



Durham E-Theses

Higher Order QCD Corrections to Electroweak Boson Production at Colliders

WALKER, DUNCAN,MARTIN

How to cite:

WALKER, DUNCAN,MARTIN (2019) *Higher Order QCD Corrections to Electroweak Boson Production at Colliders*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/13284/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Higher Order QCD Corrections to Electroweak Boson Production at Colliders

Duncan Martin Walker

A Thesis presented for the degree of
Doctor of Philosophy



Institute for Particle Physics Phenomenology
Department of Physics
Durham University
United Kingdom

July 2019

Higher Order QCD Corrections to Electroweak Boson Production at Colliders

Duncan Martin Walker

Submitted for the degree of Doctor of Philosophy

July 2019

Abstract: In this thesis we consider the Next-to-Next-to-Leading Order (NNLO) corrections to single charged electroweak boson production with associated QCD radiation in hadron-hadron collisions, calculated using the antenna subtraction method to regulate infrared (IR) divergences. Results are presented alongside the neutral current case for the inclusive transverse momentum spectrum and subsequent ratios both with and without the addition of state-of-the-art resummation results. In the former case a comparison to CMS data is also provided. We also discuss the phenomenological implications of the results when one or more jets are reconstructed from the QCD radiation. Particular attention is given to the impact on valence quark content of Parton Distribution Functions (PDFs) through a comparison with experimental results from the ATLAS, CMS and LHCb collaborations.

We then discuss the use of fixed-order QCD predictions for inclusive Drell-Yan production in the context of an effective Weinberg angle extraction using triple-differential data taken by the ATLAS collaboration at $\sqrt{s} = 8$ TeV, using the kinematics to extend the predictions to Next-to-Next-to-Next-to-Leading Order (N³LO) for certain parts of the measurement.

Finally, using the antenna subtraction method we derive the NNLO QCD corrections

to di-jet production in charged-current Deep Inelastic Scattering (DIS), allowing the first comparison to ZEUS data at this order. These results are then combined with inclusive structure functions using the method of projection-to-Born (p2B) in order to derive the first exclusive fiducial predictions for single jet inclusive production in charged-current DIS to N³LO. A comparison to data is performed, where we observe reasonable agreement with the experimental results from ZEUS.

List of Figures

1.1	The “Mexican Hat” Higgs potential for a single complex scalar ϕ . The orange line shows the degenerate set of states for which $\phi = \phi_{\min}$, and thus the potential is minimised.	12
1.2	Representative Feynman diagrams at LO, NLO and NNLO for boson production in quark-antiquark collisions.	20
1.3	The one quark-gluon vertex diagram, where the quark line couples to an off-shell external photon with invariant mass q . The incoming quarks are massless, with $p_1^2 = p_2^2 = 0$	21
1.4	The running of the strong coupling constant alongside a selection of experimental measurements used in the extraction of the world average value of $\alpha_s(M_Z)$. Image taken from [8].	25
1.5	A gluonic real emission from a quark line becoming unresolved in either the soft $ k \rightarrow 0$ or the collinear $k \parallel p$ limit.	27
1.6	The individual parton distributions (left) and valence and sea quark distributions (right) as a function of the proton momentum fraction x at $Q^2 = M_Z^2$ contained within the MMHT14 PDF set [38] and produced using the LHAPDF package [39].	30
1.7	Two interpretations corresponding to the ambiguity in PDF definition in the presence of unresolved initial state radiation ($k \parallel Pz$).	31

1.8	Summary of cross-section measurements made by the ATLAS collaboration at the LHC as of March 2019, separated by process. The data are compared to benchmark theory results at $\sqrt{s} = 5.02, 7, 8$ and 13 TeV. Image taken from [61].	35
2.1	An example 4-quark interference squared Feynman diagram present at $\mathcal{O}(\alpha\alpha_s^2)$ in W boson production through hadron-hadron collisions, which contains no double unresolved final state collinear limits. Quark momenta 1 and 2 are initial-state, with i and j final-state.	68
2.2	An example 4-quark interference squared Feynman diagram present at $\mathcal{O}(\alpha\alpha_s^2)$ in W boson production through hadron-hadron collisions, which contains double unresolved final state collinear limits. Quark momenta 1 and 2 are initial-state, with i and j final-state.	70
3.1	Born level Z/γ^* boson production through the Drell-Yan process.	76
3.2	Born level W^+ boson production through the Drell-Yan process.	77
3.3	Normalised p_T^W distribution for $W = W^+ + W^-$ production. The left hand panel shows predictions for the electron decay channel and the right hand panel shows the muon channel. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text.	83
3.4	Normalised p_T^Z distribution for Z production with subsequent decay into muons. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text.	84

- 3.5 Ratio of normalised p_T^V distributions for decays in the muon channel. The left hand panel shows the ratio between W^- and W^+ production, and the right hand panel shows the ratio between Z and $W = W^+ + W^-$ production. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text. 85
- 3.6 W^+ (left) and W^- (right) differential p_T^V distributions normalised to the inclusive cross-section at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. The lower panel shows the ratio of the results to NNLL+NLO. The uncertainty prescriptions for the theory uncertainty are as described in the text. 93
- 3.7 The Z differential p_T^V distributions normalised to the inclusive cross-section at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. The lower panel shows the ratio of the results to NNLL+NLO. The uncertainty prescriptions for the theory uncertainty are as described in the text. 94
- 3.8 W^+ (left) and W^- (right) differential p_T^V distributions normalised to the inclusive cross-section at NNLO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the `AZ` tune [218]. The lower panel shows the ratio of the results to the aforementioned `Pythia` results. The uncertainty prescriptions for the theory uncertainty are as described in the text. 95

- 3.9 The Z differential p_T^Y distributions normalised to the inclusive cross-section at NNLO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the `AZ` tune [218]. The lower panel shows the ratio of the results to the aforementioned Pythia results. The uncertainty prescriptions for the theory uncertainty are as described in the text. 96
- 3.10 Ratios of Z/W⁺ (left) and W⁻/W⁺ (right) normalised differential distributions at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volumes described in the text. The three lower panels show different prescriptions for the theory uncertainty, as described in the text and normalised to the NNLL+NLO results. 97
- 3.11 Ratios of Z/W⁺ (left) and W⁻/W⁺ (right) normalised differential distributions at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volumes described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the `AZ` tune [218]. The three lower panels show different prescriptions for the theory uncertainty, as described in the text and normalised to the aforementioned Pythia results. The lower panel shows the ratio of the 99
- 4.1 The kinematic regions of the (x_1, x_2) (left) and (x, Q^2) (right) planes accessible for the LHCb [219], ATLAS [220] ZJ and CMS [221] W[±]J selection criteria at LO ($\mathcal{O}(\alpha^2\alpha_s)$). Here x_1 and x_2 are the Bjorken- x values from beams 1 and 2 respectively, Q^2 is the invariant mass of the full final state including both charged leptons and QCD radiation and x is the Bjorken- x from either of the incoming beams. 104

- 4.2 Cross section differential in the leading jet transverse momentum for W^- , W^+ and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text. 109
- 4.3 Cross section differential in the leading jet pseudorapidity for W^- , W^+ and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text. 111
- 4.4 Cross section differential in the lepton pseudorapidity for W^- , W^+ , and the Z boson rapidity and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text. 112
- 4.5 The charge asymmetry A^\pm for W^\pm production as a function of lepton pseudorapidity. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the difference with respect to NLO is shown in the lower panel of the asymmetry plot, with the ratio to NLO is shown in the lower panel of the ratio plot. The bands correspond to scale uncertainties estimated as described in the main text. 114

- 4.6 The W^\pm charge asymmetry A^\pm in WJ final states differential in the pseudorapidity η of the decay muon, evaluated with the central member of the NNPDF3.1 (red), MMHT14 (yellow), CT14 (green) NNLO PDFs for the LHCb data of [219]. The NNPDF3.1 curve corresponds to a full NNLO calculation with scale uncertainties as described in the main text, and is used to determine a differential NNLO/NLO k -factor. The other two predictions are calculated at NLO and then rescaled by this k -factor. 115
- 4.7 Leading jet kinematic distributions for the LHCb data of [219]. The upper row shows plots for the leading jet pseudorapidity in W^- (left) and W^+ (right) production, and the lower row shows plots for the leading jet transverse momentum in W^- (left) and W^+ (right) production. Each of these distributions is evaluated with the central member of the NNPDF3.1 (red), MMHT14 (yellow), CT14 (green) NNLO PDFs. The NNPDF3.1 curve corresponds to a full NNLO calculation with scale uncertainties as described in the main text, and is used to determine a differential NNLO/NLO k -factor. The other two predictions are calculated at NLO and then rescaled by this k -factor. 116
- 4.8 WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with exactly one associated jet ($N = 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226, 228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. 119

- 4.9 WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226,228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. 120
- 4.10 Ratio of exclusive over inclusive ($N = 1/N \geq 1$) distributions for WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226,228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. Errors on the ATLAS data are approximated using uncertainties from the $N = 1$ distribution normalised to the $N \geq 1$ results. 122
- 4.11 WJ/ZJ ratio differential in the exclusive p_T of the leading jet ($N = 1$) (left) and inclusive p_T of the leading jet ($N \geq 1$) (right) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [220], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. 122

- 4.12 WJ (left) and ZJ (right) cross sections differential in the absolute rapidity $|y_j|$ of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226, 228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. 124
- 4.13 WJ/ZJ ratio differential in the absolute rapidity $|y_j|$ of the leading jet. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [220] in the ATLAS fiducial region from Eq. 4.2.1, and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. 125
- 4.14 Leading jet transverse momentum (left) and pseudorapidity (right) distributions for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text. 127
- 4.15 The differential cross section (left) and average number of jets $\langle N_{\text{jets}} \rangle$ (right) distributions as a function of H_T for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text. 129

- 4.16 Angular separation between leading jet and muon $\Delta\Phi(j_1, \mu)$ (left) and jet multiplicity (right) distributions for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text. 130
- 5.1 A comparison of the current world best measurements of $\sin^2 \theta_W^{\text{eff}}$. LEP/SLC results are taken from [242] for combined measurements taken on the Z pole, from b -quark forward-backward asymmetries and in the Z boson left-right asymmetry. CDF [243] and DØ [244] results are presented alongside their combination in [245], and LHC results from the ATLAS collaboration at 7 TeV [116] along with LHCb [115] and CMS [246] 7 and 8 TeV combinations are given. The horizontal bars represent the total uncertainties on each measurement, and the shaded vertical band give the uncertainty on the LEP/SLD measurement. 134
- 5.2 The definition of Collins-Soper frame and associated lepton decay angles θ and ϕ . p_1 and p_2 are the directions of the incoming partonic momenta in the lepton rest frame, k_1 is the negative lepton momentum and k_2 is the positive lepton. Image taken from [194]. . 137
- 5.3 Pure Z, γ^* and Z/ γ^* interference contributions to the total Drell-Yan cross section as a function of m_{ll} , calculated at LO in QCD. The upper panel shows the absolute values of the three contributions, and the lower panel shows the ratio of the Z/ γ^* interference term to the pure γ^* contribution. The black dotted vertical line at $m_{ll} \sim M_Z$ marks the change in sign of the interference term. 138

-
- 5.4 A_{FB} as a function of the di-lepton invariant mass m_{ll} at LO in QCD.
The dotted vertical red line denotes the position of the Z peak, at
which point the Z/γ^* interference term is identically zero. 140
- 5.5 Phase space constraints in the $(|y_{ll}|, \cos \theta^*)$ plane at Born level (LO)
for the central-central Z3D Drell-Yan fiducial region. Overlaid are
the measurement bins, integrated over m_{ll} 146
- 5.6 Bin classifications at LO for the central-central Z3D Drell-Yan fiducial
region in the $(|y_{ll}|, \cos \theta^*)$ plane. Overlaid are the measurement bins,
integrated over m_{ll} 147
- 5.7 Bin classifications at LO for the central-forward Z3D Drell-Yan fidu-
cial region in the $(|y_{ll}|, \cos \theta^*)$ plane. Overlaid are the measurement
bins, integrated over m_{ll} 149
- 5.8 Minimum Q_T values required for the different regions of the $(|y_{ll}|,$
 $\cos \theta^*)$ plane in the central-central region. Overlaid are the Z3D meas-
urement bins, integrated over m_{ll} 150
- 5.9 The upper panel shows the minimum values of y_1 and maximum val-
ues of $\Delta y_{ll}^{\text{max}}$ as a function of the lepton transverse momentum ratio r .
The lower panel shows the resultant minimum transverse momentum
induced by the momentum ratio r . For this we consider the forbidden
phase-space point at $y_{ll} = 2$, $\cos \theta^* = 0.4$ in the Z3D central-central
fiducial region, and take $y_2 = y_{\text{max}}^2 = 2.4$ in order to maximise the
rapidity difference. 152

- 5.10 The normalised Q_T spectrum for the $0.4 < \cos \theta^* < 0.7$, $46 < m_{ll} < 200$ region for each rapidity bin of the Z3D central-central measurement region between $y_{ll} = 1.2$ and $y_{ll} = 2.4$. The results are produced to $\mathcal{O}(\alpha_s)$, with the colours labelling as before the allowed (green), mixed (orange) and forbidden (red) bins. The kink observed at $1/\sigma \cdot d\sigma/dQ_T = 10^{-4}$ is a consequence of the linearisation of the axes between $Q_T \pm 10^{-4}$ to allow the negative contribution at $Q_T = 0$ to be shown. 154
- 5.11 Minimum Q_T values required for the different regions of the $(|y_{ll}|, \cos \theta^*)$ plane in the central-forward region. Overlaid are the Z3D measurement bins, integrated over m_{ll} 158
- 5.12 Acceptances for the CC Z3D fiducial region. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos \theta^*$ values are denoted by the central colour of each point, and the exterior colour and shape label the bin as allowed, partially allowed or forbidden. . . 159
- 5.13 NNLO/NLO ($\mathcal{O}(\alpha_S^2)/\mathcal{O}(\alpha_S)$) k -factors for the CC Z3D fiducial region. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos \theta^*$ values are denoted by the central colour of each point, and the exterior colour and shape label the bin as allowed, partially allowed or forbidden. The majority of k -factors for the forbidden region are outside the bounds of the plot. 162

- 5.14 The ratio of the NNLO QCD predictions to ATLAS data in the central-central region of the Z3D analysis. The upper plot shows the $\mathcal{O}(\alpha_s^2)$ theory predictions, and the lower plot includes the $\mathcal{O}(\alpha_s^3)$ contributions for the forbidden bins. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right). Luminosity uncertainties of $\sim 1.8\%$ are not included. 163
- 5.15 EW (NLO+partial higher order)/LO k -factors for the CC Z3D fiducial region as produced for $M_W = 79.939$ GeV using Powheg [262–265]. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos\theta^*$ values are denoted by the colour of each point. . 165
- 5.16 EW (NLO+partial higher order)/LO k -factors for the CF Z3D fiducial region as produced for $M_W = 79.939$ GeV using Powheg [262–265]. The bin number is as defined in (5.5.1), with appropriate modifications for the CF region, such that the major m_{ll} bins are divided into 5 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos\theta^*$ values are denoted by the colour of each point. 165

- 5.17 The ratio of the NNLO QCD predictions to ATLAS data in the central-forward region of the Z3D analysis. The upper plot shows the $\mathcal{O}(\alpha_s^2)$ theory predictions, and the lower plot includes the $\mathcal{O}(\alpha_s^3)$ contributions for the forbidden bins. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.3), such that the major m_{ll} bins are divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 6 y_{ll} sub-bins from 0 to 2.4 (left to right). Luminosity uncertainties of $\sim 1.8\%$ are not included. 167
- 5.18 The ratio of the central predictions of the MMHT14 (top panel) and CT14 (bottom panel) PDF sets to the predictions of NNPDF3.1 data in the central-central region of the Z3D analysis. Both plots show the predictions in all bins to $\mathcal{O}(\alpha_s^2)$. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right) which are in turn divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right). . . . 168
- 5.19 The ratio of MMHT14 and CT14 PDF sets to the central member of the NNPDF3.1 set as a function of Bjorken- x . From the top left panel clockwise, the panels correspond to the d valence, u valence, u sea quark and d sea quark contributions to the various PDF sets at $Q^2 = M_Z^2$. The uncertainties given are the PDF uncertainties. 170
- 5.20 The kinematic regions in the (x, Q^2) plane associated with the central-central (yellow) and central-forward (green) fiducial selections of the Z3D measurement. The total kinematic reach of the ATLAS detector at $\sqrt{s} = 8$ TeV is shown in red. 171

- 5.21 The ratio of the $\mathcal{O}(\alpha_s^2)$ predictions for the Z3D measurement made at $M_W = 79.864$ GeV and $M_W = 80.014$ GeV, corresponding to $\sin^2 \theta_W^{\text{eff}} = 0.23294$ and $\sin^2 \theta_W^{\text{eff}} = 0.23005$ for the CC (top) and CF (bottom) measurement regions. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty. The CC bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right). The CF bin number is defined analogously. 174
- 5.22 The “step effect” in the ratio of theory to data across the bin edge at $m_{ll} = 91$ GeV in the CC Z3D fiducial region, using the on-shell mass with a fixed propagator width. 175
- 5.23 LO predictions for A_4 as a function of m_{ll} in the running and fixed width schemes, in both cases using OS input values. The different plots show the same data with a different binning in m_{ll} 179
- 5.24 (NNLO + partial $\mathcal{O}(\alpha_s^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-central region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos \theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass m_{ll} 182

- 5.25 (NNLO + partial $\mathcal{O}(\alpha_S^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-central region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos \theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass m_{ll} 183

- 5.26 (NNLO + partial $\mathcal{O}(\alpha_S^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-forward region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos \theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass m_{ll} 184

- 5.27 The ratio of NNLO + partial $\mathcal{O}(\alpha_S^3)$ QCD predictions reweighted with NLO + partial higher order EW k -factors to ATLAS data in the central-central (upper) and central-forward (lower) region of the Z3D analysis. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right) for the central-central case. The CF binning is analogously defined. Luminosity uncertainties of $\sim 1.8\%$ are not included. 185
- 6.1 Born level kinematics for a single jet leptonic CC DIS event. . . . 190
- 6.2 Predictions at LO (green right-hatched), NLO (orange left-hatched), and NNLO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_{12} , E_{12}^T and M_{12} distributions for inclusive di-jet production in $e^- - P$ collisions. The bands correspond to scale uncertainties as described in the main text. 192
- 6.3 Predictions at LO (green right-hatched), NLO (orange left-hatched), and NNLO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_{12} , E_{12}^T and M_{12} distributions for inclusive di-jet production in $e^+ - P$ collisions. The bands correspond to scale uncertainties as described in the main text. 193

- 6.4 Schematic diagrams of the separate contributions to inclusive single jet production in CC DIS at N³LO using the method of p2B . The green points correspond to multiplicities at which the contribution fully describes the event kinematics. The orange points give the integrated out contributions, which act as a counterterm for the di-jet calculation in the divergent phase space regions where it approaches Born kinematics, and are naturally present in the inclusive structure function calculation. The red points correspond to the regions with no access to the associated event kinematics. 196
- 6.5 Comparison of the $\mathcal{O}(\alpha_S^2)$ term in single-jet inclusive e^+p CC DIS calculated via antenna subtraction and p2B differential in Q^2 , η_j , E_j^T , and x . The shaded regions correspond to the 7 point scale variation band calculated about central scale $\mu_F = \mu_R = Q$, and the vertical error bands correspond to the statistical error. The upper panel of each plot is the absolute correction to the cross section from the $\mathcal{O}(\alpha_S^2)$ terms, and the lower panel is the ratio to the central scale of the antenna calculation. 198
- 6.6 Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^-p collisions. 200
- 6.7 Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^+p collisions. 201

- 6.8 Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^-p collisions. The upper panel shows the differential cross sections for $P = -0.27$ and $P = 0.30$, and the lower panels show the individual ratios of data and experiment to the NNLO result. 203
- 6.9 Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^+p collisions. The upper panel shows the differential cross sections for $P = -0.37$ and $P = 0.32$, and the lower panels show the individual ratios of data and experiment to the NNLO result. 205
- 6.10 Cross sections as a function of polarisation in single jet leptonic CC DIS at N³LO in QCD. Uncertainty bands from scale variation, albeit small, are included as a shaded band around the central prediction. Results are compared to ZEUS data from Ref. [279]. 207
- A.1 An example spikeplot for a sub-leading colour $qg \rightarrow gg\bar{q} W^+$ double real matrix element in a limit where one final state gluon goes collinear with the incoming quark and another goes soft. As the x parameter approaches 0, the phase space points are driven further into the divergent region, and the subtraction term more closely matches the divergence of the matrix element. 214
- A.2 Cancellations of separate terms between different phase space multiplicities in the antenna subtraction method at NLO and NNLO in QCD. 217

- A.3 A comparison of the scale variation of the $\mathcal{O}(\alpha_s^2)$ term in single-jet inclusive CC DIS between NNLOJET and the result obtained analytically through running of the renormalisation scale μ_R using the renormalisation group equations. We take $\mu_R = 80\text{GeV}$ as a reference scale, and vary in the range $20 - 400\text{GeV}$ 220
- A.4 A summary of the NNLOJET regression tests between December 2017 and April 2019. The upper panel shows the test runtime excluding compilation time, and the lower panel shows the total number of tests alongside the number of passing tests as a function of time. The tests are run on a multi-use desktop containing a 4 core, 8 thread Intel i7-4790 CPU with clock speed of 3.60GHz, meaning that large fluctuations in the test runtimes are generally caused by external (non-regression test) load on the system. 222
- B.1 Representative Feynman diagram of the fl_{02} flavour class, which is forbidden for $V = W^\pm$ due to charge conservation in the quark loop. 228

List of Tables

1.1	Particle content of the standard model.	2
1.2	Tabulated quantum numbers of the fermions in GWS theory. . . .	9
1.3	Current world average values for the elements of the CKM matrix, using the parameterisation for V_{CKM} used in (1.2.32). Data taken from [8].	17
2.1	The classes of X_3^0 and X_3^1 antenna functions, broken down by hard ra- diator and radiative process. For the loop antenna functions, the tilde variants correspond to sub-leading colour functions and the hatted variants correspond to antennas containing a closed quark loop [48].	50
2.2	The different classes of unintegrated X_0^4 antenna functions, broken down by hard radiator and radiative process. The tilde variants cor- respond to sub-leading colour contributions [48].	52
2.3	Tabulated colour-connected singular limits of the amplitudes contrib- uting to $D(1_u, 2_u, i_d, j_u)$. This corresponds to the matrix element crossing shown in Fig. 2.1.	69
2.4	Tabulated colour-connected singular limits of the amplitudes contrib- uting to $D(1_u, i_u, 2_d, j_u)$. This corresponds to the matrix element crossing shown in Fig. 2.2.	71

4.1	Fiducial cross sections in picobarns for fixed-order theory predictions and LHCb results from Ref. [219]. The errors quoted for NNLO-JET correspond to the scale uncertainty and the reported LHCb errors are statistical, systematic and luminosity respectively. . . .	108
5.1	Selection criteria for the central-central and central-forward fiducial regions in the ATLAS measurement of [260].	143
5.2	Binnings for the central-central and central-forward fiducial regions in the ATLAS measurement of [260].	143
5.3	Maximum/minimum values of y_{ll} permitted in different regions of phase space for the CF Z3D selection.	148
5.4	Tabulated A_{FB} values for the $80 < m_{ll} < 100$, $1.6 < y_{ll} < 2.5$ bin of [261], performed using different EW input parameter schemes. .	178
6.1	Comparison of inclusive cross sections calculated independently through antenna subtraction and structure functions for the CC DIS e^+p $\mathcal{O}(\alpha_s^2)$ term, broken down by incoming parton channel.	197
6.2	Total cross sections at each perturbative order compared to ZEUS data from Ref. [279] for each polarisations in the experimental data. The uncertainties quoted for the theoretical predictions are from scale variations only, and statistical uncertainties are negligible.	204
B.1	C_2 coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case.	227
B.2	C_L coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case.	228

- B.3 C_3 coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case. Here, $A = 2C_F(1 + x)$, with colour factor C_F normalised such that $C_F = (N^2 - 1)/2N$ 228

Acknowledgements

First and foremost, I would like to thank my supervisor Nigel Glover for his insight over the last four years. It has been a privilege, and his support, encouragement, and knowledge on matters physics and otherwise have been immensely valuable.

It has been an absolute pleasure to be a part of the top floor NNLOJET crew alongside Juan, James and Jan, working out optimal Grid strategies with PYHEP-GRID, monitoring commits amongst so many other miscellaneous schemes.

I have also been very fortunate to have worked with such a variety of exceptional people, past and present, within the rest of the collaboration: Xuan Chen, James Currie, Rhorry Gauld, Aude Gehrmann-De Ridder, Thomas Gehrmann, Marius Höfer, Alex Huss, Imre Majer, Jonathan Mo, Tom Morgan, and Joao Pires. It's been a lot of fun, and I've greatly appreciated the many chats about QCD, computing and beyond.

I am also very grateful to my experimental collaborators from the Z3D ATLAS team, and in particular Sasha, Simone, and Eram for the guidance and encouragement. Getting to give a talk at CERN is something I'll never forget!

Special thanks must be extended those who have facilitated my (extensive) consumption of computing resources, namely Jeppe, Adam, Paul, George, and Ollie. I also have to thank Joanne, Linda and Trudy, the unsung heroes of the IPPP office, for making sure everything ran smoothly and always knowing what to do.

I am very grateful to all those who have made my time at Durham so enjoyable, in particular Stephen, Marian, and Alan for the countless pub quiz failures, David, Jamie, and Ollie for sharing the Python drop-in duties, and Robin for the much-needed perspective.

Last, but by no means least, I must thank Laura for her unwavering and unmitigated support and my parents, brother and family for their belief and for making everything possible. It is greatly appreciated.

This research was supported by the UK Science and Technology Facilities Council and by the Research Executive Agency (REA) of the European Union through the ERC Advanced Grant MC@NNLO (340983).

Declaration

The work in this thesis is based on research carried out in the Department of Physics at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification. All work is my own in collaboration with my supervisor Professor Nigel Glover unless referenced as otherwise within the main text, with the exception of Chapter 6.

Chapter 3 is based upon work done in collaboration with Aude Gehrmann-De Ridder, Thomas Gehrmann, Alexander Huss, Pier Francesco Monni, Emanuele Re, and Luca Rottoli and presented in the following publications:

- A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *Next-to-Next-to-Leading-Order QCD Corrections to the Transverse Momentum Distribution of Weak Gauge Bosons*, *Phys. Rev. Lett.* **120** (2018) 122001, [1712.07543]
- W. Bizon, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, P. F. Monni et al., *The transverse momentum spectrum of weak gauge bosons at $N^3LL+NNLO$* , [1905.05171]

Chapter 4 is based upon work done in collaboration with Aude Gehrmann-De Ridder, Thomas Gehrmann, Alexander Huss and presented in the following publications and conference proceedings:

- A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *NNLO QCD Corrections to W +jet Production in NNLOJET*, *PoS LL2018* (2018) 041, [1807.09113]

- A. Gehrmann–De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *Vector Boson Production in Association with a Jet at Forward Rapidities*, *Eur. Phys. J.* **C79** (2019) 526, [1901.11041]

Chapter 5 is based upon work done in collaboration with Tony Kwan, Pavel Shvydkin, Simone Amoroso, Elzbieta Richter-Was, Stefano Camarda, Andrey Saponov, Sasha Glazov, Richard Keeler and Eram Rizvi as an Analysis Consultant Expert (ACE) to the ATLAS collaboration. A number of the NNLOJET results contained in this thesis are presented in the conference note

- ATLAS collaboration, *Measurement of the effective leptonic weak mixing angle using electron and muon pairs from Z-boson decay in the ATLAS experiment at $\sqrt{s} = 8$ TeV*, [ATLAS-CONF-2018-037]

although the author of this thesis is not an author of the note. A joint publication is anticipated in the future.

Chapter 6 is based upon work done in collaboration with Thomas Gehrmann, Alexander Huss, Jan Niehues, and Andreas Vogt and presented in the following publications:

- J. Niehues and D. M. Walker, *NNLO QCD Corrections to Jet Production in Charged Current Deep Inelastic Scattering*, *Phys. Lett.* **B788** (2019) 243–248, [1807.02529]
- T. Gehrmann, A. Huss, J. Niehues, A. Vogt and D. M. Walker, *Jet production in charged-current deep-inelastic scattering to third order in QCD*, *Phys. Lett.* **B792** (2019) 182–186, [1812.06104]

Copyright © 2019 Duncan Martin Walker.

The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.

Contents

Abstract	i
List of Figures	iii
List of Tables	xxiii
1 Introduction to the Standard Model	1
1.1 Lagrangian Mechanics	2
1.2 Electroweak Theory	8
1.2.1 Fermions	8
1.2.2 Gauge Bosons	10
1.2.3 The Higgs Mechanism	10
1.2.4 Properties of the Higgs Boson	15
1.2.5 Quark Mixing and the CKM Matrix	16
1.3 Quantum Chromodynamics (QCD)	17
1.3.1 QCD Cross Sections	18
1.3.2 Divergences in QCD	20
1.3.3 UV Divergences and Renormalisation	22
1.3.4 IR Divergences and Cancellation - a Heuristic Overview	26

1.3.5	Slicing Methods	31
1.3.6	Subtraction Methods	32
1.4	High Precision QCD Observables at the LHC	34
2	Antenna Subtraction	39
2.1	Colour Ordering	39
2.2	Behaviour of Real Emission Contributions in Singular Limits	42
2.3	Antenna Functions	49
2.3.1	Unintegrated Antenna Functions	50
2.3.2	Integrated Antenna Functions	53
2.4	Subtraction Term Construction	54
2.4.1	Mass Factorisation Terms	54
2.4.2	NLO Subtraction	57
2.4.3	NNLO Subtraction	58
2.5	Antenna Subtraction for Charged-Current processes	66
2.6	Implementation of W^\pm processes in NNLOJET	71
3	Phenomenology of Neutral- and Charged-Current Electroweak Gauge Boson Production at the LHC	75
3.1	Overview	75
3.2	Transverse Momentum Distributions of Vector Bosons in the Drell-Yan Process	80
3.3	Fixed Order Predictions	81
3.4	Fixed Order Matched to N^3LL p_T^V Resummation	87
3.5	Summary	100

4 Phenomenology of Vector Boson Production in Association with a Jet	103
4.1 Forward Production of Vector Bosons at LHCb	106
4.2 Ratios of VJ Production at ATLAS	117
4.3 W+J Production at CMS	126
4.4 Summary	131
5 Determination of $\sin^2 \theta_W^{\text{eff}}$ from Triple Differential Inclusive Z Boson Production	133
5.1 The Effective Weak Mixing Angle, $\sin^2 \theta_W^{\text{eff}}$	134
5.2 Measurements of $\sin^2 \theta_W^{\text{eff}}$ in Drell-Yan Processes	136
5.3 ATLAS Drell-Yan Triple Differential (Z3D) Measurement	142
5.4 Kinematics of the Z3D Measurement	144
5.4.1 Born Level ($\mathcal{O}(\alpha^2 \alpha_s^0)$) Kinematics	145
5.4.2 Constraints Beyond Born Level	149
5.5 Acceptances and k -factors	158
5.6 PDF Variation	166
5.7 $\sin^2 \theta_W^{\text{eff}}$ Variation	171
5.8 Running Width Corrections	173
5.8.1 On-Shell (OS) scheme	175
5.8.2 Pole Mass (PM) scheme	177
5.9 Combined EW+QCD Predictions using Powheg and NNLOJET	180
5.10 Summary	183

6	Jet Production in Charged-Current Deep Inelastic Scattering	187
6.1	DIS Kinematics	190
6.2	Di-Jet Production in CC DIS at NNLO	191
6.3	Projection to Born Subtraction	194
6.3.1	Implementation in the NNLOJET Framework	197
6.4	Single Jet Production in CC DIS at N ³ LO	199
6.5	Summary	206
7	Conclusions	209
	Appendices	211
A	Validation and Testing Procedures in NNLOJET	211
A.1	Point-wise Matrix Element Testing	212
A.2	Subtraction Term Tests	213
A.2.1	Unresolved Limits Of Real Matrix Elements	213
A.2.2	Pole Cancellation Of Virtual Matrix Elements	215
A.2.3	Finite Term Consistency Checks	216
A.3	Integrated Cross Section Tests	218
A.3.1	Technical Cut Dependence	218
A.3.2	Renormalisation Scale Variation	219
A.3.3	3 rd Party Validation	219
A.4	Regression Tests	221

B DIS Structure Functions	225
B.1 Inclusive Cross Section Decomposition	225
B.2 Scale Variation	229
B.2.1 Renormalisation Scale Dependence	229
B.2.2 Factorisation Scale Dependence	231
B.3 Numerical Implementation	234
Bibliography	237

Chapter 1

Introduction to the Standard Model

The development of the standard model (SM) of particle physics is one of the most important achievements in fundamental physics over the last century. It has been remarkably successful in describing the phenomena of subatomic processes, and has remained robust under a vast amount of scrutiny. It is not, however, a complete description of our observed reality, absent descriptions of dark matter, neutrino masses and gravity being notable examples. Attempts to extend the SM to incorporate these known deficiencies, such as supersymmetry, have so far met with a surprising lack of experimental evidence. In many ways it is elegant in its simplicity, with the imposition of symmetries on a generic Lagrangian being the defining characteristic. The most famous example of this elegance was the prediction of a Higgs boson in the early 1960s, the first observation of which was made at the Large Hadron Collider (LHC) after an extensive experimental search almost 50 years later.

There is a rich particle content within the SM, as summarised in Table. 1.1, with the fermion content divisible into two sectors each containing three distinct generations. The fermion content comprises the quark sector, which interacts both through the electroweak and strong forces, and the electroweakly interacting lepton sector, which contains the electron, muon and tau particles alongside their associated neutrinos.

Quarks	u	c	t	Gauge Bosons	W^\pm	γ
	d	s	b		Z	g
Leptons	e^-	μ^-	τ^-	Scalar Bosons	H	
	ν_e	ν_μ	ν_τ			

Table 1.1: Particle content of the standard model.

The gauge bosons mediate the interactions between the fermions and are naturally generated as a consequence of the standard model local gauge group symmetry, $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$, which is then broken through the introduction of the Higgs field, all within the framework of quantum field theory.

In this chapter we will describe the framework of the standard model, starting with the general properties of the quantum field theory (QFT) Lagrangian in Section 1.1 after which we introduce the electroweak Lagrangian with spontaneous symmetry breaking in Section 1.2. The strong sector and the nature of calculations to higher orders in QCD is discussed in Section 1.3. These calculations require the regularisation of multiple types of divergences, and we will outline the various techniques required to produce predictions for observables in high energy scattering processes in Section 1.4.

1.1 Lagrangian Mechanics

Any local quantum field theory, including that of the SM, can be defined through the Lagrangian density¹ \mathcal{L} , which is a generalised function of the associated fields $\{\phi_i\}$ and their space-time derivatives $\{\partial_\mu \phi_i\}$. By minimising the action S ,

$$S[\{\phi_i\}] = i \int d^4x \mathcal{L}(\{\phi_i\}, \{\partial_\mu \phi_i\}) \quad (1.1.1)$$

using functional integration, one can derive the Euler-Lagrange equations of motion for the fields:

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi_i)} \right) - \frac{\partial \mathcal{L}}{\partial \phi_i} = 0. \quad (1.1.2)$$

¹For brevity, we will refer to the Lagrangian density as the Lagrangian.

In order for a Lagrangian to provide a physical description of our reality, it must be invariant under transformations belonging to the Poincaré group, containing rotations, translations, and Lorentz boosts, such that results derived are frame independent and obey the principles of relativity. As this defines an invariance of the Lagrangian, from Noether's theorem there are associated conserved quantities, namely angular momentum and four-momentum.

\mathcal{L} must also be Hermitian, in order to ensure that the action is real, as well as being renormalisable. Renormalisation provides a means of absorbing ultra-violet (UV) divergences into the definition of physical constants to render predictions finite, and will be discussed in more detail in Section 1.3.3. This is only possible if all coupling constants have mass dimension ≥ 0 , which is the relevant point for the current discussion as it considerably restricts the form of possible terms in \mathcal{L} .

Beyond these basic conditions, one can also further specify a given theory through the imposition of symmetries on \mathcal{L} . By requiring that the Lagrangian remain invariant under symmetric transformations, one can greatly constrain the possible terms in \mathcal{L} . To elucidate this point, we consider the application of a general local $SU(N)$ symmetry on the Dirac Lagrangian for a free spin- $\frac{1}{2}$ field ψ :

$$\mathcal{L} = \bar{\psi}(i\not{\partial} - m)\psi. \quad (1.1.3)$$

Here we use Dirac slash notation, where for covariant vector A_μ we define $\not{A} = \gamma^\mu A_\mu$, with γ^μ the Dirac matrices which form a matrix representation of the Clifford algebra². ψ is a Dirac spinor, which in the chiral (Weyl) representation can be separated into left- and right-handed components

$$\psi = \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix}. \quad (1.1.4)$$

In our interpretation, we consider ψ as a fermion field, $\bar{\psi} = \psi^\dagger \gamma^0$ and m as the fermion mass.

²The Clifford algebra is defined through the anti-commutation relation $\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}$ where $\eta^{\mu\nu}$ is the Minkowski metric.

If we require a local $SU(N)$ symmetry with generators T^a , under which the fermion and anti-fermion fields transform in the fundamental and anti-fundamental representations respectively (denoting the fundamental $SU(N)$ indices as i, j),

$$\begin{aligned}\psi_i(x) &\rightarrow U_{ij}(x)\psi_j(x) \\ \bar{\psi}_i(x) &\rightarrow \bar{\psi}_j(x)U_{ij}^{-1}(x) \\ U &= e^{i\alpha_a(x)T_{ij}^a},\end{aligned}\tag{1.1.5}$$

it becomes apparent that (1.1.3) is not invariant under the transformation:

$$\begin{aligned}\mathcal{L} \rightarrow \mathcal{L}' &= \bar{\psi}_j(x)U_{ij}^{-1}(x)(\not{\partial} - m)U_{ik}(x)\psi_k(x) \\ &= \bar{\psi}_j(x)(\not{\partial} - m)\psi_k(x) + \bar{\psi}_j(x)\psi_k(x)U_{ij}^{-1}(x)\not{\partial}(U_{ik}(x)) \\ &\neq \mathcal{L}.\end{aligned}\tag{1.1.6}$$

This can be remedied by promoting the standard derivative $\not{\partial}$ to the covariant derivative \not{D} , defined by its transformation as (dropping the $SU(N)$ indices)

$$\not{D}\psi(x) \rightarrow U(x)\not{D}\psi(x)\tag{1.1.7}$$

such that the term $\bar{\psi}(x)\not{D}\psi(x)$ is gauge invariant. To achieve this, one can introduce a gauge field $A_\mu^a(x)$, $a \in [1, N^2 - 1]$ in the adjoint representation of $SU(N)$, alongside coupling strength g in order to define \not{D} :

$$\not{D} \equiv \gamma^\mu(\partial_\mu - igA_\mu^a T^a),\tag{1.1.8}$$

where $A_\mu^a(x)$ transforms as

$$A_\mu^a(x)T^a \rightarrow U(x)A_\mu^a(x)T^a U^{-1}(x) + \frac{i}{g}U^{-1}(x)\partial_\mu U(x).\tag{1.1.9}$$

Under contraction with $\bar{\psi}\psi$, the second term of (1.1.9) exactly cancels the second, gauge dependent term in the final line of (1.1.6), leaving the new Lagrangian,

$$\mathcal{L} = \bar{\psi}(i\not{D} - m)\psi,\tag{1.1.10}$$

gauge invariant.

