Z and W^\pm signal strength parameters. The points show to the pull $(\hat{\nu} - \nu)/\sigma_\nu$, corresponding to the difference between post-fit value $\hat{\nu}$ and pre-fit value ν , relative to the pre-fit uncertainty σ_ν of the nuisance parameters ν , while the uncertainty bars show the ratio of post-fit to pre-fit uncertainties associated with ν . The bar histograms indicate the relative impact of a variation by one post-fit standard deviation on the signal strength parameter.

provide more and different information than supplied in the anti-isolated regions. This is not surprising due to the precision limitation in the determination of the processes in the anti-isolated regions, which is caused by the contradicting isolation criteria: in these regions, the muon is required to be isolated on the trigger-level and non-isolated offline. To mitigate systematic effects from this contradiction in the selection, dedicated muon scale factors have been calculated for the anti-isolated regions. These scale factors are subject to large uncertainties up to the order of 10 % and are effectively included in the 10 % uncertainty on the signal contamination. As a consequence, the precision in the normalization in the non-isolated region is limited and deviations from the normalization in the signal region, where all processes are better-understood, are not unexpected. In summary, the difference between the pre- and post-fit values of this nuisance parameter is not considered problematic. During the unblinding process, the robustness of the fit was studied by repeating it with the input value of the W^\pm boson cross section in the anti-isolated region blindly updated to the post-fit value. The second fit was found to be consistent with the first iteration. Despite having the largest impact on the W^\pm boson signal strength parameters, the effect is rather small because of the relatively small share of signal events in the non-isolated region in contrast to the isolated signal regions. More specifically, a 6 % variation (corresponding to one post-fit uncertainty) leads to a signal strength variation in the order of 3 – 4 %. Thus, even if this variation would iteratively increase the impact, the effect would still decrease by an order of magnitude for each iteration.

In addition to the signal contamination, several high- m_T bin uncertainties from the [QCD](#) templates are among the leading systematics. Their pre-fit uncertainties are relatively large, allowing for significant constraint in the fit. Interestingly, the direction of their impact on the signal strength parameters is not uniform. While most bins follow the intuitive trend—more [QCD](#) reduces the space available for signal, thus lowering the signal strength—the nuisance parameter for the 19th bin shows a positive impact. This can be explained with a strong correlation of roughly 84 % with the signal contamination nuisance parameter, illustrated by Figure D.11. In this case, an increased [QCD](#) yield in the 19th bin is compensated by an increase in signal contamination, which in turn decreases the [QCD](#) yield globally and thus raises the signal yield. Nuisance parameters connected to other [QCD](#) bins generally show a smaller (or even negative) correlation with the signal contamination nuisance parameter, explaining their different behavior.

Finally, the trigger uncertainty has a larger impact in the W^\pm regions than in the Z region due to the presence of only one muon. Similar to the other scale factors, the sign of its impact is determined by the effect on the normalization. Unlike the other scale factors, the corresponding nuisance parameter is pulled by almost one standard deviation. While there is no clear indication why this is the case, the pull is still covered by the uncertainty and is therefore not alarming.

A summary of the systematic impacts, where different nuisance parameters are grouped together for better overview, is provided in Table 7.6. All evaluated uncertainties lie well below the 1.4 % luminosity uncertainty, and statistical uncertainties in the data are generally negligible.

7.3.2.2 Comparison with Theoretical Predictions

Figure 7.19 presents comparison plots between the measured cross sections and theoretical predictions obtained with [DYTURBO](#) at [NNLO](#) in perturbative [QCD](#), using three different [PDF](#) sets. Overall, the theoretical predictions are compatible with the measurements. The agreement is particularly good for the Z boson cross sections, both in the fiducial and total

Table 7.6: Post-fit uncertainties in percent for the fiducial cross section measurement. For completeness, also the integrated luminosity and statistical uncertainty are given.

Uncertainty source	Uncertainty in $\sigma_{\text{fid}}\mathcal{B}$ (in %) for			
	W^+	W^-	W^\pm	Z
Muon efficiency	0.28	0.29	0.29	0.40
Finite size of MC samples (bin-by-bin)	0.27	0.27	0.25	0.08
QCD background	0.53	0.49	0.49	0.07
PDF, scales, and parton shower	0.25	0.25	0.25	0.06
Muon momentum correction	0.01	0.02	0.01	0.03
Recoil correction	0.09	0.08	0.08	0.02
EWK background normalization	0.05	0.05	0.05	0.02
Z boson p_T correction	0.03	0.04	0.03	0.01
$t\bar{t}$ background normalization	0.01	0.03	0.02	0.01
Pileup	0.01	0.02	<0.01	<0.01
Total	0.68	0.66	0.65	0.42
Integrated luminosity	1.4	1.4	1.4	1.4
Statistical uncertainty	0.03	0.03	0.02	0.06

phase space. In contrast, the predictions for the W^\pm boson cross sections show a slight tendency to overestimate the data, especially in the total phase space.

Among the theoretical uncertainties, those from the NNPDF sets are generally smaller compared to the other predictions. In particular, the CT18 uncertainties are larger and highly asymmetric in some cases. These asymmetries stem mainly from the [PDF](#) uncertainties, which are evaluated based on the corresponding procedure (Hessian or [MC](#)).

7.3.3 Cross Section Ratios

Nominal results for the fiducial cross section ratios are obtained by dividing the corresponding measured cross sections by each other. The associated uncertainties are derived analogously to those of the individual cross sections, i.e., by fixing each nuisance parameter at its $\pm 1\sigma$ variation and evaluating the resulting effect on the ratios. As a consequence, uncertainties that affect both processes in a similar way tend to cancel (partially), which enhances the precision of the ratio measurement, in particular for the W^\pm boson ratio.

As in the cross section measurement, fiducial ratios are extrapolated to the total phase space using theoretical acceptance values computed with DYTURBO. Table 7.7 summarizes the fiducial and total cross section ratios, along with the acceptance values used for the extrapolation. The uncertainty in the luminosity measurement cancels completely, due to its 100 % correlation between the different processes. In contrast, statistical uncertainties are uncorrelated across the different data regions—except for the small contributions of the Z boson process in the W^\pm regions, which has negligible impact on the precision of the Z boson cross section measurement. Consequently, the statistical uncertainties do not cancel and therefore become relevant, although they are still notably smaller than the systematic uncertainties.

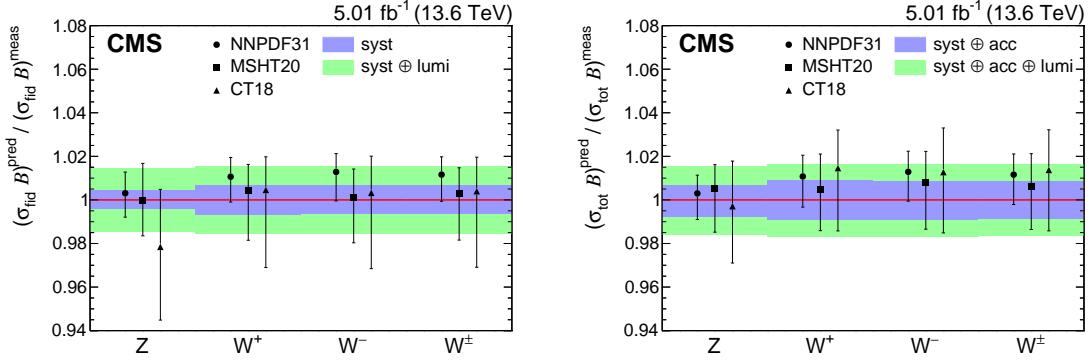


Figure 7.19: Comparison between the measured cross sections and theoretical predictions generated with DYTURBO and different PDF sets. On the left, the fiducial cross section comparisons are shown, while on the right, the total cross section comparisons are presented. All values are normalized to the measured cross sections, which are represented as a red line. The systematic uncertainties considered in the fit are represented by the blue area, while the total uncertainty—including the uncertainty in the luminosity measurement—is indicated by the green area. Since the acceptance calculated with the NNPDF set is used for the extrapolation from the fiducial to the total phase space, the corresponding entries are consistent, while the remaining theory entries vary from the fiducial to the total result. All theoretical values are found to be consistent with the measurement.

Table 7.7: Ratios of the measured product of the fiducial and total inclusive cross sections and branching fractions along with the corresponding acceptance predictions. Since some contributions of the systematic uncertainty, most prominently the luminosity uncertainty, cancel out in the ratios, the statistical component becomes relevant. For the predictions of the acceptance ratios, the first uncertainty is the PDF uncertainty, the second is the scale uncertainty, and the third is the integration uncertainty of the calculation.

Process	$R_{\text{fid}} \pm (\text{stat}) \pm (\text{syst})$	Acceptance ratio		$R_{\text{tot}} \pm (\text{stat}) \pm (\text{syst})$	$\pm (\text{acceptance})$
		$\pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$	$\pm (\text{PDF}) \pm (\text{scale}) \pm (\text{int})$		
W^+/Z	$7.105 \pm 0.005 \pm 0.048$	$1.203 \pm 0.003^{+0.004}_{-0.002} \pm 0.002$	$5.906 \pm 0.004 \pm 0.040^{+0.022}_{-0.028}$		
W^-/Z	$5.455 \pm 0.004 \pm 0.036$	$1.245 \pm 0.004^{+0.001}_{-0.005} \pm 0.002$	$4.382 \pm 0.003 \pm 0.029^{+0.023}_{-0.016}$		
W^\pm/Z	$12.559 \pm 0.008 \pm 0.081$	$1.221 \pm 0.003^{+0.003}_{-0.003} \pm 0.002$	$10.288 \pm 0.007 \pm 0.066^{+0.037}_{-0.037}$		
W^+/W^-	$1.303 \pm 0.001 \pm 0.005$	$0.966 \pm 0.003^{+0.003}_{-0.002} \pm 0.002$	$1.348 \pm 0.001 \pm 0.005^{+0.006}_{-0.007}$		

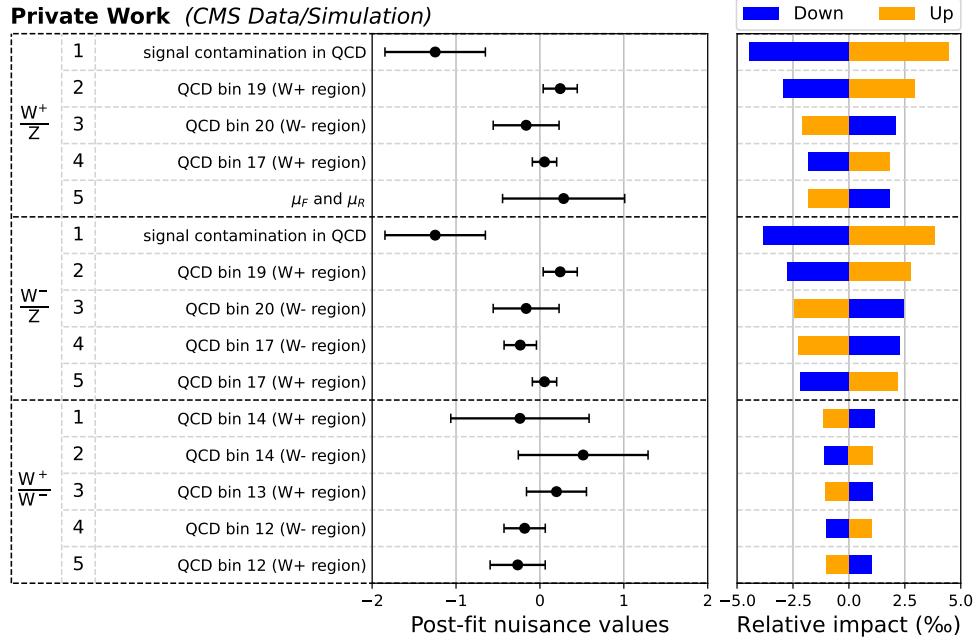


Figure 7.20: Distribution of the pull and impact values of the leading nuisance parameters with respect to the W^+/Z , W^-/Z and W^+/W^- ratios. The points show the pull ($\hat{\nu} - \nu$) / σ_ν , corresponding to the difference between post-fit value $\hat{\nu}$ and pre-fit value ν , relative to the pre-fit uncertainty σ_ν of the nuisance parameters ν , while the uncertainty bars show the ratio of post-fit to pre-fit uncertainties associated with ν . The bar histograms indicate the relative impact of a variation by one post-fit standard deviation on the cross section ratios.

7.3.3.1 Breakdown of the Leading Sources of Uncertainty

Based on the considerations for the dominant uncertainties in the cross section measurement, a few expectations can be established:

- Muon efficiencies: the impact of uncertainties in the efficiency scale factors is governed by the number of muons in the final state. As a result, their impact is expected to be negligible in the W^+/W^- ratio, and significantly reduced in the W^\pm/Z ratio.
- QCD contribution: for uncertainties related to the QCD background estimation, which primarily affect the W^\pm regions and have only second-order effects on the Z region, a dominant contribution is expected in the W^\pm/Z ratios. In contrast, a partial cancellation of these uncertainties is expected in the W^+/W^- ratio.

Figure 7.20 shows the leading sources of uncertainty and their impacts on the ratio measurements, as well as the post-fit nuisance value distributions. While the distributions of the nuisance parameters with respect to their nominal values and uncertainties remain unchanged, their impacts on the ratios differ from those on the cross sections. As expected, muon efficiency scale factors no longer appear among the top five uncertainty sources; instead, uncertainties related to the QCD modeling dominate across all ratio measurements. The W^\pm/Z ratios behave very similar to each other and to the W^\pm boson cross section uncertainties, while the W^+/W^- ratio shows notable differences. Here, uncertainty sources that were dominant in the cross section measurement are significantly reduced, leaving the QCD uncertainties in the bins close to the W^\pm boson m_T peak as the leading ones.