Given that a new field $A_\mu^a(x)$ has been introduced, it makes sense to interpret it as particle content in our Lagrangian, and as it transforms as a vector under Lorentz transformations, it can be viewed as bosonic in nature. However, for consistency purposes one must also introduce a gauge invariant kinetic term for A , as otherwise the Euler-Lagrange equations of motion for A require $\bar{\psi}\psi = 0$.

This kinetic term can be achieved through definition of the field strength tensor $F_{\mu\nu}^a$ as

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - gf^{abc}A_{\mu,b}A_{\nu,c}, \quad (1.1.11)$$

where f^{abc} are the structure constants of the $SU(N)$ theory. The structure constants are present for non-Abelian theories where the generators do not commute, and are indeed defined through the commutation relation

$$[T^b, T^c] = if^{abc}T_a. \quad (1.1.12)$$

A kinetic term for the Lagrangian can then be constructed as

$$-\frac{1}{4}F_{\mu\nu}^a F_a^{\mu\nu}, \quad (1.1.13)$$

which gives rise to the dynamics of the gauge field.

It is the terms proportional to f^{abc} contained within the gauge kinetic term that give rise to one of the defining characteristics of non-Abelian gauge theories, *gauge boson self-interactions*. In the case of the SM Lagrangian, one has two distinct sectors, the electroweak (EW) $SU(2)_L \otimes U(1)_Y$ and the QCD $SU(3)_C$, where these self-interactions are observed. While we will consider these sectors in more detail in forthcoming sections, it is instructive to point out their relationship to the present discussion.

In EW theory, one can identify the gauge fields (after mixing) as the W and Z bosons alongside the photon, and the fermion field ψ as either quarks or leptons. The gauge boson self-interactions correspond to the interactions between the photon and the W and Z bosons. For the case of QCD, the quarks are the only fermions carrying

a colour charge, so they are naturally associated with the field ψ . The gauge fields of the theory are identified as gluons, which indeed strongly self interact. In the simplest case of an Abelian $U(1)$ theory, QED, $f^{abc} = 0$ as the generators of the theory commute and the gauge fields, photons, do not interact with one another. This is as one would naïvely expect as photons only directly couple to fields charged under $U(1)_{\text{EM}}$, whilst themselves remaining charge neutral.

Mass terms for the gauge fields of the form

$$\frac{1}{2}m^2 A_\mu^a A_a^\mu, \quad (1.1.14)$$

are forbidden as they violate gauge invariance of the Lagrangian. This does not pose an issue for the cases of QED and QCD, where the gauge bosons (photons and gluons respectively) are observed to be massless, but conflicts with experimental observations of the W and Z bosons with masses of 80.379 ± 0.012 GeV and 91.1876 ± 0.0021 GeV respectively [8]. The procedure through which their mass is derived is known as the Higgs mechanism, and will be discussed in Section 1.2.3. We will also see that fermion mass terms of the form we have seen so far are not compatible with $SU(2)_L \otimes U(1)_Y$ symmetry. In a demonstration of the elegance of the Higgs mechanism, it will be shown that these can also be generated through the introduction of a scalar Higgs field.

At this point, it may seem that we have a Lagrangian from which we can begin to directly calculate physical results using perturbation theory. However, were one to use directly the subsequent action in a path integral of the form

$$\int \mathcal{D}A e^{iS[A_\mu, \{\psi_i\}, \{\partial_\mu \psi_i\}]}, \quad (1.1.15)$$

it would be found to diverge due to degeneracies caused by gauge symmetries. These symmetries generate an infinite number of physically equivalent states related by gauge transformations, which must be removed by the Fadeev-Popov procedure. Whilst we refrain from a full treatment of this procedure here, we will briefly summarise the important points (dropping the dependence of the action on ψ_i , $\partial_\mu \psi_i$ for

ease).

One must first fix the gauge of the action using some arbitrary, general condition on the gauge field,

$$G[A_\mu(x)] = 0. \quad (1.1.16)$$

If we consider the set of fields A_μ^α related by a generic gauge transformation parameterised by $\alpha(x)$, one can isolate the gauge chosen in (1.1.16) using a delta function $\delta(G[A_\mu^\alpha(x)])$. Including the appropriate Jacobian, this can be used to construct the following identity

$$1 = \int \mathcal{D}\alpha(x) \delta(G[A_\mu^\alpha(x)]) \det \left(\frac{\delta G[A_\mu^\alpha(x)]}{\delta \alpha} \right) \quad (1.1.17)$$

to be inserted directly into the path integral. One can then perform a change of variables from A to A^α and utilise the gauge invariance of the action ($S[A^\alpha] = S[A]$) to isolate the divergent integral over $\alpha(x)$:

$$\int \mathcal{D}A e^{iS[A]} = \int \mathcal{D}\alpha(x) \int \mathcal{D}A e^{iS[A]} \delta(G[A_\mu^\alpha(x)]) \det \left(\frac{\delta G[A_\mu^\alpha(x)]}{\delta \alpha} \right). \quad (1.1.18)$$

In order to factor out the dependence on A in the functional determinant, one can write it as a functional integral over Grassmannian scalar fields, known as Fadeev-Popov ghosts. Written in this form, one can choose a gauge before constructing the appropriate correlation function, in which the divergent part of the path integral cancels between numerator and denominator at the expense of a new term in the integrand of (1.1.15). This term can be expressed in the form $\exp(i \int d^4x G)$, and thus can be interpreted as an unphysical gauge-fixing term in the Lagrangian. The field content of this term includes Fadeev-Popov ghosts, scalar fields which anti-commute and therefore violate the spin-statistics theorem.

All dependence on these terms will cancel in calculations of physical observables, as long as the calculation is performed consistently within the same gauge. One can even choose the gauge such that no such ghosts arise; these choices are known collectively as axial gauges. In practice the gauge is usually determined by calculational convenience, with common choices being the Feynman and Landau gauges.

1.2 Electroweak Theory

At this point, we now have the machinery available to consider the EW sector of the SM, governed by the theory introduced by Glashow, Weinberg and Salam (GWS) [9–12]. In GWS theory, the EW sector is governed by the gauge symmetry $SU(2)_L \otimes U(1)_Y$, which is subsequently broken through the Higgs mechanism to leave a residual $U(1)_{EM}$ symmetry.

1.2.1 Fermions

We start by considering the fermion content of the theory. The first step is to define weak isospin doublets for each generation of left-handed leptons and quarks, alongside the corresponding right-handed $SU(2)$ singlets:

$$L_L = \begin{pmatrix} \nu_L \\ l_L \end{pmatrix}, \quad Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \quad (l_R), \quad (u_R), \quad (d_R). \quad (1.2.1)$$

Here u , d , l and ν correspond to up-type quarks, down-type quarks, leptons and lepton neutrinos respectively, and we exclude right-handed neutrinos ν_R as only the left-handed counterpart has been directly observed.

From this content, one can construct the term in the Lagrangian for massless EW fermions alongside the appropriate covariant derivative:

$$\mathcal{L}_{\text{fermion}} = \sum_{i=1}^{N_{\text{gen}}} \left[\bar{Q}_L \not{D} Q_L + \bar{u}_R \not{D} u_R + \bar{d}_R \not{D} d_R + \bar{L}_L \not{D} L_L + \bar{l}_R \not{D} l_R \right], \quad (1.2.2)$$

$$\not{D} = \gamma^\mu \left(\partial_\mu + ig_1 \frac{Y}{2} B_\mu + ig_2 T_W \sigma^i W_\mu^i \right), \quad (1.2.3)$$

where the sum over i corresponds to each of the three generations of fermions, g_1 and g_2 are the coupling constants to the $U(1)_Y$ B_μ and $SU(2)_L$ W_μ^i fields respectively, Y and T_W are the hypercharge and weak isospin operators with eigenvalues given in Table 1.2 and σ^i are the $(N^2 - 1) = 3$ generators of $SU(2)$ in the adjoint representations, the Pauli matrices.

	T_W	T_{W_3}	Y	Q
ν_L	$+\frac{1}{2}$	$+\frac{1}{2}$	-1	0
l_L	$+\frac{1}{2}$	$-\frac{1}{2}$	-1	-1
l_R	0	0	-2	-1
u_L	$+\frac{1}{2}$	$+\frac{1}{2}$	$+\frac{1}{3}$	$+\frac{2}{3}$
d_L	$+\frac{1}{2}$	$-\frac{1}{2}$	$+\frac{1}{3}$	$-\frac{1}{3}$
u_R	0	0	$+\frac{4}{3}$	$+\frac{2}{3}$
d_R	0	0	$-\frac{2}{3}$	$-\frac{1}{3}$

Table 1.2: Tabulated quantum numbers of the fermions in GWS theory.

Each fermion field carries three conserved quantum numbers, Y , T_W and the $U(1)_{\text{EM}}$ charge Q , which are related by

$$Q = T_{W_3} + \frac{Y}{2}, \quad (1.2.4)$$

and the value T_{W_3} is the projection of T_W along the W_μ^3 direction (the third direction is chosen by convention). The appearance of T_W in the covariant derivative ensures that only the left-handed $SU(2)_L$ doublet fields couple to the W_μ^3 gauge fields as these are the only fields with non-zero eigenvalues under the T_W operator, chosen to be $\frac{1}{2}$ by convention. Q is the usual electromagnetic charge which appears after spontaneous symmetry breaking of the $SU(2)_L \otimes U(1)_Y$ theory. It is important to note that the different quantum numbers between the left-handed and right-handed components of the theory induce parity violating effects.

To this point, we have not introduced any mass terms for the fermions in the Lagrangian, which take the form

$$\begin{aligned} m\bar{\psi}\psi &= m(\psi_L^\dagger, \psi_R^\dagger) \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix} \\ &= m(\psi_L^\dagger \psi_R + \psi_R^\dagger \psi_L). \end{aligned} \quad (1.2.5)$$

Under $SU(2)_L$ symmetry, terms of this form are no longer gauge invariant, as the left- and right-handed components of the spinor fields have different quantum numbers and therefore behave differently under gauge transformations. As we will see, the

symmetry breaking induced through Higgs mechanism will give gauge invariant mass terms to the fermions and resolve this issue³.

1.2.2 Gauge Bosons

Having introduced the quark and lepton content of GWS theory, we now consider the gauge boson fields which have thus far appeared only in the covariant derivative \mathcal{D} . As in the generic $SU(N)$ case, we must now introduce kinetic terms for these fields to form a consistent picture. This is relatively straightforward, and proceeds in the manner we have seen previously.

The unbroken $SU(2)_L \otimes U(1)_Y$ EW gauge boson kinetic term in the GWS Lagrangian is given by

$$\mathcal{L}_{\text{gauge, kinetic}} = -\frac{1}{4}W^{i,\mu\nu}W_{\mu\nu}^i - \frac{1}{4}B^{\mu\nu}B_{\mu\nu}, \quad (1.2.6)$$

where $B^{\mu\nu}$ and $W^{i,\mu\nu}$ are the $U(1)_Y$ and $SU(2)_L$ antisymmetric field strength tensors,

$$W_{\mu\nu}^i = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i - g_2 \varepsilon^{ijk} W_\mu^j W_\nu^k \quad (1.2.7)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu. \quad (1.2.8)$$

Comparing these tensors to the generic case in (1.1.11), we can see that for $SU(2)$ the structure constants are given by the completely antisymmetric Levi-Civita tensor ε_{ijk} , whereas for the Abelian $U(1)$ symmetry we have $f^{abc} = 0$.

1.2.3 The Higgs Mechanism

At this point, we return to the subject of gauge boson masses. As we have seen already, there is naïvely no way to construct a mass term for these which is gauge invariant. However, this is in direct conflict with experimental observations, where we see three massive gauge bosons; two charged W bosons and the neutral Z boson.

³For the calculations performed in this thesis we will consider all active fermions to be massless; this assumption is valid at high energies $\sqrt{s} \gg m$, where finite mass corrections are extremely small. However such terms are still required for the consistency of the theory across energy regimes.

Such terms, as will be shown, can be constructed through the Higgs mechanism [13–15], where the $SU(2)_L \otimes U(1)_Y$ symmetry is spontaneously broken, giving a residual $U(1)_{EM}$ symmetry. The spontaneous breaking of the symmetry results in a vacuum state that does not have to respect the gauge symmetry and predicts the existence of a scalar spin-0 particle, known as the Higgs boson. In a landmark achievement, the existence of a Higgs-like boson was verified experimentally by the ATLAS and CMS collaborations in 2012 [16, 17] and efforts since have focussed on a comparison of its properties with the predictions of the SM.

We start by introducing a single complex $SU(2)_L$ doublet of scalar Higgs fields:

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.2.9)$$

which has weak isospin $T_{W_3} = \frac{1}{2}$ and weak hypercharge $Y = 1$. Using this we can define an associated Lagrangian, invariant under $SU(2)_L \otimes U(1)_Y$ gauge transformations

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi) + \mathcal{L}_{\text{Yukawa}}, \quad (1.2.10)$$

where $\mathcal{L}_{\text{Yukawa}}$ contains all interactions between the massive fermions and the new Higgs field, such that

$$\mathcal{L}_{\text{Yukawa}} = -\sum_{i,j=1}^{N_{gen}} [y_{ij}^u Q_L^i \sigma^2 \Phi u_R^j + y_{ij}^d Q_L^i \Phi d_R^j + y_{ij}^l L_L^i \Phi e_R^j + \text{h.c.}]. \quad (1.2.11)$$

The term $V(\Phi)$ is the scalar Higgs potential, and is constrained to be of the form

$$V(\Phi) = \mu^2 (\Phi^\dagger \Phi) - \lambda (\Phi^\dagger \Phi)^2, \quad (1.2.12)$$

where μ and λ are arbitrary parameters. We can infer that $\lambda < 0$ if the potential is to be bounded from below and vacuum stability is to be retained.

It is then instructive to consider where the minima of this potential occur. For $\mu^2 > 0$ we have

$$\frac{\partial V(\Phi)}{\partial \Phi} = 0; \quad V(\Phi) = \frac{\mu^4}{2\lambda}; \quad |\Phi^\dagger \Phi|_{\min} = 0, \quad (1.2.13)$$




Figure 1.1: The “Mexican Hat” Higgs potential for a single complex scalar ϕ . The orange line shows the degenerate set of states for which $\phi = \phi_{\min}$, and thus the potential is minimised.

which allows one to consider $\mu^2(\Phi^\dagger\Phi)$ as a standard mass term for the field Φ . More interesting, however, is the case where $\mu^2 < 0$, which gives minima of the potential at

$$\frac{\partial V(\Phi)}{\partial \Phi} = 0; \quad V(\Phi) = \frac{\mu^4}{2\lambda}; \quad |\Phi^\dagger\Phi|_{\min} = \frac{\mu^2}{2\lambda} \left(\equiv \frac{v^2}{2} \right). \quad (1.2.14)$$

This gives rise to the “Mexican Hat” potential, shown for a single complex scalar in Fig. 1.1, and it is this configuration which leads to the spontaneous breaking of the $SU(2)_L \otimes U(1)_Y$ symmetry. At this point, one can choose one of the complex components of (1.2.9) to take this non-zero expectation value v and break the symmetry. In order to preserve $U(1)_{\text{EM}}$ gauge invariance, we choose the real part of ϕ^0 to take a non-zero expectation value⁴, such that:

$$\begin{aligned} \langle \Re(\phi^+) \rangle &= \langle \Im(\phi^+) \rangle = \langle \Im(\phi^0) \rangle = 0 \\ \langle \Re(\phi^0) \rangle &= v. \end{aligned} \quad (1.2.15)$$

⁴Giving this non-zero expectation value to either of the charged degrees of freedom breaks gauge invariance, as the resulting field $h(x)$ is real-valued, with mass terms $\propto h(x)^2$. Were h charged under $U(1)_{\text{EM}}$, the resulting mass term would transform as $h^2 \rightarrow e^{2i\alpha(x)}h^2$, breaking the $U(1)_{\text{EM}}$ symmetry.

We can now consider small perturbations about the vacuum expectation value of each field

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \chi_1 + i\chi_2 \\ v + h(x) + i\chi_3 \end{pmatrix}, \quad (1.2.16)$$

where $h(x)$ labels the real-valued Higgs field and the χ_i are Nambu-Goldstone bosons. By performing a suitable SU(2) transformation (into the unitary gauge), the χ_i can be absorbed into a redefinition of the fields:

$$\Phi = \frac{1}{\sqrt{2}} e^{iT_W^j \chi^j/v} \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} + \mathcal{O}(h\chi_i) \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \quad (1.2.17)$$

Taking the vacuum expectation value term in this redefined Φ , and substituting back into the kinetic term of $\mathcal{L}_{\text{Higgs}}$ in (1.2.10), we extract

$$\begin{aligned} \mathcal{L}_{\text{Higgs, kinetic}} &= \frac{1}{2} \begin{pmatrix} 0 & v \end{pmatrix} \begin{pmatrix} i\frac{g_1}{2} B_\mu + ig_2 \frac{\sigma^i}{2} W_\mu^i \\ i\frac{g_1}{2} B^\mu + ig_2 \frac{\sigma^i}{2} W^{i,\mu} \end{pmatrix} \begin{pmatrix} 0 \\ v \end{pmatrix} \\ &= \frac{-v^2}{8} \left(g_1^2 B_\mu B^\mu + g_2^2 \sigma^i W_\mu^i \sigma^j W^{j,\mu} + g_1 g_2 B_\mu \sigma^i W^{i,\mu} \right) \\ &= \frac{-v^2}{8} \left[g_2^2 W_\mu^1 W^{1,\mu} + g_2^2 W_\mu^2 W^{2,\mu} + (g_1 B^\mu - g_2 W^{3,\mu})^2 \right]. \end{aligned} \quad (1.2.18)$$

At this point we can see the emergence of terms $\propto W_\mu^i W_i^\mu$, $B_\mu B^\mu$ which have the correct form for gauge boson mass terms. However the current picture is not complete as we now appear to have four massive bosons. This can be remedied through a mixing between the weak and mass eigenstates to recover the massless photon. That we have 3 massive and one massless boson is a consequence of Goldstone's theorem, in which one gains one scalar degree of freedom (the *Nambu-Goldstone* boson) for each generator of the broken symmetry, in this case three. These degrees of freedom under transformation become the longitudinal degrees of freedom of each of the massive bosons.

First, we need to rewrite the $W_\mu^{1,2}$ states as eigenstates of $U(1)_{\text{EM}}$ using the identity in (1.2.4). By definition their hypercharge Y is 0 (they are $SU(2)_L$ bosons), so the charge is given by the eigenvalues of the isospin operator, T_W using the algebra

$[\sigma_3, \sigma_i] = T_{W_3} \sigma_i$. The end result gives the linear combinations of gauge fields and generators:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \quad (1.2.19)$$

$$\sigma^\pm = \sigma^1 \pm i\sigma^2. \quad (1.2.20)$$

Substituting these into (1.2.18) we obtain equal masses for the W^\pm bosons, with mass

$$M_W = \frac{g_2 v}{2}. \quad (1.2.21)$$

Next we consider the electrically neutral W_μ^3 and B_μ fields. This requires diagonalisation of the mass matrix given by the final term in (1.2.18) in order to isolate the mass eigenstates. The mass matrix takes the form

$$M = g_1^2 \begin{pmatrix} 1 & -\frac{g_2}{g_1} \\ -\frac{g_2}{g_1} & \frac{g_2^2}{g_1^2} \end{pmatrix} \quad (1.2.22)$$

with eigenvalues 0 and $g_1^2 + g_2^2$. Thus one recovers one massless boson, the photon, and the massive Z boson, with mass

$$M_Z = \frac{\sqrt{(g_1^2 + g_2^2)} v}{2}. \quad (1.2.23)$$

This amounts to a mixing of W_μ^3 and B_μ through an angle $\sin \theta_W$:

$$\begin{pmatrix} A^\mu \\ Z^\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B^\mu \\ W_3^\mu \end{pmatrix}, \quad (1.2.24)$$

where the Weinberg angle θ_W ⁵ is defined through the relative strengths of the coupling constants:

$$\sin^2 \theta_W = \frac{g_1^2}{g_1^2 + g_2^2} (\sim 0.23). \quad (1.2.25)$$

This angle can in turn be used to define the relationship between the masses of the W and Z bosons, as

$$\cos \theta_W = \frac{m_W}{m_Z}. \quad (1.2.26)$$

⁵See Chapter 5 for an in-depth discussion of an experimental determination of $\sin^2 \theta_W$.

If one instead considers the $h(x)$ terms in the expansion of Φ , one similarly recovers direct interactions of the Higgs field with the massive EW gauge bosons, with a coupling strength proportional to the mass.

Having given a mass to each of the requisite gauge bosons, we now return to the fermions which interact with the Higgs field in the Yukawa sector of the Lagrangian. As demonstrated previously these cannot directly be given mass terms before spontaneous symmetry breaking. However, the expansion of Φ about the vacuum expectation value introduces gauge-invariant mass terms in a similar manner to the gauge bosons. Expanding $\mathcal{L}_{\text{Yukawa}}$ about v , we obtain

$$\mathcal{L}_{\text{Yukawa}} = -\frac{v}{\sqrt{2}} \sum_{i=1}^{N_{\text{gen}}} [y_{ii}^u u_L^i u_R^i + y_{ii}^d d_L^i d_R^i + y_{ii}^l e_L^i e_R^i + \text{h.c.}]. \quad (1.2.27)$$

This gives mass terms to all fermions with right-handed terms present in the Lagrangian⁶, proportional to their Yukawa coupling y^f ,

$$M_f = \frac{y^f v}{\sqrt{2}}. \quad (1.2.28)$$

As in the case of the gauge bosons, the Higgs couples to the fermions with a strength proportional to their mass. Ongoing experimental efforts are underway to measure this proportionality for the fermions, with measurements which are so far in agreement with the predictions of the standard model [18, 19].

1.2.4 Properties of the Higgs Boson

The final content of the EW sector that we have so far neglected is that of the Higgs itself. If we substitute the symmetry-broken form of Φ from (1.2.17) into the potential term $V(\Phi)$, we recover

$$V(\Phi) = \frac{\mu^2}{2} (v + h(x))^2 - \frac{\lambda}{4} (v + h(x))^4, \quad (1.2.29)$$

⁶The observation of neutrino oscillation implies they carry a mass; the generation of this mass without right-handed neutrinos remains an outstanding issue in the standard model.

which when expanded (dropping constant terms) gives a term for the Higgs mass of

$$M_H = v\sqrt{2\lambda}, \quad (1.2.30)$$

which is currently measured at 125.10 ± 0.14 GeV [8], alongside Higgs self-coupling terms of the form

$$\mathcal{L}_{\text{Higgs, self-int.}} = \frac{M_H^2}{2v} h^3(x) + \frac{M_H^2}{8v^2} h^4(x). \quad (1.2.31)$$

These induce Higgs boson self-interaction terms, in the form of trilinear and quartic couplings. Unfortunately, the production rates of multiple Higgs bosons required to measure these couplings are beyond the reach of the current generation of collider experiments, however there are hopes that future experiments will be able to directly measure these effects [20–22].

1.2.5 Quark Mixing and the CKM Matrix

Analogously to the rotation required between weak and mass eigenstates in the gauge boson sector after spontaneous symmetry breaking, a similar effect is seen in the quark sector in charged current W^\pm interactions. When one diagonalises the Yukawa mass matrix⁷, it rotates away from the weak eigenbasis in which the charged current interactions sit.

This rotation is described by a 3×3 unitary matrix, the Cabibbo-Kobayashi-Maskawa (CKM) matrix [23, 24] $V_{\text{CKM}} \equiv V_L^u V_L^{d\dagger}$, which mixes the quark generations through terms in the Lagrangian of the form

$$\frac{-g_2}{\sqrt{2}} (\bar{u}_L, \bar{s}_L, \bar{t}_L) \gamma^\mu W_\mu^+ V_{\text{CKM}} \begin{pmatrix} d_L \\ s_L \\ b_L \end{pmatrix} + \text{h.c.}, \quad V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (1.2.32)$$

For 3 generations of quarks, V_{CKM} can be parameterised by three mixing angles alongside a complex CP-violating phase δ_{CP} , and the current world-average values are given in Table 1.3. CP violation is induced in the presence of at least 3 quark

⁷This was performed implicitly in (1.2.27) where only terms of the form y_{ii} were considered.

$$V_{\text{CKM}} = \begin{pmatrix} 0.97420 \pm 0.00021 & 0.2243 \pm 0.0005 & (3.94 \pm 0.36) \times 10^{-3} \\ 0.218 \pm 0.004 & 0.997 \pm 0.017 & (42.2 \pm 0.8) \times 10^{-3} \\ (8.1 \pm 0.5) \times 10^{-3} & (39.4 \pm 2.3) \times 10^{-3} & 1.019 \pm 0.025 \end{pmatrix}$$

Table 1.3: Current world average values for the elements of the CKM matrix, using the parameterisation for V_{CKM} used in (1.2.32). Data taken from [8].

generations by permitting a phase difference between a process and the equivalent CP conjugated process.

As an example, we take the processes $u \rightarrow dW^+$ and $\bar{u} \rightarrow \bar{d}W^-$, which can proceed either directly, $\propto |V_{ud}|$, or indirectly through a chain, e.g. $\propto |V_{us}||V_{sc}||V_{cd}|$ (at the amplitude level). If we permit a relative complex phase ϕ in the CKM matrix between the two mechanisms, we see

$$\begin{aligned} u \rightarrow dW^+ &\sim ||V_{ud}| + |V_{us}||V_{sc}||V_{cd}|e^{+i\phi}|^2 \\ \bar{u} \rightarrow \bar{d}W^- &\sim ||V_{ud}| + |V_{us}||V_{sc}||V_{cd}|e^{-i\phi}|^2, \end{aligned}$$

which induces a CP-violation through the rate difference parameterised by the phase

$$||V_{ud}| + |V_{us}||V_{sc}||V_{cd}|e^{+i\phi}|^2 - ||V_{ud}| + |V_{us}||V_{sc}||V_{cd}|e^{-i\phi}|^2 \propto \sin(\phi). \quad (1.2.33)$$

CP violation in the weak sector arising from the CKM matrix was first observed indirectly in neutral K -meson decays [25] in which the two weak eigenstates of the mesons were found to have different lifetimes.

1.3 Quantum Chromodynamics (QCD)

Having discussed the EW content of the SM, we now turn our attention to the strong sector, governed by the theory of quantum chromodynamics (QCD) under an $\text{SU}(3)_C$ gauge symmetry and to which the only fermions that couple are the quarks. As no symmetry breaking occurs in QCD theory and the gauge bosons, the gluons, are massless, the Lagrangian itself is relatively straightforward in comparison to its

EW counterpart. Here we neglect mass terms for the quarks, which are generated through Higgs couplings in the EW sector.

As usual, we start with the QCD Lagrangian:

$$\mathcal{L}_{\text{QCD}} = \sum_{\text{quarks}} \bar{\psi}_q i \not{D} \psi_q - \frac{1}{4} G_{\mu\nu}^a G_{\mu\nu}^a, \quad (1.3.1)$$

with the covariant derivative

$$\not{D} = \gamma^\mu \left(\partial_\mu - i g_s T^a G_\mu^a \right). \quad (1.3.2)$$

Here T^a are the $N^2 - 1 = 8$ generators of $\text{SU}(3)$ in the adjoint representation, the Gell-Mann matrices, and G_μ^a are the gluon fields. g_s is the QCD coupling constant, related by convention to α_s by $\alpha_s = \frac{g_s^2}{4\pi}$ (~ 0.12 at M_Z). At hadron colliders, g^2 is an order of magnitude greater than g_1^2 , meaning that QCD contributions generally dominate for processes that can proceed through either EW or QCD channels. Calculations in QCD are usually performed through an asymptotic perturbative expansion of transition correlation functions in α_s , valid in the regime⁸ $\alpha_s \ll 1$.

The expansion of correlation functions in a coupling constant can be used to derive Feynman rules for a theory through Wick's theorem, which allow construction of transition amplitudes to a given perturbative order without recourse to path integrals. These amplitudes can then be used to construct predictions for production rates of given final states, as we will see later on.

1.3.1 QCD Cross Sections

QCD predictions for scattering processes are conventionally given in terms of the cross section, a measure of the rate at which a scattering process proceeds, normalised to the flux of incident particles. Predictions in QFT are generally performed at the particle level, which poses a problem for calculations of scattering processes

⁸ α_s (and indeed α) is not constant, and adopts a value that is dependent on the energy scale at which it is probed, an effect known as the “running of the coupling”. This will be considered in more detail in Section. 1.3.3.

with incoming coloured particles (partons). Due to colour confinement, incoming states can only be known at hadronic level, and one must find a way to relate the partonic quantities (such as the differential partonic cross section $d\hat{\sigma}$) to measurable hadronic quantities $d\sigma$. This can be done using the cross section “master formula” for hadron-hadron collisions:

$$d\sigma(p_1, p_2) = \sum_{i,j} \int_0^1 dx_1 dx_2 f_{i/h_1}(x_1, \mu_F) f_{j/h_2}(x_2, \mu_F) d\hat{\sigma}_{ij}(x_1 p_1, x_2 p_2, \mu_F, \mu_R), \quad (1.3.3)$$

where i, j are the partons inside their parent hadrons h_1 and h_2 , and the values $x_1 p_1$ and $x_2 p_2$ are rescalings of the hadronic momenta p_1 and p_2 to give the incoming partonic momenta. The objects f_i are non-perturbative parton distribution functions (PDFs) and μ_F and μ_R are characteristic energy scales of the process, to be introduced in the following sections. The above equation relies on the principle of factorisation, whereby the soft non-perturbative physics of the hadron factorise from the hard perturbative cross section and are parameterised by the PDFs. If one wants to replace a hadronic state h with a leptonic initial state l , as is the case in deep inelastic scattering (DIS) where one has lepton-hadron collisions, the above can be altered by the substitution

$$f_{i/h_1} \rightarrow \delta_{il} \delta(1 - x_1). \quad (1.3.4)$$

The partonic cross section $d\hat{\sigma}_{ij}$ is then usually treated perturbatively as an expansion in the coupling constant, α_S in the case of QCD:

$$\begin{aligned} d\hat{\sigma}_{ij} &= \sum_{n=0}^{\infty} \left(\frac{\alpha_s}{2\pi} \right)^n d\hat{\sigma}_{ij}^{(n)} \\ &= d\hat{\sigma}_{ij,LO} + \left(\frac{\alpha_s}{2\pi} \right) d\hat{\sigma}_{ij,NLO} + \left(\frac{\alpha_s}{2\pi} \right)^2 d\hat{\sigma}_{ij,NNLO} + \dots, \end{aligned} \quad (1.3.5)$$

where all factors of α_s present at lowest order (which we call leading order, or LO) have been absorbed into $d\hat{\sigma}$ and the following terms are labelled next-to-leading order (NLO), next-to-next-to-leading order (NNLO) and so on. It is this perturbative series that is truncated for practical purposes in calculations in the region $\alpha_s \ll 1$,




Figure 1.2: Representative Feynman diagrams at LO, NLO and NNLO for boson production in quark-antiquark collisions.

with each successive order improving the approximation⁹.

Each order in the coupling constant introduces new interaction topologies, which are suppressed by factors of α_s from LO and take the form of real radiative emissions and virtual loop corrections. Examples of Feynman diagrams for these are shown in Figure. 1.2 for the case of boson production from a quark line. In both real and virtual diagrams, one must integrate over the new degrees of freedom introduced. For virtual diagrams this becomes an integration over all possible values of the internal loop momenta, and for real emissions the degrees of freedom are absorbed by an increase in the dimensionality of the phase space integral contained within $d\hat{\sigma}_{ij}$.

1.3.2 Divergences in QCD

Beyond LO, one encounters two separate forms of divergences in QCD calculations that must be handled appropriately. In order to demonstrate this, we consider the 3-leg massless one-loop integral with external 4-momenta p_1 , p_2 and $q = p_1 - p_2$ alongside the undetermined internal 4-momentum k , as occurs in the quark-gluon vertex diagram (see Fig. 1.3)

$$I_3 = \int_{-\infty}^{\infty} \frac{d^4 k}{(2\pi)^4} \frac{(p_1 + k)(-ie\gamma_\mu)(p_2 + k)}{[k^2 + i0][k^2 + i0][(k + p_1)^2 + i0][(k + p_2)^2 + i0]}. \quad (1.3.6)$$

⁹This does not apply to all orders, as the series is asymptotic. See e.g. [26] for the QED case.




Figure 1.3: The one quark-gluon vertex diagram, where the quark line couples to an off-shell external photon with invariant mass q . The incoming quarks are massless, with $p_1^2 = p_2^2 = 0$.

We have introduced the term $i0$ to maintain causal ordering, but this will henceforth be omitted. The first divergence we observe in this integral occurs in the high energy *ultra-violet* (UV) regime, where $k^2 \gg p^2$ and the integral behaves as

$$\begin{aligned}
 I_3 &\sim \int_{-\infty}^{\infty} \frac{d^4 k}{(2\pi)^4} \frac{k^2}{k^6} \\
 &\sim \int d\Omega_3 \int_0^{\Lambda} \frac{d|k|}{(2\pi)^4} \frac{|k|^5}{|k|^6} \\
 &\sim \lim_{\Lambda \rightarrow \infty} \log(\Lambda),
 \end{aligned} \tag{1.3.7}$$

where we have introduced parameter Λ to regulate the divergence. As we take $\Lambda \rightarrow \infty$, we observe singular behaviour.

Secondly, the integral is also divergent in the low energy *infrared* (IR) regime. This can be seen by introducing a small, unphysical mass parameter μ in order to regulate the divergence

$$I_3 = \int_{-\infty}^{\infty} \frac{d^4 k}{(2\pi)^4} \frac{(p_1 + k)(-ie\gamma_\mu)(p_2 + k)}{[k^2 - \mu^2][(k + p_1)^2 - \mu^2][(k + p_2)^2 - \mu^2]}. \tag{1.3.8}$$

When evaluated, this integral contains terms of the form

$$\log^2 \left(\frac{-q^2}{\mu^2} \right), \tag{1.3.9}$$

the well-known Sudakov double logarithm, which displays singular behaviour for

massless partons ($\mu \rightarrow 0$). This divergence is associated with kinematic configurations where momenta in the denominator vanish and the integrand becomes singular up to the regulator μ^2 . These divergences have a fundamentally different origin to the UV singularities, and therefore must be handled differently.

Whilst in the above example we have introduced two different *ad-hoc* regulators of the divergences in the form of μ and Λ , they can both in practice be isolated simultaneously using dimensional regularisation (DR) [27]. Compared to other regularisation schemes, DR has the important property of respecting gauge and Lorentz invariance and is implemented by analytically continuing the number of space-time dimensions from 4 to D , where $D = 4 - 2\varepsilon$. Both IR and UV divergences present themselves as poles in ε , and the $\varepsilon \rightarrow 0$ limit is taken after the calculation is performed in order to recover the intended result in 4 dimensions.

In order to take this limit in ε however, one must first have calculational techniques available in order to tackle the poles in both IR and UV regimes. In the following sections we will discuss the use of renormalisation to absorb the UV divergences of the theory before moving on to the cancellation of IR poles between radiative and virtual contributions for well-defined observables. Whilst we restrict the present discussion to QCD, it should be noted that these methods can also be applied to the EW sector where similar issues arise.

1.3.3 UV Divergences and Renormalisation

We start with the case of UV divergences which are handled by renormalisation. In general, the method of renormalisation amounts to a redefinition of the input parameters and fields within the SM Lagrangian. Thus far these objects have simply been parameterisations of the theory rather than experimentally measurable quantities and as a result there is some amount of freedom in both their interpretation and measurement.

For renormalisable theories such as QCD, one can rescale the fields and coupling

constant by overall multiplicative factors:

$$\begin{aligned}\psi_0(x) &= Z_2^{\frac{1}{2}} \psi(x); & G_0^\mu(x) &= Z_3^{\frac{1}{2}} G^\mu(x); \\ \eta_0(x) &= Z_\eta^{\frac{1}{2}} \eta(x); & g_{s,0}^2(x) &= Z_g g_s^2,\end{aligned}\tag{1.3.10}$$

where the bare quantities on the left hand side of each definition are the parameters as appearing in the Lagrangian to this point. If we let each of these scaling factors Z_i be divergent, they can be used to define a finite number of UV counterterms in the Lagrangian at each order in perturbation theory to remove the UV divergences, as

$$Z_i = 1 + \delta_{Z_i}.\tag{1.3.11}$$

This has the effect of decoupling the bare Lagrangian into a finite renormalised contribution $\mathcal{L}_{\text{renorm}}$ and a counterterm contribution $\mathcal{L}_{\text{c.t.}}$:

$$\mathcal{L}_0 = \mathcal{L}_{\text{renorm}} + \mathcal{L}_{\text{c.t.}}.\tag{1.3.12}$$

Taking our QCD Lagrangian from Eqn. (1.3.1), this explicitly gives:

$$\begin{aligned}\mathcal{L}_0 &= \sum_{\text{quarks}} \bar{\psi}_{q,0} i \not{D} \psi_{q,0} - \frac{1}{4} G_{\mu\nu,0}^a G_{a,0}^{\mu\nu} \\ &= \overbrace{\sum_{\text{quarks}} \bar{\psi}_q i \not{D} \psi_q - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}}^{\mathcal{L}_{\text{renorm}}} \\ \mathcal{L}_{\text{c.t.}} &\left\{ \begin{aligned} &+ (Z_2 - 1) \sum_{\text{quarks}} \bar{\psi}_q i \not{D} \psi_q + g(Z_g Z_2 Z_3^{1/2} - 1) \sum_{\text{quarks}} \bar{\psi}_q G^a T^a \psi_q \\ &- \frac{1}{4} (Z_3 - 1) (\partial_\mu G_\nu^a - \partial_\nu G_\mu^a)^2 - g(Z_g Z_3^{1/2} - 1) \partial_\mu G_\nu^a \varepsilon^{abc} G_{\mu,b} G_{\nu,c} \\ &+ \frac{1}{4} g^2 (Z_g^2 Z_3^2 - 1) (\varepsilon^{abc} G_{\mu,b} G_{\nu,c})^2. \end{aligned} \right.\end{aligned}\tag{1.3.13}$$

If we now systematically perform calculations and interpret experimental measurements using $\mathcal{L}_{\text{renorm}}$, these UV divergences are no longer present. This results in some loss of predictive power as the renormalised quantities cannot be directly predicted from the Lagrangian. However with the diversity of experimental measurements available one can still predict event rates and distributions for large classes of observables.