Interestingly, these specific bins were not leading contributions in the individual cross section fits. Their prominence arises from anti-correlated behavior between the W^+ and W^- regions: a variation in one bin tends to shift the signal strength in the same region negatively while shifting the other positively. Moreover, correlations between different bins (of up to 29 %) increase this effect by propagating shape variations across neighboring bins (see Figure D.15). The resulting shifts are counteracted by other systematic uncertainties that exhibit opposite behavior across the two regions. While the exact mechanisms for the translation from one region into the other are difficult to spot, this interpretation provides a conceptual understanding of the fit behavior.

To provide an overview of all uncertainties, they are grouped together and their impact is listed in Table 7.8. Across all three regions, uncertainties related to the [QCD](#) modeling are dominant. Their relative impact is reduced to some extent in the W^+/W^- ratio due to the partial cancellation of correlated dominant effects.

In contrast to the naive expectation, the muon efficiency uncertainties are not completely canceled in the ratios. In the W^+/W^- ratio, residual contributions remain due to differences in the trigger efficiency for differently charged muons, in particular at high $|\eta|$ values. In the W^\pm/Z ratios, the muon efficiency uncertainties remain at the level observed in the W^\pm boson cross section measurements, which is consistent given that the effective difference is the presence of one muon in the final state.

The uncertainty associated with the limited size of the simulated samples decreases by 30 % for the W^+/W^- ratio, but not significantly in the W^\pm/Z ratios. This can be explained by implicit correlations through other nuisance parameters, similarly to the bin uncertainties in the estimation of the [QCD](#) multijet background. Since the impact of the limited size of the simulated samples is significantly smaller for the Z boson process, the resulting effect on the W^\pm/Z ratios is small.

Other uncertainty groups, such as those associated with [PDF](#), scale variations, and parton showers, behave similarly to their contributions in the cross section measurements. However, as the overall impact of the remaining uncertainty groups is small, a detailed discussion is omitted.

Overall, the total systematic uncertainty on the W^\pm/Z ratios does not differ significantly from that of the W^\pm cross section measurement and remains around 0.7 %, while the W^+/W^- ratio benefits from larger cancellations and achieves a reduced uncertainty of approximately 0.4 %. The dominant trait in the ratio measurement lies in the full cancellation of the luminosity uncertainty, allowing the intrinsic precision of the analysis to be fully exploited.

7.3.3.2 Comparison with Theoretical Predictions

Figure 7.21 presents comparison plots between the measured cross section ratios and the corresponding theoretical predictions, obtained with DYTURBO at [NNLO](#) in perturbative [QCD](#) using three different [PDF](#) sets. Due to the cancellation of the luminosity uncertainty, the total measurement uncertainty is significantly reduced. Similarly, the uncertainties in the theoretical predictions are notably smaller, primarily due to reduced uncertainties in the [PDFs](#), as well as the variations of the renormalization and factorization scales.

In general, the theoretical predictions are in good agreement with the measurements, with deviations well below the 2σ level. However, the spread among the different predictions is comparable in size to the experimental uncertainty, highlighting the potential of future, high-precision differential measurements to further constrain the [PDFs](#).

Table 7.8: Post-fit uncertainties in percent for the fiducial cross section ratio measurement. For completeness, also the statistical uncertainty is given.

Uncertainty source	Uncertainty in R_{fid} (in %) for			
	W^+/Z	W^-/Z	W^\pm/Z	W^+/W^-
Muon efficiency	0.27	0.29	0.28	0.02
Finite size of MC samples (bin-by-bin)	0.26	0.26	0.25	0.19
QCD background	0.52	0.47	0.48	0.29
PDF, scales, and parton shower	0.24	0.24	0.24	0.04
Muon momentum correction	0.04	0.05	0.04	0.01
Recoil correction	0.10	0.09	0.09	0.04
EWK background normalization	0.07	0.07	0.07	<0.01
Z boson p_T correction	0.04	0.05	0.05	0.01
$t\bar{t}$ background normalization	0.02	0.04	0.03	0.02
Pileup	0.01	0.02	<0.01	0.03
Total	0.68	0.66	0.64	0.38
Statistical uncertainty	0.07	0.07	0.06	0.04

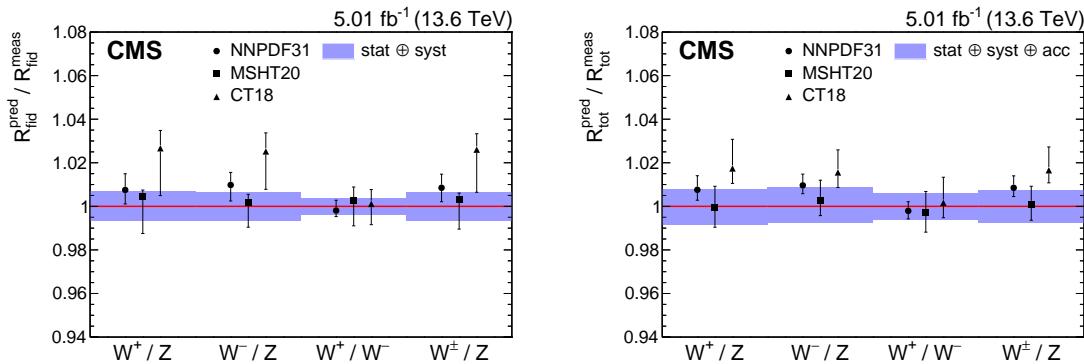


Figure 7.21: Comparison between the measured cross section ratios and theoretical predictions generated with DYTURBO and different PDF sets. On the left, comparisons of the fiducial cross section ratios are shown, while on the right, the total cross section ratios are compared. All values are normalized to the measured cross section ratios, which are represented as a red line. The systematic uncertainties in the measurement are indicated by the blue area. Since the acceptance calculated with the NNPDF set is used for the extrapolation from the fiducial to the total phase space, the corresponding entries are consistent, while the remaining theory entries vary from the fiducial to the total result. All theoretical values are found to be consistent with the measurement.

7.4 Discussion

The cross section measurements presented in this thesis are among the most precise determinations of inclusive W^\pm and Z boson production to date. Main driver of this exceptional precision is the luminosity uncertainty of 1.4 %, which is unprecedented for the CMS experiment and has been achieved following extensive studies and excellent work of the corresponding CMS subgroup. Although the remaining uncertainties are significantly smaller than the luminosity uncertainty, they play a central role in the cross section ratios, where many common uncertainties cancel.

A key factor influencing the precision is the level of pileup in the data, which has a direct impact on the accuracy of the p_T^{miss} reconstruction. With fewer simultaneous pp interactions, fewer unrelated particles contribute to the p_T^{miss} calculation, thereby improving its resolution. This, in turn, enables a more precise reconstruction of the W^\pm boson m_T , yielding a sharper peak around its mass. As a result, the fraction of signal events in the low- m_T region is significantly reduced, allowing for the definition of an additional control region to better constrain the QCD background.

For these reasons, certain previous measurements have been performed using low-pileup datasets that—despite their reduced statistical size—offer improved precision, in particular in the estimation of the QCD background. Figure 7.22 provides a comparison of the total inclusive cross section measurements at various center-of-mass energies conducted with the CMS experiment.

Already in 2024, the ATLAS Collaboration published results for vector boson production cross sections at the new center-of-mass energy of 13.6 %TeV [106]. While the nominal values reported by ATLAS and those presented in this thesis are compatible within uncertainties, the precision achieved in this thesis exceeds that of the ATLAS measurement substantially. In particular, for the Z boson production cross section, the uncertainty quoted by ATLAS is larger by a factor of 1.8; for all other measured quantities, the difference in precision is even more pronounced. To understand this difference in precision, it is necessary to highlight differences in the analysis strategy.

While the measurement presented in this analysis uses an early subset of the data recorded in 2022, the ATLAS measurement is based on the full dataset. In addition, ATLAS includes final states with both electrons and muons, further increasing the effective size of the dataset. However, the resulting statistical gain by a factor of $\sqrt{12}$ (six times more luminosity, two times the final states) is practically negligible, since the statistical uncertainties are already tiny in comparison with the systematic uncertainties—even for the smaller dataset used in this analysis. Furthermore, the early dataset used in this thesis contains a smaller average pileup value, which generally allows for a higher resolution of the m_T spectrum in the W^\pm region.

Beyond these differences in the dataset, the ATLAS measurement uses a simplified approach, in which each signal region is represented by a single bin in the fit. As a result, the measurement does not fully exploit the discriminating power of the shape information, unlike the multi-bin fit used in this analysis. The ATLAS analysis also follows a different strategy for estimating the QCD background in the W^\pm regions, which ultimately affects the precision via an implicit threshold in the p_T^{miss} variable. In this thesis, the measurement is inclusive in both p_T^{miss} and m_T , and the normalization of the QCD background is extracted directly in the statistical fit. While the control-region setup is similar in the ATLAS analysis—using non-isolated regions to estimate the QCD contribution in the isolated signal region—the normalization is constrained externally, not in the fit. The external normalization is obtained by dividing the p_T^{miss} (or m_T) distribution into two regions, where

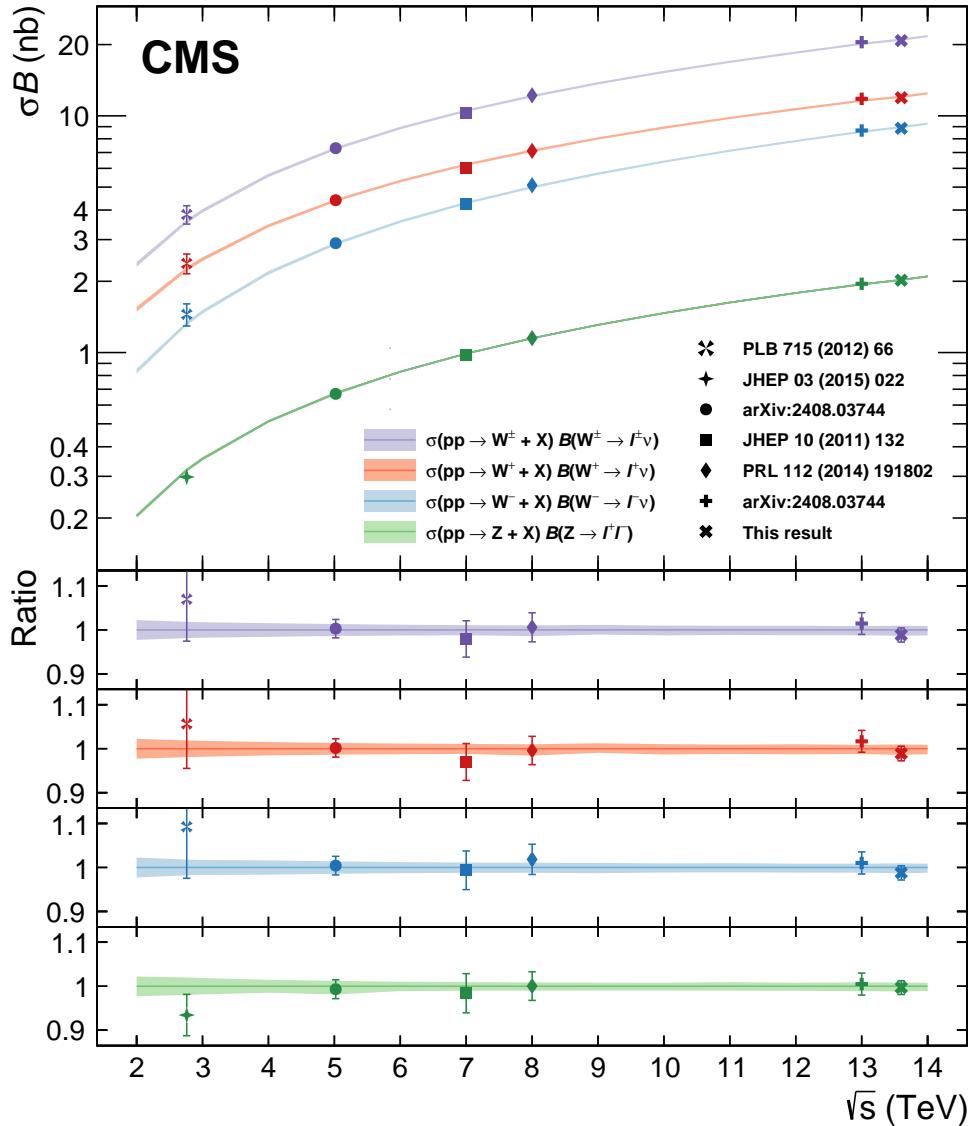


Figure 7.22: Comparison of measured products of total cross section and branching fractions for W^\pm and Z boson production at different center-of-mass energies with the corresponding theoretical prediction. All quoted measurements were conducted by the CMS Collaboration, while the predictions are derived consistently from DYTURBO at NNLO+NNLL accuracy in QCD. The uncertainties in the theoretical prediction include variations of the renormalization and factorization scales, as well as the PDF uncertainty evaluated with the NNPDF 3.1 set. The vertical uncertainty bars on the markers represent the total uncertainty of the measurement.

the low- m_T region with limited signal contamination is used to extract transfer factors carrying the normalization information. However, the relatively large signal contamination in this low- m_T region limits the precision of the transfer factors. The implicit p_T^{miss} threshold introduced in the ATLAS approach leads to additional uncertainties, particularly from jet-related sources at the threshold.

Furthermore, the muon efficiency uncertainties reported in the ATLAS measurement are substantial—above 2% for the $Z \rightarrow \mu\mu$ category, compared to 0.4% in this thesis—which is also atypical for the ATLAS experiment, based on results in Run 2 [131].

Finally, the uncertainty in the luminosity measurement is also notably larger in the ATLAS analysis—at 2.2–2.5% for the ATLAS measurement compared to the 1.4% used in this analysis. Combined, these factors result in a considerably lower total uncertainty in the measurement presented in this thesis.

In summary, the measurements presented in this thesis rank among the most precise cross section determinations ever achieved in hadron collider experiments. This level of precision is achieved primarily due to the highly accurate luminosity measurement—with 1.4% uncertainty the most precise achieved with the CMS experiment so far—and due to the use of a multi-bin fit approach that fully exploits the shape information. While the luminosity remains the dominant uncertainty source in the absolute cross section measurements, the precision of their ratios is limited by the modeling of the QCD background, which in turn is affected by pileup. Low-pileup conditions improve the resolution of p_T^{miss} and sharpen the m_T distribution, reducing the signal contamination in the low- m_T region and allowing for a more precise background estimation.

Assuming the validity of the theoretical results, the excellent agreement between the measurement and the predictions—in combination with the precision of the results—highlights the impressive performance of the CMS detector in the latest LHC Run 3. Apart from this demonstration of the potential of the new data, the processes studied in this analysis serve as important backgrounds in many searches and measurements. Thus, the results help to validate the modeling in such analyses. Moreover, the results emphasize the potential to further constrain the proton PDFs, for example through differential measurements of Z boson production in association with jets. Since the uncertainty in the luminosity dominates the total uncertainty, these measurements could also serve as input for alternative luminosity determinations, effectively inverting the measurement.

Future improvements in the precision of the inclusive cross sections of Z and W^\pm boson production will mainly depend on further reducing the luminosity uncertainty. However, with higher pileup expected in the high-luminosity LHC era, improvements in the p_T^{miss} reconstruction will be necessary to match the precision reached, in particular in the cross section ratios. Alternatively, reducing the reliance on p_T^{miss} , for example by focusing on the lepton momentum or using combined observables, could help mitigating the impact of pileup while still maintaining sensitivity.

8 Conclusions

In 2022, the [Large Hadron Collider \(LHC\)](#) started into its latest data taking era, colliding protons at the unprecedented center-of-mass energy of 13.6 TeV. After the pause in operations during the preceding [long shutdown \(LS\) 2](#) (2018–2022), checking the early-recorded data was an essential step to ensure optimal performance later on.

This thesis is a strong contribution to the successful start of [LHC](#) Run 3 from the perspective of the [Compact Muon Solenoid \(CMS\)](#) Collaboration. A dataset recorded early in Run 3 and corresponding to an integrated luminosity of 5.01 fb^{-1} is analyzed to measure the cross sections of W^\pm and Z boson production in muonic final states. Achieving high precision in this measurement required a detailed calibration of the simulated samples, particularly in the modeling of the muon momentum and the missing transverse momentum. A method for reducing residual differences in scale and resolution of muon momenta was developed, generalized, and made available for use in future [CMS](#) analyses.

Due to their large production rates at the [LHC](#), W^\pm and Z boson cross sections can be measured with high statistical precision. This sensitivity makes them ideal for early studies, as even small systematic discrepancies can become visible. These measurements also provide important tests of theoretical predictions at high orders in perturbative [quantum chromodynamics \(QCD\)](#), and serve to validate background models for analyses, in which W^\pm and Z bosons appear as backgrounds.

Systematic differences between data and simulation are reduced through a sequential procedure: simulated events are corrected to match the pileup distribution in data, and to account for missing higher-order corrections in the description of the boson transverse momenta. Furthermore, scale factors are applied to the simulation to improve the description of muon selection efficiencies. Then, both muons and the hadronic recoil are corrected in terms of scale and resolution of their transverse momentum.

After having addressed these effects in the description of the final-state particles, the representation of the theoretical prediction in the statistical fit can be constructed. For the W^\pm boson, this requires the estimation of a background contribution that arises from events in which muons are not produced in the main interaction, but in the hadronization process of quarks and gluons. Unlike the other contributions, this background is estimated in a data-driven approach by extrapolating the shape from a control region enriched in such [QCD](#) multijet events. The normalization, on the other hand, is obtained through the [maximum likelihood \(ML\)](#) fit.

The measured values for the product of total cross sections and muonic branching fractions in the total phase space are:

$$\begin{aligned}\sigma(W^+) \mathcal{B}(W^+ \rightarrow \mu^+ \nu_\mu) &= 11.93 \pm 0.08 \text{ (syst)} \pm 0.17 \text{ (lumi)} \begin{array}{l} +0.07 \\ -0.07 \end{array} \text{ (acc) nb,} \\ \sigma(W^-) \mathcal{B}(W^- \rightarrow \mu^- \bar{\nu}_\mu) &= 8.86 \pm 0.06 \text{ (syst)} \pm 0.12 \text{ (lumi)} \begin{array}{l} +0.05 \\ -0.06 \end{array} \text{ (acc) nb,} \\ \sigma(Z) \mathcal{B}(Z \rightarrow \mu^+ \mu^-) &= 2.021 \pm 0.009 \text{ (syst)} \pm 0.028 \text{ (lumi)} \begin{array}{l} +0.011 \\ -0.013 \end{array} \text{ (acc) nb.}\end{aligned}$$

In these results, statistical uncertainties are negligible, the systematic uncertainty is denoted with ‘syst’, the uncertainty in the integrated luminosity with ‘lumi’, and the uncertainty in the acceptance with ‘acc’. Based on these values, cross section ratios are calculated, reaching relative uncertainties below the level of 1%, as the uncertainty in the luminosity measurement cancels along with further correlated uncertainties. All measured values are in good agreement with theoretical predictions at [next-to-next-to-leading order \(NNLO\)](#)+[next-to-next-to-leading logarithmic \(NNLL\)](#) precision in [QCD](#).

In summary, the results represent some of the most precise cross section measurements ever conducted at hadron colliders. Despite the high precision, the results are still in good agreement with theoretical predictions. Due to the high trust in the theoretical predictions, building upon the manifold of previous tests, this combination indicates an excellent quality of the data and great potential for future analyses. The analysis has successfully finished the internal review process and has been submitted to the [Journal of High Energy Physics](#), a preprint being already available [5].

Following the detailed studies of muons in the scope of the cross section measurement, the effort in the calibration of muon momenta was further intensified and the results were ultimately provided internally to the [CMS](#) Collaboration. In the calibration procedure, residual differences in the muon description are divided into two categories: scale and resolution. The correction parameters are derived in four consecutive steps, where first a rough parametrization is extracted based on single muon distributions. These initial rough corrections are then fine tuned by using a dimuon resonance, such as the Z boson peak or the J/ψ peak. Here, the peak position gives rise to the momentum scale of the muons while the width of the peak contains information on the resolution. The correction procedure has been made publicly available in a [Detector Performance Summary](#) [83].

The muon scale and resolution corrections are a key piece of the puzzle in conducting high-precision analyses with the [CMS](#) experiment. Along with the efficiency corrections they belong to the fundamental corrections applied to almost any muon-related [CMS](#) analysis. For this reasons a fast and reliable alternative to calculate the uncertainties of the correction was developed, tested and provided to the [CMS](#) Collaboration in the scope of this thesis. This method simplifies the uncertainty calculation drastically, decreasing the computing demand for the uncertainty calculation by a factor of 25 without losing significant description power.

Although the method leads to highly precise results, there are still possible ways to further advance the accuracy in the correction. High potential lies in the combination of information from the different muon momentum regions, fully exploiting the different sensitivities of each region. Furthermore, additional resonances, such as the Υ mesons, which are located in the mass region between the J/ψ meson and the Z boson, can be added either to cross-check the results or to add to the sensitivity.

Bibliography

- [1] CMS Collaboration, “Luminosity measurement in proton-proton collisions at 13.6 TeV in 2022 at CMS”, CMS Physics Analysis Summary CMS-PAS-LUM-22-001, CERN, 2024.
- [2] CERN, “The CERN Experimental Programme - Grey Book database”, 2025. Accessed: 2025-04-18. <https://greybook.cern.ch/experiment/detail?id=CMS>.
- [3] A. Bodek et al., “Extracting Muon Momentum Scale Corrections for Hadron Collider Experiments”, *Eur. Phys. J. C* **72** (2012) 2194, doi:10.1140/epjc/s10052-012-2194-8, arXiv:1208.3710.
- [4] CMS Collaboration, “Measurement of the inclusive cross section of Z boson production in pp collisions at $\sqrt{s} = 13.6$ TeV”, CMS Physics Analysis Summary CMS-PAS-SMP-22-017, CERN, Geneva, 2023.
- [5] CMS Collaboration, “Measurements of the inclusive W and Z boson production cross sections and their ratios in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, March, 2025. arXiv:2503.09742. Submitted to *JHEP*.
- [6] ERA Forum stakeholders, “Living Guidelines on the Responsible Use of Generative AI in Research”, April, 2025. Second edition. https://research-and-innovation.ec.europa.eu/document/download/2b6cf7e5-36ac-41cb-aab5-0d32050143dc_en?filename=ec_rtd_ai-guidelines.pdf.
- [7] Particle Data Group Collaboration, “Review of particle physics”, *Phys. Rev. D* **110** (2024), no. 3, 030001, doi:10.1103/PhysRevD.110.030001.
- [8] M. E. Peskin and D. V. Schroeder, “An Introduction to Quantum Field Theory”. Westview Press, 1995. ISBN 978-0-201-50397-5.
- [9] J. Ellis, “Higgs Physics”, in *2013 European School of High-Energy Physics*, pp. 117–168. 2015. arXiv:1312.5672. doi:10.5170/CERN-2015-004.117.
- [10] K. J. Barnes, “Group Theory for the Standard Model of Particle Physics and Beyond”. CRC Press, 1st edition, 2010. doi:10.1201/9781439895207.
- [11] J. Campbell, J. Huston, and F. Krauss, “The Black Book of Quantum Chromodynamics: A Primer for the LHC Era”. Oxford University Press, December, 2017. doi:10.1093/oso/9780199652747.001.0001, ISBN 9780199652747.
- [12] Symmetry Magazine, “The standard model”, 2025. Accessed: 2025-02-25. <https://www.symmetrymagazine.org/standard-model/>.
- [13] E. Noether, “Invariant Variation Problems”, *Gott. Nachr.* **1918** (1918) 235–257, doi:10.1080/00411457108231446, arXiv:physics/0503066.

- [14] M. Gell-Mann, “Symmetries of Baryons and Mesons”, *Phys. Rev.* **125** (1962) 1067–1084, doi:10.1103/PhysRev.125.1067.
- [15] M. Gell-Mann, “A Schematic Model of Baryons and Mesons”, *Phys. Lett.* **8** (1964) 214–215, doi:10.1016/S0031-9163(64)92001-3.
- [16] G. Zweig, “An SU_3 model for strong interaction symmetry and its breaking; Version 2”, 1964. CERN preprint CERN-TH-412.
- [17] D. J. Gross and F. Wilczek, “Ultraviolet Behavior of Nonabelian Gauge Theories”, *Phys. Rev. Lett.* **30** (1973) 1343–1346, doi:10.1103/PhysRevLett.30.1343.
- [18] D. J. Gross and F. Wilczek, “Asymptotically Free Gauge Theories - I”, *Phys. Rev. D* **8** (1973) 3633–3652, doi:10.1103/PhysRevD.8.3633.
- [19] D. J. Gross and F. Wilczek, “Asymptotically Free Gauge Theories. II”, *Phys. Rev. D* **9** (1974) 980–993, doi:10.1103/PhysRevD.9.980.
- [20] S. L. Glashow, “Partial Symmetries of Weak Interactions”, *Nucl. Phys.* **22** (1961) 579–588, doi:10.1016/0029-5582(61)90469-2.
- [21] A. Salam and J. C. Ward, “Weak and electromagnetic interactions”, *Nuovo Cim.* **11** (1959) 568–577, doi:10.1007/BF02726525.
- [22] S. Weinberg, “A Model of Leptons”, *Phys. Rev. Lett.* **19** (1967) 1264–1266, doi:10.1103/PhysRevLett.19.1264.
- [23] J. S. Schwinger, “The Theory of quantized fields. 1.”, *Phys. Rev.* **82** (1951) 914–927, doi:10.1103/PhysRev.82.914.
- [24] G. Luders, “On the Equivalence of Invariance under Time Reversal and under Particle-Antiparticle Conjugation for Relativistic Field Theories”, *Kong. Dan. Vid. Sel. Mat. Fys. Medd.* **28N5** (1954), no. 5, 1–17.
- [25] R. Jost, “A remark on the C.T.P. theorem”, *Helv. Phys. Acta* **30** (1957) 409–416.
- [26] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons”, *Phys. Rev. Lett.* **13** (1964) 321–323, doi:10.1103/PhysRevLett.13.321.
- [27] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Phys. Rev. Lett.* **13** (1964) 508–509, doi:10.1103/PhysRevLett.13.508.
- [28] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, “Global Conservation Laws and Massless Particles”, *Phys. Rev. Lett.* **13** (1964) 585–587, doi:10.1103/PhysRevLett.13.585.
- [29] P. W. Higgs, “Spontaneous Symmetry Breakdown without Massless Bosons”, *Phys. Rev.* **145** (1966) 1156–1163, doi:10.1103/PhysRev.145.1156.
- [30] H. Yukawa, “On the Interaction of Elementary Particles I”, *Proc. Phys. Math. Soc. Jap.* **17** (1935) 48–57, doi:10.1143/PTPS.1.1.
- [31] N. Cabibbo, “Unitary Symmetry and Leptonic Decays”, *Phys. Rev. Lett.* **10** (1963) 531–533, doi:10.1103/PhysRevLett.10.531.
- [32] M. Kobayashi and T. Maskawa, “CP Violation in the Renormalizable Theory of Weak Interaction”, *Prog. Theor. Phys.* **49** (1973) 652–657, doi:10.1143/PTP.49.652.