This loss of predictivity is related to an ambiguity in the definition of $\mathcal{L}_{\text{renorm}}$. When partitioning the Lagrangian into its renormalised contribution and counterterm there is a generic freedom as to where one chooses to place finite terms, which is generally fixed by a choice of *renormalisation scheme*. These finite terms are then absorbed into a redefinition of the input parameters. Care has to be taken to ensure that the input parameters to the Lagrangian are accounted for consistently in the same scheme to ensure the correctness of the calculation. Failing to do so can result in large differences between schemes¹⁰.

The simplest scheme, known as the Minimal Subtraction (MS) scheme, is the one which allocates no finite parts to the counterterm, rendering it purely divergent. However, in practice the Modified Minimal Subtraction ($\overline{\text{MS}}$) scheme, in which one also absorbs finite corrections that always accompany poles in the dimensional regulator:

$$\frac{1}{\varepsilon} + \log(4\pi) - \gamma_E + \mathcal{O}(\varepsilon), \quad (1.3.14)$$

where γ_E is the Euler-Mascheroni constant ~ 0.578 is most commonly used for QCD calculations. This can be rewritten as a multiplicative factor in a similar manner to the renormalisation factors Z_i , as

$$\frac{1}{\varepsilon} \rightarrow (4\pi)^\varepsilon e^{-\varepsilon\gamma_E} \frac{1}{\varepsilon} \equiv \bar{C}(\varepsilon) \frac{1}{\varepsilon}. \quad (1.3.15)$$

In practice, when one performs loop calculations in D -dimensions, the poles in ε become explicit as Gamma (Γ) functions of the regulator, at which point an expansion around $\varepsilon = 0$ can be performed and the UV divergences systematically renormalised. For an l -loop calculation, one will find UV poles at most of order ε^l .

There is, however, one important consideration that has been neglected so far, the energy scale μ_R (*renormalisation scale*) at which the theory is renormalised. When moving from $D = 4$ to $D = 4 - 2\varepsilon$, one must introduce this unphysical quantity in

¹⁰As an example, the mass of the top quark between the $\overline{\text{MS}}$ and pole mass schemes is $\mathcal{O}(10)$ GeV, amounting to a $\mathcal{O}(5)\%$ effect on the total top quark mass [28]. It is obviously imperative to use the correct definitions for precision measurements.




Figure 1.4: The running of the strong coupling constant alongside a selection of experimental measurements used in the extraction of the world average value of $\alpha_s(M_Z)$. Image taken from [8].

order to maintain the dimensionality of the Lagrangian as e.g. $\alpha_s \rightarrow \mu_R^{2\varepsilon} \alpha_s$ ¹¹. The choice of μ_R is in principle arbitrary, as physical observables will be independent of its value, and observables calculated to all orders will have no dependence on μ_R . However, for observables computed at finite orders in perturbation theory this is not the case and they will retain some residual dependence.

Given that μ_R is arbitrary, it is also possible to renormalise the theory at a new scale, μ'_R . By considering a general observable \mathcal{O} dependent on momentum scale q^2 , one can write this independence as

$$\begin{aligned} 0 &= \mu_R^2 \frac{d}{d\mu_R^2} \mathcal{O} \left(\frac{q^2}{\mu_R^2}, \alpha_s(\mu_R^2) \right) \\ &= \left[\frac{-\partial}{\partial t} + \beta(\alpha_s(\mu_R^2)) \frac{\partial}{\partial \alpha_s(\mu_R^2)} \right] \mathcal{O} \left(e^t, \alpha_s(\mu_R^2) \right), \end{aligned} \quad (1.3.16)$$

where in the second line we have substituted $t = \log(q^2/\mu_R^2)$, and have defined the QCD β -function as

$$\beta(\alpha_s(\mu_R^2)) = \mu_R^2 \frac{\partial \alpha_s(\mu_R^2)}{\partial \mu_R^2}. \quad (1.3.17)$$

¹¹Through dimensional analysis, the same procedure is not required of the fields ψ , G^μ and η .

For $\alpha_s \ll 1$, β can be perturbatively expanded with coefficients β_i as

$$\frac{\beta(\alpha_s(\mu_R^2))}{2\pi} = - \sum_{i=0} \beta_i \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^{i+2}, \quad (1.3.18)$$

and considering only the leading β_0 term¹² one can solve to obtain

$$\alpha_s(\mu_R^2) = \frac{\alpha_s(\mu_R'^2)}{1 - \frac{\beta_0}{2\pi} \alpha_s(\mu_R'^2) \log \left(\frac{\mu_R^2}{\mu_R'^2} \right)} \quad (1.3.19)$$

with some secondary scale $\mu_R'^2$. From this we can deduce that there is a dependence of the renormalised parameter α_s on the energy scale at which it is measured - it is a *running coupling*. This running has been extracted from an extensive range of experimental data, of which a selection of results are summarised in Fig. 1.4.

One particularly important point to mention here is that the sign of the β_0 coefficient determines the behaviour of the coupling at both low and high energies. For QCD, $\beta_0 > 0$, and the coupling strength decreases at higher energies, whilst becoming non-perturbative in the low energy regime, a phenomenon known as *asymptotic freedom*. At low energies, the coupling strength increases. This is in contrast to QED which has a negative β_0 coefficient; this leads to a coupling strength that grows with energy.

1.3.4 IR Divergences and Cancellation - a Heuristic

Overview

Having dealt with the UV divergences arising from the Lagrangian, it now remains to consider the low energy IR singularities in calculations. We provide only an overview here, and postpone a detailed discussion to Chapter 2 where we will consider the subject in the context of antenna subtraction.

Within a perturbative calculation beyond the leading order, IR divergences arise in two different sectors; through loop diagrams where internal momenta become arbitrarily small, as we have seen already, and for particular kinematic configurations

¹²For results to higher orders including the full cross section dependence on μ_R , see Appendix. B.2.1.




Figure 1.5: A gluonic real emission from a quark line becoming unresolved in either the soft $|k| \rightarrow 0$ or the collinear $k \parallel p$ limit.

of real emission contributions. These configurations occur when massless external partons (quarks and gluons) become either soft or collinear with one another and thus become degenerate with the kinematics of the process in which the real emission does not occur.

To illustrate such a real IR divergence, we take the example of real gluon emission from a massless quark line as shown in Fig. 1.5, where the external quark and gluons are taken on-shell ($k^2 = p^2 = 0$). In this case, the internal quark propagator with momentum $p + k$ takes the form

$$\begin{aligned} P_q &\sim \frac{\not{p} + \not{k}}{(p^\mu + k^\mu)^2 + i0} \\ &\sim \frac{\not{p} + \not{k}}{2|p||k|(1 - \cos(\theta_{qg}))}, \end{aligned} \quad (1.3.20)$$

where θ_{qg} is the angle between the gluon and quark 3-vectors, and $|p|$ and $|k|$ can be identified as the energies E_p and E_k of the quark and gluon respectively. At this point it becomes evident that the propagator diverges in two limits¹³:

1. $|k| (= E_k) \rightarrow 0$, the *soft gluon* limit,
2. $\cos(\theta_{qg}) \rightarrow 1$; $\theta_{qg} \rightarrow 0$, the *quark-gluon collinear* limit.

These divergent contributions lie in a higher multiplicity phase space ($N+1$ external particles) than those in the loop sector (N external particles), and do not directly include explicit poles in the dimensional regulator ε . However, when the phase space

¹³The quark cannot go soft as it is charged under $U(1)_{\text{EM}}$ and thus $E_p \rightarrow 0$ would violate the conservation of the EM current.

integration is performed, one in effect integrates over the unresolved configurations and recovers explicit ε poles.

At this point we take a slight detour to consider the impact of the fact that these singularities occur for momentum configurations that are exactly equivalent to the configurations in which no real emission occurred at all. In particular, this implies that the kinematics of the divergence map directly on to kinematics from a phase space of multiplicity N . This observation is a powerful one! The presence of such $N + 1 \rightarrow N$ mappings allows the derivation of factorisation formulae for singular limits, in which the matrix element in $(N + 1)$ -multiplicity space factorises into a contribution in N -multiplicity space convolved with term \mathcal{V}_1 encoding the universal singular structure in a 1-particle phase space, schematically

$$|\mathcal{M}_{N+1}|^2 \xrightarrow[\text{limit}]{\text{singular}} |\mathcal{M}_N|^2 \otimes \mathcal{V}_1. \quad (1.3.21)$$

There are multiple such mappings and associated phase space factorisations, including the Catani-Seymour dipole [29] and antenna [30] factorisations, which directly inform many of the techniques used to handle IR singularities in practical calculations as we will see later.

Whilst we have only briefly illustrated the case in which a single particle becomes unresolved, as occurs at NLO, this concept generalises to higher orders. In these cases the singularity structure becomes more complex due to the presence of divergent limits in which multiple external partons simultaneously become unresolved, and due to combinations of both loop and soft/collinear divergences.

Combining the IR divergences from the loop sector and real emissions, we now have two separate sources of IR singular behaviour in our fixed-order cross-section calculations. It has been proven, first for the case of QED by Bloch and Nordsieck [31] and later for the SM in general by Kinoshita, Lee and Nauenberg [32, 33] in the famous KLN theorem, that for well-defined observables¹⁴ these contributions must

¹⁴In this context, a well-defined observable is taken to mean that the observable is not directly sensitive to the IR singular contributions, and is therefore *IR-safe*.

cancel to give a finite result. Ensuring that this cancellation occurs is non-trivial as it occurs between phase spaces of different dimensionalities, making a direct numerical integration impossible.

Before we outline the techniques developed to facilitate this cancellation, we must first address an additional complication that occurs when one or more of the initial states is hadronic. In this case one does not have a single coloured incoming particle, rather a non-perturbative bound state of partons due to colour confinement, and one must incorporate the non-perturbative contributions to the cross-section through the PDFs. The PDFs, as shown in Fig. 1.6, model low energy strongly-coupled physics and at leading order can be understood as providing the probability of finding a given parton within the hadron at a given energy scale μ_F with a given fraction of the hadronic momentum z . These PDFs are in general determined from fits to experimental data and lattice QCD results, and cannot be derived from first principles.

The energy scale μ_F , known as the *factorisation scale*, can be thought of as the energy scale which partitions the soft physics of the PDF from the perturbative hard scattering cross section, and is necessary to derive the cross-section master formula in (1.3.3). The factorisation of the two contributions has only been proven for a limited number of scattering processes, including deep inelastic scattering (DIS) [34] in lepton-hadron collisions and Drell-Yan production in hadron-hadron collisions [35, 36], and its violation has even been demonstrated in pure QCD processes at the 3-loop level [37]. Despite this, factorisation is generally assumed within perturbative calculations as both experimental and theoretical technology are not yet in a position to probe the violation.

The introduction of PDFs leads to an ambiguity in initial state real emission contributions as they can be interpreted either as part of the hard scattering process or as part of the PDF. This ambiguity is made manifest as a residual dependence of the cross-section on the factorisation scale μ_F . These emissions can change the incoming momentum fraction z and generate divergences when they become collinear with the

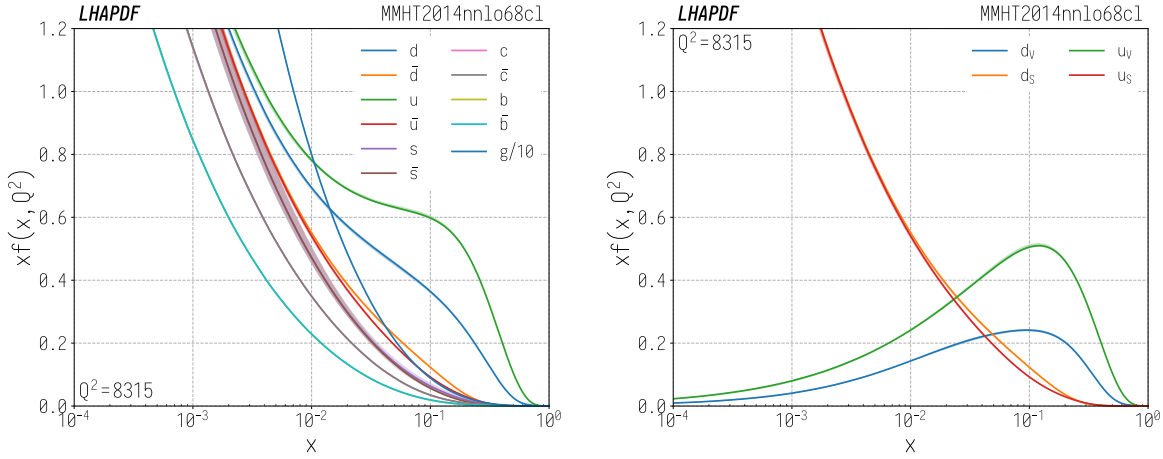


Figure 1.6: The individual parton distributions (left) and valence and sea quark distributions (right) as a function of the proton momentum fraction x at $Q^2 = M_Z^2$ contained within the MMHT14 PDF set [38] and produced using the LHAPDF package [39].

initial state. In analogy with renormalisation, these singularities can be absorbed into a redefinition of the PDFs in order to ensure finiteness, shown schematically in Fig. 1.7. Once again this is a scheme dependent procedure, with the finite contribution absorbed into most modern PDF sets defined in the $\overline{\text{MS}}$ scheme, allowing one to view μ_F as the “renormalisation scale” of the PDF. The full dependence of observables on μ_F can be reconstructed analytically and is conventionally used alongside the renormalisation scale dependence to parameterise the uncertainties from missing higher orders in theoretical predictions. The full dependence of the hard scattering cross section on the factorisation scale is discussed and reconstructed in Appendix B.2.2, and a more detailed discussion of the “mass-factorisation” terms absorbed into the PDF definition is given in Section 2.4.1.

This PDF renormalisation procedure allow us to fully cancel IR singularities and derive IR-finite predictions for QCD observables in processes with hadronic initial states. For simple processes at low perturbative orders it is indeed possible to perform analytic phase space integrations in order to extract the implicit poles from the real emission graphs; this is not true in general, particularly so in the presence of fiducial cuts. As a scheme for providing predictions for the vast majority of collider processes, this evidently is not sufficient. This inadequacy comes mainly from the




Figure 1.7: Two interpretations corresponding to the ambiguity in PDF definition in the presence of unresolved initial state radiation ($k \parallel Pz$).

inability to perform arbitrary phase space integrals over divergent integrands in D -dimensions either numerically or analytically. However, there exist two primary classes of techniques, *slicing* and *subtraction* which have been developed in order to facilitate the numerical evaluation of such phase space integrals.

1.3.5 Slicing Methods

Slicing methods, first developed in the early ‘90s [40, 41], deal with the issue of numerical integration over divergent regions of real emission phase space by introducing a cut on observable \mathcal{O} . This observable approaches some limit \mathcal{O}_0 as the integrand approaches a singularity, allowing it to be used to define “singular” and “non-singular” regions of phase space. In the non-singular region the integrand remains finite and the integration can be performed numerically in 4 dimensions using standard techniques. In the singular region, one exploits the factorisation properties of phase space to give an integral in N -particle space over some analytic function \mathcal{F} that approximates the behaviour of the real corrections as the singular limit is approached. In combination with the virtual contributions which also reside in the N -particle phase space this can be used to construct the total higher order cross section.

At NLO, using value \mathcal{O}_{cut} to divide the singular and non-singular regions, this can be written schematically as

$$\hat{\sigma}_{ij}^{NLO} = \int_{\Phi_{n+1}} d\hat{\sigma}_{ij}^R \Theta(\mathcal{O} - \mathcal{O}_{\text{cut}}) + \int_{\Phi_n} d\mathcal{F}_{ij} + \int_{\Phi_n} d\hat{\sigma}_{ij}^V \Big|_{\mathcal{O}=\mathcal{O}_0} \quad (1.3.22)$$

where

$$\int_{\Phi_{n+1}} d\hat{\sigma}_{ij} \Theta(\mathcal{O}_{\text{cut}} - \mathcal{O}) \sim \int_{\Phi_n} d\mathcal{F}_{ij}. \quad (1.3.23)$$

Typical observables used for slicing include Q_T , the transverse momentum of a colour singlet [42], and N -jettiness, an event shape variable [43], both of which have associated factorisation theorems and formalisms extended up to NNLO. One must ensure that the final result is independent of \mathcal{O}_{cut} . In other words, Eqn. (1.3.23) must hold to a sufficient accuracy. This is a highly non-trivial constraint, as computations typically become significantly more time intensive as one probes the singular $\mathcal{O}_{\text{cut}} \rightarrow \mathcal{O}_0$ region and encounters growing numerical instabilities. However this is relatively straightforward to implement if equipped with function \mathcal{F} , as the $\mathcal{O} > \mathcal{O}_0$ real contribution can be produced using computations of $\hat{\sigma}$ in association with one extra partonic emission at one order lower in perturbation theory.

Slicing methods are a form of non-local singularity cancellation, where the divergences do not cancel for individual points in phase space, rather only after integration. This means that they are not sufficient for use with tools such as parton showers, which procedurally generate soft emissions on top of phase space points in order to perform all-orders resummations.

1.3.6 Subtraction Methods

The second class of methods used to keep divergent phase space integrals finite are those that rely on subtraction. In these methods, one introduces a counterterm for the real emission parts of the calculation which exactly mimics the divergent structure of the integrand and can be analytically integrated over the extra degrees of freedom not present at Born level. This counterterm ensures that the real emission parts are finite in the singular limits, making it amenable to numerical integration. The counterterm can then be integrated over the unresolved phase space to give explicit poles in ε , and can then be reintroduced in the lower multiplicity phase spaces to directly cancel the explicit poles in the virtual contributions. These integrated

counterterms exactly compensate for the contributions removed in the real part, and therefore leave the total cross section unchanged.

Taking the example of NNLO corrections to a given process, and neglecting the mass factorisation contributions to show the schematics:

$$d\hat{\sigma}_{ij}^{NNLO} = \int_{\Phi_n} d\hat{\sigma}_{ij}^{VV} + \int_{\Phi_{n+1}} d\hat{\sigma}_{ij}^{RV} + \int_{\Phi_{n+2}} d\hat{\sigma}_{ij}^{RR}, \quad (1.3.24)$$

if we construct subtraction counterterms $\hat{\sigma}_{ij}^S$, $\hat{\sigma}_{ij}^T$ and $\hat{\sigma}_{ij}^U$ such that each term in

$$d\hat{\sigma}_{ij}^{NNLO} = \int_{\Phi_n} (d\hat{\sigma}_{ij}^{VV} - d\hat{\sigma}_{ij}^U) + \int_{\Phi_{n+1}} (d\hat{\sigma}_{ij}^{RV} - d\hat{\sigma}_{ij}^T) + \int_{\Phi_{n+2}} (d\hat{\sigma}_{ij}^{RR} - d\hat{\sigma}_{ij}^S), \quad (1.3.25)$$

is IR finite, we can then see that each term is integrable in an integer number of dimensions. The condition that this leaves the overall cross section unaltered is equivalent to

$$\int_{\Phi_n} d\hat{\sigma}_{ij}^U + \int_{\Phi_{n+1}} d\hat{\sigma}_{ij}^T + \int_{\Phi_{n+2}} d\hat{\sigma}_{ij}^S = 0, \quad (1.3.26)$$

such that the subtraction terms exist purely to redistribute divergent quantities between different phase space multiplicities.

When undertaking this procedure there is a freedom in the finite contribution accompanying the divergent structure incorporated into the subtraction counterterm. As a result, one can define a subtraction scheme in a multitude of ways, which is reflected in the number of methods available for use in calculation. The procedure has been automated at NLO for the FKS [44] and Catani-Seymour [45] schemes, whilst at NNLO automation has not yet been achieved due to the substantial increase in complexity. However many NNLO schemes have been developed including antenna subtraction [46–48] (see Chapter 2), projection to Born (p2B) [49, 50], (see Section 6.3), COLORFULNNLO [51, 52], nested soft-collinear subtraction [53, 54] and local analytic sector subtraction [55].

Whilst these calculations are generally more time-intensive to construct compared to slicing methods due to the intricate nature of singularity cancellation at NNLO and beyond, they have multiple advantages. They are usually much more numerically

stable which allows their use deep into regions of phase space that approach singularities, such as the low- p_T region of colour singlets, and in many cases are local (or quasi-local). Computational efficiency is usually improved as well as one constructs a point-by-point cancellation of divergences rather than relying on cancellation of divergences post-integration. Both slicing and subtraction techniques have lead to a proliferation of NNLO calculations for low multiplicity processes where 2-loop matrix elements are available, dubbed the “NNLO revolution”, and as a result NNLO has become the de-facto standard for high-precision QCD.

There is a third class of methods, based on sector decomposition, which we only briefly mention here. In these methods, the integrand is divided into separate sectors which are amenable to numerical integration in order to extract the poles in ε . The method was first developed in [56, 57], and has since been implemented in the form of sector-improved residue subtraction in the program STRIPPER [58–60].

1.4 High Precision QCD Observables at the LHC

A wide variety of scattering processes have been measured experimentally at the LHC since data collection began in 2009, with measured cross-sections spanning 14 orders of magnitude. Figure 1.8 summarises the latest set of published results from the ATLAS collaboration [61]. For the majority of processes, higher order calculations in QCD are mandatory and provide the dominant corrections beyond LO; the increasing accuracy of the experimental results provides an essential testing ground for the predictions of the standard model. Whilst the canonical example of the necessity of QCD corrections is inclusive Higgs production through gluon fusion, where NNLO QCD corrections are $\mathcal{O}(20\%)$ over the NLO calculation, it is also invaluable for e.g. vector boson production, where experimental precision is approaching the




Figure 1.8: Summary of cross-section measurements made by the ATLAS collaboration at the LHC as of March 2019, separated by process. The data are compared to benchmark theory results at $\sqrt{s} = 5.02, 7, 8$ and 13 TeV. Image taken from [61].

per-mille level¹⁵. As we will see in later chapters, such processes can be exploited for a large number of important measurements, from PDF determinations to measurements of the fundamental EW and QCD parameters of the standard model.

Beyond inclusive measurements of EW boson production, in which the dynamics of the boson are probed through its decay products, one can also use outgoing hadronic radiation to provide a direct handle on the QCD dynamics of processes. One typically accomplishes this through the definition of final state objects called *jets*, which provide a theoretical description of collimated outgoing QCD radiation. As a result of colour confinement, one can not observe the outgoing partons directly in detectors, but only the hadronic decay products produced by the hard scattering process. A procedure is required in order to match the outgoing hadronic signal

¹⁵If one excludes luminosity uncertainty which can be removed through distribution normalisation.

to theory, which is accomplished by *jet algorithms*, exploiting the fact that the kinematics of hard jets are largely described by partonic emissions at the hard cross-section level. The resultant jets must be defined in an IR-safe manner in order to avoid ill-defined theory predictions arising from soft or collinear final state radiation which cannot be resolved by experiment.

The jet algorithms which see the most widespread use at the LHC and beyond belong to the class of sequential recombination algorithms. These use an iterative method to experimentally “undo” partonic fragmentation and reconstruct the partonic emissions from the hard scattering process. They rely on two geometric distance measures in order to define jets. These distances are d_{ij} , the distance between final state entities i and j , which may be partons or pseudojets, and d_{iB} , the distance between a final state entity i and the beam direction. The most common definition of these distance measures is [62]:

$$d_{ij} = \min(k_{T,i}^{2p}, k_{T,j}^{2p}) \frac{\Delta_{ij}^2}{R^2}, \quad d_{iB} = k_{T,i}^{2p} \quad \Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2, \quad (1.4.1)$$

where y_i , ϕ_i , and $k_{T,i}$ are the rapidity, azimuthal angle and transverse momentum of object i , R is a radius parameter of the jet algorithm and p is a parameter which gives the relative clustering power of the energy with respect to Δ_{ij}^2 . All of the above quantities are manifestly invariant under boosts along the beam axis.

These two distance measures can be used to define an iterative procedure, which identifies the smallest distance measure $d \in \{d_{ij}\} \cup \{d_{iB}\}$ considering all pairs of final state objects i, j . If $d \in \{d_{ij}\}$, then one clusters objects i and j to form a protojet, whereas if $d \in \{d_{iB}\}$, i is labelled as a jet and removed from the set of final state objects. This procedure is then repeated until there are no remaining objects that have not been identified as jets. There are 3 common choices of p which define algorithms in widespread use:

- $p = -1$, the anti- k_T algorithm [62],
- $p = 0$, the Cambridge/Aachen algorithm [63, 64],

- $p = +1$, the Durham, or k_T algorithm [65, 66].

Of these the anti- k_T algorithm is the most common, having been adopted by the experiments at the LHC. The choice $p = -1$ preferentially clusters soft partons before hard partons, resulting in conical jets with a circular profile in the y, ϕ plane when hard partons are widely separated, which is useful experimentally for the understanding of detector effects. The effective radius of these cones is determined by the R parameter, with a large value $R \sim 1$ resulting in fewer, larger jets and small values $R \rightarrow 0$ giving more, smaller jets with a greater resolution of soft effects.

Once a jet algorithm has been chosen, one can then use requirements on the number of jets present to define hard scattering processes with partonic final states in an IR safe manner. At LO, an outgoing parton in the hard process generally corresponds to a jet¹⁶, whereas more partons become available for clustering through real emissions at higher orders.

¹⁶This can be bypassed for certain cases where fiducial cuts on non-QCD final states can be used to render the total cross section finite. One example, which we consider in Section 3.2, is the use of a minimum transverse momentum of a vector boson, which implicitly requires at least one partonic emission, even if not defined by a jet algorithm.

Chapter 2

Antenna Subtraction

In this chapter we will discuss the method of antenna subtraction and its application to next-to-next-to-leading order (NNLO) calculations in QCD. We begin in Section 2.1 with an introduction to colour ordered amplitudes and matrix elements, before discussing their behaviour in unresolved infrared (IR) limits in Section 2.2. We introduce antenna functions in Section 2.3 and their use in subtraction term construction in Section 2.4. Finally in Sections 2.5 and 2.6 we discuss particular points relating to the application of antenna subtraction to charged-current processes and their implementation within the NNLOJET framework, in particular W^\pm boson production both with and without a jet in hadron-hadron collisions, and single/dijet production in charged-current lepton-hadron scattering.

2.1 Colour Ordering

In order to discuss the cancellation of IR divergences used in antenna subtraction, we must first introduce colour ordering. This defines a decomposition of the matrix elements into subamplitudes on to which higher multiplicity amplitudes factorise in IR singular limits [67–72].

The colour dependence of any given QCD amplitude can be decomposed into a linear combination of colour factors $\mathcal{F}(\{T, f\})$, dependent only on the $SU(3)$ generators T

and structure constants f multiplied by colour-ordered partial amplitudes $|M^i\rangle$:

$$|\mathcal{M}\rangle = \sum_i \mathcal{F}_i(\{T, f\}) |M^i\rangle, \quad (2.1.1)$$

where $|M^i\rangle$ is a purely kinematic object. This allows us to decompose the squared matrix element

$$|\mathcal{M}|^2 = \langle \mathcal{M} | \mathcal{M} \rangle \quad (2.1.2)$$

into sums of products of colour-ordered amplitudes [73]:

$$|\mathcal{M}|^2 = \sum_{i,j} \mathcal{F}_i^\dagger(\{T, f\}) \mathcal{F}_j(\{T, f\}) \langle M^i | M^j \rangle. \quad (2.1.3)$$

The colour algebra within the factor $\mathcal{F}_i^\dagger(\{T, f\}) \mathcal{F}_j(\{T, f\})$ can be evaluated straightforwardly through iterative use of the Fierz identity:

$$T_{ij}^a T_{kl}^a = \frac{1}{2} \left(\delta_{il} \delta_{jk} - \frac{1}{N} \delta_{ij} \delta_{kl} \right), \quad (2.1.4)$$

alongside the rearrangement of (1.1.12),

$$f^{abc} = -2i \left(T_{ij}^a T_{jk}^b T_{ki}^c - T_{ij}^a T_{jk}^c T_{ki}^b \right). \quad (2.1.5)$$

Within these colour-decomposed amplitudes there exists a strict ordering of partons determined by the colour prefactor and resultant colour flow which only allows inverse powers of momentum invariants between partons adjacent in the ordering. As a result, unresolved singular IR limits of real amplitudes can only occur between these colour-connected partons, allowing the treatment of singular behaviour to be decomposed into that of colour-adjacent parton pairs.

If we take $i = j$ in (2.1.3), we isolate only the dominant leading colour (LC) terms and the matrix elements become sums of squares of colour-ordered partial amplitudes:

$$\begin{aligned} |\mathcal{M}_n|^2 \Big|_{\text{LC}} &= \sum_i \mathcal{F}_i^\dagger(\{T, f\}) \mathcal{F}_i(\{T, f\}) \langle M^i | M^i \rangle \\ &= \sum_i |\mathcal{F}_i(\{T, f\})|^2 |M^i|^2. \end{aligned} \quad (2.1.6)$$

For states with up to 6 external partons, the sub-leading colour (SLC) terms can also be written exclusively as squares of coherent sums of amplitudes $\widetilde{M}_n = \sum_{\text{perm}} |M^i\rangle$, where the sum over permutations is constructed such that the three- and four- gluon vertex contributions cancel for one or more gluons using the identity from (2.1.5). Removing these self-interactions means that the gluon becomes effectively Abelian and no longer exhibits collinear limits with other gluons in the squared sub-leading matrix elements, analogous to the IR divergence structure of the photon.

As an example, we take the case of a matrix element containing a quark pair and three gluons at tree level. Here we have the colour-ordered amplitude decomposition

$$|\mathcal{M}\rangle = \sum_{i,j,k \in S_3} (T^{a_i} T^{a_j} T^{a_k}) M_5^0(1_q, i_g, j_g, k_g, 2_{\bar{q}}). \quad (2.1.7)$$

The LC term can be expressed as

$$\langle \mathcal{M} | \mathcal{M} \rangle \Big|_{LC} = N^2 \sum_{i,j,k \in S_3} |M_5^0(1_q, i_g, j_g, k_g, 2_{\bar{q}})|^2, \quad (2.1.8)$$

the SLC term with Abelian gluon j can be expressed as

$$\begin{aligned} \langle \mathcal{M} | \mathcal{M} \rangle \Big|_{SLC} = & - \sum_{i,j,k \in S_3} |M_5^0(1_q, i_g, j_g, k_g, 2_{\bar{q}}) \\ & + M_5^0(1_q, i_g, k_g, j_g, 2_{\bar{q}}) \\ & + M_5^0(1_q, j_g, i_g, k_g, 2_{\bar{q}})|^2 \end{aligned} \quad (2.1.9)$$

and the SSLC term with three Abelian gluons can be expressed as

$$\langle \mathcal{M} | \mathcal{M} \rangle \Big|_{SSLC} = \frac{N^2 + 1}{N^2} \left| \sum_{i,j,k \in S_3} M_5^0(1_q, i_g, j_g, k_g, 2_{\bar{q}}) \right|^2. \quad (2.1.10)$$

In the SLC colour term gluon j is Abelian, having been commuted through the non-Abelian gluons i and k . This contribution comes from terms in the colour algebra where two of the generators are contracted to become the identity operator in colour space through the Fierz identity, such that they then commute. This is in direct analogy to the generators in Abelian theories such as QED in which we have photons rather than gluons. For the SSLC term, all of the generators in the colour

algebra are contracted to form identity matrices in colour space and thus all gluons are effectively Abelian, or photon-like.

This discussion can be repeated in the case of loop matrix elements, with the exception that at one-loop the squared matrix elements are interferences of one-loop and tree-level amplitudes,

$$\langle \mathcal{M} | \mathcal{M} \rangle_{1\text{-loop}} = \langle \mathcal{M}^{(1)} | \mathcal{M}^{(0)} \rangle + \langle \mathcal{M}^{(0)} | \mathcal{M}^{(1)} \rangle \quad (2.1.11)$$

and at two loops we have both two-loop \times tree and one-loop squared terms:

$$\langle \mathcal{M} | \mathcal{M} \rangle_{2\text{-loop}} = \langle \mathcal{M}^{(2)} | \mathcal{M}^{(0)} \rangle + \langle \mathcal{M}^{(0)} | \mathcal{M}^{(2)} \rangle + \langle \mathcal{M}^{(1)} | \mathcal{M}^{(1)} \rangle. \quad (2.1.12)$$

For each of these loop contributions, one can also understand the associated explicit poles in ε through the colour-decomposed Catani operators, which can be applied to derive the singularity structure at the virtual level for arbitrary processes at up to two-loop level [74]. This colour decomposition is a particularly natural form for these explicit poles, and serves to reflect the singularity structure in colour-ordered real-emission matrix elements as required by the KLN theorem.

2.2 Behaviour of Real Emission Contributions in Singular Limits

We are now equipped to discuss in more detail the behaviour of squared colour-ordered matrix elements in singular limits. From here it will be assumed that the arguments of each matrix element are colour-ordered, adopting notation in which the ordering of the arguments of the matrix element squared matches their colour ordering.

In the IR divergent limits of the cross section integrand, we observe that the phase space itself factorises, to give

$$d\Phi_n(\{p_n\}) \xrightarrow{\text{div}} d\Phi_m(\{p_m\}) \cdot d\Phi_{n-m}(\{p_{n-m}\}) \quad (2.2.1)$$

when $n - m$ external partons go unresolved. To construct $\{p_m\}$ from $\{p_n\}$ requires a set of momenta in this reduced phase space, produced through a generic set of momentum mappings which relate the momentum set of the full n -particle phase space point to the momenta in the divergent limit in which the phase space factorises. There is a freedom of choice in these mappings, however they must satisfy several criteria, principally conservation of momentum, preservation of on-shell particles and Lorentz symmetries.

The antenna mappings [30, 47] provide such a set of generic mappings that simultaneously parameterise soft and collinear limits. Each mapping contains two hard radiators adjacent in the colour ordering of the unresolved partons and where mapped momenta compensate for any recoil of the unresolved parton(s) for momentum conservation purposes. If a hard radiator is in the initial state, this recoil adjustment takes the form of a rescaling along the beam axis. We adopt a notation for $n + 1 \rightarrow n$ mappings when parton j goes unresolved between hard radiators i, k such that the resultant mapped momenta are labelled as

$$(i, j, k) \rightarrow (\widetilde{(ij)}, \widetilde{(jk)}). \quad (2.2.2)$$

The two composite momenta $\widetilde{(ij)}, \widetilde{(jk)}$ reside in the n -parton phase space. Similarly for the $n + 2 \rightarrow n$ mappings used for double-real contributions:

$$(i, j, k, l) \rightarrow (\widetilde{(ijk)}, \widetilde{(jkl)}). \quad (2.2.3)$$

In the same manner as the phase space itself, the matrix elements also factorise in the divergent limits, to give a matrix element lying in the reduced multiplicity phase space $d\Phi(\{p_m\})$. This reduced matrix element is multiplied by a parameterisation of the divergence in terms of soft and/or splitting functions, lying in the remainder of the phase space, $d\Phi(\{p_{n-m}\})$. The exact factorisation behaviour is specific to the form of IR singularity being considered, as we will now see for the singular structures relevant to the calculation of QCD corrections to NNLO.

Single Soft

We begin with the limits that occur for the first time at NLO as outlined in Section 1.3.4, where at most one final state parton can become unresolved. The first of these limits is the soft limit, in which the momentum p_j of an external gluon¹ $\rightarrow 0$. In this region the matrix element factorises into a reduced squared matrix element multiplied by an Eikonal factor S_{ijk} :

$$|M_{n+1}^0|^2(\dots, i, j, k, \dots) \xrightarrow{j \text{ soft}} S_{ijk} |M_n^0|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots), \quad (2.2.4)$$

where S_{ijk} is given by

$$S_{ijk} = \frac{2s_{ik}}{s_{ij}s_{jk}}. \quad (2.2.5)$$

Here $|M_n^0|^2$ is now a function of mapped composite momenta \widetilde{ij} , \widetilde{jk} , which in the perfectly soft limit $|p_j| = 0$ become

$$\widetilde{ij} = i; \quad \widetilde{jk} = k, \quad (2.2.6)$$

where the composite momenta \widetilde{ij} , \widetilde{jk} inherit the partonic identities of their parent radiator (i.e. if i is a quark, then \widetilde{ij} corresponds to a quark momentum in the reduced matrix element since it is j which goes unresolved).

Single Collinear

The second IR-divergent configuration that can occur at NLO is the configuration in which two colour-connected external partons are emitted parallel to one another. In this case the real emission matrix element factorises into a reduced matrix element multiplied through by an object P known as a splitting function:

$$|M_{n+1}^0|^2(\dots, i, j, k, \dots) \xrightarrow{i \parallel j} \frac{P^{(0)}_{ij \rightarrow \widetilde{ij}}}{s_{ij}} |M_n^0|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots) + \text{azimuthal terms}, \quad (2.2.7)$$

¹Recall that there is no divergence when a single external (anti-)quark becomes soft for the conservation of EM current.

where we have partons i, j collinear, with hard radiator k absorbing the momentum recoil. Here $P_{ij}^{(l)}(z)$ is the flavour dependent l -loop *spin averaged* splitting function, where the flavour of collinear partons determines the exact form of P as well as the flavour of the outgoing composite parton. The parameter z is the fraction of the momentum of (\widetilde{ij}) carried by parton i , the explicit dependence on which we have dropped above for brevity. At LO ($l = 0$), the splitting function can be interpreted as the probability of finding parton i in parton j with fraction z of the longitudinal momentum of parent parton j , assuming the transverse momentum of i with respect to j is much smaller than the factorisation scale.