- [33] H. Pilkuhn, “The Interactions of Hadrons”. North-Holland, Amsterdam, 1967.
- [34] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (October, 2017) doi:10.1140/epjc/s10052-017-5199-5.
- [35] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand, “Parton Fragmentation and String Dynamics”, *Phys. Rep.* **97** (1983) 31–145, doi:10.1016/0370-1573(83)90080-7.
- [36] B. Andersson, “The Lund Model”, volume 7. Cambridge University Press, 1998. doi:10.1017/9781009401296, ISBN 978-1-009-40129-6, 978-1-009-40125-8, 978-1-009-40128-9, 978-0-521-01734-3, 978-0-521-42094-5, 978-0-511-88149-7.
- [37] G. C. Fox and S. Wolfram, “A Model for Parton Showers in QCD”, *Nucl. Phys. B* **168** (1980) 285–295, doi:10.1016/0550-3213(80)90111-X.
- [38] R. D. Field and S. Wolfram, “A QCD Model for $e^+ e^-$ Annihilation”, *Nucl. Phys. B* **213** (1983) 65–84, doi:10.1016/0550-3213(83)90175-X.
- [39] B. R. Webber, “A QCD Model for Jet Fragmentation Including Soft Gluon Interference”, *Nucl. Phys. B* **238** (1984) 492–528, doi:10.1016/0550-3213(84)90333-X.
- [40] KATRIN Collaboration, “Direct neutrino-mass measurement based on 259 days of KATRIN data”, *Science* **388** (2025), no. 6743, 180–185, doi:10.1126/science.adq9592.
- [41] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr., “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4\text{kpc}$) to UGC 2885 ($R = 122\text{kpc}$)”, *Astrophys. J.* **238** (1980) 471, doi:10.1086/158003.
- [42] E. Corbelli and P. Salucci, “The Extended Rotation Curve and the Dark Matter Halo of M33”, *Mon. Not. R. Astron. Soc.* **311** (2000) 441–447, doi:10.1046/j.1365-8711.2000.03075.x, arXiv:astro-ph/9909252.
- [43] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln”, *Helv. Phys. Acta* **6** (1933) 110–127, doi:10.1007/s10714-008-0707-4.
- [44] D. Clowe et al., “A direct empirical proof of the existence of dark matter”, *Astrophys. J. Lett.* **648** (2006) L109–L113, doi:10.1086/508162, arXiv:astro-ph/0608407.
- [45] WMAP Collaboration, “Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation”, *Astrophys. J. Suppl.* **192** (2011) 18, doi:10.1088/0067-0049/192/2/18, arXiv:1001.4538.
- [46] Planck Collaboration, “Planck 2018 results. V. CMB power spectra and likelihoods”, *Astron. Astrophys.* **641** (2020) A5, doi:10.1051/0004-6361/201936386, arXiv:1907.12875.
- [47] L. de Broglie, “Waves and quanta”, *Nature* **112** (Oct, 1923) 540–540, doi:10.1038/112540a0.
- [48] CMS Collaboration, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30–61, doi:10.1016/j.physletb.2012.08.021, arXiv:1207.7235.

[49] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1–29, [doi:10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020), [arXiv:1207.7214](https://arxiv.org/abs/1207.7214).

[50] O. S. Bruning et al., “LHC Design Report Vol.1: The LHC Main Ring”, Technical Report CERN-2004-003-V-1, CERN, June, 2004.
[doi:10.5170/CERN-2004-003-V-1](https://doi.org/10.5170/CERN-2004-003-V-1).

[51] CMS Collaboration, “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”, Technical Report CERN-LHCC-2006-001, CMS-TDR-8-1, CERN, 2006. [doi:10.2172/2510878](https://doi.org/10.2172/2510878).

[52] O. Aberle et al., “High-Luminosity Large Hadron Collider (HL-LHC): Technical design report”. CERN Yellow Reports: Monographs. CERN, Geneva, 2020.
[doi:10.23731/CYRM-2020-0010](https://doi.org/10.23731/CYRM-2020-0010).

[53] CMS Collaboration, “The Phase-2 Upgrade of the CMS Tracker”, Technical Report CERN-LHCC-2017-009, CERN, Geneva, 2017. [doi:10.17181/CERN.QZ28.FLHW](https://doi.org/10.17181/CERN.QZ28.FLHW).

[54] CMS Collaboration, “The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger”, Technical Report CERN-LHCC-2021-007, CMS-TDR-022, CERN, Geneva, 2021. This is the final version of the document, approved by the LHCC.

[55] CMS Collaboration, “The Phase-2 Upgrade of the CMS Muon Detectors”, Technical Report CERN-LHCC-2017-012, CMS-TDR-016, CERN, Geneva, 2017.
[doi:10.17181/CERN.5T9S.VPMI](https://doi.org/10.17181/CERN.5T9S.VPMI), This is the final version, approved by the LHCC.

[56] R. Bailey and P. Collier, “Standard Filling Schemes for Various LHC Operation Modes”, Technical Report LHC-PROJECT-NOTE-323, CERN, Geneva, 2003.

[57] D. Abbaneo et al., “Studies on the upgrade of the muon system in the forward region of the CMS experiment at LHC with GEMs”, *JINST* **9** (January, 2014) [doi:10.1088/1748-0221/9/01/C01053](https://doi.org/10.1088/1748-0221/9/01/C01053).

[58] CMS Collaboration, “Performance and quality control of the first CMS GE2/1 muon production chambers”, Technical Report CMS-DP-2024-075, CERN, 2024.

[59] Y. Hong, “GEM detectors for the CMS Endcap Muon System: status of three detector stations”, *JINST* **20** (April, 2025) C04022,
[doi:10.1088/1748-0221/20/04/C04022](https://doi.org/10.1088/1748-0221/20/04/C04022).

[60] R. L. Gluckstern, “Uncertainties in track momentum and direction, due to multiple scattering and measurement errors”, *Nucl. Instrum. Methods* **24** (1963) 381–389, [doi:10.1016/0029-554X\(63\)90347-1](https://doi.org/10.1016/0029-554X(63)90347-1).

[61] CMS Tracker Group Collaboration, “The CMS Phase-1 Pixel Detector Upgrade”, *JINST* **16** (2021), no. 02, P02027, [doi:10.1088/1748-0221/16/02/P02027](https://doi.org/10.1088/1748-0221/16/02/P02027), [arXiv:2012.14304](https://arxiv.org/abs/2012.14304).

[62] CMS Collaboration, “CMS Tracker Material Budget Plots”. Accessed: 2025-03-10.
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/TrackerMaterialBudgetplots>.

[63] CMS Collaboration, “The CMS Experiment at the CERN LHC”, *JINST* **3** (2008) S08004, [doi:10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004).

- [64] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *JINST* **9** (2014), no. 10, P10009, [doi:10.1088/1748-0221/9/10/P10009](https://doi.org/10.1088/1748-0221/9/10/P10009), [arXiv:1405.6569](https://arxiv.org/abs/1405.6569).
- [65] CMS Collaboration, “The CMS electromagnetic calorimeter project: Technical Design Report”, Technical Report CERN-LHCC-97-033, CERN, 1997.
- [66] CMS Collaboration, “Prompt and Non-Prompt J/ψ Production in pp Collisions at $\sqrt{s} = 7$ TeV”, *Eur. Phys. J. C* **71** (2011) 1575, [doi:10.1140/epjc/s10052-011-1575-8](https://doi.org/10.1140/epjc/s10052-011-1575-8), [arXiv:1011.4193](https://arxiv.org/abs/1011.4193).
- [67] CMS Collaboration, “Development of the CMS detector for the CERN LHC Run 3”, *JINST* **19** (2024), no. 05, P05064, [doi:10.1088/1748-0221/19/05/P05064](https://doi.org/10.1088/1748-0221/19/05/P05064), [arXiv:2309.05466](https://arxiv.org/abs/2309.05466).
- [68] R. E. Kalman, “A new approach to linear filtering and prediction problems”, *J. Basic Eng.* **82** (March, 1960) 35–45, [doi:10.1115/1.3662552](https://doi.org/10.1115/1.3662552).
- [69] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Methods Phys. Res.* **262** (1987), no. 2, 444–450, [doi:10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4).
- [70] W. Adam, B. Mangano, T. Speer, and T. Todorov, “Track reconstruction in the CMS tracker”, Technical Report CERN-CMS-NOTE-2006-041, CMS-NOTE-2006-041, CERN, December, 2005.
- [71] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017), no. 10, P10003, [doi:10.1088/1748-0221/12/10/P10003](https://doi.org/10.1088/1748-0221/12/10/P10003), [arXiv:1706.04965](https://arxiv.org/abs/1706.04965).
- [72] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_t jet clustering algorithm”, *JHEP* **04** (2008) 063, [doi:10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063), [arXiv:0802.1189](https://arxiv.org/abs/0802.1189).
- [73] W. Heisenberg, “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik”, *Zeitschrift für Physik* **43** (1927), no. 3, 172–198, [doi:10.1007/BF01397280](https://doi.org/10.1007/BF01397280).
- [74] CMS Collaboration, “Muon ID and Isolation efficiencies with muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, Technical Report CMS-DP-2024-067, CERN, 2024.
- [75] CMS Collaboration, “Performance of the reconstruction and identification of high-momentum muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 02, P02027, [doi:10.1088/1748-0221/15/02/P02027](https://doi.org/10.1088/1748-0221/15/02/P02027), [arXiv:1912.03516](https://arxiv.org/abs/1912.03516).
- [76] CMS Collaboration, “Muon identification using multivariate techniques in the CMS experiment in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **19** (2024), no. 02, P02031, [doi:10.1088/1748-0221/19/02/P02031](https://doi.org/10.1088/1748-0221/19/02/P02031), [arXiv:2310.03844](https://arxiv.org/abs/2310.03844).
- [77] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018), no. 06, P06015, [doi:10.1088/1748-0221/13/06/P06015](https://doi.org/10.1088/1748-0221/13/06/P06015), [arXiv:1804.04528](https://arxiv.org/abs/1804.04528).
- [78] CMS Collaboration, “Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV”, *JINST* **7** (2012) P10002, [doi:10.1088/1748-0221/7/10/P10002](https://doi.org/10.1088/1748-0221/7/10/P10002), [arXiv:1206.4071](https://arxiv.org/abs/1206.4071).

[79] G. Breit and E. Wigner, “Capture of slow neutrons”, *Phys. Rev.* **49** (April, 1936) 519–531, doi:10.1103/PhysRev.49.519.

[80] D. Guthmann, “Muon p_T Scale and Resolution Calibration for the CMS Experiment”, Master’s thesis, Karlsruhe Institute of Technology (KIT), 2024.

[81] M. J. Oreglia, “A study of the reactions $\psi' \rightarrow \gamma\gamma\psi$ ”. PhD thesis, Stanford University, 1980. SLAC-R-236.

[82] J. E. Gaiser, “Charmonium spectroscopy from radiative decays of the J/ψ and ψ' ”. PhD thesis, Stanford University, 1982. SLAC-R-255.

[83] CMS Collaboration, “Muon momentum calibration with proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, Technical Report CMS-DP-2024-065, CERN, 2024.

[84] J. Floßmann, “Uncertainty calculations on Muon p_T Scale and Resolution Calibration for the CMS Experiment”, Bachelor’s thesis, Karlsruhe Institute of Technology (KIT), 2024.

[85] CMS Collaboration, “High-precision measurement of the W boson mass with the CMS experiment at the LHC”, December, 2024. arXiv:2412.13872. Submitted to *Nature*.

[86] UA1 Collaboration, “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV”, *Phys. Lett. B* **122** (1983) 103–116, doi:10.1016/0370-2693(83)91177-2.

[87] UA2 Collaboration, “Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider”, *Phys. Lett. B* **122** (1983) 476–485, doi:10.1016/0370-2693(83)91605-2.

[88] UA1 Collaboration, “Experimental Observation of Lepton Pairs of Invariant Mass Around $95\text{ GeV}/c^2$ at the CERN SPS Collider”, *Phys. Lett. B* **126** (1983) 398–410, doi:10.1016/0370-2693(83)90188-0.

[89] UA2 Collaboration, “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider”, *Phys. Lett. B* **129** (1983) 130–140, doi:10.1016/0370-2693(83)90744-X.

[90] ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group Collaboration, “Precision electroweak measurements on the Z resonance”, *Phys. Rep.* **427** (2006) 257–454, doi:10.1016/j.physrep.2005.12.006, arXiv:hep-ex/0509008.

[91] ATLAS Collaboration, “Measurement of W^\pm -boson and Z -boson production cross-sections in pp collisions at $\sqrt{s} = 2.76$ TeV with the ATLAS detector”, *Eur. Phys. J. C* **79** (2019), no. 11, 901, doi:10.1140/epjc/s10052-019-7399-7, arXiv:1907.03567.

[92] CMS Collaboration, “Study of Z production in $PbPb$ and pp collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV in the dimuon and dielectron decay channels”, *JHEP* **03** (2015) 022, doi:10.1007/JHEP03(2015)022, arXiv:1410.4825.

[93] ATLAS Collaboration, “Measurements of W and Z boson production in pp collisions at $\sqrt{s} = 5.02$ TeV with the ATLAS detector”, *Eur. Phys. J. C* **79** (2019), no. 2, 128, doi:10.1140/epjc/s10052-019-6622-x, arXiv:1810.08424. [Erratum: Eur.Phys.J.C 79, 374 (2019)].

[94] CMS Collaboration, “Measurement of the inclusive cross sections for W and Z boson production in proton-proton collisions at $\sqrt{s} = 5.02$ and 13 TeV”, August, 2024. [arXiv:2408.03744](https://arxiv.org/abs/2408.03744). Submitted to *JHEP*.

[95] ATLAS Collaboration, “Precision measurement and interpretation of inclusive W^+ , W^- and Z/γ^* production cross sections with the ATLAS detector”, *Eur. Phys. J. C* **77** (2017), no. 6, 367, [doi:10.1140/epjc/s10052-017-4911-9](https://doi.org/10.1140/epjc/s10052-017-4911-9), [arXiv:1612.03016](https://arxiv.org/abs/1612.03016).

[96] CMS Collaboration, “Measurements of Inclusive W and Z Cross Sections in pp Collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **01** (2011) 080, [doi:10.1007/JHEP01\(2011\)080](https://doi.org/10.1007/JHEP01(2011)080), [arXiv:1012.2466](https://arxiv.org/abs/1012.2466).

[97] CMS Collaboration, “Measurement of the Inclusive W and Z Production Cross Sections in pp Collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **10** (2011) 132, [doi:10.1007/JHEP10\(2011\)132](https://doi.org/10.1007/JHEP10(2011)132), [arXiv:1107.4789](https://arxiv.org/abs/1107.4789).

[98] LHCb Collaboration, “Measurement of the forward W boson cross-section in pp collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **12** (2014) 079, [doi:10.1007/JHEP12\(2014\)079](https://doi.org/10.1007/JHEP12(2014)079), [arXiv:1408.4354](https://arxiv.org/abs/1408.4354).

[99] LHCb Collaboration, “Measurement of the cross-section for $Z \rightarrow e^+e^-$ production in pp collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **02** (2013) 106, [doi:10.1007/JHEP02\(2013\)106](https://doi.org/10.1007/JHEP02(2013)106), [arXiv:1212.4620](https://arxiv.org/abs/1212.4620).

[100] LHCb Collaboration, “Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV”, *JHEP* **08** (2015) 039, [doi:10.1007/JHEP08\(2015\)039](https://doi.org/10.1007/JHEP08(2015)039), [arXiv:1505.07024](https://arxiv.org/abs/1505.07024).

[101] LHCb Collaboration, “Measurement of forward $W \rightarrow e\nu$ production in pp collisions at $\sqrt{s} = 8$ TeV”, *JHEP* **10** (2016) 030, [doi:10.1007/JHEP10\(2016\)030](https://doi.org/10.1007/JHEP10(2016)030), [arXiv:1608.01484](https://arxiv.org/abs/1608.01484).

[102] LHCb Collaboration, “Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **09** (2016) 136, [doi:10.1007/JHEP09\(2016\)136](https://doi.org/10.1007/JHEP09(2016)136), [arXiv:1607.06495](https://arxiv.org/abs/1607.06495).

[103] LHCb Collaboration, “Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV”, *JHEP* **01** (2016) 155, [doi:10.1007/JHEP01\(2016\)155](https://doi.org/10.1007/JHEP01(2016)155), [arXiv:1511.08039](https://arxiv.org/abs/1511.08039).

[104] CMS Collaboration, “Measurement of inclusive W and Z boson production cross sections in pp collisions at $\sqrt{s} = 8$ TeV”, *Phys. Rev. Lett.* **112** (2014) 191802, [doi:10.1103/PhysRevLett.112.191802](https://doi.org/10.1103/PhysRevLett.112.191802), [arXiv:1402.0923](https://arxiv.org/abs/1402.0923).

[105] ATLAS Collaboration, “Measurement of W^\pm and Z-boson production cross sections in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Phys. Lett. B* **759** (2016) 601–621, [doi:10.1016/j.physletb.2016.06.023](https://doi.org/10.1016/j.physletb.2016.06.023), [arXiv:1603.09222](https://arxiv.org/abs/1603.09222).

[106] ATLAS Collaboration, “Measurement of vector boson production cross sections and their ratios using pp collisions at $\sqrt{s} = 13.6$ TeV with the ATLAS detector”, *Phys. Lett. B* **854** (2024) 138725, [doi:10.1016/j.physletb.2024.138725](https://doi.org/10.1016/j.physletb.2024.138725), [arXiv:2403.12902](https://arxiv.org/abs/2403.12902).

[107] CMS Collaboration, “Performance of CMS muon reconstruction in cosmic-ray events”, *JINST* **5** (March, 2010) T03022, [doi:10.1088/1748-0221/5/03/T03022](https://doi.org/10.1088/1748-0221/5/03/T03022).

[108] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (2017), no. 10, 663, [doi:10.1140/epjc/s10052-017-5199-5](https://doi.org/10.1140/epjc/s10052-017-5199-5), [arXiv:1706.00428](https://arxiv.org/abs/1706.00428).

[109] C. Bierlich et al., “A comprehensive guide to the physics and usage of PYTHIA 8.3”, *SciPost Phys. Codeb.* **2022** 8, [doi:10.21468/SciPostPhysCodeb.8](https://doi.org/10.21468/SciPostPhysCodeb.8), [arXiv:2203.11601](https://arxiv.org/abs/2203.11601).

[110] CMS Collaboration, “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”, *Eur. Phys. J. C* **80** (2020), no. 1, 4, [doi:10.1140/epjc/s10052-019-7499-4](https://doi.org/10.1140/epjc/s10052-019-7499-4), [arXiv:1903.12179](https://arxiv.org/abs/1903.12179).

[111] GEANT4 Collaboration, “Geant4—a simulation toolkit”, *Nucl. Instrum. Methods Phys. Res.* **506** (2003), no. 3, 250–303, [doi:10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).

[112] J. Alwall et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *JHEP* **07** (2014) 079, [doi:10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079), [arXiv:1405.0301](https://arxiv.org/abs/1405.0301).

[113] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO”, *JHEP* **12** (2012) 061, [doi:10.1007/JHEP12\(2012\)061](https://doi.org/10.1007/JHEP12(2012)061), [arXiv:1209.6215](https://arxiv.org/abs/1209.6215).

[114] P. Nason, “A New method for combining NLO QCD with shower Monte Carlo algorithms”, *JHEP* **11** (2004) 040, [doi:10.1088/1126-6708/2004/11/040](https://doi.org/10.1088/1126-6708/2004/11/040), [arXiv:hep-ph/0409146](https://arxiv.org/abs/hep-ph/0409146).

[115] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method”, *JHEP* **11** (2007) 070, [doi:10.1088/1126-6708/2007/11/070](https://doi.org/10.1088/1126-6708/2007/11/070), [arXiv:0709.2092](https://arxiv.org/abs/0709.2092).

[116] E. Re, “Single-top Wt-channel production matched with parton showers using the POWHEG method”, *Eur. Phys. J. C* **71** (2011) 1547, [doi:10.1140/epjc/s10052-011-1547-z](https://doi.org/10.1140/epjc/s10052-011-1547-z), [arXiv:1009.2450](https://arxiv.org/abs/1009.2450).

[117] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”, *JHEP* **06** (2010) 043, [doi:10.1007/JHEP06\(2010\)043](https://doi.org/10.1007/JHEP06(2010)043), [arXiv:1002.2581](https://arxiv.org/abs/1002.2581).

[118] R. Frederix, E. Re, and P. Torrielli, “Single-top t-channel hadroproduction in the four-flavour scheme with POWHEG and aMC@NLO”, *JHEP* **09** (2012) 130, [doi:10.1007/JHEP09\(2012\)130](https://doi.org/10.1007/JHEP09(2012)130), [arXiv:1207.5391](https://arxiv.org/abs/1207.5391).

[119] P. Nason and G. Zanderighi, “ $W^+ W^-$, WZ and ZZ production in the POWHEG-BOX-V2”, *Eur. Phys. J. C* **74** (2014), no. 1, 2702, [doi:10.1140/epjc/s10052-013-2702-5](https://doi.org/10.1140/epjc/s10052-013-2702-5), [arXiv:1311.1365](https://arxiv.org/abs/1311.1365).

[120] P. Artoisenet, R. Frederix, O. Mattelaer, and R. Rietkerk, “Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations”, *JHEP* **03** (2013) 015, [doi:10.1007/JHEP03\(2013\)015](https://doi.org/10.1007/JHEP03(2013)015), [arXiv:1212.3460](https://arxiv.org/abs/1212.3460).

[121] S. Camarda et al., “DYTurbo: Fast predictions for Drell-Yan processes”, *Eur. Phys. J. C* **80** (2020), no. 3, 251, [doi:10.1140/epjc/s10052-020-7757-5](https://doi.org/10.1140/epjc/s10052-020-7757-5), [arXiv:1910.07049](https://arxiv.org/abs/1910.07049). [Erratum: *Eur.Phys.J.C* 80, 440 (2020)].

[122] S. Camarda, L. Cieri, and G. Ferrera, “Drell–Yan lepton-pair production: q_T resummation at N^3LL accuracy and fiducial cross sections at N^3LO ”, *Phys. Rev. D* **104** (2021), no. 11, L111503, [doi:10.1103/PhysRevD.104.L111503](https://doi.org/10.1103/PhysRevD.104.L111503), [arXiv:2103.04974](https://arxiv.org/abs/2103.04974).

- [123] S. Camarda, L. Cieri, and G. Ferrera, “Fiducial perturbative power corrections within the q_T subtraction formalism”, *Eur. Phys. J. C* **82** (2022), no. 6, 575, doi:10.1140/epjc/s10052-022-10510-x, arXiv:2111.14509.
- [124] S. Camarda, L. Cieri, and G. Ferrera, “Drell–Yan lepton-pair production: q_T resummation at N^4LL accuracy”, *Phys. Lett. B* **845** (2023) 138125, doi:10.1016/j.physletb.2023.138125, arXiv:2303.12781.
- [125] S. Bailey et al., “Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs”, *Eur. Phys. J. C* **81** (2021), no. 4, 341, doi:10.1140/epjc/s10052-021-09057-0, arXiv:2012.04684.
- [126] T.-J. Hou et al., “New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC”, *Phys. Rev. D* **103** (2021), no. 1, 014013, doi:10.1103/PhysRevD.103.014013, arXiv:1912.10053.
- [127] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020), no. 10, P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165.
- [128] C. Collaboration, “CMS Public MET Analysis Information”. Accessed: 2025-04-02. https://twiki.cern.ch/twiki/bin/view/CMSPublic/WorkBookMetAnalysis#xy_Shift_Correction.
- [129] R. Barlow and C. Beeston, “Fitting using finite Monte Carlo samples”, *Comput. Phys. Commun.* **77** (1993) 219, doi:10.1016/0010-4655(93)90005-W.
- [130] CMS Collaboration, “The CMS Statistical Analysis and Combination Tool: Combine”, *Comput. Softw. Big Sci.* **8** (2024), no. 1, 19, doi:10.1007/s41781-024-00121-4, arXiv:2404.06614.
- [131] ATLAS Collaboration, “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** (2021), no. 7, 578, doi:10.1140/epjc/s10052-021-09233-2, arXiv:2012.00578.
- [132] D. E. Groom, N. V. Mokhov, and S. I. Striganov, “Muon Stopping Power and Range Tables: 10 MeV–100 TeV”, *At. Data Nucl. Data Tables* **76** (July, 2001) 183–356, doi:10.1006/adnd.2001.0861.

Appendix

A Estimations for Relevant Final-State Particles

In this thesis, the most important final-state particles are muons and neutrinos. To provide a deeper understanding, this appendix makes a few rough estimations on the behavior of these particles in the CMS detector.

A.1 Muons in the CMS Detector

Usually, muons in the medium- p_T range are simply assumed to consistently reach the muon chambers. Since this assumption is crucial for the muon reconstruction and therefore has a strong impact on the different studies in this thesis, it is examined briefly.

The assumption rests on two main claims:

1. **Energy content:** muons retain sufficient energy to traverse the detector material without significant energy loss.
2. **Lifetime considerations:** muons live long enough to reach the muon chambers before decaying.

If a muon does not fulfill either of these aspects, then it is not expected to reach the muon chambers.