The flavour of the resultant composite parton (\widetilde{ij}) is mapped for $i, j \rightarrow (\widetilde{ij})$ as follows:

$$q\bar{q} \rightarrow g \quad \bar{q}g \rightarrow \bar{q} \quad gg \rightarrow g, \quad (2.2.8)$$

which enforces the conservation of quantum numbers. The spin-averaging of the splitting function leads to residual azimuthal terms which are absent if one works with spin-dependent splitting functions that directly interact with the spin structure of the matrix elements. We use the spin-averaged version here as it is the form of the limit reproduced by antenna functions within the framework of antenna subtraction. If one averages over the azimuthal angle about the collinear axis, these contributions exactly cancel.

Double Soft

We now turn our attention to the limits which first occur in the double real contribution at NNLO, in which two partons both go unresolved. For each of the possible combinations of soft or collinear limits, the factorisation pattern depends on the colour-connections of the two unresolved partons. When the two unresolved partons are adjacent in the colour ordering (colour-connected), the factorisation behaviour in the limits is fundamentally different to that at NLO, and requires the introduction of new soft and splitting functions. Where the two unresolved partons

are colour separated, they can be viewed as iterations of NLO unresolved limits and the treatment simplifies.

The first of these limits we consider is the double soft, in which two partons go soft simultaneously. Where these two partons are colour connected, a new four-parton soft S_{ijkl} function is required to describe the IR behaviour [75]:

$$|M_{n+2}^0|^2(\dots, i, j, k, l, \dots) \xrightarrow{j, k \text{ soft}} S_{ijkl} |M_n^0|^2(\dots, \widetilde{(ijk)}, \widetilde{(jkl)}, \dots). \quad (2.2.9)$$

For the first time, one can encounter configurations where j and k are a colour-connected quark-antiquark pair, which can go unresolved to form a composite gluon. All of the other flavour identifications can be found through iterations of the NLO case.

For the remaining double soft limits, where the two soft gluons are not colour connected, the factorisation pattern is that of two single unresolved limits:

$$|M_{n+2}^0|^2(\dots, i, j, k, \dots, l, m, n, \dots) \xrightarrow{j, m \text{ soft}} S_{ijk} S_{lmn} |M_n^0|^2(\dots, \widetilde{(ij)}, \widetilde{(jk)}, \dots, \widetilde{(lm)}, \widetilde{(mn)}, \dots). \quad (2.2.10)$$

Soft Collinear

Alongside the double soft limit, it is also possible at NNLO to generate soft collinear limits in which one parton goes soft at the same time that a pair goes collinear. For the case in which the soft parton is colour connected to the collinear pair, the double real matrix element factorises into the product of a soft-collinear $S_{i;jkl}$ factor and a single collinear splitting function [76]

$$|M_{n+2}^0(\dots, i, j, k, l, \dots)|^2 \xrightarrow{j \text{ soft}, k \parallel l} S_{i;jkl} \frac{P_{kl \rightarrow (kl)}^{(0)}}{s_{kl}} |M_n^0|^2(\dots, \widetilde{(ijk)}, \widetilde{(jkl)}, \dots) + \text{azimuthal terms}. \quad (2.2.11)$$

As before, in the non-colour connected case, we have a set of iterated single-soft and

single-collinear limits

$$\begin{aligned}
 & |M_{n+2}^0(\dots, i, j, k, \dots, l, m, n, \dots)|^2 \\
 & \xrightarrow{j \text{ soft}, l \parallel m} S_{ijk} \frac{P_{lm \rightarrow (\widetilde{lm})}^{(0)}}{s_{lm}} |M_n^0|^2 \left(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots, (\widetilde{lm}), (\widetilde{mn}), \dots \right) \\
 & + \text{azimuthal terms.}
 \end{aligned} \tag{2.2.12}$$

As we only encounter a single-soft and a single-collinear limit, the flavour identification of the composite partons can be treated as iterations of the NLO case.

Multi-collinear

For three colour-connected unresolved partons which are all simultaneously collinear, one observes a “triple collinear” limit where the three partons are produced along the same axis in phase space. In this instance, one must introduce triple collinear splitting functions $P_{ijk \rightarrow (\widetilde{ijk})}^{(0)}$ to describe the IR behaviour in the singular limit. Like the single collinear splitting functions, these have different forms dependent on the unresolved partons’ flavours. The factorisation in this case behaves as [76]

$$\begin{aligned}
 & |M_{n+2}^0|^2(\dots, i, j, k, l, \dots) \xrightarrow{i \parallel j \parallel k} P_{ijk \rightarrow (\widetilde{ijk})}^{(0)} |M_n^0|^2 \left(\dots, (\widetilde{ijk}), (\widetilde{kl}), \dots \right) \\
 & + \text{azimuthal terms,}
 \end{aligned} \tag{2.2.13}$$

and the flavour of the composite partons as

$$\begin{aligned}
 & ggg \rightarrow g \quad q\bar{q}g \rightarrow g \quad q\bar{q}\widetilde{g} \rightarrow \widetilde{g} \quad \overleftarrow{q}gg \rightarrow \overleftarrow{q} \\
 & \overleftarrow{q}\widetilde{g}\widetilde{g} \rightarrow \overleftarrow{q} \quad q\bar{q}q \rightarrow q \quad q\bar{q}'q' \rightarrow q,
 \end{aligned} \tag{2.2.14}$$

where we identify Abelian-like gluons that occur at SLC as \widetilde{g} . The explicit forms of these triple-collinear splitting functions can be found in [46].

When the unresolved partons go collinear with different hard radiators, the “double collinear” limit, the factorisation pattern behaves as an iterated form of the single

unresolved limit:

$$|M_{n+2}^0|^2(\dots, i, j, \dots, k, l, \dots) \xrightarrow{i||j \quad k||l} \frac{P^{(0)}_{ij \rightarrow \widetilde{ij}}}{s_{ij}} \frac{P^{(0)}_{kl \rightarrow \widetilde{kl}}}{s_{kl}} |M_n^0|^2(\dots, \widetilde{ij}, \dots, \widetilde{kl}, \dots) + \text{azimuthal terms.} \quad (2.2.15)$$

One Loop

Having introduced the new factorisation properties that occur in the double real contributions at NNLO, we now briefly discuss the factorisation of the one-loop matrix elements in the real-virtual contribution. As at NLO, we only have two possible IR limits, the single soft and the single collinear. However, due to the structure of the one-loop contributions (see (2.1.11)) this requires the introduction of new universal singular functions. One must do this in order to match the (loop \times tree) + (tree \times loop) structure of the matrix element, where one can have contributions proportional to either the loop or tree part.

For the single-soft limit of a one-loop matrix element, the one-loop soft function $S_{ijk}^{(1)}$ is required, which provides a contribution containing explicit ε poles factorising on to a tree-level matrix element:

$$|M_{n+1}^1|^2(\dots, i, j, k, \dots) \xrightarrow{j \text{ soft}} S_{ijk}^{(0)} |M_n^1|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots) + S_{ijk}^{(1)}(\varepsilon) |M_n^0|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots). \quad (2.2.16)$$

For the collinear limit, the one-loop splitting function $P_{ij \rightarrow \widetilde{ij}}^{(1)}$ is introduced in the same manner:

$$|M_{n+1}^1|^2(\dots, i, j, k, \dots) \xrightarrow{i||j} \frac{P^{(0)}_{ij \rightarrow \widetilde{ij}}}{s_{ij}} |M_n^1|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots) + \frac{P^{(1)}_{ij \rightarrow \widetilde{ij}}(\varepsilon)}{s_{ij}} |M_n^0|^2(\dots, \widetilde{ij}, \widetilde{jk}, \dots) + \text{azimuthal terms,} \quad (2.2.17)$$

where we once again see the emergence of azimuthal terms, and the $P^{(1)}$ functions

contain explicit poles in the dimensional regulator similar to the one-loop soft functions.

2.3 Antenna Functions

Having established the behaviour of matrix elements in IR limits for all contributions up to NNLO, one can now define functions that can replicate this divergent structure for the real contributions. If this can be done, one can construct local counterterms, ‘subtraction terms’, in order to keep the integrand finite. These functions must also be integrable over the phase space of the unresolved parton(s) $d\Phi_{(n-m)\in[1,2]}$ such that the divergent contributions can be reintroduced in the virtual contributions for closure of the subtraction method. For the method of antenna subtraction, these functions are (somewhat unsurprisingly) called antenna functions, and are constructed using ratios of n -parton matrix elements to $(n - m)$ -parton reduced matrix elements. These matrix elements are calculated using the simplest processes displaying the required divergent structures.

Breaking down the possible pairs of hard radiators, we have (alongside the process from which the antenna functions are derived):

- $q\bar{q}$, derived using matrix elements from the decay of a virtual photon into a quark-antiquark pair alongside additional QCD radiation [77],
- qg , derived using matrix elements from the decay of a heavy neutralino into a gluon and gluino alongside additional QCD radiation in SUSY [78],
- gg , derived using matrix elements from the decay of a Higgs boson into a gluon pair alongside additional QCD radiation in Higgs effective field theory [79].

The qg terms must be taken from non-SM processes, as no $2 \rightarrow 3$ -, 4 -type processes with the QCD content qgg , $qggg$ occur within the SM.

Hard Radiators	M_0^3 Parton Content	X_3^0 class	X_3^1 class
$q\bar{q}$	$qg\bar{q}$	A_3^0	$A_3^1, \tilde{A}_3^1, \hat{A}_3^1$
qg	qgg $qq'\bar{q}'$	D_3^0 E_3^0	D_3^1, \widehat{D}_3^1 $E_3^1, \tilde{E}_3^1, \hat{E}_3^1$
gg	ggg $gq\bar{q}$	F_3^0 G_3^0	F_3^1, \hat{F}_3^1 $G_3^1, \tilde{G}_3^1, \hat{G}_3^1$

Table 2.1: The classes of X_3^0 and X_3^1 antenna functions, broken down by hard radiator and radiative process. For the loop antenna functions, the tilde variants correspond to sub-leading colour functions and the hatted variants correspond to antennas containing a closed quark loop [48].

2.3.1 Unintegrated Antenna Functions

Single Unresolved X_3^0

At NLO there is only one class of antenna functions, the X_3^0 , which contains all tree level single-unresolved limits. For each crossing of hard radiators between initial and final states, and for each flavour of hard radiators, these are defined as [46]

$$X_3^0(i, j, k) = \mathcal{S} \frac{|M_3^0(i, j, k)|^2}{|M_2^0(\widetilde{(ij)}, \widetilde{(jk)})|^2}, \quad (2.3.1)$$

where \mathcal{S} is a symmetry factor to account for differences in the number of identical final state partons between $|M_3^0|$ and $|M_2^0|$. The full set of three-parton antenna functions is listed in Table 2.1 [46, 48].

We can show schematically how this can be used to replicate the divergences in some arbitrary tree-level matrix element $|\bar{M}_{n+1}^0|^2$. If we consider some single-unresolved IR limit contained in $|\bar{M}_{n+1}^0|^2$, we have in the divergent limit:

$$|\bar{M}_{n+1}^0|^2 \xrightarrow{\text{div}} A^0 \times |\bar{M}_n^0|^2, \quad (2.3.2)$$

where $A^0 \in \{P^{(0)}, S\}$ is either a tree-level splitting function or a tree-level soft factor (neglecting azimuthal terms). One can then construct a term with the same behaviour in the singular limit if X_3^0 contains the same divergence. This term takes

the form $X_3^0 |\bar{M}_n^0|^2$, and in the divergent limit behaves as

$$\begin{aligned} X_3^0 |\bar{M}_n^0|^2 &= \frac{|M_3^0|^2}{|M_2^0|^2} |\bar{M}_n^0|^2 \\ &\xrightarrow{\text{div}} \frac{A^0 \times |M_2^0|^2}{|M_2^0|^2} |\bar{M}_n^0|^2 = A^0 \times |\bar{M}_n^0|^2. \end{aligned} \quad (2.3.3)$$

Due to the universality of the IR divergences summarised in A^0 , $X_3^0 |\bar{M}_n^0|^2$ exhibits exactly the same behaviour as $|\bar{M}_{n+1}^0|^2$ and illustrates that X_3^0 functions can be used to regulate the limits associated with $|\bar{M}_{n+1}^0|$. Outside of these limits, the subtraction term remains finite, and so through combinations of these X_3^0 functions, the full single unresolved divergence structure of any given matrix element can be constructed from process-independent functions. Azimuthal terms can be handled through an angular averaging of the $n + 1$ particle phase space, under which they cancel.

Double Unresolved X_4^0

At NNLO, one must introduce a new class of antenna functions which contain the divergent structures associated with doubly unresolved limits. These are defined analogously to the X_3^0 functions, as [46]

$$X_4^0(i, j, k, l) = \mathcal{S} \frac{|M_4^0(i, j, k, l)|^2}{|M_2^0(\widetilde{(ijk)}, \widetilde{(jkl)})|^2}. \quad (2.3.4)$$

The full set of four-parton antenna functions are listed in Table 2.2. As in the NLO case, we can show schematically how the unresolved behaviour of an arbitrary matrix element $|\bar{M}_{n+2}^0|^2$ can be reconstructed using terms of the form $X_4^0 |\bar{M}_n^0|^2$:

$$\begin{aligned} X_4^0 |\bar{M}_n^0|^2 &= \frac{|M_4^0|^2}{|M_2^0|^2} |\bar{M}_n^0|^2 \\ &\xrightarrow{\text{div}} \frac{A^0 \times |M_2^0|^2}{|M_2^0|^2} |\bar{M}_n^0|^2 = A^0 \times |\bar{M}_n^0|^2. \end{aligned} \quad (2.3.5)$$

Combinations of these terms can be used to regulate the doubly divergent behaviour of an integrand in a similar manner to the X_3^0 for the single unresolved limits [46].

Hard Radiators	$ M_0^4 $ Parton Content	X_0^4 class
$q\bar{q}$	$qgg\bar{q}$ $qq'\bar{q}'\bar{q}$ $qq\bar{q}\bar{q}$	A_0^4, \tilde{A}_0^4 B_0^4 C_0^4
qg	$qggg$ $qq'\bar{q}'g$	D_0^4 E_0^4, \tilde{E}_0^4
gg	$gggg$ $gq\bar{q}g$ $q\bar{q}q'\bar{q}'$	F_0^4 G_0^4, \tilde{G}_0^4 H_0^4

Table 2.2: The different classes of unintegrated X_0^4 antenna functions, broken down by hard radiator and radiative process. The tilde variants correspond to sub-leading colour contributions [48].

One Loop X_3^1

The only class of antenna functions it now remains to define are those which recreate the singular limits of one-loop matrix elements. Due to the extra term in the one-loop factorisation formulae in Eqns 2.2.16 and 2.2.17, these are defined in a different manner to the X_3^0 and X_4^0 functions [46]:

$$X_3^1(i, j, k) = \mathcal{S} \frac{|M_3^1(i, j, k)|^2}{|M_2^0(\widetilde{(ij)}, \widetilde{(jk)})|^2} - X_3^0 \frac{|M_2^1(\widetilde{(ij)}, \widetilde{(jk)})|^2}{|M_2^0(\widetilde{(ij)}, \widetilde{(jk)})|^2}, \quad (2.3.6)$$

which comes naturally from the definition

$$|M_3^1(i, j, k)|^2 = X_3^1(i, j, k) |M_2^0(\widetilde{(ij)}, \widetilde{(jk)})|^2 + X_3^0(i, j, k) |M_2^1(\widetilde{(ij)}, \widetilde{(jk)})|^2, \quad (2.3.7)$$

analogous to the factorisation of the one-loop matrix element in IR limits. These one-loop antenna functions are listed in Table 2.1 [48].

If we denote the one-loop splitting/soft functions as A^1 , which these antenna functions are designed to capture, we can once again show schematically how the X_3^1 functions can be used to reconstruct the second term in both (2.2.16) and (2.2.17):

$$\begin{aligned} X_3^1 |\bar{M}_n^0|^2 &= \left(\frac{|M_3^1|^2}{|M_2^0|^2} - \frac{|M_3^0|^2}{|M_2^0|^2} \frac{|M_2^1|^2}{|M_2^0|^2} \right) |\bar{M}_n^0|^2 \\ &= \frac{1}{|M_2^0|^2} \left(|M_3^1|^2 - |M_3^0|^2 \frac{|M_2^1|^2}{|M_2^0|^2} \right) |\bar{M}_n^0|^2 \end{aligned}$$

$$\begin{aligned}
& \xrightarrow{\text{div}} \frac{1}{|M_2^0|^2} \left(A^0 \times |M_2^1|^2 + A^1 \times |M_2^0|^2 - A^0 \times |M_2^0|^2 \frac{|M_2^1|^2}{|M_2^0|^2} \right) |\bar{M}_n^0|^2 \\
& = A^1 \times |\bar{M}_n^0|^2.
\end{aligned} \tag{2.3.8}$$

In combination with X_3^0 functions to subtract the first term in (2.2.16) and (2.2.17), one can then reconstruct the divergent behaviour of one-loop matrix elements.

Soft Function

We will also encounter the need to subtract explicit wide-angle single-soft limits at NNLO where 5 or more partons are present at double real level. These cannot be directly replicated using antenna functions alone. In this case one also requires the three-parton Eikonal factor S_{ijk} from Eqn. (2.2.5), which can be used to reconstruct directly single soft limits using Eqn. (2.2.4).

2.3.2 Integrated Antenna Functions

For each initial-final crossing of the unintegrated antenna functions listed in the previous sections, the analytic integration has been performed over the unresolved phase space [47, 80–83] for the case of massless partons. As the phase space mappings are different for each crossing of the hard radiators within the initial and final states, the integration over the antenna phase space must be performed separately for each crossing. For each unintegrated antenna, we adopt the standard notation in which the integrated antennas are labelled with the calligraphic form of the unintegrated antenna, with further flavour labelling if there is remaining ambiguity. The crossing can be seen from the function arguments, in which mapped momenta are labelled with a bar, and initial state momenta are labelled with hats. Whilst we refrain from listing all integrals separately here, we give the example of the initial-final integration of a generic antenna function $X_n^{(l)}$, which proceeds as [47]:

$$\mathcal{X}_n^{(l)}(\hat{1}, i) = \left[\frac{8\pi^2}{\bar{C}(\varepsilon)} \right]^{n-2} \int d\Phi_{n-1} \delta(x_1 - \hat{x}_1) \frac{Q^2}{2\pi} X_n^{(l)}(\hat{1}, i, \dots), \tag{2.3.9}$$

where the ellipses correspond to the remaining final state parton content of the unintegrated antenna, $Q = \sum p_i$, the sum of all momenta, x_1 is the Bjorken- x of the initial state parton and \hat{x}_1 is the rescaling factor along the beam axis for the initial state composite parton involved in the reduced momentum set. Using this technique the soft Eikonal functions have also been integrated in the final-final [84] and initial-final [81] configurations, which is sufficient for all parton multiplicities.

2.4 Subtraction Term Construction

As each antenna function only accounts for a subset of the possible IR limits in a general process, one must assemble subtraction terms containing multiple antenna functions in order to keep the integrand finite across the entirety of $d\Phi$. In this section we will outline the assembly of these terms at cross-section level, which requires particular care at NNLO due to the appearance of spurious IR divergences.

2.4.1 Mass Factorisation Terms

We begin by more formally defining the initial-state IR terms which are absorbed into the PDF, as outlined in Section 1.3.4, working as usual in the $\overline{\text{MS}}$ scheme.

The modification of the PDFs beyond LO takes place by treating the PDFs in (1.3.3) as ‘bare PDFs’ denoted \mathbf{f}^0 , analogous to bare quantities before renormalisation. However, in this case, as the physical PDFs \mathbf{f}_i can be treated as a vector in flavour space with indices $i \in \{g, q, \bar{q}\}$, $q \in \{u, d, c, s, b\}$, we use a factorisation kernel $\mathbf{\Gamma}^{-1}$, with inverse $\mathbf{\Gamma}$ in matrix form taking the place of the multiplicative Z factor in renormalisation:

$$\begin{aligned}\mathbf{f}^0(x) &= [\mathbf{f}(x, \mu_F) \otimes \mathbf{\Gamma}^{-1}(\mu_F)] \\ \mathbf{f}(x, \mu_F) &= [\mathbf{f}^0(x) \otimes \mathbf{\Gamma}(\mu_F)].\end{aligned}\tag{2.4.1}$$

The convolution operator \otimes is defined as

$$\left[\mathbf{f}(x, \mu_F) \otimes \mathbf{\Gamma}^{-1}(\mu_F) \right] = f_i^0(z, \mu_F) = \int dx dy \delta(z - xy) f_j(x, \mu_F) \Gamma_{ji}^{-1}(y, \mu_F). \quad (2.4.2)$$

At leading order, these factorisation kernels are simply the identity matrix. However, they can be expanded perturbatively in α_s alongside the inverse:

$$\begin{aligned} \mathbf{\Gamma} &= \mathbf{I} + \left(\frac{\alpha_s}{2\pi} \right) \mathbf{\Gamma}^1 + \left(\frac{\alpha_s}{2\pi} \right)^2 \mathbf{\Gamma}^2 + \mathcal{O}(\alpha_s^3) \\ \mathbf{\Gamma}^{-1} &= \mathbf{I} - \left(\frac{\alpha_s}{2\pi} \right) \mathbf{\Gamma}^1 - \left(\frac{\alpha_s}{2\pi} \right)^2 \left[\mathbf{\Gamma}^2 - [\mathbf{\Gamma}^1 \otimes \mathbf{\Gamma}^1] \right] + \mathcal{O}(\alpha_s^3). \end{aligned} \quad (2.4.3)$$

When these are substituted into the cross section master formula, they generate extra contributions to the cross section beyond LO in α_s . These contain explicit poles in ε , and appear in addition to the contributions in the expansion of $d\hat{\sigma}_{ij}$. One can rewrite these extra terms as contributions to the partonic cross section by considering the cross section in terms of the renormalised PDFs:

$$\begin{aligned} d\sigma &= \mathbf{f}_a^0 \cdot d\hat{\sigma} \cdot \mathbf{f}_b^0 \\ &= \mathbf{f}_a \otimes \mathbf{\Gamma}^{-1} \cdot d\hat{\sigma} \cdot \mathbf{\Gamma}^{-1} \otimes \mathbf{f}_b \\ &= \mathbf{f}_a \cdot d\hat{\sigma}^{FAC} \cdot \mathbf{f}_b, \end{aligned} \quad (2.4.4)$$

where the factorised cross section is denoted

$$d\hat{\sigma}^{FAC} = \mathbf{\Gamma}^{-1} \cdot d\hat{\sigma} \cdot \mathbf{\Gamma}^{-1}. \quad (2.4.5)$$

In full, this gives

$$\begin{aligned} d\sigma(x_1, x_2) &= \int da_1 db_1 \delta(x_1 - a_1 b_1) f_j(a_1, \mu_F) \Gamma_{ji}^{-1}(b_1, \mu_F) \\ &\quad d\hat{\sigma}_{ik} \int da_2 db_2 \delta(x_2 - a_2 b_2) \Gamma_{kl}^{-1}(b_2, \mu_F) f_l(a_2, \mu_F), \end{aligned} \quad (2.4.6)$$

where x_1 and x_2 are the aforementioned momentum fractions of the two partons at LO. The factorised cross section can then be written as the sum of its component contributions:

$$d\hat{\sigma}^{FAC} = d\hat{\sigma} + d\hat{\sigma}^{MF}, \quad (2.4.7)$$

where $d\hat{\sigma}^{MF}$ is the mass factorisation contribution that appears at NLO and beyond, which can itself be expanded order by order in α_s :

$$d\hat{\sigma}^{MF} = \left(\frac{\alpha_s}{2\pi}\right) d\hat{\sigma}_{NLO}^{MF} + \left(\frac{\alpha_s}{2\pi}\right)^2 d\hat{\sigma}_{NNLO}^{MF} + \mathcal{O}(\alpha_s^3), \quad (2.4.8)$$

with the NLO and NNLO contributions given by

$$d\hat{\sigma}_{ij}^{MF,NLO} = - \left[\Gamma_{ik}^1 \otimes d\sigma_{kj}^{LO} + d\sigma_{ik}^{LO} \otimes \Gamma_{kj}^1 \right] \quad (2.4.9)$$

$$\begin{aligned} d\hat{\sigma}_{ij}^{MF,NNLO} = & - \left[\Gamma_{ik}^2 \otimes d\sigma_{kj}^{LO} + d\sigma_{ik}^{LO} \otimes \Gamma_{kj}^2 - \Gamma_{ik}^1 \otimes d\sigma_{kl}^{LO} \otimes \Gamma_{lj}^1 \right. \\ & - \Gamma_{ik}^1 \otimes \Gamma_{kl}^1 \otimes d\sigma_{lj}^{LO} - d\sigma_{ik}^{LO} \otimes \Gamma_{kl}^1 \otimes \Gamma_{lj}^1 \\ & \left. + \Gamma_{ik}^1 \otimes d\sigma_{kj}^{NLO} + d\sigma_{ik}^{NLO} \otimes \Gamma_{kj}^1 \right]. \end{aligned} \quad (2.4.10)$$

Through this procedure, we obtain corrections to the cross section that occur in hadronic processes which should be included at each perturbative order. $d\hat{\sigma}^{MF,NLO}$ lies in the same phase space as $d\hat{\sigma}^V$, and $d\hat{\sigma}^{MF,NNLO}$ has components in both $d\hat{\sigma}^{VV}$ and $d\hat{\sigma}^{RV}$, and so should be included in the integrands of the respective phase space integrals.

As discussed previously, these mass factorisation contributions are necessary for the cancellation of IR divergences in scattering processes with hadronic initial states. At NLO, the relationship between the mass factorisation contributions and initial state collinear divergences is made manifest as one can directly relate the kernels to the splitting functions:

$$\Gamma_{ij}^1(x) = \frac{-1}{\varepsilon} P_{ij}^{(0)}(x). \quad (2.4.11)$$

These are exactly the splitting functions that one observes at NLO in initial state collinear divergences as we saw in Section 2.2, and serve to cancel exactly those contributions. The explicit forms of the splitting kernels in the $\overline{\text{MS}}$ scheme can be found in [48].

2.4.2 NLO Subtraction

Having defined the mass factorisation terms we can now construct our subtraction counterterm at NLO, to give the $d\hat{\sigma}_{ij}^{\text{S,NLO}}$ and $d\hat{\sigma}_{ij}^{\text{T,NLO}}$ terms in our NLO corrections:

$$d\hat{\sigma}_{ij}^{\text{NLO}} = \int_{\Phi_{n+1}} [d\hat{\sigma}_{ij}^{\text{R}} - d\hat{\sigma}_{ij}^{\text{S,NLO}}] + \int_{\Phi_n} [d\hat{\sigma}_{ij}^{\text{V}} - d\hat{\sigma}_{ij}^{\text{T,NLO}} + d\hat{\sigma}_{ij}^{\text{MF,NLO}}]. \quad (2.4.12)$$

The NLO case is the most straightforward contribution to put together, and forms a subset of the terms at NNLO that we will consider later on. The only class of antenna functions that we observe at NLO is the X_3^0 , which will fully account for all of the single unresolved IR divergences.

In this case, the real subtraction term $d\hat{\sigma}_{ij}^{\text{S,NLO}}$ for a process with n final state partons at LO and m initial-state partons is constructed for each level in the colour decomposition as:

$$d\hat{\sigma}_{ij}^{\text{S,NLO}} \sim \sum_{\text{perms}} \sum_{j \in \text{FS}} d\Phi_{n+1} \frac{1}{S_{n+1}} X_3^0(i, j, k) |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2, \quad (2.4.13)$$

where we neglect overall prefactors. Here the sum over permutations simply indicates the sum over the permutations of colour-orderings that contribute at a given colour level. The exact choice of each X_3^0 is determined by the parton content of the unresolved parton j and the hard radiators which accompany it on either side in the colour ordering. This will serve to remove any divergences in the real phase space integral, leaving it finite.

This real subtraction term is compensated by the equivalent integrated contribution in the single (1-loop) virtual matrix element counterterm again decomposed by colour level, where the integrated \mathcal{X}_3^0 match exactly the X_3^0 content in the reals:

$$d\hat{\sigma}_{ij}^{\text{T,NLO}} \sim \sum_{\text{perms}} \sum_{i,j} d\Phi_n \frac{1}{S_n} \mathcal{X}_3^0(i, j) |M_{n+m}^0(\dots, i, j, \dots)|^2. \quad (2.4.14)$$

Once these are introduced, they will cancel exactly the explicit poles of the virtual matrix element, alongside any mass factorisation contribution if initial state partons are present.

One can also systematically write the NLO virtual subtraction term in combination with mass factorisation terms. This is performed by combining the mass factorisation kernels and integrated antenna functions \mathcal{X}_3^0 which must always appear in combination into objects

$$\mathbf{J}_2^{(1)} \sim \mathcal{X}_3^0 + \sum \Gamma_{ij}^{(1)} \quad (2.4.15)$$

known as *integrated antenna dipoles*, where the sum is dependent on the dipole².

The composition of these dipoles depends on the flavour content, colour level and initial state configuration of the integrated antenna functions. For flavour-changing antenna functions, which cancel exclusively against mass factorisation counterterms, the corresponding dipoles are finite and contain no explicit ε poles, whereas for flavour-preserving terms they reconstruct exactly the poles of the corresponding loop matrix elements. The full composition of each integrated antenna dipole can be found in [48]. In this formalism, the NLO subtraction and mass factorisation terms take the simple form

$$d\hat{\sigma}_{ij}^{\text{T,NLO}} - d\hat{\sigma}_{ij}^{\text{MF,V}} \sim \sum_{\text{perms}} \sum_{i,j} d\Phi_n \frac{1}{S_n} \mathbf{J}_2^{(1)} |M_{n+m}^0(\dots, i, j, \dots)|^2. \quad (2.4.16)$$

Whilst not strictly necessary at NLO, these integrated dipoles simplify the construction of NNLO subtraction terms considerably.

2.4.3 NNLO Subtraction

Now that the NLO subtraction terms have been assembled, our focus moves to the regulation of the IR divergences at NNLO, where the construction of antenna subtraction terms becomes considerably more complex. The overall structure at NNLO is given by

$$\begin{aligned} d\hat{\sigma}_{ij}^{\text{NNLO}} = & \int_{\Phi_{n+2}} \left[d\hat{\sigma}_{ij}^{\text{RR}} - d\hat{\sigma}_{ij}^{\text{S}} \right] \\ & + \int_{\Phi_{n+1}} \left[d\hat{\sigma}_{ij}^{\text{RV}} - d\hat{\sigma}_{ij}^{\text{T}} + d\hat{\sigma}_{ij}^{\text{MF,RV}} \right] \end{aligned}$$

²For final-final dipoles there are no corresponding mass factorisation terms and the sum over splitting kernels is not present.

$$+ \int_{\Phi_n} \left[d\hat{\sigma}_{ij}^{\text{VV}} - d\hat{\sigma}_{ij}^{\text{U}} + d\hat{\sigma}_{ij}^{\text{MF,VV}} \right], \quad (2.4.17)$$

where we have explicitly included the mass factorisation terms. Whilst we will not provide an in depth discussion here, we will give an overview of each subtraction term and the divergent behaviours they regulate. For an extensive consideration of the subject, see e.g. [48].

NNLO Subtraction - Double Reals

We start with the double real (RR) subtraction term $d\hat{\sigma}_{ij}^{\text{S}}$, which regulates the highest multiplicity matrix elements. This can be decomposed into multiple contributions which account for all different unresolved configurations (and correct for any oversubtractions):

$$\hat{\sigma}_{ij}^{\text{S}} = \hat{\sigma}_{ij}^{\text{S},a} + \hat{\sigma}_{ij}^{\text{S},b_1} + \hat{\sigma}_{ij}^{\text{S},b_2} + \hat{\sigma}_{ij}^{\text{S},c} + \hat{\sigma}_{ij}^{\text{S},d}, \quad (2.4.18)$$

each term of which we will consider in the following.

The first term, $d\hat{\sigma}_{ij}^{\text{S},a}$, accounts for the single unresolved pieces of the double real integrand, and is directly analogous to the NLO counterterm $d\hat{\sigma}_{ij}^{\text{S,NLO}}$:

$$d\hat{\sigma}_{ij}^{\text{S},a} \sim \sum_{\text{perms}} \sum_{j \in \text{FS}} d\Phi_{n+2} \frac{1}{S_{n+2}} X_3^0(i, j, k) |M_{n+m+1}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2. \quad (2.4.19)$$

The second term, $d\hat{\sigma}_{ij}^{\text{S},b_1}$, accounts for the colour-connected double unresolved emissions, using the double unresolved limits of the X_4^0 antenna functions:

$$d\hat{\sigma}_{ij}^{\text{S},b_1} \sim \sum_{\text{perms}} \sum_{j,k \in \text{FS}} d\Phi_{n+2} \frac{1}{S_{n+2}} X_4^0(i, j, k, l) |M_{n+m}^0(\dots, (\widetilde{ijk}), (\widetilde{jkl}), \dots)|^2, \quad (2.4.20)$$

where j, k are colour connected. However, these X_4^0 functions also contain single unresolved limits, which must themselves be subtracted using X_3^0 antennas.

This issue of residual divergences is also present in $d\hat{\sigma}_{ij}^{\text{S},a}$, where the matrix elements $|M_{n+m+1}^0|^2$ can also still go singly unresolved, which must again be regulated by further X_3^0 terms. The remaining divergent behaviour in $d\hat{\sigma}_{ij}^{\text{S},a}$ and $d\hat{\sigma}_{ij}^{\text{S},b_1}$ can be

resolved simultaneously through the addition of a single term $d\hat{\sigma}_{ij}^{S,b_2}$, which explicitly removes single unresolved divergences in the X_4^0 antenna whilst implicitly removing the single limits within $|M_{n+m+1}^0|^2$.

This is made simple by the construction of antennas as ratios of matrix elements, which allows undesired single unresolved limits of X_4^0 functions to be regulated directly using X_3^0 antennas:

$$\begin{aligned} d\hat{\sigma}_{ij}^{S,b_2} &\sim \sum_{\text{perms}} \sum_{j \in X_4^0} d\Phi_{n+2} \frac{1}{S_{n+2}} X_3^0(i, j, k) X_3^0(\widetilde{(ij)}, \widetilde{(jk)}, l) \\ &\quad \times |M_{n+m}^0(\dots, \widetilde{(ij)(jk)}, \widetilde{(jk)l}, \dots)|^2. \end{aligned} \quad (2.4.21)$$

In the single unresolved limits of the X_4^0 , the universal divergent behaviour is mimicked by the first X_3^0 function, whilst the remaining finite part reduces onto a second X_3^0 . This compensates exactly for the oversubtractions contained within $d\hat{\sigma}_{ij}^{S,b_1}$ and the remaining colour-connected double unresolved limits in $d\hat{\sigma}_{ij}^{S,a}$.

At this point, all terms necessary for complete IR subtraction of processes containing up to 4 partons at the double real level are complete, as no divergences can occur between colour-disconnected partons. This is the case for inclusive Higgs production in gluon-gluon fusion, inclusive Drell-Yan production of vector bosons, single jet production in lepton-hadron scattering and di-jet production in lepton-antilepton scattering.

However, for processes where one wants to probe further the QCD kinematics at NNLO in hadron-hadron collisions, one must have at least one final state jet at Born level, which requires a minimum of 5 partons at the double real level. In this case, there is a double counting of single unresolved divergences between colour almost-connected partons arising from $d\hat{\sigma}_{ij}^{S,a}$ and $d\hat{\sigma}_{ij}^{S,b_2}$. These can be removed for each $\widetilde{X}_4^0(i, j, k, l)$ antenna function with unresolved j, k using terms of the form

$$\begin{aligned} d\Phi_{n+2} &\left[+ \frac{1}{2} X_3^0(i, j, l) X_3^0(\widetilde{(ij)}, k, \widetilde{(jl)}) |M_{n+m}^0(\dots, a, \widetilde{(ij)k}, \widetilde{(kjl)}, b, \dots)|^2 \right. \\ &\quad \left. - \frac{1}{2} X_3^0(a, j, i) X_3^0(\widetilde{(ij)}, k, l) |M_{n+m}^0(\dots, \widetilde{(aj)}, \widetilde{(ij)k}, \widetilde{(kl)}, b, \dots)|^2 \right] \end{aligned}$$

$$-\frac{1}{2}X_3^0(b, j, l)X_3^0(i, k, (\widetilde{jl}))|M_{n+m}^0(\dots, a, (\widetilde{ik}), (\widetilde{k(jl)}), (\widetilde{j\bar{b}}), \dots)|^2 \Big] \quad (2.4.22)$$

which systematically account for the oversubtraction. However, these structures themselves introduce new wide-angle single soft divergences which require the use of explicit Eikonal factors to remove [85, 86]. To subtract these requires terms

$$\begin{aligned} d\Phi_{n+2} \Bigg[\frac{1}{2} \Big[(S_{(\widetilde{ij}), j, (\widetilde{jl})} - S_{((\widetilde{ij})k)j(\widetilde{k(jl)})}) - (S_{aj(\widetilde{ij})} - S_{aj((\widetilde{ij})k)}) - (S_{bj(\widetilde{jl})} - S_{bj((\widetilde{jl})k)}) \Big] \\ \times X_3^0((\widetilde{ij}), k, (\widetilde{jl}))|M_{n+m}^0(\dots, a, ((\widetilde{ij})k), (\widetilde{k(jl)}), b, \dots)|^2 \Big], \end{aligned} \quad (2.4.23)$$

where the integral of the Eikonal function S over the single unresolved phase space is known, and can be explicitly reintroduced at the RV level. Each term in (2.4.23) can be directly associated with one in (2.4.22). In combination, these two terms along with the equivalent terms with a $(j \leftrightarrow k)$ substitution applied give the contribution which we label $d\hat{\sigma}_{ij}^{S,c}$.

The final term to be mentioned is $d\hat{\sigma}_{ij}^{S,d}$, which contains the contribution from colour disconnected divergent limits. Whilst this is not used in any of the work in this thesis, as it requires six or more partons at double-real level, this can be constructed as two independent NLO-like limits:

$$\begin{aligned} d\hat{\sigma}_{ij}^{S,d} \sim \sum_{\text{perms}} \sum_{j \in \text{FS}} d\Phi_{n+2} \frac{1}{S_{n+2}} X_3^0(i, j, k) X_3^0(l, m, n) \\ \times |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots, (\widetilde{lm}), (\widetilde{mn}), \dots)|^2. \end{aligned} \quad (2.4.24)$$

This completes the set of terms required to construct double real subtraction terms.