A.1.1 Limitations Due to Energy Loss

A muon can reach the muon chambers only if the energy it loses while traversing the detector is smaller than its initial energy. Since its mass is approximately 200 times larger than that of the electron, the critical energy, at which bremsstrahlung becomes dominant over ionization, is significantly higher for muons than for electrons. In copper, for example, muons between the energies of approximately 100 MeV and 10 GeV primarily lose energy through ionization, with a stopping power of just below $2 \text{ MeVcm}^2/\text{g}$. As the stopping power $-\text{d}E/\text{d}X$ is roughly constant across different materials, this value is considered representative for the entire CMS detector. The average density of the detector material excluding the muon system can be estimated as follows:

$$\rho = \frac{14\,000 \text{ t} - 12\,500 \text{ t}}{\pi \cdot (3 \text{ m})^2 \cdot 12.5 \text{ m}} \approx 4.24 \text{ g/cm}^3, \quad (8.1)$$

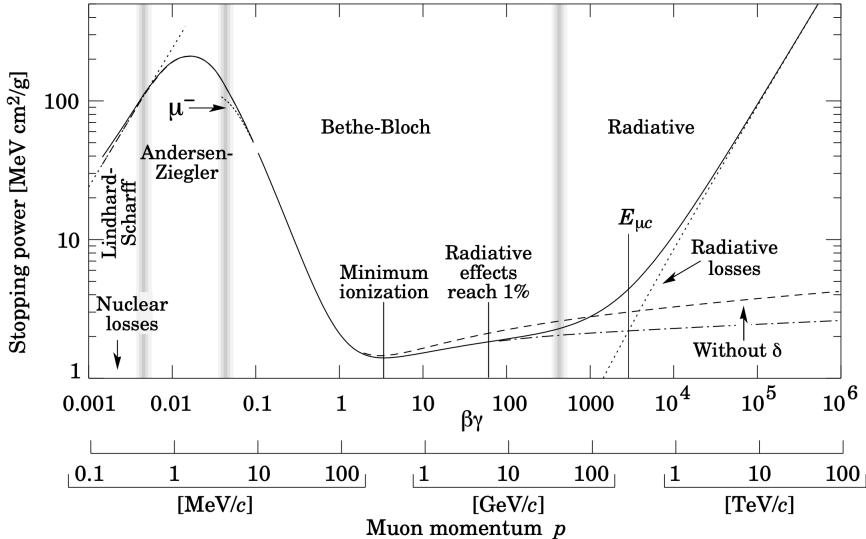


Figure A.1: Muon energy loss in copper, taken from Reference [132].

where the numerator represents the difference between the total detector mass and the mass of the iron return yoke, and the denominator is derived from the solenoid dimensions. Using this estimate, the average energy loss for muons per meter of travel within the detector is approximately 0.85 GeV/m. This implies that a 10 GeV muon could traverse over 11 m of detector material before being stopped, which corresponds to nearly half a revolution in the transverse plane of the detector while propagating along the z axis. Since this estimation is conservative, all muons with energies above 20 GeV—which are the focus of this study—should reach the muon chambers with high probability, ensuring a high muon reconstruction efficiency.

A.1.2 Limitations Due to Lifetime Effects

While muons are unstable, they have a relatively long lifetime of $\tau_\mu = 2.20 \mu\text{s}$. This finite lifetime means that some muons may decay before reaching the muon chambers, reducing the efficiency of muon reconstruction. To estimate an upper limit on this effect, consider a low-energy muon reaching the muon chambers with $\beta\gamma = 1$, corresponding to a momentum of approximately 100 MeV. Assuming conservatively that the muon has this final momentum throughout its entire journey in the central part of the detector, the probability that it decays before reaching the muon chambers (assumed to be at 3 m from the interaction point) is given by:

$$P(x < 3 \text{ m}) = 1 - \exp\left(-\frac{3 \text{ m}}{\gamma \tau_\mu \beta c}\right) \approx 0.5 \%. \quad (8.2)$$

Since the muons analyzed in this thesis typically have much higher momentum, their probability of decaying before reaching the muon chamber is negligible. Moreover, effects like these are fully accounted for in the detailed detector simulations, meaning they only contribute as second-order corrections.

A.2 Neutrinos

As neutrinos carry neither electromagnetic nor color charge, they only interact via the weak interaction. Consequently, the cross sections σ_ν of neutrino interactions are very small. The number of interactions of a neutrino stream within a material can be estimated using the mean free path λ . This quantity is defined as the average distance a neutrino can cover in the material before taking part in an interaction with the material:

$$\lambda \equiv \frac{1}{n\sigma_\nu}, \quad (8.3)$$

where n is the number density of the interactive particles in the material, which can be calculated from the mass density ρ and the atomic mass M of the material, as well as the Avogadro constant N_A :

$$\lambda = \frac{M}{\rho N_A \sigma_\nu}. \quad (8.4)$$

Using this mean free path, the interaction probability for a neutrino to decay in a material of maximum depth L can be calculated by evaluating the formula:

$$P = 1 - \exp\left\{ \left(-\frac{L}{\lambda} \right) \right\} \quad (8.5)$$

For neutrinos at typical collider energies at the GeV scale, the cross section for interaction with material is roughly $10 \times 10^{-38} \text{ cm}^2$. Using the material properties of lead, $M = 207 \text{ u}$, $\rho = 11.3 \text{ g cm}^{-3}$ and the Avogadro constant of $N_A = 6.02 \times 10^{23} \text{ mol}^{-1}$, the mean free path is of the order $\lambda = 3 \times 10^{13} \text{ m}$. This value corresponds to an interaction probability in 10 m of lead of roughly $P = 3 \times 10^{-13}$.

In conclusion, the scattering of single neutrinos can not be detected with the [CMS](#) experiment at a sufficient statistical significance.

B Muon Selection Efficiencies

This appendix provides a detailed overview of the different efficiencies in the muon reconstruction and selection procedure. All efficiencies are calculated using the **tag-and-probe (TnP)** technique as outlined in Section 6.2.

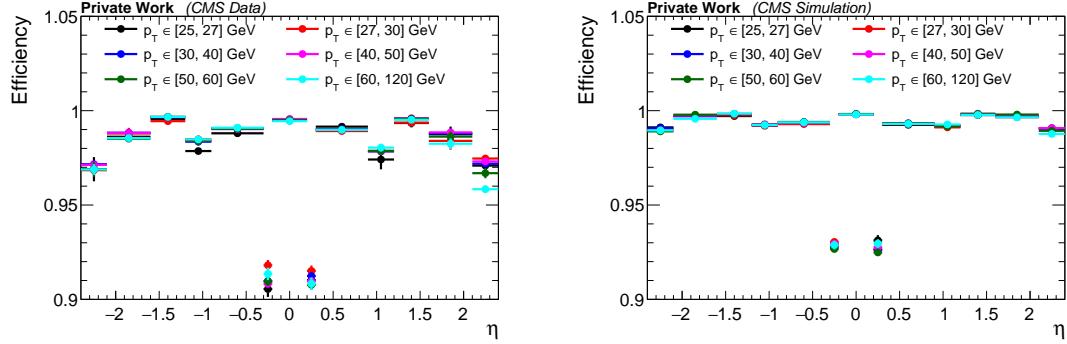


Figure B.2: Distributions of the standalone muon reconstruction efficiency in the data (left) and in the simulation (right). The standalone muon reconstruction efficiency is calculated based on tracker muons. The drop in efficiency between $0.2 < |\eta| < 0.3$ is a result of a gap between the muon wheels in the geometry of the CMS experiment, which is illustrated by Figure 4.1.

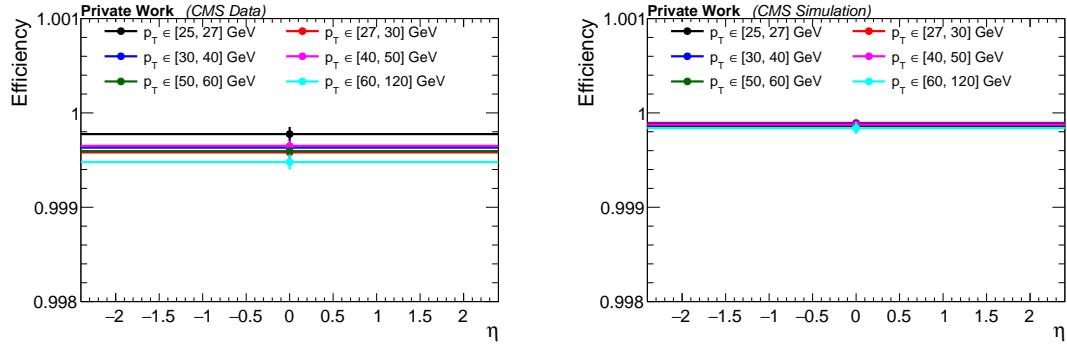


Figure B.3: Distributions of the global muon reconstruction efficiency in the data (left) and in the simulation (right). The global muon reconstruction efficiency is calculated based on standalone muons.

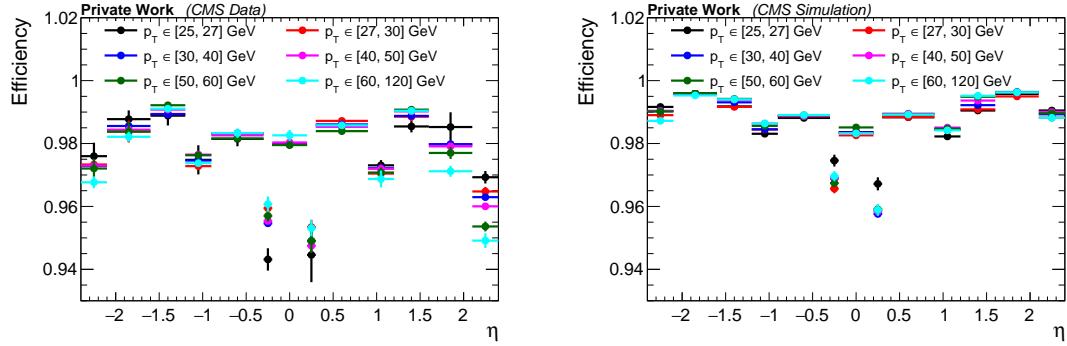


Figure B.4: Distributions of the tight identification selection efficiency in the data (left) and in the simulation (right). The tight identification selection efficiency is calculated based on global muons.

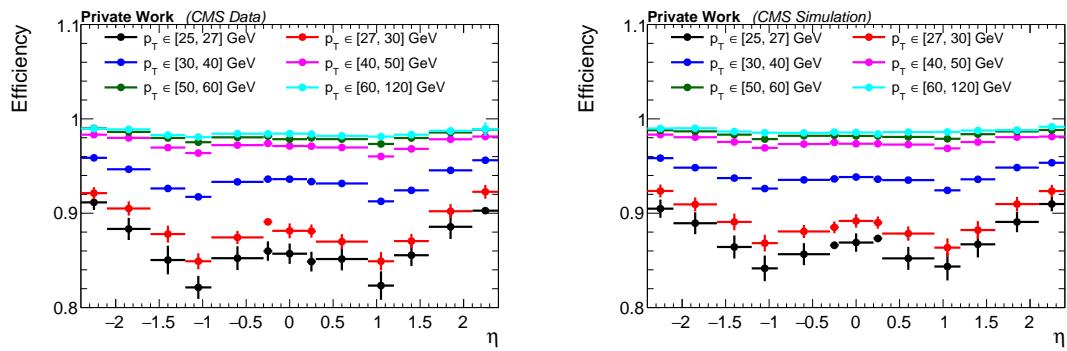


Figure B.5: Distributions of the tight isolation selection efficiency in the data (left) and in the simulation (right). The tight isolation selection efficiency is calculated based on tight identified muons.

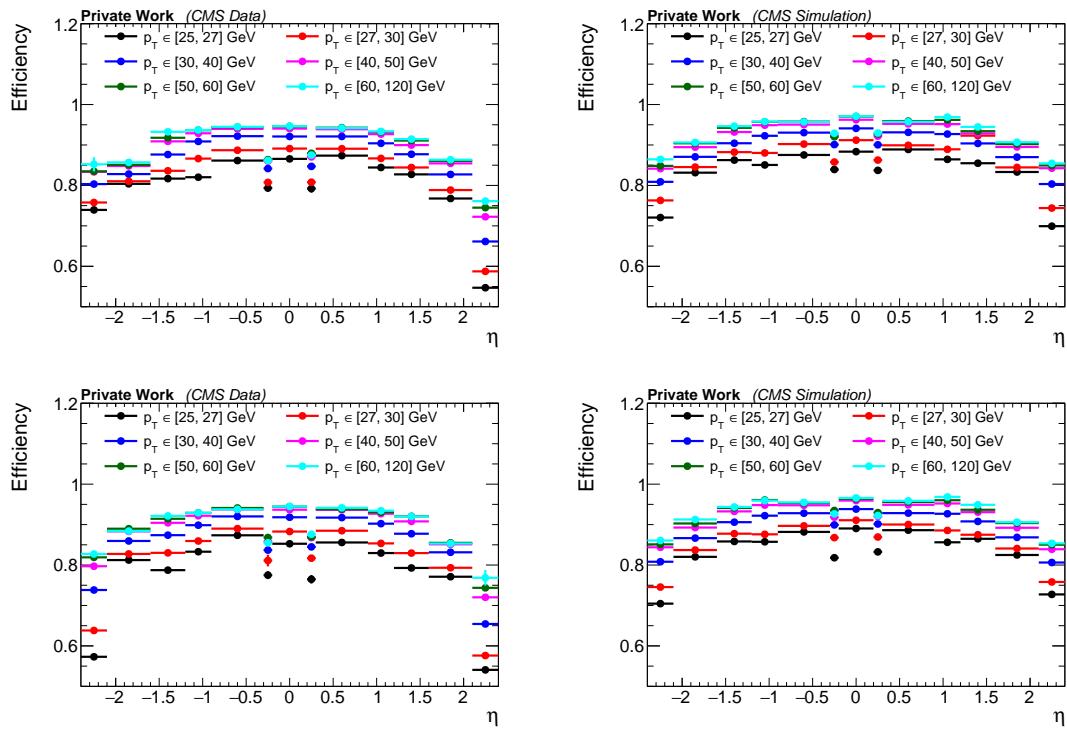


Figure B.6: Distributions of the muon trigger efficiency in the data (left) and in the simulation (right). The efficiency is calculated separately for negatively (top) and positively (bottom) charged muons, due to slight asymmetries across the whole η - p_T spectrum, in particular in the endcap regions at $|\eta| > 0.9$. The muon trigger efficiency is calculated based on tight identified and isolated muons.