NNLO Subtraction - Real-Virtuals

The real virtual subtraction term $d\hat{\sigma}_{ij}^T$ can be decomposed in a similar way to its counterpart in the double reals:

$$d\hat{\sigma}_{ij}^T - d\hat{\sigma}_{ij}^{\text{MF,RV}} = d\hat{\sigma}_{ij}^{T,a} + d\hat{\sigma}_{ij}^{T,b_1} + d\hat{\sigma}_{ij}^{T,b_2} + d\hat{\sigma}_{ij}^{T,c}. \quad (2.4.25)$$

We now have for the first time a mixture of contributions which contain the correct explicit pole structure in ε , which regulate implicit poles, and which must compensate correctly for finite parts in the double real and double virtual terms. These contributions lie in the $(n+1)$ -parton phase space, and contain only single unresolved limits which factorise on to reduced terms in n -parton phase space.

The first term, $d\hat{\sigma}_{ij}^{\text{T},a}$, is an NLO-like counterterm and takes exactly the same form as $d\hat{\sigma}_{ij}^{\text{T,NLO}}$ with the addition of an extra final state parton:

$$d\hat{\sigma}_{ij}^{\text{T},a} \sim \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} \sum_{i,j} d\Phi_{n+1} \frac{1}{S_{n+1}} \mathbf{J}_2^{(1)}(i,j) |M_{n+m+1}^0(\dots, i, j, \dots)|^2. \quad (2.4.26)$$

This term completely accounts for the explicit poles contained in the RV matrix elements and mass factorisation terms residing, and is exactly equivalent to the NLO subtraction term for the process with an additional radiated parton. The integrals over dx_a are present in the case that initial state partons are present, and form the required convolutions with the incoming momentum fractions x_a .

The second and third terms in $d\hat{\sigma}_{ij}^{\text{T}}$ concern the implicit singularities present at the real-virtual level. As seen in Section 2.2, there are two terms required to cancel divergent one-loop contributions containing one-loop and tree-level reduced matrix elements respectively, which we identify as $d\hat{\sigma}_{ij}^{\text{T},b_1}$ and $d\hat{\sigma}_{ij}^{\text{T},b_2}$.

The first of these contains the one-loop reduced matrix element M_{n+m}^1 , which itself has an explicit singularity structure:

$$\begin{aligned} d\hat{\sigma}_{ij}^{\text{T},b_1} \sim & \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} \sum_{j \in \text{FS}} d\Phi_{n+1} \frac{1}{S_{n+1}} X_3^0(i, j, k) \\ & \left\{ \left[\prod_{a=1}^m \delta(1 - x_a) \right] |M_{n+m}^1(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2 \right. \\ & \left. + \sum_{(\widetilde{ij}), (\widetilde{jk})} \mathbf{J}_2^{(1)}((\widetilde{ij}), (\widetilde{jk})) |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2 \right\}. \end{aligned} \quad (2.4.27)$$

Within this contribution, the final line containing $\mathbf{J}_2^{(1)}((\widetilde{ij}), (\widetilde{jk}))$ is constructed to cancel the explicit poles in the one-loop matrix element such that $d\hat{\sigma}_{ij}^{\text{T},b_1}$ is finite.

This cancellation can be viewed heuristically by considering $M_{n+m}^1(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)$

as the virtual matrix element for an $(n+m)$ -parton process at NLO. A corresponding virtual subtraction term can then be constructed to compensate a hypothetical $M_{n+m+1}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)$ real contribution in the same manner as $d\hat{\sigma}_{ij}^{\text{T,NLO}}$ in (2.4.16). This term replicates exactly the pole structure of M_{n+m}^1 and can therefore be used to ensure that $d\hat{\sigma}_{ij}^{\text{T},b_1}$ is free from explicit poles.

The second term maps on to the tree level matrix element M_{n+m}^0 and takes the form

$$\begin{aligned} d\hat{\sigma}_{ij}^{\text{T},b_2} \sim & \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} \sum_{j \in \text{FS}} d\Phi_{n+1} \frac{1}{S_{n+1}} \\ & \left\{ \left[\prod_{a=1}^m \delta(1-x_a) \right] X_3^1(i, j, k) |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2 \right. \\ & + \sum_{i,j} \mathbf{J}_2^{(1)}(i, j) X_3^0(i, j, k) |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2 \\ & \left. - A_{X_3^1} X_3^0(i, j, k) \mathbf{J}_2^{(1)}((\widetilde{ij}), (\widetilde{jk})) |M_{n+m}^0(\dots, (\widetilde{ij}), (\widetilde{jk}), \dots)|^2 \right\}. \end{aligned} \quad (2.4.28)$$

Here we see for the first time the appearance of one-loop X_3^1 antenna functions which exist to regulate the divergences in the second term of (2.2.17). These X_3^1 themselves contain explicit poles which are accounted for in the final two lines of (2.4.28) in a similar manner to the M_{n+m}^1 term in $d\hat{\sigma}_{ij}^{\text{T},b_1}$. Neglecting function arguments, one can schematically rewrite (2.3.7) as:

$$\begin{aligned} X_3^1 &= \frac{|M_3^1|^2}{|M_2^0|^2} - X_3^0 \frac{|M_2^1|^2}{|M_2^0|^2} \\ &= X_3^0 \left(\frac{|M_3^1|^2}{|M_3^0|^2} - \frac{|M_2^1|^2}{|M_2^0|^2} \right). \end{aligned} \quad (2.4.29)$$

Here an explicit pole structure is contained within M_2^1 and M_3^1 which can be treated as virtual contributions to some 2- and 3-parton processes. As before, one can then regulate the explicit poles using integrated dipoles, giving terms

$$\begin{aligned} & + \sum_{i,j} \mathbf{J}_2^{(1)}(i, j) X_3^0(i, j, k) |M_{n+m}^0|^2 \\ & - A_{X_3^1} X_3^0(i, j, k) \mathbf{J}_2^{(1)}((\widetilde{ij}), (\widetilde{jk})) |M_{n+m}^0|^2, \end{aligned} \quad (2.4.30)$$

where the first line corresponds to the first term in (2.4.29) and the second line to the second term. We have also introduced a constant of the X_3^1 antenna, $A_{X_3^1}$, which

is zero when there is no $|M_2^1|^2$ matrix element onto which the antenna maps, and otherwise absorbs relative symmetry factors.

Here one must also account for the fact that the one-loop antenna X_3^1 is renormalised at the mass scale of the antenna s_{ijk} , rather than the renormalisation scale of the calculation, μ_R^2 . This can be systematically corrected for by the substitution

$$X_3^1(i, j, k) \rightarrow X_3^1(i, j, k) + \frac{\beta_0}{\varepsilon} C(\varepsilon) X_3^0(i, j, k) \left[\left(\frac{s_{ijk}}{\mu_R^2} \right)^{-\varepsilon} - 1 \right] \quad (2.4.31)$$

everywhere that X_3^1 occurs, where

$$C(\varepsilon) = \frac{(4\pi)^\varepsilon e^{-\gamma_E \varepsilon}}{8\pi^2}. \quad (2.4.32)$$

The final contribution to the real-virtual subtraction term required to complete $d\hat{\sigma}_{ij}^T$ is the counterterm to $d\hat{\sigma}_{ij}^{S,c}$, $d\hat{\sigma}_{ij}^{T,c}$. This contains terms of the form

$$\begin{aligned} d\Phi_{n+1} \frac{1}{S_{n+1}} \frac{1}{2} \sum_j \Bigg[& + \left(\mathcal{X}_3^0(i, k) - \mathcal{X}_3^0(\widetilde{(ij)}, \widetilde{(jk)}) \right) \\ & - \left(\mathcal{X}_3^0(a, i) - \mathcal{X}_3^0(a, \widetilde{(ij)}) \right) - \left(\mathcal{X}_3^0(b, k) - \mathcal{X}_3^0(b, \widetilde{(jk)}) \right) \Bigg] \\ & \times X_3^0(i, j, k) |M_{n+m}^0(\dots, \widetilde{(ij)}, \widetilde{(jk)}, \dots)|^2 \end{aligned} \quad (2.4.33)$$

and

$$\begin{aligned} d\Phi_{n+1} \frac{1}{S_{n+1}} \delta(1-x_1) \delta(1-x_2) \frac{1}{2} \sum_j \Bigg[& + \left(\mathcal{S}(i, k) - \mathcal{S}(\widetilde{(ij)}, \widetilde{(jk)}) \right) \\ & - \left(\mathcal{S}(a, i) - \mathcal{S}(a, \widetilde{(ij)}) \right) - \left(\mathcal{S}(b, k) - \mathcal{S}(b, \widetilde{(jk)}) \right) \Bigg] \\ & \times X_3^0(i, j, k) |M_{n+m}^0(\dots, \widetilde{(ij)}, \widetilde{(jk)}, \dots)|^2 \end{aligned} \quad (2.4.34)$$

to systematically account for the terms (2.4.22) and (2.4.23) introduced in $d\hat{\sigma}_{ij}^{S,c}$, where \mathcal{S} are the integrated Eikonal factors [81, 84].

NNLO Subtraction - Double Virtuals

We move now to the final component of the NNLO subtraction, the double virtual counterterm $d\hat{\sigma}_{ij}^U$. Once again we fold in the mass factorisation counterterm before decomposing into separate contributions:

$$d\hat{\sigma}_{ij}^U - d\hat{\sigma}_{ij}^{\text{MF,RV}} = d\hat{\sigma}_{ij}^{U,a} + d\hat{\sigma}_{ij}^{U,b} + d\hat{\sigma}_{ij}^{U,c}. \quad (2.4.35)$$

We no longer have to deal with implicit divergences at this stage, and the task amounts to systematically constructing the integrated counterparts to all contributions from the double-real and real-virtual subtraction terms which contain reduced matrix elements in the Born level phase space Φ_n . The poles of these terms will exactly cancel against the two-loop matrix elements, ensuring that the double virtual integrand is finite.

The first term is the counterpart of the second line of $d\hat{\sigma}_{ij}^{\text{T},b_1}$ in (2.4.27):

$$d\hat{\sigma}_{ij}^{U,a} \sim \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} d\Phi_n \frac{1}{S_n} \sum_{i,j} \mathbf{J}_2^{(1)}(i,j) \left[|M_{n+m}^1(\dots, i, j, \dots)|^2 - \frac{\beta_0}{\varepsilon} |M_{n+m}^0(\dots, i, j, \dots)|^2 \right]. \quad (2.4.36)$$

Here the second term proportional to β_0 is once again included in order to account for the systematic redefinition of X_3^1 for renormalisation purposes in the real-virtual. The component proportional to $(s_{ij}/\mu^2)^{-\varepsilon}$ is absorbed elsewhere, in $d\hat{\sigma}_{ij}^{U,c}$.

The second term in the decomposition of $d\hat{\sigma}_{ij}^U$ includes the remaining integrated counterpart to $d\hat{\sigma}_{ij}^{\text{T},b_1}$, as well as $d\hat{\sigma}_{ij}^{\text{T},c}$ and $d\hat{\sigma}_{ij}^{\text{S},d}$ if present:

$$d\hat{\sigma}_{ij}^{U,b} \sim \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} d\Phi_n \frac{1}{S_n} \frac{1}{2} \left[\sum_{i,j} \mathbf{J}_2^{(1)}(i,j) \otimes \sum_{k,l} \mathbf{J}_2^{(1)}(k,l) \right] |M_{n+m}^0(\dots, i, j, \dots)|^2. \quad (2.4.37)$$

This convolution structure naturally generates the (one-loop) \otimes (one-loop) mass factorisation terms we see in (2.4.10).

The final term contains the remaining contributions, which have been written here

in terms of the 2-loop integrated dipole function $\mathbf{J}_2^{(2)}(i, j)$:

$$d\hat{\sigma}_{ij}^{\text{U},c} \sim \int \left(\prod_{a=1}^m \frac{dx_a}{x_a} \right) \sum_{\text{perms}} d\Phi_n \frac{1}{S_n} \sum_{i,j} \mathbf{J}_2^{(2)}(i, j) |M_{n+m}^0(\dots, i, j, \dots)|^2. \quad (2.4.38)$$

The $\mathbf{J}_2^{(2)}(i, j)$ contain the two loop factorisation kernel contributions, as well as the integrated \mathcal{X}_4^0 , \mathcal{X}_3^1 and any remaining $\mathcal{X}_3^0 \otimes \mathcal{X}_3^0$ terms [48]. The most general form is given by

$$\begin{aligned} \mathbf{J}_2^{(2)}(i, j) = & c_1 \mathcal{X}_4^0 + c_2 \widetilde{\mathcal{X}}_4^0 + c_3 \mathcal{X}_3^1 + c_4 \frac{\beta_0}{\varepsilon} \left(\frac{s_{ij}}{\mu^2} \right)^{-\varepsilon} \mathcal{X}_3^0 + c_5 \mathcal{X}_3^0 \otimes \mathcal{X}_3^0 \\ & - c_6 \Gamma_{ik}^{(2)}(z_1) \delta(1 - z_2) - c_7 \Gamma_{kj}^{(2)}(z_2) \delta(1 - z_1) \end{aligned} \quad (2.4.39)$$

where the c_i are constants specific to each integrated dipole and initial-final crossing, and $\Gamma_{ij}^{(2)}$ are the two-loop mass factorisation terms in the first line of (2.4.10). For the case in which there are no hadronic initial states in the dipole, $c_6 = c_7 = 0$.

2.5 Antenna Subtraction for Charged-Current processes

We now turn to the application of antenna subtraction to charged current processes which we consider in much of the remainder of this thesis. In particular, this relates to inclusive W boson and W boson production in association with a jet in hadron-hadron collisions, alongside single and di-jet production in lepton-hadron collisions to NNLO. These processes are very closely related to the neutral current equivalents, first performed using antenna subtraction in [87–94]. However, there are some important differences, primarily in contributions containing four quark interference terms. Here we observe flavour structures that are not present in the neutral current case, and require special consideration. Examples of these are shown in Figs 2.1 and 2.2, which first occur at $\mathcal{O}(\alpha_s^2)$ in the perturbative series and which we label the D -type matrix elements.

We adopt the following notation for the colour ordered four quark amplitude:

$$\mathcal{C}(i_q^a, j_q^b, k_{q'}^b, l_q^a), \quad (2.5.1)$$

where i, j, k, l are the quark momenta, and the a, b indices give the colour connected quarks which share the same SU(3) index. These \mathcal{C} -type amplitudes form two separate squared matrix elements dependent on whether they interfere (neglecting the momenta of the W boson):

$$\begin{aligned} C(i_q^a, j_q^b, k_{q'}^b, l_q^a) &= \mathcal{C}(i_q^a, j_q^b, k_{q'}^b, l_q^a) \mathcal{C}^\dagger(i_q^a, j_q^b, k_{q'}^b, l_q^a) \\ D(i_q^a, j_q^b, k_{q'}^b, l_q^a) &= \mathcal{C}(i_q^a, j_q^b, k_{q'}^b, l_q^a) \mathcal{C}^\dagger(i_q^a, l_q^b, k_{q'}^b, j_q^a). \end{aligned} \quad (2.5.2)$$

The C -type matrix elements behave in a similar manner to the neutral current counterpart, as the boson couples to the same quark line in both the amplitude and its Hermitian conjugate. This one quark line must remain resolved for reduced matrix elements to exist, which means that the singularity structure of the matrix element is much the same in both cases³.

The D -type matrix elements display different behaviour, as the quark interference structure is directly modified by the presence of a flavour changing current. These matrix elements are noteworthy in that they do not contain any single unresolved limits between the quarks, as no appropriate reduced matrix elements exist. The interfering quarks can only participate in double-collinear limits, in which two quarks go collinear, forming a composite gluon which then goes soft or collinear with the remaining quark line. This factorises on to a squared matrix element with a single quark line, which we label $B(1_q^a, 2_{q'}^a)$.

To elucidate this, we take as an example the interference terms for the hadron-

³The only caveat being that one must subdivide the neutral current singularities across multiple crossings in the charged current case as the change in flavour means that one can always identify the quark line from which the W was radiated. Schematically this follows:

$$C(q, Q, Q, q) \Big|_{\text{NC}} \rightarrow [C(Q, u, d, Q) + C(u, Q, Q, d)] \Big|_{\text{CC}}, \quad (2.5.3)$$

where the quarks which can go unresolved are labelled as Q and q , and those known to couple to the W boson are labelled u, d .




Figure 2.1: An example 4-quark interference squared Feynman diagram present at $\mathcal{O}(\alpha\alpha_s^2)$ in W boson production through hadron-hadron collisions, which contains no double unresolved final state collinear limits. Quark momenta 1 and 2 are initial-state, with i and j final-state.

hadron process $u(1)u(2) \rightarrow d(i)u(j)$, corresponding to $D(1_u^a, 2_u^b, i_d^b, j_u^a)$ as shown in Fig. 2.1. If we consider the possible limits of colour-connected partons, one might naïvely believe that there is a single unresolved limit:

$$D(1_u^a, 2_u^b, i_d^b, j_u^a) \xrightarrow{(1||j)_g || i_u} B(1_u^a, 2_d^a). \quad (2.5.4)$$

This would occur when two u -type quarks form a collinear gluon, which in turn becomes collinear with the d quark, reducing on to a $B(1_u^a, 2_d^a)$ matrix element. However, this is not the case as both of the colour ordered \mathcal{C} amplitudes which constitute the D matrix element must simultaneously display divergent behaviour for the phase space integral itself to diverge. The integration carries factors of the vanishing invariants, such that the volume of the phase space vanishes in the exactly singular limit. This means that the matrix element must diverge faster than the associated decrease in phase space volume for a singular limit to occur. When only a single amplitude diverges the singular behaviour of the matrix element is not enough to cause a divergence under integration.

That only one amplitude is singular in (2.5.4) can be seen by tabulating the possible divergences, as is performed in Table 2.3 by considering that collinear partons must

$D(1_u, 2_u, i_d, j_u)$	$\mathcal{C}(1_u^a, 2_u^b, i_d^b, j_u^a)$	$\mathcal{C}^\dagger(1_u^a, j_u^b, i_d^b, 2_u^a)$
Double Limits:	$1_u^a \parallel j_u^a$ ✓	$1_u^a \parallel 2_u^a$ ✗
	$2_u^b \parallel i_d^b$ ✗	$j_u^b \parallel i_d^b$ ✗
Triple Limits:	$(1_u^a \parallel j_u^a)_g \parallel i_d$ ✓	$(1_u^a \parallel 2_u^a)_g \parallel i_d$ ✗
	$(1_u^a \parallel j_u^a)_g \parallel 2_u$ ✗	$(1_u^a \parallel 2_u^a)_g \parallel j_u$ ✗
	$(2_u^b \parallel i_d^b)_g \parallel 1_u$ ✗	$(j_u^b \parallel i_d^b)_g \parallel 1_u$ ✗
	$(2_u^b \parallel i_d^b)_g \parallel j_u$ ✗	$(j_u^b \parallel i_d^b)_g \parallel 2_u$ ✗

Table 2.3: Tabulated colour-connected singular limits of the amplitudes contributing to $D(1_u, 2_u, i_d, j_u)$. This corresponds to the matrix element crossing shown in Fig. 2.1.

be colour-connected and of the correct flavour/initial state configuration. The singular limits are labelled with a green tick, and the non-singular limits are labelled with a red cross. One can immediately see that in \mathcal{C} , $(1_u \parallel j_u)$ forms a valid limit, whereas $(2_u \parallel i_d)$ does not due to the difference in quark flavour. The composite gluon formed from quark pair $(1_u, j_u)$ can then go collinear with i_d , forming a $(1_u \parallel j_u \parallel i_d)$ triple collinear limit within the amplitude.

However, in \mathcal{C}^\dagger there are no valid single unresolved limits, as $1_u \parallel 2_u$ would contain two initial state partons that by definition cannot go collinear, and $j_u \parallel i_d$ contains quarks of different flavours and is finite. There can therefore be no triple collinear limits in \mathcal{C}^\dagger . Combining the two contributions to form $D = \mathcal{C}\mathcal{C}^\dagger$, one can see that there are no singular limits that occur simultaneously in both amplitudes, and therefore the contribution from the given matrix element crossing is finite. Note that were i and j the same flavour of quark, as is the case for the neutral current, \mathcal{C}^\dagger would display a $(j \parallel i)$ single unresolved limit and a $(1 \parallel j \parallel i)$ triple collinear limit as is observed in \mathcal{C} , and the matrix element would diverge.

The finiteness of charged-current D -type matrix elements is not however a generic behaviour. If we take a different crossing $(u(1) d(2) \rightarrow u(i) u(j))$

$$D(1_u, i_u, 2_d, j_u) = \mathcal{C}(1_u^a, i_u^b, 2_d^b, j_u^a) \mathcal{C}^\dagger(1_u^a, j_u^b, 2_d^b, i_u^a) \quad (2.5.5)$$

shown in Fig. 2.2, and perform the same procedure, we see that we indeed observe $(1 \parallel j \parallel i)$ limits in both amplitudes and the D -type matrix element is divergent (see




Figure 2.2: An example 4-quark interference squared Feynman diagram present at $\mathcal{O}(\alpha\alpha_s^2)$ in W boson production through hadron-hadron collisions, which contains double unresolved final state collinear limits. Quark momenta 1 and 2 are initial-state, with i and j final-state.

Table 2.4). These limits can be subtracted straightforwardly using the C_4^0 antenna function. That there are finite crossings of a matrix element that also contains IR-divergent contributions at the double-real level is noteworthy, and a distinctive feature of charged-current processes where the EW flavour-changing behaviour directly interacts with the QCD limits.

There is a second distinction with the neutral current case for these four-quark interference terms that should also be mentioned. In the neutral current case, all relevant crossings can be constructed directly from $D(i_q, j_q, k_q, l_q)$. This ceases to be true for the charged current where one can have two separate contributions related by charge conjugation. Taking W^+ production at $\mathcal{O}(\alpha_s^2)$ in hadron-hadron collisions as an example, we have the two sub-processes (and associated crossings)

$$\text{a) } u \bar{d} \rightarrow u \bar{u} W^+$$

$$\text{b) } u \bar{d} \rightarrow d \bar{d} W^+$$

which cannot be related through crossing symmetry alone. The D -type matrix elements for the second sub-process can however be generated from a line reversal symmetry applied to the first (or vice versa), which allows the two matrix elements

$D(1_u, i_u, 2_d, j_u)$	$\mathcal{C}(1_u^a, i_u^b, 2_d^b, j_u^a)$	$\mathcal{C}^\dagger(1_u^a, j_u^b, 2_d^b, i_u^a)$
Double Limits:	$1_u^a \parallel j_u^a$ ✓	$1_u^a \parallel i_u^a$ ✓
	$i_u^b \parallel 2_d^b$ ✗	$j_u^b \parallel 2_d^b$ ✗
Triple Limits:	$(1_u^a \parallel j_u^a)_g \parallel i_u$ ✓	$(1_u^a \parallel i_u^a)_g \parallel 2_d$ ✗
	$(1_u^a \parallel j_u^a)_g \parallel 2_d$ ✗	$(1_u^a \parallel i_u^a)_g \parallel j_u$ ✓
	$(i_u^b \parallel 2_d^b)_g \parallel 1_u$ ✗	$(j_u^b \parallel 2_d^b)_g \parallel 1_u$ ✗
	$(i_u^b \parallel 2_d^b)_g \parallel j_u$ ✗	$(j_u^b \parallel 2_d^b)_g \parallel i_u$ ✗

Table 2.4: Tabulated colour-connected singular limits of the amplitudes contributing to $D(1_u, i_u, 2_d, j_u)$. This corresponds to the matrix element crossing shown in Fig. 2.2.

to be directly related.⁴ This relationship holds independent of the number of gluons present in the amplitude, as it is purely an artefact of the quark-interference effects.

2.6 Implementation of W^\pm processes in NNLOJET

Having established the primary differences between the neutral and charged boson production it becomes possible to fully implement calculations of charged current processes to NNLO within the NNLOJET framework using the existing calculations of inclusive Drell-Yan, and Z boson plus jet (ZJ) production in hadron-hadron collisions, and single jet inclusive production and di-jet production in deep inelastic scattering (DIS) [87–94]. These calculations use the matrix elements first calculated in [95–102].

Once these matrix elements are implemented and antenna subtraction terms constructed to regulate them in the IR-divergent limits, one must then perform a phase space integration to extract physical cross-sections. To do this integration analytically is impossible for all but the most inclusive processes due to the complex nature of fiducial cuts, the structure of the integrand, and the high dimensionalities involved. Because of this, one must resort to numerical methods in order to evaluate the result up to some sufficient statistical precision. In particular, Monte Carlo

⁴This applies for all tree level matrix elements. For 4-quark one-loop matrix elements as occur in e.g. the real-virtual contribution to W plus jet production in hadron-hadron scattering, this line reversal must instead take place at amplitude level due to renormalisation terms in the loop part.

methods are favoured due to the independence of the variance of the result on the dimensionality of the integral, ensuring that they scale well with the number of final state particles. These methods require repeated evaluation of the IR-regulated integrand for large numbers of randomly generated phase space points respecting the usual momentum conservation constraints, with the standard deviation on the final result decreasing as $\sigma_{\text{STD}} \propto N^{-1/2}$ for N evaluations. For integral dimensionality $D \geq 8$, Monte Carlo methods become favourable over quadrature methods such as the extended Simpson's rule, which scales as $\sigma_{\text{STD}} \propto N^{-4/D}$. They also have the advantage of simultaneously being able to construct differential distributions through point-by-point bin classification, and can be interfaced with tools such as parton showers to extend physical descriptions beyond fixed order perturbation theory.

The Monte Carlo algorithm implemented within NNLOJET is VEGAS [103], which during a warmup phase adapts to the integrand in order to improve sampling and therefore reduce the standard deviation with respect to a uniform distribution of random numbers. For intricate, high-dimensionality integrands, VEGAS typically requires a large number of phase space points in order to both adapt to the integrand and then to converge to the correct result. One must also increase N substantially when probing regions of phase space where large cancellations are observed within the integrand, i.e. when evaluating deep into soft and/or collinear limits where one relies on antenna functions to regulate divergences. Numerical stability is a particular concern in these limits if the cancellation between matrix element and subtraction term exceeds the floating point accuracy of the computing architecture, and a technical phase space cut is generally employed to control miscancellations (see also Appendix A.3.1).

Representative values of N required for the double real integrations in WJ production are still in the hundreds of millions however; one must typically resort to parallel grid computing techniques in order to provide the necessary CPU resources to obtain results in a reasonable timeframe. When one includes the evaluation time of the integrand, the total (unnormalised) CPU requirements can reach $\mathcal{O}(10^7)$ hrs

for the most pathological cases⁵ which is considerable even with currently available grid resources. Numerous techniques, including amplitude caching and multichannel integration have been implemented within NNLOJET in order to reduce this computational load, however it still remains a sizeable constraint on the production of results.

The output of this Monte Carlo integration constitutes the physical cross sections and differential distributions, and these results will form the majority of the work in the remainder of this thesis. In particular, the calculation of WJ production in hadron-hadron scattering will be used for work in Chapters 3 and 4, with the inclusive W cross section used also in Chapter 3 for the normalisation of distributions. In Chapter 6 we consider the subject of DIS, where we present results obtained using calculations of single- and di-jet production in charged current scattering which are again performed with antenna subtraction and share matrix elements with the above hadron-hadron processes. Each of these has been rigorously validated, and whilst for brevity we do not present those results in this thesis, the testing procedures applied to each calculation are discussed in Appendix A.

⁵Of the results presented in this thesis, the most computationally expensive are those contained in Section 3.4, which across all boson production channels reached the order of magnitude quoted above.

Chapter 3

Phenomenology of Neutral- and Charged-Current Electroweak Gauge Boson Production at the LHC

3.1 Overview

The production of electroweak (EW) gauge bosons decaying leptonically, known as the Drell–Yan (DY) process [104], is one of the most important processes at hadron colliders such as the LHC. Not only does it provide a testing ground for precision theoretical predictions and precision measurements of the Standard Model, it also forms a dominant background for many searches for new physics beyond the Standard Model (BSM). As a consequence, a robust understanding of DY W^\pm and Z boson production is mandatory and has been a subject of considerable theoretical and experimental attention over the last half-century.

Precision experimental measurements where the decay leptons are either muons or electrons have been performed for a wide variety of observables since the first obser-




Figure 3.1: Born level Z/γ^* boson production through the Drell-Yan process.

vation of the Z and W bosons at the UA1 and UA2 Super Proton Synchrotron (SPS) experiments at CERN in 1983 [105–108]. For neutral-current (NC) DY production mediated by a virtual photon or Z boson (see Fig. 3.1),

$$p + p \rightarrow Z/\gamma^*(\rightarrow \ell^+ + \ell^-) + X, \quad (3.1.1)$$

where X is some final state containing zero or more hadronic jets, the characteristic experimental signature of two oppositely charged lepton tracks allows exceptionally precise measurements to be taken across a large volume of phase space. A high production cross-section ensures that large sample sizes are relatively straightforward to collect, and that statistical uncertainties are generally relatively small. Experimentally, these properties allow NC DY production to fulfil valuable roles in terms of detector calibration (through measurements of previously well-known quantities such as the Z-boson mass M_Z and width Γ_Z) and luminosity determination (through the total cross section), which are crucial in order to understand all measurements made at hadron colliders such as the Tevatron and the LHC.

Charged-current (CC) DY production (see Fig. 3.2),

$$p + p \rightarrow W^\pm(\rightarrow \ell + \nu_\ell) + X, \quad (3.1.2)$$

mediated through either of the charged W^\pm bosons occurs at a rate approximately an order of magnitude larger than that of the NC DY production as the EW coupling constant α is not suppressed by the Weinberg angle, $\sin^2 \theta_W$. Unlike the NC




Figure 3.2: Born level W^+ boson production through the Drell-Yan process.

case where the full final state can be directly probed, the decay neutrino cannot be measured and instead presents as a missing transverse energy which can be reconstructed using momentum conservation. Due to errors introduced through this procedure, CC DY production measurements typically have larger systematic uncertainties than their NC counterparts, usually offset by lower statistical uncertainties.

Both of these DY production mechanisms can be accompanied by associated QCD radiation, usually measured in the form of jets defined by some IR-safe jet algorithm (see Section 1.4). Selection criteria on the multiplicity of associated jets control the sensitivity of measurements to different sectors of the SM by altering the power of the strong coupling constant present at Born level. Where one permits events in which no additional jets are produced, one probes directly the EW content of the final state, and as one increases the number of required jets, one increases the sensitivity to QCD radiation and jet dynamics. In the following we label the case in which some minimum constraint $N_{\text{jets}} \geq 1$ is present on the number of associated jets as V + jet(s) (VJ) production and refer to the case where $N_{\text{jets}} \geq 0$ as DY production.

The extraordinary experimental precision that these processes afford enables extremely high precision measurements of SM EW parameters at hadron colliders. In particular, the most precise determination of the W boson mass, $M_W = 80.387 \pm 0.016$ GeV is given by the combination of the CDF [109] and DØ [110] measurements at the Tevatron [111], made using fits to lepton transverse momentum and

W transverse mass distributions which contain Jacobian peaks at M_W and $M_W/2$ respectively. Comparable precision has since been reached at the LHC, with a 7 TeV M_W measurement performed by the ATLAS experiment giving a measured value of $M_W = 80.370 \pm 0.019$ GeV [112], and work is underway on equivalent extractions using the larger 8 and 13 TeV data samples, as well as a combined Tevatron-ATLAS extraction. There are similar prospects for the W-boson width, Γ_W , where the most precise measurement of $\Gamma_W = 2.046 \pm 0.049$ GeV was made at the Tevatron as a combination of CDF and DØ results [113].

Precision measurements of the EW sector at hadron-hadron colliders are not limited to properties of the W boson. Extraction of NC DY angular coefficients including the Lam-Tung relation, as well as forward-backward asymmetry measurements in NC DY production have already been made at the LHC [114–116], which permit extractions of the effective weak mixing angle $\sin^2 \theta_W^{\text{eff}}$. We defer a discussion of this to Chapter 5, where we consider in detail elements of such an extraction using ATLAS 8 TeV data.

There is also a strong phenomenological importance to EW gauge boson production at colliders, where both NC and CC production provide complementary constraints in PDF determinations. Charge asymmetry measurements in CC DY production alongside NC DY rapidity measurements give strong bounds on the quark flavour decomposition of the proton across a wide range of Bjorken- x values (see e.g. [117]), with high invariant mass NC DY giving some sensitivity to the photon content of the PDF [118]. If one instead considers VJ production, the gluon distribution can also be probed starting at LO in QCD, particularly in the less well determined intermediate and high x regions. If flavour tagging is also used on final state hadronic radiation, sensitivity to specific quark flavour content increases further; this is particularly notable for $W^\pm + \text{charm}$ production which gives important constraints on strange quark distributions. An extensive review of modern PDF frameworks and relevant experimental measurements can be found in [119].

Searches for BSM physics, such as dark matter, supersymmetry and invisible Higgs

decays also benefit from our understanding of EW gauge boson production, most notably as an irreducible background to new physics searches in mono-jet events. The experimental signature of these events consists of large missing transverse energy accompanied by hadronic radiation, the dominant contributions to which are produced from WJ events and ZJ events in which the Z decays to neutrinos. Understanding of the background contributions is crucial, and comes most prominently from independent control regions in leptonically decaying VJ measurements which can be used to evaluate efficiencies and acceptances [120].

Given the considerable interest in precision measurements of EW gauge boson production at hadron-hadron colliders, it is only natural that comparable effort has been made in reducing the uncertainty on theoretical predictions. Calculations of DY cross sections to $\mathcal{O}(\alpha_s^2)$ have been known in total since 1991 [121], inclusively since 2003 [122] and fully exclusively since 2006 [123], which have been followed by implementations in multiple codes using a wide variety of calculational methodologies [1, 124–128]. These fixed-order DY predictions have been supplemented by parton showers [126], Q_T [129–132] and threshold [133] resummation to various logarithmic accuracies in order to extend the range of their kinematic validity. Steps have also been made towards evaluation of the inclusive DY cross section at $\mathcal{O}(\alpha_s^3)$, including calculation of the 3-loop quark form factors [134, 135], threshold approximation [136] and N -jettiness beam function contributions [137, 138], with considerable further progress anticipated in the near future.

Fixed order calculations of VJ production through to $\mathcal{O}(\alpha_s^3)$ have been performed using the methods of N -jettiness slicing and antenna subtraction [1, 87–89, 139–143], and are now being matched to resummation results [2, 144, 145]. NLO calculations for weak boson production with up to 5 jets have been performed [146–148], with parton shower matching and multi-jet merging [149–152] applied for lower multiplicity cases, and all-orders logarithmic corrections have been derived in the case of jets well separated in rapidity [153, 154].

QED and EW corrections to DY production are similarly well developed, and a

variety of implementations of different higher order corrections are available [124, 155–171], with benchmarking between different approaches and approximations an area of continued effort [172]. NLO EW results have been produced for jet multiplicities up to and including 3 [173], and the calculation of full NLO EW corrections has recently been automated by two separate collaborations [174, 175]. The calculation of mixed QCD+EW corrections has also been performed in the pole approximation [176, 177], amounting to a determination of cross terms in the α and α_s perturbative series for the first time.

In this chapter we will study the phenomenology of both NC and CC DY using the VJ calculations implemented within the NNLOJET framework as outlined in Section. 2.6. We begin in Section 3.2 by discussing the use of fixed-order predictions for evaluation of the p_T^V spectrum of gauge bosons. In Section 3.3, we show a comparison of these fixed-order NNLO results with data from the CMS collaboration before incorporating state-of-the-art resummation to extend the kinematic validity of the results in Section 3.4.

3.2 Transverse Momentum Distributions of Vector Bosons in the Drell-Yan Process

The transverse momentum distributions of vector bosons play a particularly important role in our understanding of SM physics at hadron-hadron colliders. Different kinematic regimes of this observable probe various aspects of the predictions, including resummation and non-perturbative effects at low p_T^V , fixed-order predictions at intermediate to high p_T^V , and electroweak Sudakov logarithms at very high p_T^V . As such, detailed theory-data comparisons of this observable are useful probes for testing Standard Model predictions and the regions of their validity. The p_T^V distribution can also provide important constraints in the fit of parton distribution functions (PDFs), as was studied in Ref. [178] for the case of the Z-boson spectrum.

Ratios between the p_T^V spectra of different processes, such as W^-/W^+ and Z/W , shed further light on the composition of the proton and are also important inputs to precision measurements. Most notably, a precise theoretical understanding of the ratio between Z- and W-boson production cross sections is of critical importance in the measurement of the W-boson mass [111, 112], where the modelling of the W-boson p_T^V spectrum is obtained indirectly through p_T^Z as it can't be measured directly.

The p_T^V spectra of weak gauge bosons have been studied by the CDF [179] and DØ [180–186] collaborations at the Tevatron collider, and corresponding measurements at the LHC have been performed by the ATLAS [187, 188], CMS [189–191], and LHCb [192, 193] experiments. These measurements constitute an integral part of the current and future LHC programme. The corresponding ratios of distributions have also been studied, and the CMS collaboration have produced results for the W^-/W^+ and Z/W ratios using 8 TeV data [191].

3.3 Fixed Order Predictions

Using the calculation of $W^\pm J$ production through to $\mathcal{O}(\alpha_s^3)$ from Section 2.6, we can calculate the $\mathcal{O}(\alpha_s^3)$ NNLO QCD corrections to DY production at finite transverse momentum regulated by minimum value $p_T^{V,\text{cut}}$ including leptonic decay,

$$p + p \rightarrow W^\pm (\rightarrow \ell + \nu_\ell) \Big|_{p_T^V > p_T^{V,\text{cut}}} + X, \quad (3.3.1)$$

which is closely related to the Z transverse momentum distribution as discussed in [88, 89, 194]. This is made possible as the EW bosons produced at Born level in DY have $p_T^V = 0$ due to momentum conservation. The first non-trivial order for the p_T^V distribution occurs at $\mathcal{O}(\alpha_s)$, where the recoil of the QCD radiation allows the boson to move away from the beam axis. Consequently one can use the calculation of $V + \text{jet}$ production if the IR divergence at $p_T^V = 0$ is regulated by some minimum transverse momentum cut $p_T^{V,\text{cut}}$.

In order to test these predictions, we adopt the setup of the CMS measurement of Ref. [191] and perform a comparison of the predictions to data for the normalised p_T^V distributions for W- and Z-boson production and their ratios. The measurement is performed independently for electrons and muons in the fiducial volumes defined by

$$\begin{aligned} p_T^e &> 25 \text{ GeV}, \quad p_T^\mu > 20 \text{ GeV}, \\ |\eta^e| &< 2.5, \quad |\eta^\mu| < 2.1 \end{aligned} \tag{3.3.2}$$

For the neutral-current process, an additional invariant-mass cut of

$$60 \text{ GeV} < m_{ll} < 120 \text{ GeV} \tag{3.3.3}$$

is imposed on the lepton-pairs, and for the charged current process, p_T^W is reconstructed from the vector sum of the lepton and neutrino momenta. No cut is placed on the neutrino momentum, in line with the treatment of the missing transverse momentum in the experimental analysis [191]. The transverse-momentum distributions we consider here are $\mathcal{O}(\alpha_s^3)$, where final-state QCD emissions are treated fully inclusively whilst using a transverse-momentum cut of $p_T^{V,\text{cut}} > 7.5 \text{ GeV}$ on the vector bosons to regulate the cross-section. This cut value is chosen in order to align with the upper edge of the first p_T^V bin of the charged-current and ratio measurements.