C QCD Multijet Background Estimation

This appendix shows example plots in support of the QCD multijet background estimation.

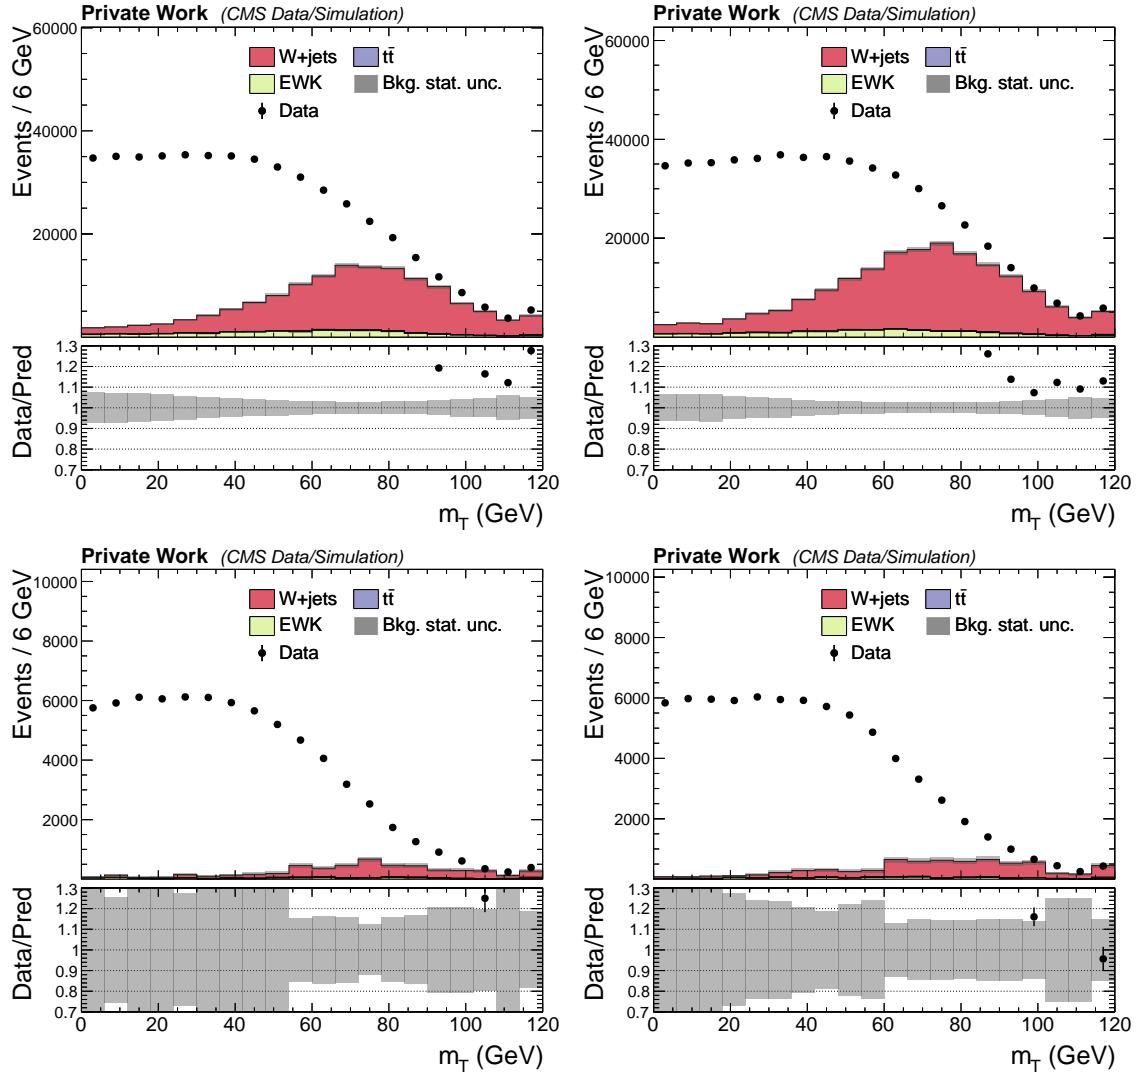


Figure C.7: Example plots for the first step of the QCD extrapolation procedure in the W^- region (left) and in the W^+ region (right)—corresponding to isolation slices from Figure 7.16. The first step consists of the QCD shape extraction in the non-isolated region. Here, the top plots correspond to muon isolation values between 0.2 and 0.25, and the bottom plots correspond to muon isolation values between 0.45 and 0.5. The difference between the m_T distributions in the data and the combined simulation is assumed to come from QCD multijet background and is therefore extracted by subtracting the simulated background contributions from the data. As the contributions from Z boson production are very small in the non-isolated regions, the contributions are counted in the [electroweak \(EWK\)](#) category.

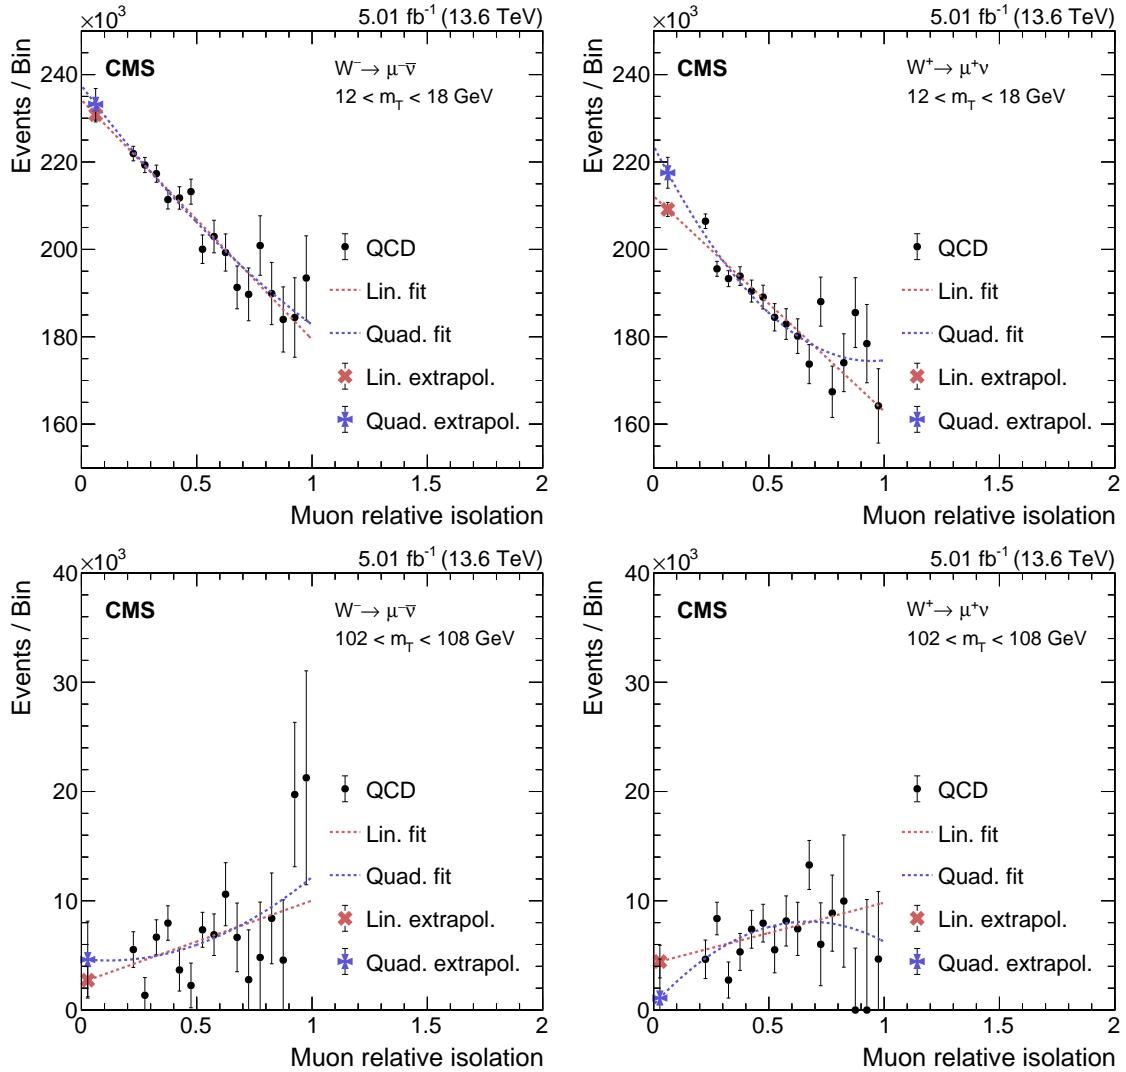


Figure C.8: Example plots for the second step of the **QCD** extrapolation procedure in the W^\pm regions—corresponding to m_T slices from Figure 7.16. The second step consists of the extrapolation of the **QCD** shape from the non-isolated into the isolated W^\pm signal region and is displayed here for two different bins of m_T in the W^- region (left) and in the W^+ region (right). At the top, the m_T bin with values in the range of 12 to 18 GeV is considered, while the right plot shows the extrapolation for the bin with m_T values between 102 and 108 GeV. The linear extrapolations are used to construct the nominal **QCD** shape, while the quadratic extrapolation and the statistical fit uncertainties are used as an estimate of the uncertainty.

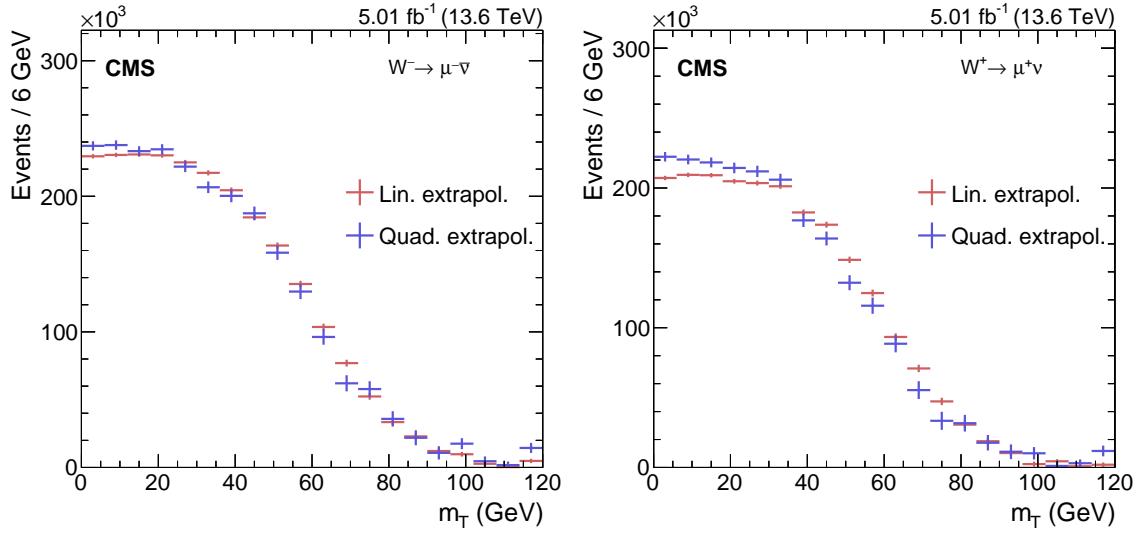


Figure C.9: Extrapolated QCD templates in the W^- region (left) and in the W^+ region (right), corresponding to the colored distributions from Figure 7.16. The inputs to this distribution come from the linear or quadratic extrapolation from the non-isolated control regions, such as those in Figure C.8.

D Correlation Plots of Maximum Likelihood Fit

This appendix provides details on the correlations between the leading systematic uncertainties with respect to the different points of interest.

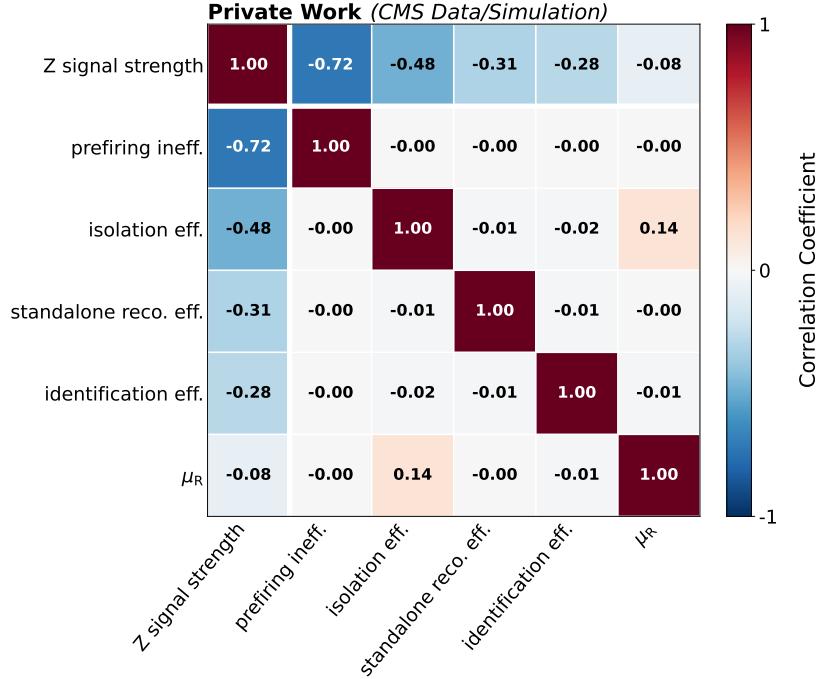


Figure D.10: Correlation matrix of signal strength parameter and leading uncertainty contributions in the Z cross section measurement.