The normalisation of the distributions is obtained from the calculation of inclusive vector boson production to $\mathcal{O}(\alpha_s^2)$. For the PDFs, we employ the central member of the NNPDF31_nnlo [117] set with $\alpha_s(M_Z) = 0.118$ for all predictions at LO, NLO, and NNLO. We use the G_μ scheme, with the electroweak parameters set to the most recent PDG values [8], equal to

$$\begin{aligned} M_Z &= 91.1876 \text{ GeV}, \quad M_W = 80.379 \text{ GeV}, \\ \Gamma_Z &= 2.4952 \text{ GeV}, \quad \Gamma_W = 2.085 \text{ GeV}, \\ G_F &= 1.1663787 \times 10^{-5} \text{ GeV}^{-2}. \end{aligned} \tag{3.3.4}$$

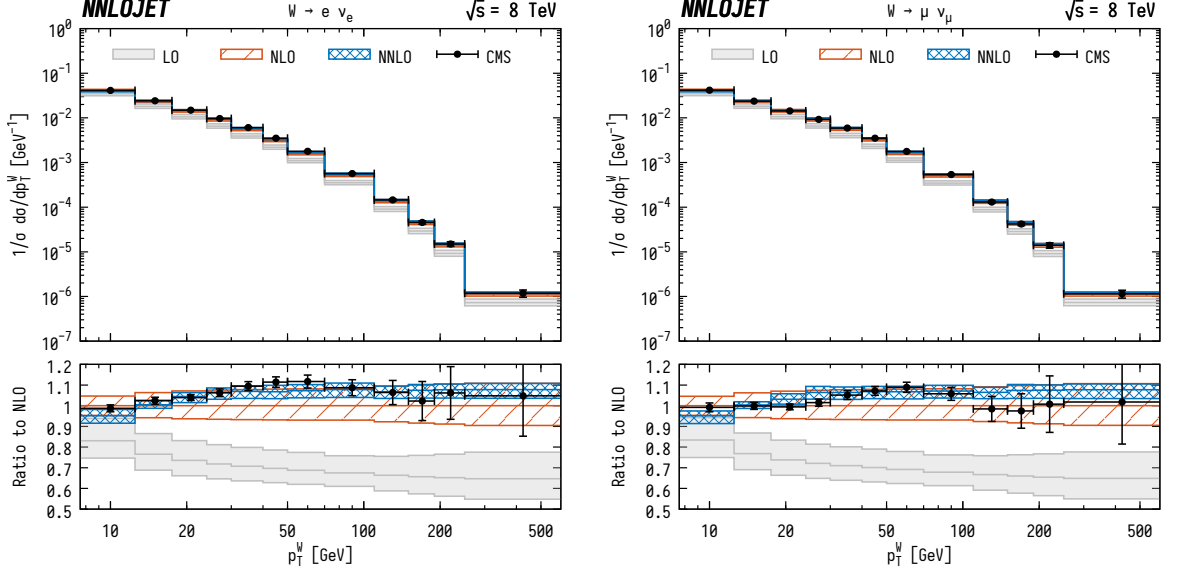


Figure 3.3: Normalised p_T^W distribution for $W = W^+ + W^-$ production. The left hand panel shows predictions for the electron decay channel and the right hand panel shows the muon channel. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text.

We use a diagonal CKM matrix, which for the distributions shown have been verified as accurate to the per-mille level at NLO.

In order to assess the theory uncertainties, we independently vary the factorisation (μ_F) and renormalisation (μ_R) scales by factors of $\frac{1}{2}$ and 2 around the central scale μ_0 , while imposing the restriction $\frac{1}{2} \leq \mu_F/\mu_R \leq 2$. The central scale choice is given by the transverse energy

$$\mu_0 = E_T \equiv \sqrt{M_V^2 + (p_T^Y)^2}, \quad (3.3.5)$$

where M_V and p_T^Y denote the invariant mass and transverse momentum of the final-state lepton pair. For the ratios and double-ratios encountered in the normalised distributions and their ratios, we generalise this procedure and consider the uncorrelated variation of all scales appearing inside the different parts while imposing $\frac{1}{2} \leq \mu/\mu' \leq 2$ between all pairs of scales. This prescription results in a total of 31 and 691 points in the scale variation of the normalised distributions and their ratios, respectively.




Figure 3.4: Normalised p_T^Z distribution for Z production with subsequent decay into muons. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text.

Normalised Fiducial Distributions

The left and right panels in Figure 3.3 show the normalised transverse-momentum distribution of the W boson decaying in the electron and muon channels respectively. In the following, the label “ $W \rightarrow \ell \nu_\ell$ ” denotes the sum of both the $W^- \rightarrow \ell^- \bar{\nu}_\ell$ and $W^+ \rightarrow \ell^+ \nu_\ell$ processes. The NLO corrections are between 10–40% with residual scale uncertainties at the level of around $\pm 10\%$. Although the scale-uncertainty bands at NLO mostly cover the experimental data points, systematic differences in the shape between data and the central theory prediction are visible, which in view of the experimental precision clearly demonstrates the necessity of higher-order predictions with smaller uncertainties in order to discriminate such behaviours. The NNLO corrections are positive and between 5–10% in the intermediate- to high- p_T^W region. Towards lower p_T^W , the NNLO corrections become smaller and turn negative in the lowest- p_T^W bin. The residual scale uncertainties reduce to the level of about $\pm 2\%$ and overlap with the NLO scale bands, exhibiting good perturbative convergence. Most notably, we observe that the shape distortion induced by the NNLO corrections brings the central predictions in line with the measured distributions.

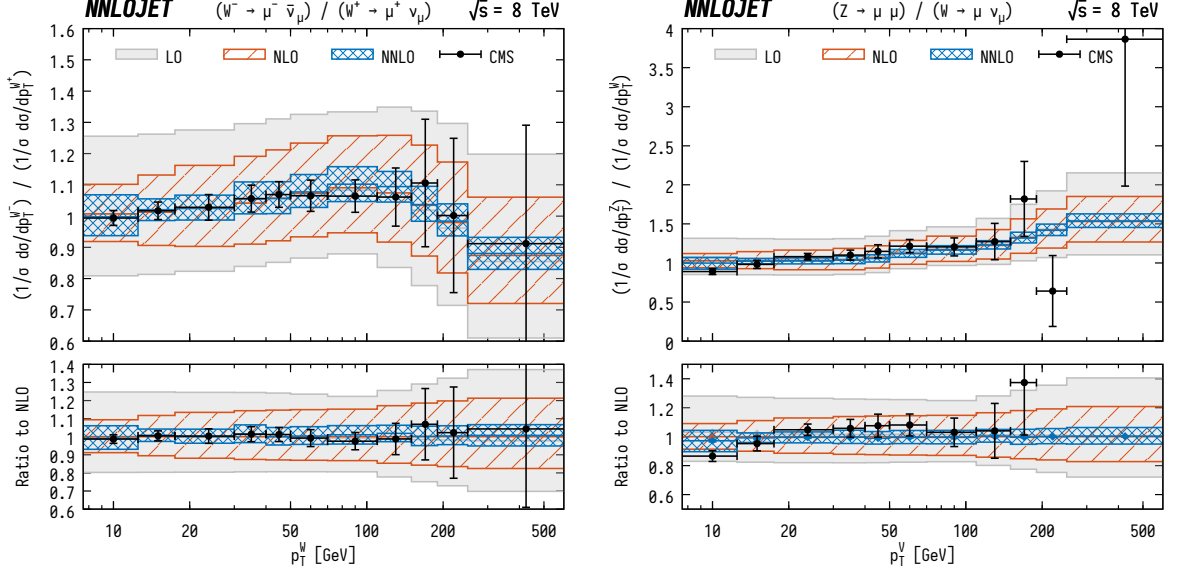


Figure 3.5: Ratio of normalised p_T^V distributions for decays in the muon channel. The left hand panel shows the ratio between W^- and W^+ production, and the right hand panel shows the ratio between Z and $W = W^+ + W^-$ production. Predictions at LO (grey fill), NLO (orange hatched), and NNLO (blue cross-hatched) are compared to CMS data from Ref. [191]. The bands correspond to scale uncertainties estimated as described in the main text.

The corresponding comparison for the Z -boson spectrum is shown in Fig. 3.4, where the measurement was only performed in the muon channel. As in the charged-current case, there is a substantial reduction in the scale uncertainties accompanied by an improved description of the shape. We note that the vertical range of the bottom panel in Fig. 3.4 has been increased compared to the respective figures of the charged-current process in order to accommodate the experimental data which exhibit larger statistical fluctuations due to the smaller total cross-section.

Ratios of Fiducial Distributions

The left hand panel of Figure 3.5 shows the ratio between the normalised distributions of the $W^- \rightarrow \ell^- \bar{\nu}_\ell$ and $W^+ \rightarrow \ell^+ \nu_\ell$ processes. The ratio is close to one in the lowest p_T^W bin and rises up to ~ 1.1 at $p_T^W \approx 150$ GeV, where it turns over and slowly decreases to 0.9 at $p_T^W = 500$ GeV, a shape which can be traced to the fact that down-type valence quarks typically carry a slightly softer distribution than the up-type counterparts in the PDFs which impacts the relative production rates of the

W^\pm bosons due to their different charges. At low p_T^V , the production is sea-quark dominated such that the W^\pm production rates converge (see Fig. 1.6). The central predictions remain remarkably stable between the perturbative orders, resulting in k -factors that are very close to one. However, the precision of the theory prediction is substantially improved by going to higher orders: While the scale uncertainties at NLO are between ± 10 – 20% , the NNLO corrections reduce the uncertainties to the level of $\pm 5\%$ across most of the p_T^W range, never exceeding $\pm 10\%$.

The ratio between the Z- and W-boson processes are shown in the right hand panel of Fig. 3.5. Here, the ratio is again close to one in the low- p_T^V bin and shows a steady increase towards higher p_T^V , reaching about 1.5 at $p_T^V \sim 500$ GeV. As was the case for the W^-/W^+ ratio, the QCD corrections are very stable and leave the central predictions largely unaffected, supporting the expected similarity of QCD corrections between Z and W production. The higher-order corrections again have a big impact on the scale uncertainties, which are reduced by more than a factor of two across almost all p_T^V -bins by going from NLO to NNLO and are at the level of ± 5 – 10% .

Throughout this section, our predictions have been based solely on fixed-order perturbation theory, which produces logarithmic terms of the form $L = \log(p_T^V/M_V)$ at each order in the expansion. In the limit of very large or very small transverse momentum, these logarithms can become potentially large, thereby affecting the convergence of the fixed-order expansion and requiring the use of techniques such as all-orders resummation.

The onset of these effects can be seen in the lowest p_T^V bin of the figures shown thus far, where the scale dependence and uncertainty is substantially larger than in the other bins. Detailed studies of the p_T^Z distribution [88, 89, 194], where multi-differential data are available for low values of the transverse momentum, have indicated that inclusion of the NNLO corrections considerably extends the range where fixed-order perturbation theory provides reliable predictions. However, in order to obtain a description valid throughout the full transverse momentum range

one must incorporate the matching of the NNLO predictions with resummation of large logarithmic corrections [132]. At large p_T^V , logarithmic QCD contributions are only very moderate in size [195], while electroweak Sudakov logarithms can become potentially important in individual distributions [196–199], although they largely cancel in the ratios. At low p_T^V , the dominant logarithmic contributions come from QCD effects and can also be resummed to allow theory predictions to probe the $p_T^V \rightarrow 0$ limit in combination with fixed-order results.

3.4 Fixed Order Matched to N³LL p_T^V Resummation

In order to extend our fixed-order results into the low p_T^V regime through combination with p_T^V resummation, we must first define the counting through which we understand the accuracy of the resummation. In the low p_T regime, the logarithmic accuracy is defined in terms of the logarithm of the cumulative cross-section Σ as

$$\begin{aligned} \ln \left(\Sigma(p_T^V) \right) &\equiv \ln \left(\int_0^{p_T^V} dp_T^{V'} \frac{d\Sigma(p_T^{V'})}{dp_T^{V'}} \right) \\ &= \sum_n \left\{ \mathcal{O} \left(\alpha_s^n L^{n+1} \right) + \mathcal{O} \left(\alpha_s^n L^n \right) + \dots \right\}. \end{aligned} \quad (3.4.1)$$

One refers to the dominant terms $\alpha_s^n L^{n+1}$ as leading logarithmic (LL), to terms $\alpha_s^n L^n$ as next-to-leading logarithmic (NLL), to $\alpha_s^n L^{n-1}$ as next-to-next-to-leading logarithmic (NNLL), and so on. These logarithms appear as an artefact of the cancellation between real and virtual contributions in divergent limits, where they appear as finite terms that accompany the cross-section as it approaches the strictly singular limit. This miscancellation occurs as the virtual diagrams lie directly in the Born phase space, unlike the real emissions against which they cancel in the total cross-section. These real emissions therefore generate a divergent p_T^V contribution close to the $p_T^V = 0$ limit which cannot be compensated for by the virtual component

at finite p_T^V values. At $p_T^V = 0$, these contributions manifest themselves as the usual singularities which directly cancel, such that the p_T^V integrated DY cross-section is finite.

At low transverse momenta, there are two real kinematic configurations that contribute to the cross-section in the low p_T^V region. The first occurs when all outgoing partons recoiling against the vector boson are soft, giving a low transverse momentum to the boson itself. This is exponentially suppressed and sub-dominant to the second contribution, in which one has multiple hard emissions which cancel vectorially to give a small transverse momentum to the recoiling system. This configuration is only power suppressed and thus provides the dominant contribution as one approaches the singular limit.

The resummation of the p_T^V spectrum of SM bosons has been studied in a multitude of theoretical formulations throughout the years [129, 130, 200–207], and the current state of the art for phenomenological studies at the LHC reaches next-to-next-to-next-to leading logarithm (N³LL) accuracy [144, 207–209]. In order to extend our fixed-order results into the low p_T^V regime through combination with Q_T resummation, we can utilise the existing resummation results which exist at N³LL as implemented in the RADISH program [205, 207], which performs the transverse momentum resummation in direct space as opposed to impact parameter space.

In this direct space approach, the factorisation properties of the QCD matrix elements in the soft and collinear limits are exploited to devise a numerical procedure to generate the radiation at all perturbative orders. This allows for the resummation of the large logarithmic terms in a fashion that is similar in spirit to a Monte Carlo generator. A detailed technical description of the method can be found in [205, 207], and the formulae up to N³LL accuracy can be found in [144].

In order to have a reliable prediction across the whole p_T^V spectrum, the fixed-order and resummed predictions must be consistently combined through a matching procedure. The matching is performed in such a way that it reduces to the resummed

calculation at small p_T^V , while reproducing the fixed-order prediction at large transverse momentum. At a given perturbative order, one can adopt various schemes that differ from one another by terms beyond the considered order. Here we adopt the multiplicative scheme formulated in [144, 210], in which the matching is performed at the level of the cumulative distribution (3.4.1) as follows:

$$\Sigma_{\text{match}}^{\text{N}^3\text{LL}}(p_T^V) = \frac{\Sigma^{\text{N}^3\text{LL}}(p_T^V)}{\Sigma_{\text{asym.}}^{\text{N}^3\text{LL}}} \left[\Sigma_{\text{asym.}}^{\text{N}^3\text{LL}} \frac{\Sigma^{\text{N}^3\text{LO}}(p_T^V)}{\Sigma^{\text{EXP}}(p_T^V)} \right]_{\text{N}^3\text{LO}}, \quad (3.4.2)$$

where Σ^{EXP} denotes the expansion of the resummation formula $\Sigma^{\text{N}^3\text{LL}}$ to $\mathcal{O}(\alpha_s^3)$ (N³LO), and the whole squared bracket in Eqn. (3.4.2) is expanded to N³LO. It should be recalled that $\mathcal{O}(\alpha_s^3)$ corresponds to N³LO in the total (i.e. p_T^V -integrated) cross-section and in any cumulative distribution, while being NNLO in the fixed-order p_T^V -differential cross-section.

In the above equation, $\Sigma^{\text{N}^3\text{LO}}$ is the cumulative fixed-order distribution at N³LO, i.e.

$$\Sigma^{\text{N}^3\text{LO}}(p_T^V) = \sigma_{\text{tot}}^{\text{N}^3\text{LO}} - \int_{p_T^V}^{\infty} dp_T^{V'} \frac{d\Sigma^{\text{NNLO}}(p_T^{V'})}{dp_T^{V'}}, \quad (3.4.3)$$

where $\sigma_{\text{tot}}^{\text{N}^3\text{LO}}$ is the total cross-section for the charged or neutral DY processes at N³LO, and $d\Sigma^{\text{NNLO}}/dp_T^{V'}$ denotes the corresponding NNLO p_T^V -differential distribution obtained with NNLOJET. Both of these quantities are accurate to $\mathcal{O}(\alpha_s^3)$. Since the N³LO inclusive cross-section for DY production is currently unknown, we approximate it with the NNLO cross-section [121–123, 125, 211–215] in the following. This approximation impacts only terms at N⁴LL order, and is thus beyond the accuracy considered in this study.

Finally, the quantity $\Sigma_{\text{asym.}}^{\text{N}^3\text{LL}}$ is defined as the asymptotic ($p_T^V \gg M$) limit of the resummed cross-section

$$\Sigma^{\text{N}^3\text{LL}}(p_T^V) \xrightarrow{p_T^V \gg M} \Sigma_{\text{asym.}}^{\text{N}^3\text{LL}}. \quad (3.4.4)$$

This prescription ensures that, in the $p_T^V \gg M$ limit, Eq. (3.4.2) reproduces by construction the fixed-order result, while in the $p_T^V \rightarrow 0$ limit it reduces to the resummed prediction. The detailed matching formulae are given in Appendix A

of [144]. This matching scheme is also applicable at lower perturbative orders, NNLL+NLO and NLL+LO, which can be obtained without any approximation to σ_{tot} as follows:

$$\Sigma_{\text{match}}^{\text{NNLL}}(p_T^V) = \frac{\Sigma^{\text{NNLL}}(p_T^V)}{\Sigma_{\text{asym.}}^{\text{NNLL}}} \left[\Sigma_{\text{asym.}}^{\text{NNLL}} \frac{\Sigma^{\text{NNLO}}(p_T^V)}{\Sigma^{\text{EXP}}(p_T^V)} \right]_{\text{NNLO}}, \quad (3.4.5)$$

$$\Sigma_{\text{match}}^{\text{NLL}}(p_T^V) = \frac{\Sigma^{\text{NLL}}(p_T^V)}{\Sigma_{\text{asym.}}^{\text{NLL}}} \left[\Sigma_{\text{asym.}}^{\text{NLL}} \frac{\Sigma^{\text{NLO}}(p_T^V)}{\Sigma^{\text{EXP}}(p_T^V)} \right]_{\text{NLO}}. \quad (3.4.6)$$

These matching schemes guarantee that in the large- p_T^V limit the matched cumulative cross-sections reproduce the total cross-sections

$$\begin{aligned} \Sigma_{\text{match}}^{\text{N}^3\text{LL}}(p_T^V) &\xrightarrow[p_T^V \gg M]{} \sigma_{\text{tot}}^{\text{NNLO}}, \\ \Sigma_{\text{match}}^{\text{NNLL}}(p_T^V) &\xrightarrow[p_T^V \gg M]{} \sigma_{\text{tot}}^{\text{NNLO}}, \\ \Sigma_{\text{match}}^{\text{NLL}}(p_T^V) &\xrightarrow[p_T^V \gg M]{} \sigma_{\text{tot}}^{\text{NLO}}. \end{aligned} \quad (3.4.7)$$

Here we note that $\Sigma_{\text{match}}^{\text{N}^3\text{LL}}$ reproduces the NNLO total cross-section at large p_T^V as the total N³LO cross-section for the inclusive DY process is currently unknown.

In order to extract normalised p_T^V distributions, one can differentiate Eqns. (3.4.2), (3.4.5) and (3.4.6), and divide by the respective total cross-sections of the right hand side of Eq. (3.4.7) to give

$$\frac{1}{\sigma} \frac{d\Sigma}{dp_T^V}. \quad (3.4.8)$$

Note that the rate of change of the cumulative cross-section with respect to p_T^V is equal to the rate of change of the usual cross-section for (3.4.3). As was the case for the CMS analysis of [191], normalising the distributions is preferred for precision experimental analyses that rely primarily on shapes to extract information. This gives a substantial reduction in systematic uncertainties, most notably for the luminosity which gives an $\mathcal{O}(2\%)$ error on the fiducial distributions [191]. The W mass extraction can be treated as one such shape-based analysis, through sensitivities of the shapes of the W transverse mass and lepton p_T spectra to M_W [112].

When probing the low p_T^V region of the spectrum, one also becomes sensitive to non-

perturbative effects. In particular, the resummed calculation contains a Landau singularity arising from configurations where the radiation is generated with transverse momentum scales

$$k_T \sim M \exp \{ -\pi / (\beta_0 \alpha_s) \} \sim \mathcal{O}(200) \text{ MeV} \quad (3.4.9)$$

(with $\alpha_s = \alpha_s(M)$ and $\beta_0 = (11 C_A - 2 N_F)/6$). This singularity comes from the running of α_s into the region where (1.3.19) diverges. Within the resummation calculation, the results are set to zero when the hardest radiation's transverse momentum reaches the singularity¹. We note that for the invariant masses studied here, this procedure acts on radiation emitted at very small transverse momentum that, due to the vectorial nature of the p_T^V observable [200, 207], gives a very moderate contribution to the spectrum. However, for a precise and accurate description of this kinematic regime, a more thorough study of the impact of non-perturbative corrections would be required.

For the fixed-order results which we match to the resummation, we adopt a similar approach to the previous section in which we integrate down to a transverse momentum cut $p_T^{V,\text{cut}}$. However due to the enhanced low p_T accuracy afforded by the resummation, this cut can be set far more aggressively. As a result, we adopt $p_T^{V,\text{cut}} = 1 \text{ GeV}$ such that we gain a much higher resolution to low p_T effects which are relevant for e.g. W boson mass extractions. Running the VJ calculation down to transverse momenta of 1 GeV is exceptionally computationally expensive, as one numerically probes the emergence of the IR divergence as the partonic emission of the VJ goes unresolved. However, the implementation of the antenna subtraction method in NNLOJET has proven to be both robust and numerically stable [208] allowing the use of the fixed-order results in this region. It remains to be seen whether the same applies to the VJ calculations performed using the N -jettiness slicing method.

¹Due to the subtraction procedure, these contributions are systematically removed from the fixed-order prediction, as they correspond to unresolved soft limits.

We now move to the results of this matched calculation. We consider pp collisions at a centre-of-mass energy of 13 TeV, and we use the NNLO NNPDF3.1 parton distribution function set [117] with $\alpha_s(M_Z) = 0.118$. We evolve the parton densities with LHAPDF [39], which implements the relevant heavy quark thresholds in the PDFs. All convolutions within RADISH are instead handled with the **Hoppet** package [216], and the EW scheme is chosen as in the previous section for the CMS analysis.

The final state for the neutral DY production is defined by applying the following set of fiducial selection cuts on the leptonic pair:

$$\begin{aligned} p_T^\ell &> 25 \text{ GeV}, \quad |\eta^\ell| < 2.5, \\ 66 \text{ GeV} &< m_{ll} < 116 \text{ GeV}, \end{aligned} \tag{3.4.10}$$

where p_T^ℓ are the transverse momenta of the two leptons, η^ℓ are their pseudo-rapidities in the hadronic centre-of-mass frame, and m_{ll} is the invariant mass of the di-lepton system. As in the previous section, the central factorisation and renormalisation scales are chosen to be $\mu_R = \mu_F = E_T$ and the central resummation scale is set to $Q = m_{ll}/2$.

In the case of charged DY production, the fiducial volume is defined as

$$\begin{aligned} p_T^\ell &> 25 \text{ GeV}, \quad E_T^{\text{miss}} > 25 \text{ GeV}, \\ |\eta^\ell| &< 2.5, \quad m_T^V > 50 \text{ GeV}, \end{aligned} \tag{3.4.11}$$

where E_T^{miss} is the missing transverse energy vector and the transverse mass

$$m_T^V = \sqrt{(p_T^V)^2 + m_V^2}. \tag{3.4.12}$$

These fiducial regions are those chosen for the upcoming 13 TeV ATLAS W mass extraction. The central factorisation and renormalisation scales are also chosen to be $\mu_R = \mu_F = E_T$ and the central resummation scale is again set to $Q = m_{l\nu}/2$.

In both processes, we assess the missing higher order uncertainty by performing a variation of the renormalisation and factorisation scales by a factor of two around

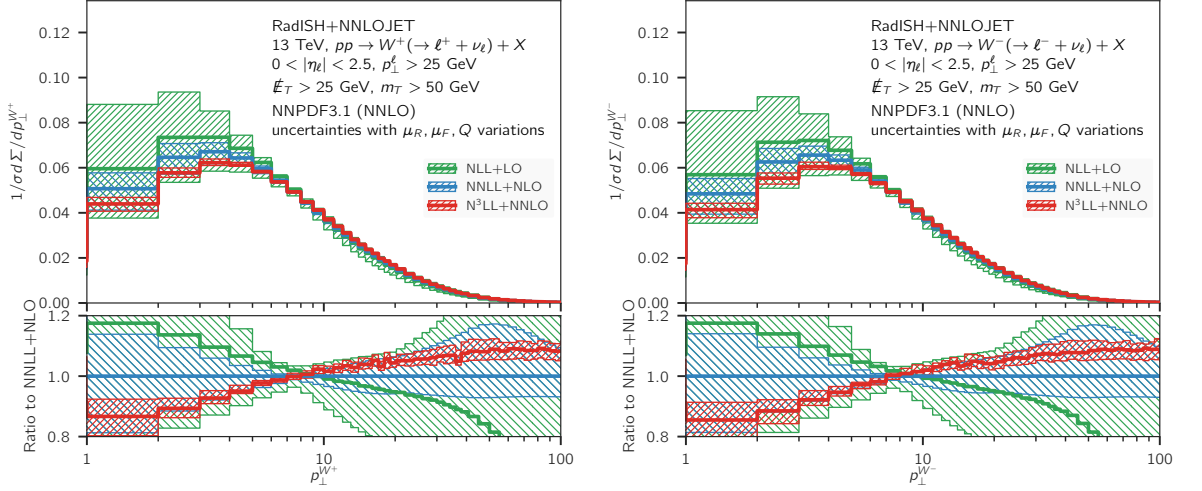


Figure 3.6: W^+ (left) and W^- (right) differential p_T^V distributions normalised to the inclusive cross-section at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. The lower panel shows the ratio of the results to NNLL+NLO. The uncertainty prescriptions for the theory uncertainty are as described in the text.

their respective central values whilst keeping $1/2 \leq \mu_R/\mu_F \leq 2$. In addition, for central factorisation and renormalisation scales, we vary the resummation scale Q by a factor of two in either direction. The final uncertainty is built as the envelope of the resulting 9 scale variations.

Normalised Fiducial Distributions

We start by showing, in Figures 3.6 and 3.7, the comparison of the Z and W^\pm normalised distributions at NLL+LO (green), NNLL+NLO (blue), and N³LL+NNLO (red) in the fiducial volumes defined above. The difference between each prediction and the next order is of $\mathcal{O}(\alpha_s)$, both in the large p_T region and in the limit $p_T \rightarrow 0$ where $\alpha_s L \sim 1$. In comparison to the NNLL+NLO result, we note that the N³LL+NNLO corrections lead to important distortions in the shape of the distributions, making the spectrum harder for $p_T \gtrsim 10$ GeV, and softer below this scale. The perturbative errors are reduced by more than a factor of two across the whole p_T range, and the residual uncertainty is at the 5% level. As expected, we see that each of the W^\pm and Z distributions displays the same behaviour, although the inclusion of the N³LL+NNLO corrections gives a $\mathcal{O}(5\%)$ smaller shape distortion




Figure 3.7: The Z differential p_T^V distributions normalised to the inclusive cross-section at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. The lower panel shows the ratio of the results to NNLL+NLO. The uncertainty prescriptions for the theory uncertainty are as described in the text.

around 50 GeV in the neutral current case.

In general, we observe a good convergence of the perturbative description when the order is increased, although in some p_T regions the N³LL+NNLO and the NNLL+NLO bands overlap only marginally. This feature can be explained by the fact that both predictions are normalised to the same NNLO total cross-section. Since at large p_T the NNLO corrections lead to an increase in the spectrum of about 10%, by unitarity of the matching procedure (that preserves the total cross-section) this must be balanced by an analogous decrease in the spectrum in the region governed by resummation, as we indeed observe. Despite this, the two orders are compatible within the quoted uncertainties and one can consider the predictions to be robust.

In Figures 3.8 and 3.9, we show the comparison among the NNLO (green), the NNLL+NLO (blue), and N³LL+NNLO (red) predictions, alongside Monte Carlo predictions obtained using the *Pythia* 8 generator [217] with the AZ tune [218], which was tuned to the Z-boson p_T distribution at 7 TeV. At 7 TeV and 8 TeV the above tune is known to correctly describe the Z transverse momentum spectrum within a few percent for $p_T \lesssim 50$ GeV [218], and has been employed in a 7 TeV




Figure 3.8: W^+ (left) and W^- (right) differential p_T^V distributions normalised to the inclusive cross-section at NNLO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the `AZ` tune [218]. The lower panel shows the ratio of the results to the aforementioned `Pythia` results. The uncertainty prescriptions for the theory uncertainty are as described in the text.

extraction of the W -boson mass by the ATLAS collaboration [112]. However, it is currently unknown how this tune performs at 13 TeV in comparison to the data. Understanding these limitations, we use the `Pythia 8` prediction as a proxy for the as yet unreleased data. The lower panel shows the ratio of each prediction to the `Pythia 8` tune.

Here we see for the first time the drastic improvement of the pure fixed-order NNLO predictions through the matching with resummation. In particular, the fixed-order and matched predictions differ below $p_T^V \lesssim 20$ GeV, where the fixed-order exhibits a much harder distribution which is offset by a negative contribution in the lowest bin as the fixed order perturbative expansion breaks down. Further evidence of this is seen in the scale variation of the fixed-order, which rapidly increases as $p_T^V \rightarrow 0$. Above $p_T \sim 20$ GeV, the NNLO results can be considered to provide a reliable theoretical prediction. We observe reasonable agreement between the N³LL+NNLO predictions and `Pythia 8` below 30 GeV, while it deteriorates for larger p_T values, a feature which is particularly visible in the case of W^\pm production.




Figure 3.9: The Z differential p_T^V distributions normalised to the inclusive cross-section at NNLO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volume described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the `AZ` tune [218]. The lower panel shows the ratio of the results to the aforementioned `Pythia` results. The uncertainty prescriptions for the theory uncertainty are as described in the text.

Ratios of Fiducial Distributions

Another set of important quantities of interest are ratios of the above distributions, which as previously discussed play a central role in recent extractions of the W boson mass at the LHC [112]. When taking ratios of perturbative quantities one has to decide how to combine the uncertainties in the numerator and denominator to obtain the final error. This becomes much more important for SM parameter determinations, where the uncertainties are direct inputs into e.g. χ^2 fits, unlike simple comparisons to data.

One option is to try to identify the possible sources of correlation in the three processes considered here. From the point of view of the perturbative (massless) QCD description we adopted, one expects that the structure of radiative corrections to such reactions is nearly identical. This is certainly the case as far as resummation is concerned, since it is governed by the same anomalous dimensions and all-order structure in W and Z production. As a consequence, the resummation scale should be varied in a correlated manner in both predictions considered in the ratio. A




Figure 3.10: Ratios of Z/W^+ (left) and W^-/W^+ (right) normalised differential distributions at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volumes described in the text. The three lower panels show different prescriptions for the theory uncertainty, as described in the text and normalised to the NNLL+NLO results.

similar argument can be made regarding the renormalisation scale μ_R . However, an important difference between Z, W^+ , and W^- production lies in the different combination of partonic channels probed by each process and, in particular, in the sensitivity to different heavy quark thresholds in the PDFs at small p_T . Therefore, it is not clear whether a fully correlated variation of the factorisation scale μ_F is physically justified.

A more conservative uncertainty prescription is to vary the scales μ_R and Q in numerator and denominator in a fully correlated way, while varying μ_F in an uncorrelated manner within the constraint [1]

$$\frac{1}{2} \leq \frac{x_{\mu_F}^{\text{num.}}}{x_{\mu_F}^{\text{den.}}} \leq 2, \quad (3.4.13)$$

where x_{μ_F} is the ratio of the factorisation scale to its central value. This corresponds to a total of 17 scale combinations.

Finally, for comparison we also consider the uncorrelated variation of μ_R and μ_F in the ratio, while imposing

$$\frac{1}{2} \leq \frac{x_{\mu}^{\text{num.}}}{x_{\mu}^{\text{den.}}} \leq 2, \quad (3.4.14)$$

where x_{μ} is the ratio of the scale μ to its central value, with $\mu \equiv \{\mu_R, \mu_F\}$, together with a correlated variation of the resummation scale Q . This recipe amounts to taking the envelope of the predictions resulting from 33 different combinations of scales in the ratio, and is equivalent to the scale prescription used in the previous section accompanied by fully correlated resummation scale variation. Note that here the scale variation of the total cross-section is neglected unlike in the previous section, which again leads to a more aggressive scale uncertainty evaluation.

To examine the reliability of the above uncertainty schemes, in Figure 3.10 we analyse the convergence of the perturbative series for the ratios of distributions, by comparing the results at NLL+LO (green), NNLL+NLO (blue), and N³LL+NNLO (red). The three lower panels in each plot show the theory uncertainties obtained according to the three prescriptions outlined above, respectively, in comparison to

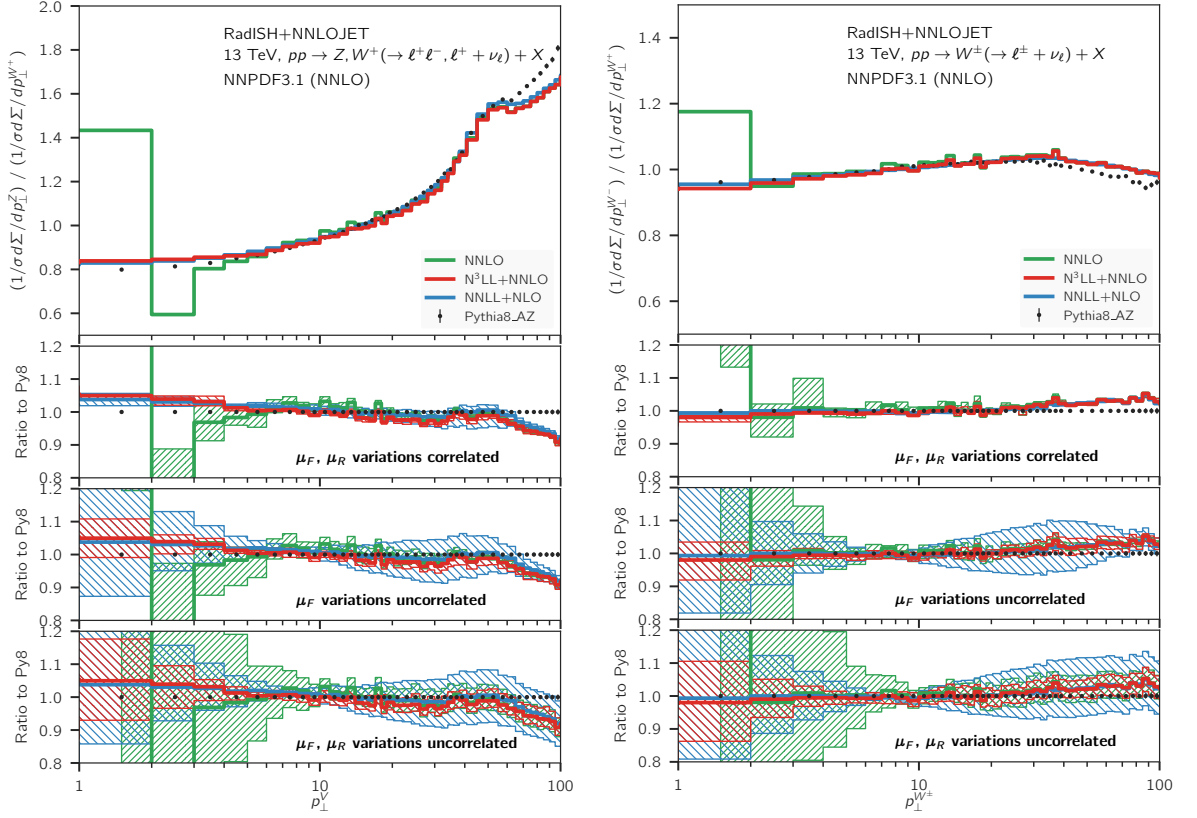


Figure 3.11: Ratios of Z/W^+ (left) and W^-/W^+ (right) normalised differential distributions at NLL+LO (green), NNLL+NLO (blue) and N³LL+NNLO (red) at $\sqrt{s} = 13$ TeV for the fiducial volumes described in the text. These are presented alongside Monte Carlo predictions obtained using the `Pythia 8` generator [217] with the AZ tune [218]. The three lower panels show different prescriptions for the theory uncertainty, as described in the text and normalised to the aforementioned Pythia results. The lower panel shows the ratio of the

the old state-of-the-art prediction at NNLL+NLO.

In the case of the Z/W^+ ratio (shown in the upper plot of Figure 3.10), we observe that the different perturbative orders are very close to one another. The results are compatible even within the uncertainty bands obtained with the more aggressive error estimate, which in some bins is sensitive to minor statistical fluctuations due to the the statistical requirements of the NNLO calculation. This feature is strikingly evident in the case of the W^-/W^+ ratio (right plot), where the excellent convergence of the series indicates that a scale variation of the type (3.4.13) is well justified.

Figure 3.11 shows the comparison of the same two ratios (Z/W^+ and W^-/W^+) to the NNLO result (green), and to `Pythia 8`. We observe that in both cases the

N³LL+NNLO calculation leads to an important reduction of the theory uncertainty. In particular, even with the most conservative estimate of the theory error, our best prediction leads to errors of the order of $\mathcal{O}(5\%)$, with the exception of the first bin where the perturbative uncertainty is at the $\mathcal{O}(10\%)$ level.