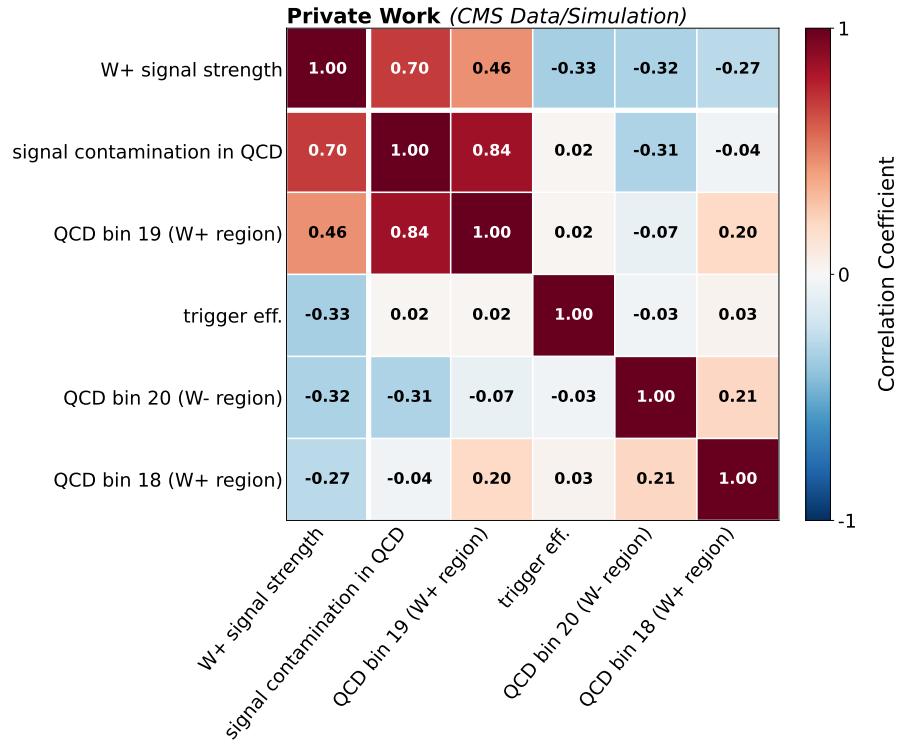


Figure D.11: Correlation matrix of signal strength parameter and leading uncertainty contributions in the W^+ cross section measurement.

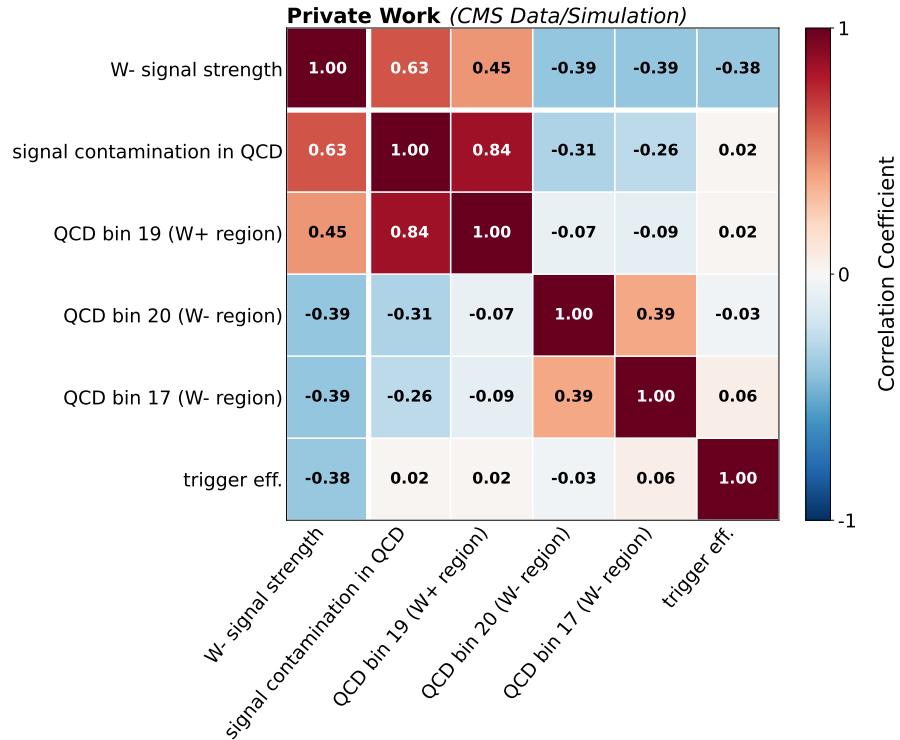


Figure D.12: Correlation matrix of signal strength parameter and leading uncertainty contributions in the W^- cross section measurement.

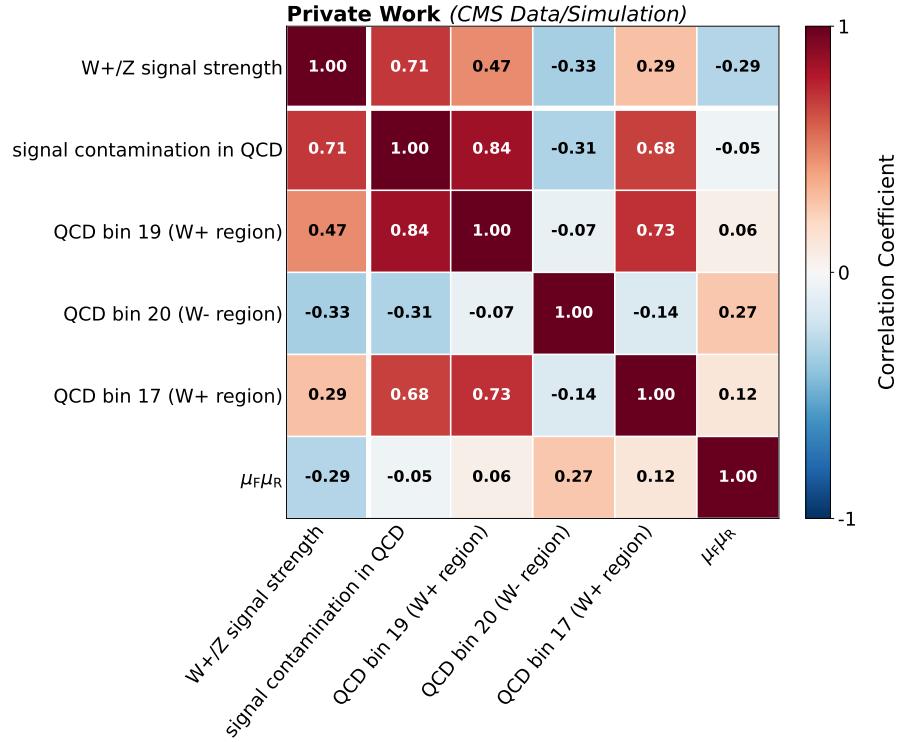


Figure D.13: Correlation matrix of signal strength parameter and leading uncertainty contributions in the measurement of the W^+/Z cross section ratio..

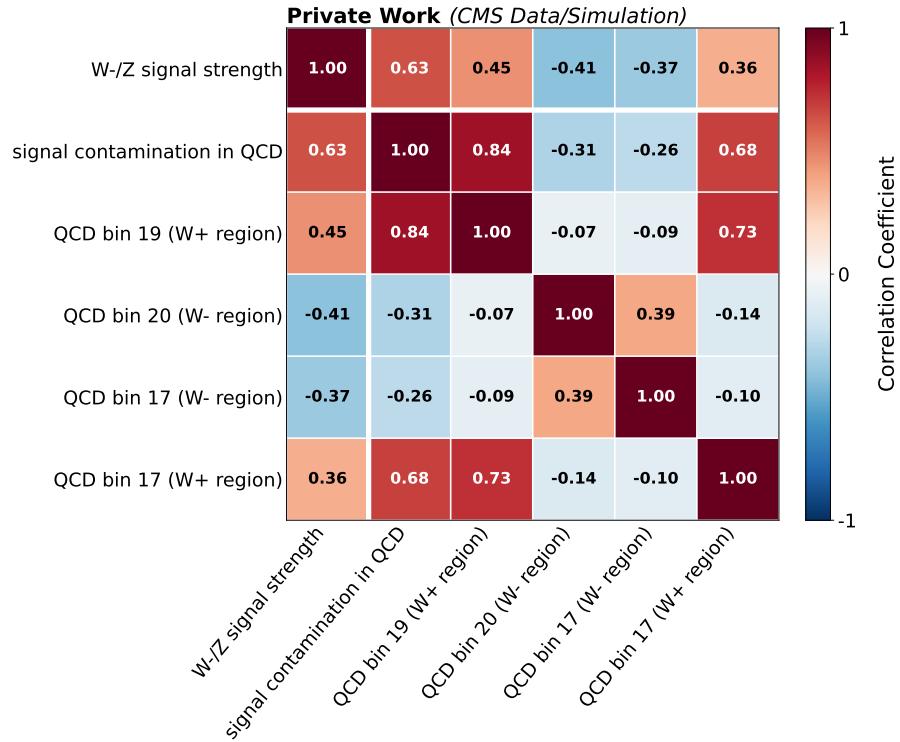


Figure D.14: Correlation matrix of signal strength parameter and leading uncertainty contributions in the measurement of the W^-/Z cross section ratio..

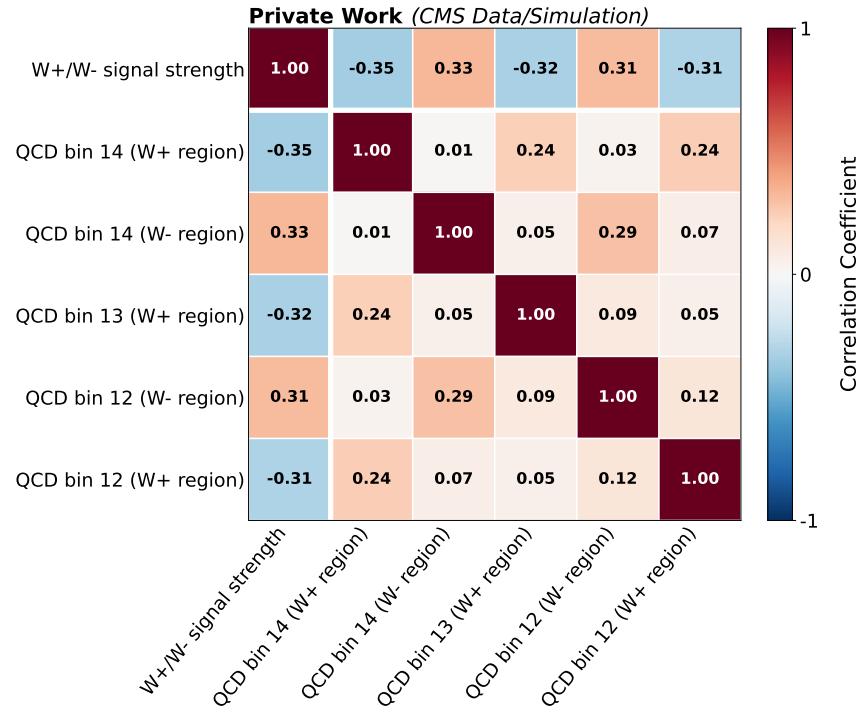


Figure D.15: Correlation matrix of signal strength parameter and leading uncertainty contributions in the measurement of the W^+/W^- cross section ratio..

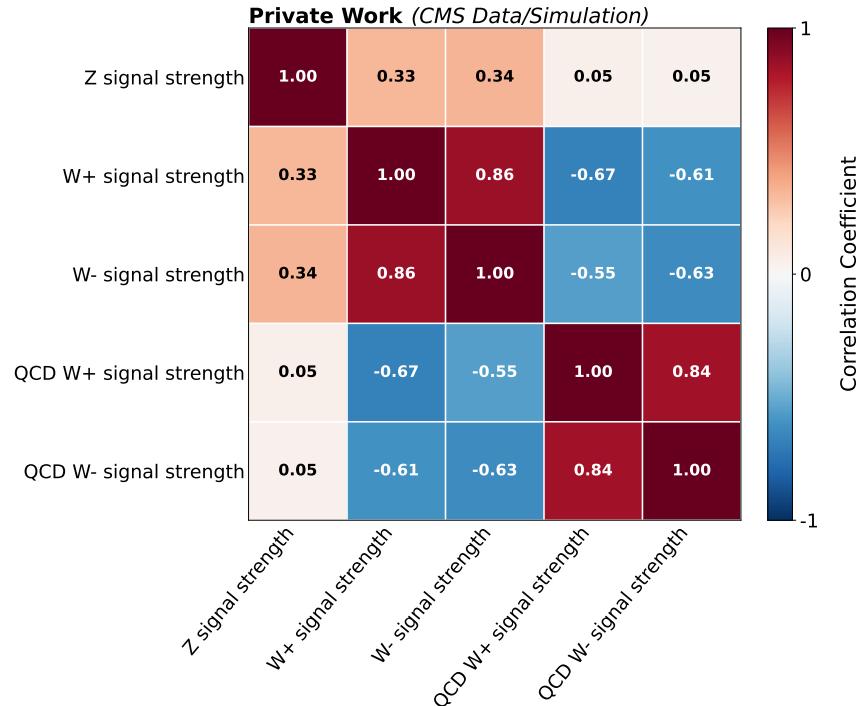


Figure D.16: Correlation matrix for the individual signal strength parameters in the combined [ML](#) fit.