The kink around $p_T \sim 50 - 60$ GeV in the Z/W^+ ratio (upper plot in Figure 3.11) is due to the difference in fiducial selection cuts between the neutral and charged current processes, in particular the $m_T^W > 50$ GeV and $m_{ll} > 66$ GeV bounds. A change in the shape of the distributions around this scale is indeed visible in both Figures 3.8 and 3.9, at slightly different p_T values for Z and W^\pm production, which is reflected in the structure observed in Figure 3.11.

We find a good agreement between our best predictions at N³LL+NNLO and the **Pythia 8** Monte Carlo in the small p_T region of the ratios, where the two predictions are compatible within the quoted theory uncertainties. On the other hand, for $p_T \gtrsim 40$ GeV, the **Pythia 8** result disagrees with the matched calculation. This behaviour is not unexpected, since the nominal perturbative accuracy of **Pythia 8** is well below any of the matched calculations, and the **AZ** tune is optimised to describe the Z spectrum in the region $p_T \leq 50$ GeV at 7 TeV. When we take into account the excellent agreement with the CMS analysis in the previous section, we can conclude that the inclusion of the matched fixed-order and resummation results is likely to describe the anticipated data well across a wide p_T^V range, from the soft region to the onset of large EW Sudakov logarithms.

3.5 Summary

In this chapter we have first provided an overview of the Drell-Yan process and its phenomenological importance, before using the fixed order calculation of V plus jet production to calculate the transverse momentum distribution and ratios. These are compared to 8 TeV CMS data from [191] where we observe good agreement with data down to $p_T^V = 7.5$ GeV.

We then combine the fixed-order results with p_T^V resummation provided by the RADISH program in order to extend the range of kinematic validity down to 1 GeV. We provide predictions for 13 TeV data, using the ATLAS fiducial region. We find that resummation effects become important for $p_T^V \lesssim 20$ GeV, and the effect of the N³LL+NNLO corrections with respect to the NNLL+NLO prediction is as large as $\sim 10\%$, giving a significant shape distortion. We find that the ratios are considerably more perturbatively stable, and one can justify a large degree of correlation between numerator and denominator in their construction.

Chapter 4

Phenomenology of Vector Boson Production in Association with a Jet

The production of a vector boson in association with a hadronic jet is perturbatively the simplest hadron-collider process that directly probes both the strong and electroweak interactions at Born level. It has been measured extensively at the Tevatron [222–225] and the LHC [219–221, 226–234], covering a large range in transverse momentum and rapidity of the final-state particles. When compared to theory predictions, these measurements provide important tests of the dynamics of the Standard Model and help to constrain the momentum distributions of partons in the proton.

The study of the forward-rapidity region for this process is particularly important for our understanding of parton distribution functions (PDFs) at extremal values of Bjorken- x , due to the different kinematic regimes that are probed compared to the inclusive case. Owing to the extended rapidity coverage of the LHC experiments, data is now available for both highly boosted leptons and jets, giving direct access to these regions in phenomenological studies.

For a given vector boson plus jet (VJ) event, one can directly infer the valid range

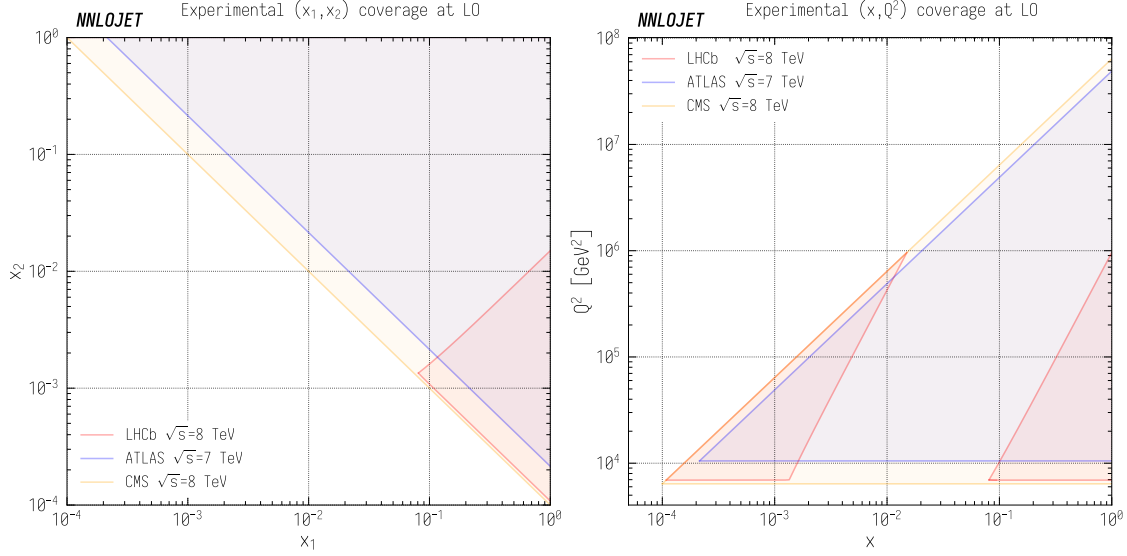


Figure 4.1: The kinematic regions of the (x_1, x_2) (left) and (x, Q^2) (right) planes accessible for the LHCb [219], ATLAS [220] ZJ and CMS [221] $W^\pm J$ selection criteria at LO ($\mathcal{O}(\alpha^2\alpha_s)$). Here x_1 and x_2 are the Bjorken- x values from beams 1 and 2 respectively, Q^2 is the invariant mass of the full final state including both charged leptons and QCD radiation and x is the Bjorken- x from either of the incoming beams.

in Bjorken- x values from the event kinematics at the hadronic centre-of-mass energy \sqrt{s} through

$$\begin{aligned} x_1 &\geq \frac{1}{\sqrt{s}} \left(m_T^V \cdot e^{+y^V} + p_T^{j1} \cdot e^{+y^{j1}} \right), \\ x_2 &\geq \frac{1}{\sqrt{s}} \left(m_T^V \cdot e^{-y^V} + p_T^{j1} \cdot e^{-y^{j1}} \right), \end{aligned} \quad (4.0.1)$$

with $m_T^V = \sqrt{(p_T^V)^2 + m_V^2}$ denoting the transverse mass. x_1 and x_2 correspond to the momentum fractions of the incoming partons present in the colliding protons, p_T^V and p_T^{j1} are the transverse momenta of the vector boson and the leading- p_T jet, m_V is the invariant mass of the combined system of the decay products of the vector boson and the rapidities of the vector boson and the leading jet are labelled y_V and y_{j1} . At Born level, the inequalities in (4.0.1) become strict equalities.

In general, the smallest x value that can be probed simultaneously ($x_1 \sim x_2$) is

$$x_{\min} = \frac{m_{V+j}^{\min}}{\sqrt{s}}, \quad (4.0.2)$$

which is relevant primarily for data where fiducial cuts are symmetric in rapidity.

Here m_{V+j} is the invariant mass of the vector boson plus jet final state at LO. We also have the combined kinematic constraint

$$x_1 x_2 \geq \frac{1}{s} \left(m_T^{V,\min} + p_T^{j1,\min} \right)^2, \quad (4.0.3)$$

where $m_T^{V,\min}$ and $p_T^{j1,\min}$ are the minimal values of the vector boson transverse mass and leading jet p_T admitted by the fiducial cuts. This constraint is particularly relevant in phase-space regions that are asymmetric in rapidity, which in turn probes more asymmetric values in x_1 , x_2 and gives rise to a more complex interplay between the kinematics and the event selection cuts.

There is a strong experimental motivation for precise predictions for VJ processes due to the high statistics and clean decay channels observed at the LHC, their relevance to determinations of Standard Model parameters, and as backgrounds for new physics searches [235]. Fitting procedures for PDFs also benefit greatly from improved predictions, due to the increased sensitivity to the gluon and quark content of the proton [117,178]. Owing to the large gluon luminosity at the LHC, the dominant initial state for vector boson plus jet production is quark–gluon scattering, with different quark flavour combinations probed by the different bosons.

In the following chapter, we perform a comparison between NNLO QCD predictions for vector boson plus jet (VJ) production and measurements by the LHCb [219], ATLAS [220] and CMS [221] experiments. These measurements are highly complementary, allowing one to probe a much larger kinematic region than if any of them were taken alone due to the different rapidity coverages of the three detectors. The region of the (x, Q^2) plane which is probed at LO in QCD in VJ production is shown in the right-hand panel of Fig. 4.1. One can see in this plot that LHCb covers two distinct sectors corresponding to the x values of the two beams as a result of its distinctly asymmetric configuration. The corresponding plot for the (x_1, x_2) plane is shown in the left-hand panel, where the asymmetry of the LHCb region preferentially probes large x_1 and small x_2 values in contrast to the symmetric (x_1, x_2) coverage of the ATLAS and CMS fiducial regions. The kinematic constraints on the

LHCb region are relaxed beyond LO as the presence of radiation permits larger Q^2 and x_2 values, unlike the ATLAS/CMS regions where LO kinematics already fully cover the kinematic region accessible at higher orders. The LO kinematics dominate in the contribution to the total cross section however, and give a good indication of where the sensitivities of the experiments lie.

Throughout this chapter, the theoretical predictions employ a diagonal CKM matrix. The electroweak parameters are set according to the G_μ scheme with the following input parameters:

$$\begin{aligned} M_Z &= 91.1876 \text{ GeV}, & M_W &= 80.385 \text{ GeV}, \\ \Gamma_Z &= 2.4952 \text{ GeV}, & \Gamma_W &= 2.085 \text{ GeV}, \\ G_F &= 1.1663787 \times 10^{-5} \text{ GeV}^{-2}, & & \end{aligned} \tag{4.0.4}$$

and the PDF set used at all perturbative orders is the central replica of NNPDF31_nnlo [117] with $\alpha_s(M_Z) = 0.118$.

4.1 Forward Production of Vector Bosons at LHCb

At a proton-proton centre-of-mass energy of 8 TeV, the LHCb experiment has measured each of $W^\pm J$ and ZJ where the vector bosons decay in the muon channel [219]. The acceptance in the forward region of the LHCb experiment allows it to reliably isolate PDFs contributions at both much higher and lower momentum fractions x than the general-purpose detectors at the LHC. This sensitivity arises from kinematic configurations that are asymmetric in x_1 and x_2 , which in turn means that the event is boosted into the forward region. PDF uncertainties at large x and Q^2 are generally driven by uncertainties in the d content of the proton, which these measurements have the capacity to constrain due to their flavour sensitivity, particularly in the charged-current channels. This provides a strong motivation to use

the NNLO QCD calculation implemented in the NNLOJET framework to test the quantitative agreement of the predictions with the experimental data.

The fiducial cuts applied to the charged leptons and the jets, which we label as the LHCb cuts for both $W^\pm J$ and ZJ production are given by

$$\begin{aligned} p_T^j &> 20 \text{ GeV}, & 2.2 < \eta^j < 4.2, \\ p_T^\mu &> 20 \text{ GeV}, & 2 < y^\mu < 4.5, \\ \Delta R_{\mu,j} &> 0.5, \end{aligned} \tag{4.1.1}$$

where p_T^j and p_T^μ are the transverse momenta of the jets and muons respectively, η^j is the jet pseudorapidity, y^μ is the muon rapidity and $\Delta R_{\mu,j}$ is the angular separation between the leading jet and the muon. In addition, the requirement $p_T^{\mu+j} > 20 \text{ GeV}$ is applied for $W^\pm J$ production, where $p_T^{\mu+j}$ is the magnitude of the transverse component of the vector sum of the charged lepton and jet momenta. For ZJ production, the invariant mass of the di-muon system $m_{\mu\mu}$ is restricted to the window $60 \text{ GeV} < m_{\mu\mu} < 120 \text{ GeV}$ around the Z-boson resonance. The anti- k_T jet algorithm [62] is used throughout, with radius parameter $R = 0.5$. In the original LHCb analysis [219], the VJ data were compared to NLO theory predictions, which were observed to overshoot the data throughout, albeit being consistent within the combined theoretical and experimental uncertainties.

For the theoretical predictions presented in this section, we set the central scale as in [219], i.e.,

$$\mu_R = \mu_F = \sqrt{m_V^2 + \sum_i (p_{T,j}^i)^2} \equiv \mu_0, \tag{4.1.2}$$

with scale variations performed independently for the factorisation and renormalisation scales μ_F , μ_R by factors of $\frac{1}{2}$ and 2 subject to the constraint $\frac{1}{2} < \mu_F/\mu_R < 2$.

The predictions for the fiducial cross section are shown in Table 4.1 for LO, NLO and NNLO QCD and compared to the results reported by the LHCb experiment for the individual VJ channels. We see large corrections when going from LO to NLO as observed in the NLO/LO k -factor of 1.34 for W^- , 1.35 for W^+ and 1.32

	W^+J	W^-J	ZJ
LO	$46.9^{+5.6}_{-2.2}$	$27.2^{+3.2}_{-2.6}$	$4.59^{+0.53}_{-0.43}$
NLO	$62.8^{+3.6}_{-3.5}$	$36.7^{+2.2}_{-2.1}$	$6.04^{+0.32}_{-0.31}$
NNLO	$63.1^{+0.4}_{-0.5}$	$36.8^{+0.3}_{-0.2}$	$6.03^{+0.02}_{-0.04}$
LHCb	$56.9 \pm 0.2 \pm 5.1 \pm 0.7$	$33.1 \pm 0.2 \pm 3.5 \pm 0.4$	$5.71 \pm 0.06 \pm 0.27 \pm 0.07$

Table 4.1: Fiducial cross sections in picobarns for fixed-order theory predictions and LHCb results from Ref. [219]. The errors quoted for NNLOJET correspond to the scale uncertainty and the reported LHCb errors are statistical, systematic and luminosity respectively.

for Z , where the difference in k -factor between the processes is largely driven by the different relative contributions of quark-gluon and quark-quark channels driven by the PDFs. On the other hand, going from NLO to NNLO produces much smaller and more stable results, with a NNLO/NLO k -factor of 1.006 for W^- , 1.003 for W^+ and 0.998 for Z , and the NNLO corrections lie within the scale bands of the NLO results. We note that the uncertainty bands overlap marginally between theory and data in Table 4.1 for W^- and Z production, but not for W^+ production, when added in quadrature.

Distributions Differential in Leading Jet p_T

Figure 4.2 shows the distributions for transverse momentum of the leading jet in W^-J , W^+J and ZJ production respectively. Similarly to the fiducial cross section, the scale dependence of the differential distributions is considerably reduced when going from NLO to NNLO. The NNLO corrections are stable with respect to NLO, indicating a good convergence of the perturbative series. In addition, these results exhibit a strong similarity in behaviour between the W^- , W^+ and Z production channels. We see that the theory overshoots the data by ~ 5 –10% over the bulk of the distribution, rising to 30% in the highest p_T bin. This closely mirrors the effects seen at NLO as well as in the total cross section. The considerable decrease in theory uncertainty from NLO to NNLO makes the tension between data and theory more pronounced.




Figure 4.2: Cross section differential in the leading jet transverse momentum for W^- , W^+ and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text.

For the cuts placed on the $W^\pm J$ final states, we are also able to associate the bins in p_T^j to lower limits on the Bjorken- x invariants. The lowest p_T bin has the loosest constraint on the forward x , with $x_1 > 4.1 \times 10^{-2}$, $x_2 > 5.4 \times 10^{-5}$. However, for the highest p_T bins, between 50 and 100 GeV, the restrictions translate to $x_1 > 7.5 \times 10^{-2}$, $x_2 > 1.1 \times 10^{-4}$. Due to the invariant mass cuts applied in the ZJ case, shown in Fig. 4.2, the smallest values in Bjorken- x that can be probed only extend down to $x_1 > 0.11$, $x_2 > 2 \times 10^{-4}$ in the highest p_T bin. As a result, one probes larger values of x for ZJ production than for $W^\pm J$ in general. At large p_T , we see that the same features are present in the neutral- and charged-current cases, where we observe that the NNLO predictions overshoot the data.

Distributions Differential in Pseudorapidity

The leading jet pseudorapidity distributions in Fig. 4.3 show a similar pattern of deviation between NNLO predictions and data to the previous p_T^j results, with theory predictions exceeding the data at the largest values of η_{j1} . The behaviour is similar for W^+ , W^- and Z, which may further indicate that the discrepancy is mainly due to the gluon distribution being overestimated at large x . Changes in individual quark or antiquark distributions would instead give a pattern of discrepancy that is more pronounced in one of the channels than in the others. In the pseudorapidity distributions, we probe simultaneously more extreme regions of x_1 and x_2 than for the p_T^j distributions as the directional dependence on y_j as given in Eq. (4.0.1) allows us to more directly discriminate the two Bjorken- x values. This can be seen most explicitly for the ZJ case, in which the forward-most bin in pseudorapidity requires implicitly $x_1 > 0.16$, $x_2 > 1.1 \times 10^{-4}$, meaning that the large $x > \mathcal{O}(0.1)$ regions are probed more efficiently than the highest p_T^j bin.

The distributions for the rapidity of the charged lepton η_ℓ are shown in Fig. 4.4 for W^- and W^+ respectively. Here the NNLO predictions lie $\sim 5\text{--}15\%$ above the central value of the data across the entire considered range in η_ℓ , with the exception of the most forward $\eta_{\ell-}$ bin where agreement is observed. Note that it would be preferable




Figure 4.3: Cross section differential in the leading jet pseudorapidity for W^- , W^+ and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text.

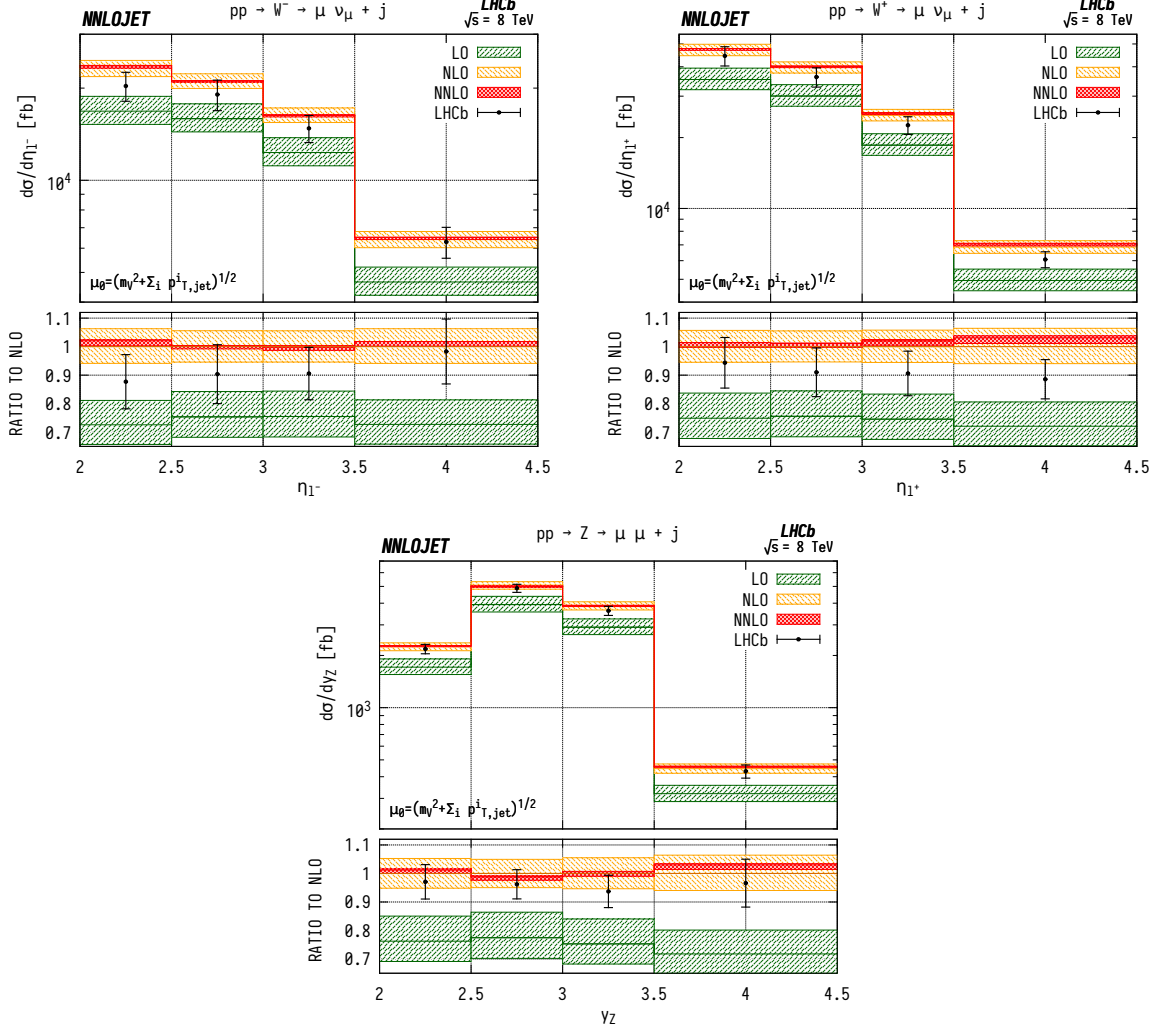


Figure 4.4: Cross section differential in the lepton pseudorapidity for W^- , W^+ , and the Z boson rapidity and Z production. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the ratio to NLO is shown in the lower panel of each plot. The bands correspond to scale uncertainties estimated as described in the main text.

to construct these distributions as a function of the W rapidity y_W , which however can not be unambiguously reconstructed experimentally due to the unknown longitudinal component of the neutrino momentum. For the case of neutral-current production, on the other hand, this is possible and is shown in the bottom panel of Fig. 4.4 differentially with respect to the rapidity of the reconstructed Z boson.

From the charged-current data one can further construct the charge asymmetry

differentially in the lepton pseudorapidity $A^\pm(\eta_\ell)$,

$$A^\pm(\eta_\ell) = \frac{d\sigma^{W^+j}/d\eta_\ell - d\sigma^{W^-j}/d\eta_\ell}{d\sigma^{W^+j}/d\eta_\ell + d\sigma^{W^-j}/d\eta_\ell}. \quad (4.1.3)$$

The charge asymmetry is a valuable input to PDF fits as many systematic experimental errors cancel due to correlations in systematic errors (including luminosity) between the measurements of W^+J and W^-J , giving a higher level of precision than for the total cross sections alone. This is also true for the theory predictions, where many higher-order contributions cancel between W^+J and W^-J , and the similarity of the two calculations justifies some correlation between scale errors. A^\pm directly provides information on the difference between the u and d quark (as well as between the \bar{d} and \bar{u} anti-quark) content of the proton.

The advantage of considering the charge asymmetry for events where a jet is produced in association with the W boson, which can be regarded as an exclusive asymmetry as opposed to A^\pm in inclusive W^\pm production, is that the implicit constraint on Bjorken- x is tightened due to the increase in partonic energy required. Before comparing our predictions with LHCb data for the exclusive charge asymmetry, it is instructive to recall the status of measurements of its inclusive analogue. The LHCb measurement of the inclusive charge asymmetry [236] probes larger values of x than at ATLAS or CMS. Currently the main constraints on u and d content at $x > 0.1$ come primarily from fixed-target DIS experiments and the DØ inclusive lepton charge asymmetry data [237]. The inclusion of the latest Tevatron results in PDF fits generally results in a harder u/d behaviour in this high- x region [238].

In Fig. 4.5, we show a comparison between our theoretical predictions for A^\pm related to WJ production and the LHCb data. Inside the numerator and the denominator expressions, we fully correlate the scales between the W^+ and W^- cross sections, which amounts to taking the sum and difference of the cross sections as independent physical quantities $[d\sigma^{W^+} \pm d\sigma^{W^-}](\mu_F, \mu_R)$ instead of the W^+ and W^- cross sections. The scale uncertainty shown is then obtained by independently varying the factorisation (μ_F) and renormalisation (μ_R) scales of both the numerator and




Figure 4.5: The charge asymmetry A^\pm for W^\pm production as a function of lepton pseudorapidity. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to LHCb data from Ref. [219], and the difference with respect to NLO is shown in the lower panel of the asymmetry plot, with the ratio to NLO is shown in the lower panel of the ratio plot. The bands correspond to scale uncertainties estimated as described in the main text.

denominator by factors of $\frac{1}{2}$ and 2 around the central scale, while imposing the restriction $\frac{1}{2} \leq \mu/\mu' \leq 2$ between all pairs of scales (μ, μ') in Eq. (4.1.3).

The shape of A^\pm as a function of η_ℓ is generally determined by two competing effects [239]. The first is the (anti-)quark content of the PDF, where the u/d ratio and q/\bar{q} asymmetry increase with momentum fraction x , and therefore with η_ℓ . This alone gives an increase in A^\pm with η_ℓ since u -initiated production is dominant in W^+ production while d -initiated production is dominant for W^- .

The second factor is due to the left-handedness of the couplings in the W^\pm production and decay process, which results in opposite preferential directions of the positive and negative decay leptons relative to the W^\pm spin. As a consequence, for the W^+ case, the lepton is preferentially produced at lower η than the W^+ , whereas for the W^- case, the lepton is preferentially produced at higher relative η . This effect causes the asymmetry to decrease with η_ℓ , and dominates over the quark PDF effects at higher x , as can be seen in Fig. 4.5.

We find that the NNLO predictions for the asymmetry describe the data reasonably well, but in general show a less steep slope with η_ℓ than the data. This may be




Figure 4.6: The W^\pm charge asymmetry A^\pm in WJ final states differential in the pseudorapidity η of the decay muon, evaluated with the central member of the NNPDF3.1 (red), MMHT14 (yellow), CT14 (green) NNLO PDFs for the LHCb data of [219]. The NNPDF3.1 curve corresponds to a full NNLO calculation with scale uncertainties as described in the main text, and is used to determine a differential NNLO/NLO k -factor. The other two predictions are calculated at NLO and then rescaled by this k -factor.

indicative of a PDF overestimate in the u/d ratio for $x \gtrsim 0.1$ which would lead to the observed overprediction of the charge asymmetry in this region. It is noted that the large u/d ratio is in particular inferred [117, 238] from the Tevatron $D\bar{O}$ lepton charge asymmetry data [237]. It will thus be very informative to combine these data with the LHCb results [219] in a global fit to determine whether they are mutually consistent.

The sensitivity of the W^\pm asymmetry in WJ final states on the PDF parameterisations is illustrated in the top plot of Figure 4.6, which shows this asymmetry at NNLO for NNPDF3.1 [117], MMHT14 [38] and CT14 [238] parton distributions. The NNPDF3.1 prediction is obtained from a full NNLO calculation of the individual cross sections entering into the ratio, which are also used to extract differential NNLO k -factors. Predictions for the other two PDF parameterisations are computed at NLO at cross section level, and then rescaled by these k -factors, before computing the ratio. The large spread of the predictions in the last bin reflects the different modelling of the quark distributions at large x in the three parameterisations, and




Figure 4.7: Leading jet kinematic distributions for the LHCb data of [219]. The upper row shows plots for the leading jet pseudorapidity in W^- (left) and W^+ (right) production, and the lower row shows plots for the leading jet transverse momentum in W^- (left) and W^+ (right) production. Each of these distributions is evaluated with the central member of the NNPDF3.1 (red), MMHT14 (yellow), CT14 (green) NNLO PDFs. The NNPDF3.1 curve corresponds to a full NNLO calculation with scale uncertainties as described in the main text, and is used to determine a differential NNLO/NLO k -factor. The other two predictions are calculated at NLO and then rescaled by this k -factor.

demonstrates the potential discriminating power of the LHCb asymmetry measurement.

The jet pseudorapidity and transverse momentum spectra for both the W^\pm cases are constructed in the same manner for the plots of Figure 4.7, where we see that there is a considerable difference between the central PDF members of each set for the W^- case which is much more suppressed for the W^+ . This is consistent with the previous observations and again points to a large uncertainty in the d -quark determination in these forward regions, further demonstrating the constraining power of these measurements for the high- x quark distributions. Whilst at first glance the systematic uncertainties may seem too large for the data to have a meaningful impact on PDF fits, the exceptionally strong correlation¹ in these uncertainties between bins means that when the data is taken as a whole it can be very competitive with other experimental data in the high- x region.

4.2 Ratios of VJ Production at ATLAS

The second set of experimental data we consider is the 7 TeV (electron and muon) measurement by the ATLAS experiment [220], which combines data from the W and Z analyses of [226] and [228] with a small modification to the lepton selection criteria applied in the Z analysis when taking ratios. This modification is applied in order to better match the W selection criteria.

The ATLAS detector has a large rapidity range, capable of measuring pseudorapidities of up to $|\eta| = 4.9$ in the endcap region for both hadronic and electromagnetic final states. Unlike the LHCb measurement region, the large pseudorapidity reach of ATLAS also allows to probe large rapidity separations between final state particles, which correspond to configurations in which the Bjorken- x of both incoming protons is relatively large. In the following, we perform a comparison of fixed-order

¹The smallest correlation coefficient between any pair of bins in any of the $W^\pm J$ pseudorapidity distributions is 0.927.

NNLO results to the individual $W^\pm J$ and ZJ distributions of [226] and [228], before constructing the ratios of WJ ($\equiv W^+J + W^-J$) and ZJ distributions and comparing those to the results of [220]. We consider leading jet p_T distributions in inclusive (at least one jet is required) and exclusive (exactly one jet is required) jet production, as well as inclusive leading jet rapidity distributions. The inclusive distributions have previously been compared to NNLO QCD predictions in [141], however exclusive distributions and ratios of distributions were not considered.

The fiducial cuts used in the ATLAS analyses are as follows:

$$\begin{aligned} p_T^j &> 30 \text{ GeV}, & |y^j| &< 4.4, \\ p_T^\ell &> 25 \text{ GeV}, & |y^\ell| &< 2.5, \\ \Delta R_{\ell,j} &> 0.5. \end{aligned} \tag{4.2.1}$$

For $W^\pm J$ production, the restrictions $E_T^{\text{miss}} > 25 \text{ GeV}$, and $m_T^W > 40 \text{ GeV}$ on the missing transverse energy and transverse mass of the W boson are imposed. For ZJ production the requirements $66 \text{ GeV} < m_T^{\ell\ell} < 116 \text{ GeV}$ and $\Delta R_{\ell\ell} > 0.2$ are applied to the transverse mass of the di-lepton system and angular separation of the leptons. In the ZJ distributions, we relax the lepton p_T cut from 25 to 20 GeV in order to compare directly with the results of [228]. However we keep the lepton p_T cut at 25 GeV when constructing ratios of WJ and ZJ distributions.

Jets are reconstructed using the anti- k_T algorithm [62] with radius parameter $R = 0.4$, and we choose the central scale of the theory predictions as

$$\mu_F = \mu_R = \frac{1}{2} H_T = \frac{1}{2} \sum_{i \in \text{jets}, \ell, \nu} |p_T^i| \equiv \mu_0, \tag{4.2.2}$$

where H_T is the scalar sum of the transverse momenta of all final state jets and leptons/neutrinos as appropriate. We denote the number of jets as N , such that in the selection criteria $N = 1$ corresponds to the exclusive case and $N \geq 1$ corresponds to the inclusive case.

The scale variation uncertainties for the ratios are obtained in a similar manner

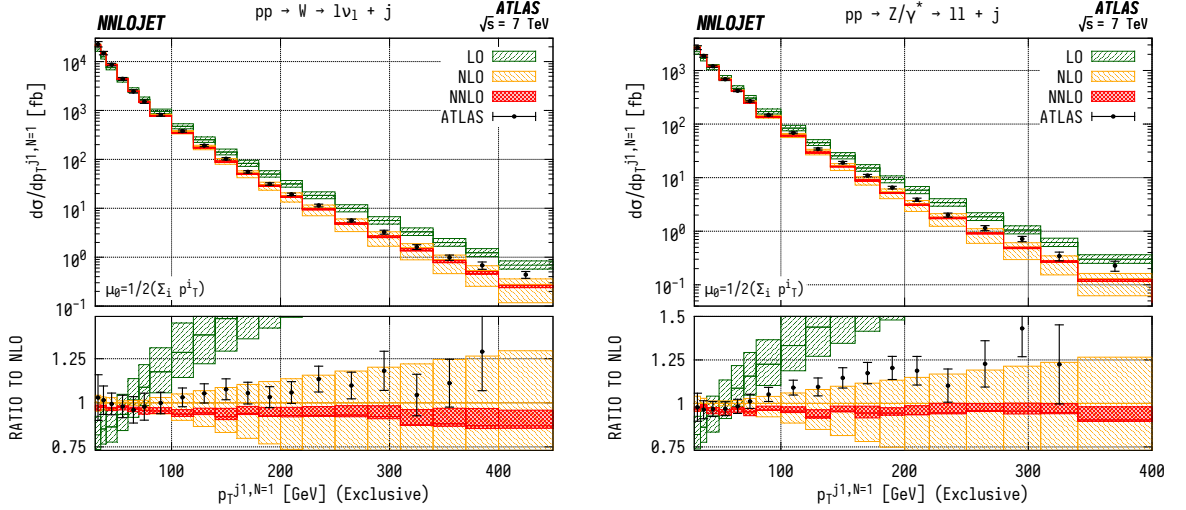


Figure 4.8: WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with exactly one associated jet ($N = 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226, 228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text.

as for LHCb W^\pm asymmetries, with fully correlated scales between the W^+ and W^- processes in the numerator, but taking the envelope of the scales when taking the ratio to the Z distributions, imposing $\frac{1}{2} \leq \mu/\mu' \leq 2$ between all pairs of scales.

Exclusive p_T^{j1} Distributions

First we consider the exclusive ($N = 1$) p_T distribution of the leading jet for WJ production using the data from [226] as shown in the left hand plot of Fig. 4.8. Here we observe agreement of the theory with data within errors up to $p_T^{j1} \sim 100$ GeV, beyond which the theoretical predictions are systematically below the data. This behaviour is closely replicated in the right hand plot, which shows the equivalent ZJ distribution. However beyond $p_T^{j1} \sim 80$ GeV, the agreement with data is noticeably worse than for the WJ distribution by $\mathcal{O}(10\%)$. While we neglect electroweak corrections which have a well understood impact on the weak boson p_T distributions [198, 199, 240] from large Sudakov logarithms, these generally give considerable reductive k -factors at large p_T^j and would further worsen the agreement with data in both cases (see e.g. [240]).

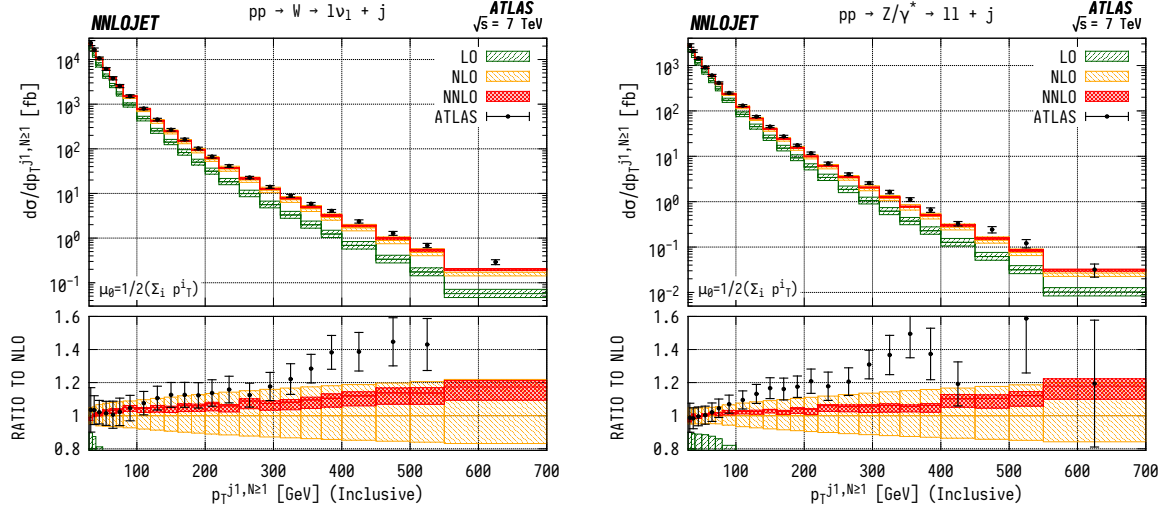


Figure 4.9: WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226, 228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text.

Inclusive p_T^{j1} Distributions

For the inclusive ($N \geq 1$) p_T^{j1} spectrum in WJ production, shown in the left hand panel of Fig. 4.9, we observe marginally improved agreement over a wider range of p_T , with overlapping uncertainty bands between data and theory up to $p_T^{j1} \sim 300$ GeV. Beyond this point, there are substantial, $\mathcal{O}(15\%)$, shape corrections when moving from NLO to NNLO which improve the agreement with data with respect to the NLO results. In ZJ production, shown in the right hand panel of Fig. 4.9, the pattern of perturbative corrections is very similar. However we do not observe the same level of improved agreement with data when moving from exclusive to inclusive jet production as for the WJ process and we again see that the theory prediction is systematically below the data from $p_T^j \sim 100$ GeV onwards.

Allowing extra QCD radiation, as in the inclusive case, entails also allowing for di-jet-type configurations where two hard jets are produced alongside a relatively soft vector boson. In the full NNLO calculation, these $\mathcal{O}(\alpha_s)$ contributions are first described at NLO, and give rise to a large QCD k -factors at high p_T^j [241].

This is the dominant cause of the distinct structure of the perturbative corrections between exclusive and inclusive production; for $N = 1$ we see a decrease in the high p_T^{j1} cross sections with the inclusion of higher orders as opposed to an increase in $N \geq 1$ production. The difference in theory-to-data agreement between the Z and W distributions persists however, and may be related to the different quark flavour combinations probed by the different processes. Whilst not as constraining as the W^+/W^- ratio, the W/Z ratio still retains some sensitivity to the u/d ratio due to different coupling strengths, alongside some dependence on the strange quark distributions, albeit suppressed compared to the inclusive Drell-Yan cross sections due to the Born-level gluon contribution. The inclusion of higher order EW terms are unlikely to describe the difference with respect to data at high p_T , as the EW corrections to the leading p_T^j distribution in vector boson plus di-jet events behave in a very similar manner for WJ and ZJ production as demonstrated in [240].

Exclusive/Inclusive Ratios

In order to better understand the description of real emission by the fixed-order predictions, one can construct the ratio between the exclusive and inclusive leading jet distributions for both the WJ and the ZJ case, shown in Figure 4.10. The experimental measurements [226, 228] do not explicitly quote the data in terms of exclusive/inclusive ratios. We have therefore reconstructed it here using the central values of the relevant distributions with the errors approximated using uncertainties from the $N = 1$ distribution normalised to the $N \geq 1$ results. For both distributions we observe similar behaviour, with a good description of the data across the range of p_T^{j1} , from which we can conclude that the extra jet radiation is well described by the fixed-order predictions.

W/Z Ratios Differential in Leading Jet p_T

The left hand panel of Fig. 4.11 shows the WJ/ZJ ratio as a function of p_T^{j1} , for the exclusive ($N = 1$) case. The large scale variation bands visible at NLO are




Figure 4.10: Ratio of exclusive over inclusive ($N = 1/N \geq 1$) distributions for WJ (left) and ZJ (right) cross sections differential in the transverse momentum p_T of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226,228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text. Errors on the ATLAS data are approximated using uncertainties from the $N = 1$ distribution normalised to the $N \geq 1$ results.

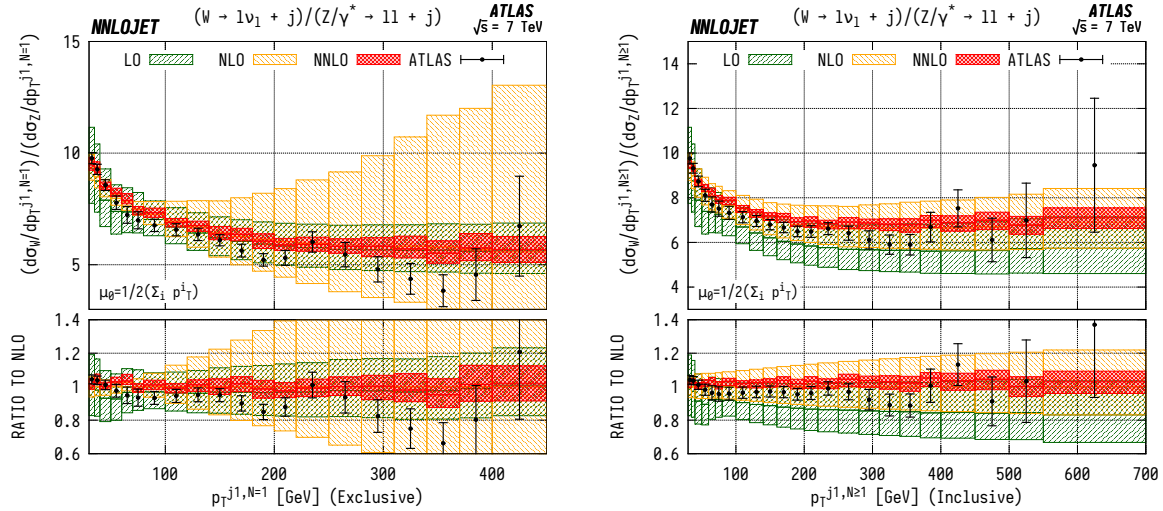


Figure 4.11: WJ/ZJ ratio differential in the exclusive p_T of the leading jet ($N = 1$) (left) and inclusive p_T of the leading jet ($N \geq 1$) (right) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [220], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text.

a result of large NLO corrections at high p_T^j that increase the scale uncertainties when propagated through ratios. In particular, as shown in Fig. 4.8, we observe large reductive NLO/LO k -factors at high p_T^{j1} for the individual ZJ and $W^\pm J$ distributions, reaching $K = 0.3$ in the highest p_T^{j1} bin, whereas the absolute size of the scale variation bands does not reduce significantly when going from LO to NLO. This has the effect of making the exclusive WJ/ZJ ratio much more sensitive to scale variation in the constituent distributions at NLO than LO, artificially inflating the scale uncertainties at this order. The inclusive ($N \geq 1$) ratio, shown in the right hand panel of Fig. 4.11, has very similar central values at LO, NLO and NNLO, but does not display the inflated NLO scale uncertainty.

When taking the ratio, the impact of the extra jet activity is strongly suppressed, while the PDF sensitivity is enhanced. As mentioned in the case of the individual distributions, the W/Z ratio can be used to provide constraints on the ratio of up and down valence quark distributions inside the PDFs, as well as on the strange distribution, due to the different couplings of the vector bosons. Taking only the dominant incoming qg partonic configurations, we can see that naïvely the ratio behaves as

$$\frac{\sigma^{\text{WJ}}}{\sigma^{\text{ZJ}}} \sim \frac{ug + dg}{0.29ug + 0.37dg}, \quad (4.2.3)$$

where the numerical factors are the appropriate sums of the squares of the vector and axial vector quark to Z couplings. Discarding the common factor of the gluon PDF, this can be used to interpret a theory-to-data excess in the W/Z ratio as an overestimate of the u/d ratio. If we look back to the individual distributions, we see that for each of the W and Z cases, the theory falls below the data. From this, it can be deduced that the most probable cause is an underestimate in the d quark content of the PDF.

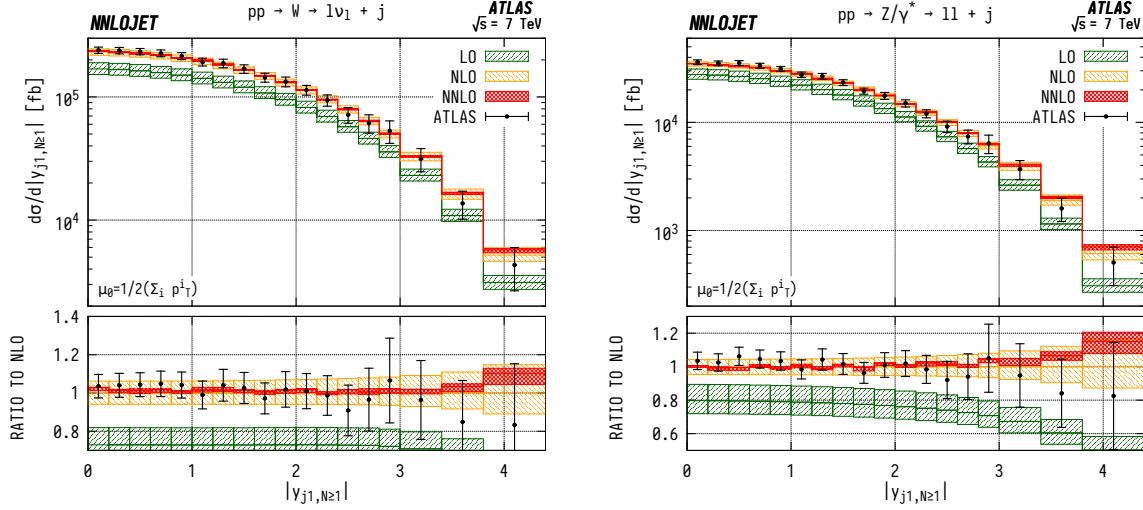


Figure 4.12: WJ (left) and ZJ (right) cross sections differential in the absolute rapidity $|y_j|$ of the leading jet for events with one or more associated jets ($N \geq 1$) in the ATLAS fiducial region from Eq. 4.2.1. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [226, 228], and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text.

Inclusive Leading Jet Rapidity Distributions

The leading jet rapidity distributions $|y_{j1}|$ for both WJ and ZJ events are shown in Fig. 4.12. Here we observe that the higher order QCD predictions are relatively stable up to $|y_{j1}| \sim 3$. Beyond this point, we see a change in shape when transitioning from LO to NLO. The shape is kept unmodified under the inclusion of the NNLO corrections. There is an increase in scale uncertainty at higher rapidities $|y_{j1}| \gtrsim 3.5$ due to large sub-leading jet contributions in this region, which are only described at lower orders for inclusive observables in the NNLO VJ calculation. In both cases, we see good agreement for all rapidities, with overlapping scale errors and experimental error bars for the entire distribution. However, the shape corrections induced at NNLO for $|y_{j1}| \gtrsim 3.5$ modify the central values of the theory predictions such that the tension with data increases compared to NLO given the correlated systematic uncertainties.

If one associates the higher-energy incoming parton with x_1 and the lower-energy incoming parton with x_2 , such that the sum of all final state momenta lies in the




Figure 4.13: WJ/ZJ ratio differential in the absolute rapidity $|y_j|$ of the leading jet. Predictions at LO (green), NLO (orange), and NNLO (red) are compared to ATLAS data from Refs. [220] in the ATLAS fiducial region from Eq. 4.2.1, and the ratios to NLO are shown in the lower panels. The bands correspond to scale uncertainties estimated as described in the main text.

same direction as parton 1, the forward-most bin ($3.8 < y_{j1} < 4.4$) in rapidity here corresponds to $x_1 > 0.19$, $x_2 > 1.2 \times 10^{-4}$ for WJ production and $x_1 > 0.19$, $x_2 > 1.9 \times 10^{-4}$ in ZJ production. One can then analyse the distributions in a similar manner to the LHCb predictions in Fig. 4.3. As is the case for the LHCb data, we see a theory excess in the jet rapidity bins corresponding to $x \gtrsim 0.1$. This is again indicative of an overestimate of the gluon contributions to the PDF in this region since this excess is present in both W and Z distributions. The central rapidity bins allow us to quantify better the PDF description at intermediate Bjorken- x , with the central-most bin in y_{j1} requiring $x_1 > 4.4 \times 10^{-3}$ and $x_2 > 3.6 \times 10^{-3}$ for both neutral- and charged-current production. Here we see good agreement with the data, indicating that the behaviour in this region is well under control.

The ratio of WJ to ZJ differential in the absolute rapidity $|y_{j1}|$ of the leading jet is shown in Fig. 4.13. Due to the cross-cancellation in the ratios, we see that these predictions display a considerably better perturbative stability than the individual distributions at high rapidities. We observe excellent agreement with the ATLAS data across the entire rapidity range. In the ratio, the PDF dependence of the predictions is in general lowered, particularly for gluonic contributions due to their

similarity between the WJ and ZJ cases. The agreement on the ratio demonstrates that the NNLO QCD description of the underlying parton-level process is reliable. It indicates that the discrepancies observed in the individual distributions are of parametric origin and can be remedied by an improved determination of the gluon distribution.

4.3 W+J Production at CMS

We now move to the final experimental comparison, this time restricted only to WJ production using 19.6 fb^{-1} of data taken by the CMS collaboration at 8 TeV published in [221]. The results provided consider only the muon decay channel of the W boson, and combine the W^+ and W^- channels in the same manner as the ATLAS data in the previous section. Whilst these results do not probe the forward region in rapidity, with a maximum jet pseudorapidity of 2.4 which results in a considerably reduced sensitivity to the PDF contributions in the high- x region, a wider range of distributions is considered. In particular this allows one to assess the usefulness and applicability of fixed-order NNLO predictions in certain regions of phase space.

We first define the phase space region for the CMS analysis, given by

$$\begin{aligned} p_T^j &> 30 \text{ GeV}, \quad |\eta|^j < 2.4, \quad |y^{\text{lep}}| < 2.1, \\ p_T^{\text{lep}} &> 25 \text{ GeV}, \quad m_T^W > 50 \text{ GeV}, \end{aligned} \quad (4.3.1)$$

where m_T^W is the transverse mass of the W boson, and with jets defined using the anti- k_T jet algorithm with radius parameter $R = 0.5$. The central scale used for the predictions is

$$\mu_R = \mu_F = \sqrt{m_{\mu\nu}^2 + \sum_i (p_{T,j}^i)^2}, \quad (4.3.2)$$

with scale variations performed in the same manner as the W ATLAS distributions. The transverse momentum and pseudorapidity distributions of the leading jet are shown in Fig. 4.14, where we see good agreement across the full ranges of both distributions. This agreement again suggests it is indeed the forward region which

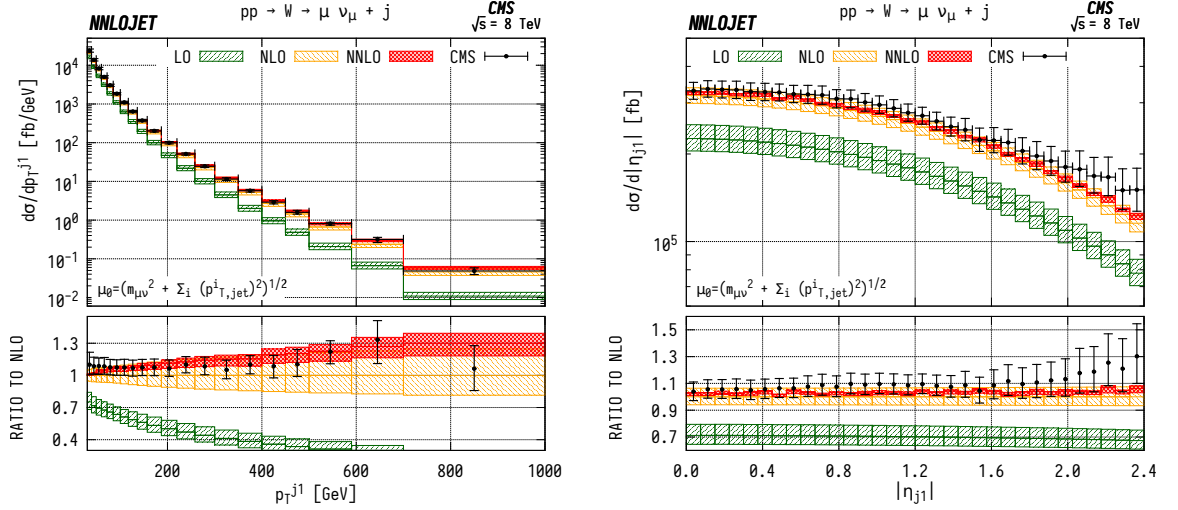


Figure 4.14: Leading jet transverse momentum (left) and pseudorapidity (right) distributions for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text.

is the cause of the discrepancies we observe with the LHCb and ATLAS analyses; once this is removed through a maximum jet rapidity of 2.4, the minimum Bjorken- x required for the final bin in the pseudorapidity region is only $x = 0.034$, an order of magnitude smaller than for the most forward ATLAS and LHCb bins. As a result, any inaccuracies induced by the high x region are generally negligible with respect to the bulk of the cross section.

There is a slight shape underestimation of the CMS data at $\eta_j \sim 2$, which is most likely an artefact of the omission of a minimum angular separation between the lepton and the jet. As $\Delta R^2 = (\Delta\Phi)^2 + (\Delta y)^2$ this cut, present in the LHCb and ATLAS analyses, amounts to an indirect rapidity cut on the jets that is not present here. This rapidity cut is active when the lepton (highly correlated with the W boson dynamics) is close to the leading jet, and so reduces the sensitivity of the measurement to soft jets against which this system recoils which are not well described in fixed-order calculations.

Perhaps more novel are the H_T distributions shown in Fig. 4.15, where H_T is now defined as the scalar sum of the jet transverse momenta unlike the previous definition

in which it also included the vector boson decay products. These distributions are highly sensitive to higher order kinematics arising from partonic radiation, and are generally not well described by NLO calculations. The left panel shows the cross section differential in H_T , where one observes good agreement with data up to $\sim 700\text{GeV}$ where effects of higher jet multiplicities > 3 than can be described in an NNLO calculation become dominant. This is a marked improvement over NLO, where the agreement of the NLO prediction begins to fail at $\sim 200\text{GeV}$.

The right hand panel shows the average number of jets $\langle N_{\text{jets}} \rangle$ as a function of H_T , which can be constructed from the exclusive n -jet distributions $d\sigma_n$ differential in H_T as:

$$\langle N_{\text{jets}} \rangle (H_T) = \frac{1}{d\sigma_{\text{TOT}}/dH_T} \sum_{n=1} n \frac{d\sigma_n}{dH_T} \quad (4.3.3)$$

In this observable, we become much more sensitive again to the soft dynamics, as the average jet multiplicity treats each jet democratically, irrespective of their transverse momentum (as long as it is enough to overcome the jet p_T threshold defined in the fiducial region). As a result, we see agreement with the fixed-order NNLO prediction only up to 200 GeV, beyond which there is substantial disagreement. This is to be expected, as fixed-order calculations have an explicit restriction on the maximum number of associated jets produced, equal to the maximum number of partonic emissions in the real contribution. This is most obvious for the NLO prediction which saturates at 2 jets, meaning that above $H_T \sim 800\text{ GeV}$ essentially all events are di-jet in nature. One trivially observes perfect agreement up to 60 GeV, as the minimum jet p_T permitted in the fiducial region only allows a single jet to be produced.

The picture is more complex at NNLO, where one does not reach a saturation point at three jets and the prediction for the average jet multiplicity actually decreases beyond 600 GeV. This can be understood as a constraint coming from the kinematics; beyond this point it becomes more viable to distribute the total energy into two hard back-to-back jets than into three jets with a more equal azimuthal separation. This is exactly the region in which one observes large QCD k -factors [241]

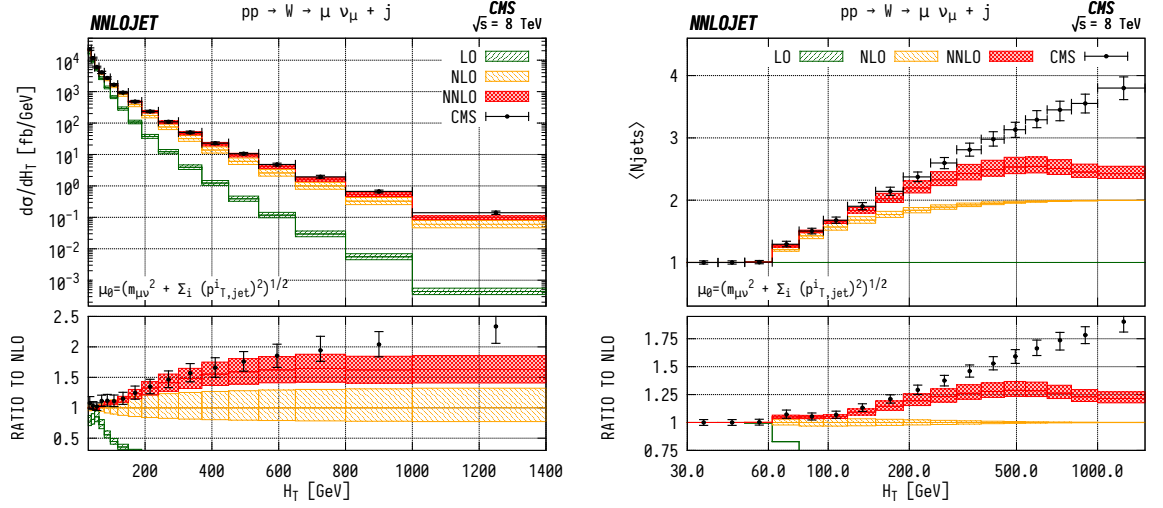


Figure 4.15: The differential cross section (left) and average number of jets $\langle N_{\text{jets}} \rangle$ (right) distributions as a function of H_T for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text.

generated by hard back-to-back di-jet kinematics coupled with the emission of a soft W boson. It therefore would be expected that the remaining discrepancy with data is largely due to the contribution of higher multiplicities of soft jets which cannot be emulated in fixed-order well and are much better modelled with e.g. parton showers. The cross section differential in H_T supports this, showing agreement to much higher H_T values ~ 800 GeV which suggests that the fixed-order calculation describes reasonably well the dominant two and three jet contributions. The remaining discrepancy is small, which implies that the contributions from extra jets are generally soft in nature. Despite this, the addition of the double real contribution in the NNLO calculation substantially increases the agreement with data over NLO in the $\langle N_{\text{jets}} \rangle$ distribution, which highlights the usefulness of the observable for understanding the behaviour of fixed-order predictions.

It must also be mentioned that scale variation does not provide a good measure of the uncertainty in the above jet multiplicity distribution as it does not directly affect the number of jets produced for a given event. Evaluating the same event at different scales, as is done to evaluate the scale uncertainty, will not alter the

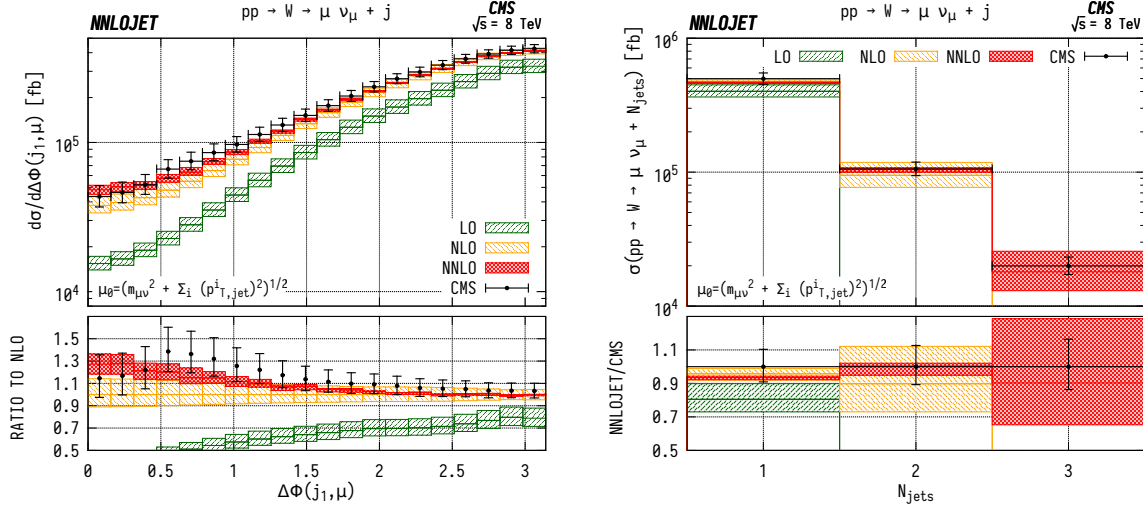


Figure 4.16: Angular separation between leading jet and muon $\Delta\Phi(j_1, \mu)$ (left) and jet multiplicity (right) distributions for W events in association with one or more jets ($N \geq 1$) in the CMS fiducial region from (4.3.1), and compared to data from [221]. The ratios to data are shown in the lower panels, and the bands correspond to scale uncertainties estimated as described in the main text.

number of jets produced which is fixed by the kinematics. The variation can in fact only alter the relative contribution of the jet multiplicity to the total cross section in (4.3.3) due to the overall normalisation factor. At LO scale variation generates a zero uncertainty as only one jet can ever be produced, and with successive orders beyond this one allows more mixing between the jet multiplicities. This in effect increases the scale uncertainty band with each perturbative term in the fixed-order expansion. This ceases to be the case in regions where the jet multiplicity is limited by the event kinematics, as can be seen up to 120 GeV where a maximum of 3 jets can be produced and the scale variation decreases at NNLO. At this point the series has begun to converge around the correct result and is stabilised by the addition of higher orders so the uncertainty band decreases.

The left panel of Fig. 4.16 shows the cross section differential in the angular separation between the muon and the leading jet $\Delta\Phi(\mu, j_1)$. We observe good agreement across the distribution, with the inclusion of NNLO corrections providing a shape correction in the region $\Delta\Phi(\mu, j_1) < 1.5$ to bring the predictions closer to the data. In the low $\Delta\Phi$ region one observes larger perturbative corrections as the muon

direction is strongly correlated with the W direction, and thus a hadronic system recoiling against the WJ system (as generated at higher order) is more often required to generate contributions. This effect is also seen in the scale uncertainty, which only begins to converge at NNLO.

Finally for completeness we show the exclusive jet multiplicity for the data in the right hand panel of Fig. 4.16. There is good agreement for the 1 jet case, and agreement within larger scale variation bands is found for 2 and 3 jets where the NNLO WJ predictions are only NLO and LO accurate respectively.

4.4 Summary

In this chapter we have performed an in-depth comparison of NNLO vector boson plus jet results to data from the LHCb, ATLAS and CMS collaborations. These data sets contain strong sensitivity to the PDF content, particularly in the forward region which can be used to derive strong constraints at high Bjorken- x , and the inclusion of NNLO QCD corrections leads to a substantial reduction of theory uncertainties.

We observe deviations between data and theory in various distributions, which are further investigated by constructing ratios between different vector bosons, and between inclusive and exclusive vector boson plus jet cross sections. The pattern of vector boson ratios and related asymmetries points to an overestimate of the PDF parameterisation in the gluon distribution for Bjorken- $x \gtrsim 0.1$ and equally to an overestimate in the u/d quark ratio in the same region.

Chapter 5

Determination of $\sin^2 \theta_W^{\text{eff}}$ from Triple Differential Inclusive Z Boson Production

As we have seen so far, Drell-Yan production of lepton pairs in hadron-hadron scattering mediated by a virtual photon or Z boson [104] is a process that can be measured experimentally to extremely high precision, and has been the focus of extensive efforts to produce theoretical predictions of competing accuracy. It has long been a benchmark process for our understanding of collider behaviour, including overall luminosity, and it performs a crucial role in determinations of PDFs and SM EW parameters, including the effective Weinberg angle $\sin^2 \theta_W^{\text{eff}}$.

In this chapter we will consider the extraction of $\sin^2 \theta_W^{\text{eff}}$ from triple-differential neutral-current Drell-Yan data taken by the ATLAS collaboration at 8 TeV. We begin with an overview of the effective weak mixing angle and its measurement in Drell-Yan processes in Sections 5.1 and 5.2. We then introduce the ATLAS measurement in Section 5.3 before discussing the kinematics of the phase space in Section 5.4. We then see the impact of these kinematics on the fixed-order QCD results in Section 5.5. In Sections 5.6 and 5.7 we show the effect of varying the PDF and $\sin^2 \theta_W^{\text{eff}}$ values respectively in the theory predictions, and then discuss




Figure 5.1: A comparison of the current world best measurements of $\sin^2 \theta_W^{\text{eff}}$. LEP/SLC results are taken from [242] for combined measurements taken on the Z pole, from b -quark forward-backward asymmetries and in the Z boson left-right asymmetry. CDF [243] and DØ [244] results are presented alongside their combination in [245], and LHC results from the ATLAS collaboration at 7 TeV [116] along with LHCb [115] and CMS [246] 7 and 8 TeV combinations are given. The horizontal bars represent the total uncertainties on each measurement, and the shaded vertical band give the uncertainty on the LEP/SLD measurement.

considerations relating to EW schemes in Section 5.8. Finally, we conclude with combined NNLOJET + Powheg EW predictions produced for the purposes of a $\sin^2 \theta_W^{\text{eff}}$ scan in Section 5.8.

5.1 The Effective Weak Mixing Angle, $\sin^2 \theta_W^{\text{eff}}$

The Weinberg angle $\sin^2 \theta_W$ was introduced in the SM Lagrangian in Section 1.2.3 as a mixing angle between the mass and weak eigenstates of the weak bosons. It therefore provides an important probe of the spontaneous symmetry breaking in-

duced by the Higgs mechanism. As one of the inputs to the SM Lagrangian, precise measurements of $\sin^2 \theta_W$ are invaluable, not only for consistency tests of the SM and searches for new physics in the EW sector, but also in reducing parametric theory uncertainties in other, related, precision measurements.

In making a theory prediction, one must make a choice as to which parameters are taken as input parameters and which are derived. This choice is known as the electroweak input scheme, and for the QCD predictions throughout this chapter, we will use the G_μ scheme [247]. This scheme takes the Z and W boson masses M_Z and M_W , Fermi constant G_F , and the Z boson width Γ_Z , as input parameters, with derived quantities $\sin^2 \theta_W$ and QED coupling constant $\alpha(M_Z)$:

$$\sin^2 \theta_W \Big|_{G_\mu} = \left(1 - \frac{M_W^2}{M_Z^2}\right); \quad \alpha(M_Z) \Big|_{G_\mu} = \frac{\sqrt{2}}{\pi} G_F M_W^2 \sin^2 \theta_W. \quad (5.1.1)$$

The above equations hold at LO in EW theory, however at higher orders the definitions depend on the renormalisation scheme. For consistency reasons, experimental determinations are generally performed within a given renormalisation scheme and then the appropriate translations between schemes are performed *a posteriori* where required. It is most common experimentally to perform measurements of the *effective weak mixing angle* [248]:

$$\sin^2 \theta_W^{\text{eff}} = \left(1 - \frac{M_W^2}{M_Z^2}\right) (1 + \Delta\kappa), \quad (5.1.2)$$

where at LO EW, $\Delta\kappa = 0$ and beyond which $\Delta\kappa \neq 0$ absorbs higher-order modifications from electroweak virtual and radiative corrections as appropriate. This definition contains a residual dependence on the process in which it is measured due to the differences in the required EW corrections. For leptonic decays, the current world average value from global electroweak fits is $\sin^2 \theta_W^{\text{eff}} = 0.23150 \pm 0.00006$ [249], with the most constraining direct measurements made at the LEP and SLD experiments [242] through precision measurements at the Z pole in e^+e^- collisions. A summary of direct experimental measurements is shown in Fig. 5.1. One notes that there is a large (3.2σ) tension between the two most precise individual measure-

ments made using combined LEP/SLD b -quark asymmetry data and SLD left-right polarisation data [242].

It is also possible to measure $\sin^2 \theta_W^{\text{eff}}$ through Drell-Yan processes at hadron-hadron colliders because the differential cross section retains a dependence on $\sin^2 \theta_W^{\text{eff}}$. Due to the hadronic initial states, the experimental measurement is considerably more challenging than in e^+e^- collisions primarily as a result of PDF uncertainties [115, 116, 243–246]. Despite this, the prospects for future hadron collider measurements are positive, and uncertainties potentially competitive to those for the current world-leading measurements could be achieved using data from the LHC [250]. Beyond the obvious use of these hadron collider measurements as an overall closure test of the EW sector, there is particular interest in the use of such measurements to resolve the current tensions in $\sin^2 \theta_W^{\text{eff}}$ between the LEP measurements.

5.2 Measurements of $\sin^2 \theta_W^{\text{eff}}$ in Drell-Yan Processes

In order to gain maximal sensitivity to $\sin^2 \theta_W^{\text{eff}}$ in Drell-Yan lepton production, it is necessary to consider the cross section differentially with respect to the variables that describe the behaviour of the inclusive cross section. The kinematics of inclusive Drell-Yan production can be described using five kinematic variables to all orders in QCD, namely m_{ll} , y_{ll} , p_T^Z and the two decay angles of the leptons in the rest frame of the Z boson, θ and ϕ . For lepton colliders such as LEP, the incoming directions of the leptons are known explicitly, defining a natural scattering angle as the angle between the negatively charged lepton and the outgoing final state. A direct equivalent, replacing the lepton with an incoming quark, is not possible at hadron colliders, where the incoming partons cannot be uniquely identified due to colour confinement. At proton-antiproton colliders the majority of quarks are produced in alignment with the proton beam due to the valence quark content.




Figure 5.2: The definition of Collins-Soper frame and associated lepton decay angles θ and ϕ . p_1 and p_2 are the directions of the incoming partonic momenta in the lepton rest frame, k_1 is the negative lepton momentum and k_2 is the positive lepton. Image taken from [194].

However this is not true for the symmetric initial states of proton-proton colliders such as the LHC.

In order to define the scattering angle θ in an unambiguous manner, it is essential to define a rest frame which facilitates the measurement. In particular, this implies working in the rest frame of the final state di-lepton system. This substantially reduces the sensitivity to initial state radiation which can give a non-zero transverse momentum to the system. To this end it is usual to employ the Collins-Soper frame [251], as shown in Fig. 5.2. The Collins-Soper frame is defined as the rest frame of the decay leptons using the bisector of the incoming beam directions as the z -axis, with the positive z direction aligned with the z -direction of the lepton pair in the lab frame. One then defines $\cos \theta^*$ as the angle between this z -axis and the negatively charged lepton. The x -axis lies in the plane defined by the incoming beams, orthogonal to the z -axis, with the remaining y direction fixed through the requirement of a right-handed Cartesian coordinate system.

In this frame, the z -direction correlates with the direction of the incoming quark due to the momentum distribution within the proton, allowing assignment of the q and \bar{q} directions on a statistical basis. As y_{ll} increases, the incoming quark direction




Figure 5.3: Pure Z, γ^* and Z/ γ^* interference contributions to the total Drell-Yan cross section as a function of m_{ll} , calculated at LO in QCD. The upper panel shows the absolute values of the three contributions, and the lower panel shows the ratio of the Z/ γ^* interference term to the pure γ^* contribution. The black dotted vertical line at $m_{ll} \sim M_Z$ marks the change in sign of the interference term.

becomes more strongly correlated with the final state longitudinal direction due to the dominance of valence quarks at high Bjorken- $x \sim \frac{p_T e^y}{\sqrt{s}}$. The Collins-Soper angle $\cos \theta^*$ can also be defined in a frame independent manner as

$$\cos \theta^* = \frac{p_{ll}^z}{|p_{ll}^z|} \frac{2(l^+ \bar{l}^- - l^- \bar{l}^+)}{Q \sqrt{Q^2 + Q_T^2}} \Big|_{\text{all frames}}, \quad (5.2.1)$$

where

$$\begin{aligned} l^\pm &= \frac{1}{\sqrt{2}}(p_l^E \pm p_l^z), \\ Q &= |m_{ll}| = \sqrt{Q^2 + Q_T^2}, \\ l, \bar{l} &= \{e^-, \mu^-\}, \{e^+, \mu^+\}, \end{aligned} \quad (5.2.2)$$

with Q the invariant mass of the di-lepton system, and Q_T the transverse momentum of the di-lepton system (which by conservation is equal to the transverse momentum of any recoiling hadronic system).

One can write the differential Drell-Yan cross-section at LO in both the QCD and EW couplings using $\cos \theta^*$ as,

$$\frac{d^3\sigma}{dm_{ll}dy_{ll}d\cos\theta^*} = \frac{\pi\alpha^2}{3m_{ll}s} \sum_q P_q [f_q(x_1, Q^2)f_{\bar{q}}(x_2, Q^2) + (q \leftrightarrow \bar{q})]$$

where P_q can be decomposed into contributions from pure virtual photon exchange, pure Z boson exchange and a parity violating Z/ γ^* interference term:

$$\begin{aligned} P_q = & P_{\gamma^*}(1 + \cos^2 \theta^*) \\ & + P_{Z/\gamma^*}[v_l v_q(1 + \cos^2 \theta^*) + 2a_l a_q \cos \theta^*] \\ & + P_Z[(a_l^2 + v_l^2)(a_q^2 + v_q^2)(1 + \cos^2 \theta^*) + 8a_l a_q v_l v_q \cos \theta^*]. \end{aligned} \quad (5.2.3)$$

The separate contributions can be themselves written in terms of the appropriate couplings and propagators:

$$\begin{aligned} \text{Pure } \gamma^* : \quad P_{\gamma^*} &= e_l^2 e_q^2 \\ Z/\gamma^* \text{ Interference} : \quad P_{Z/\gamma^*} &= e_l e_q \frac{2m_{ll}^2(m_{ll}^2 - m_Z^2)}{\sin^2 \theta_W \cos^2 \theta_W [(m_{ll}^2 - m_Z^2)^2 + \Gamma_Z^2 m_Z^2]} \\ \text{Pure } Z : \quad P_Z &= \frac{m_{ll}^4}{(\sin^2 \theta_W \cos^2 \theta_W)^2 [(m_{ll}^2 - m_Z^2)^2 + \Gamma_Z^2 m_Z^2]}. \end{aligned} \quad (5.2.4)$$

The relative contributions of each of these terms to the total cross section as a function of the di-lepton invariant mass is shown in Fig. 5.3. At low invariant m_{ll} , the photon term dominates, until the vicinity of the Breit-Wigner Z resonance where the pure Z contribution takes over. The interference term is generally the smallest contribution, and is negative up to $m_{ll} = M_Z$ where it changes in sign.

Considering the form of the differential cross section, one can see that there are two terms linear in $\cos \theta^*$, arising from the Z/ γ^* interference and Z contributions, which induce an asymmetry between positive (forward) and negative (backward) values of $\cos \theta^*$. The forward-backward asymmetry A_{FB} is defined as¹

$$A_{\text{FB}} = \frac{d^3\sigma(\cos \theta^* > 0) - d^3\sigma(\cos \theta^* < 0)}{d^3\sigma(\cos \theta^* > 0) + d^3\sigma(\cos \theta^* < 0)}. \quad (5.2.5)$$

¹This is directly related to the angular coefficient decomposition of the DY cross section, where coefficient A_4 is given by $A_{\text{FB}} = \frac{3}{8}A_4$.




Figure 5.4: A_{FB} as a function of the di-lepton invariant mass m_{ll} at LO in QCD. The dotted vertical red line denotes the position of the Z peak, at which point the Z/γ^* interference term is identically zero.

As this observable directly probes only the pure Z and Z/γ^* interference terms, it provides strong sensitivity to the axial and vector components of the Z boson coupling, and hence $\sin^2 \theta_W^{\text{eff}}$. This sensitivity is enhanced by large cancellations in systematic uncertainties between numerator and denominator, which allows extremely precise experimental measurements to be made. Despite the small contribution to the total cross-section, the interference contribution to A_{FB} dominates except at $m_{ll} = M_Z$ due to the suppression of the pure Z term by the vector coupling factor $v_l v_q$. This can be seen in the shape of the asymmetry as a function of m_{ll} , shown in Fig. 5.4, where A_{FB} changes sign at $m_{ll} \sim M_Z$, matching the behaviour of the Z/γ^* interference term.

A_{FB} also has a strong dependence on the di-lepton rapidity due to the probabilistic determination of the quark direction. At central rapidities, the incoming quark and antiquark have nearly equal momenta, substantially reducing the correlation between the z direction in the CS frame and the incoming quark direction, leading to a ‘washing out’ of A_{FB} . The opposite is true at high rapidities, with the proviso that fiducial cuts on the individual lepton rapidities can impact the asymmetry at extreme y_{ll} . This means that data taken in forward regions can have considerable

constraining power on $\sin^2 \theta_W^{\text{eff}}$ determinations even with relatively large uncertainties when compared to measurements made in more central regions.

Template/multivariate likelihood fits using A_{FB} , or equivalently A_4 are perhaps the most common method used for $\sin^2 \theta_W^{\text{eff}}$ extraction from experimental data at hadron-hadron colliders², and have been performed using both Tevatron and LHC data [114, 115, 243, 245, 246, 252–255]. The discriminating power of A_{FB} compared to bare cross sections occurs due to the cancellation of the virtual photon contributions, which dilutes the measurement through a large $\sin^2 \theta_W^{\text{eff}}$ -independent cross-section. The most constraining region of the m_{ll} spectrum for a $\sin^2 \theta_W^{\text{eff}}$ determination is normally around the Z pole, due to the large variation in A_{FB} with m_{ll} and the low statistical/systematic errors of experimental measurements performed in the peak region.

The majority of $\sin^2 \theta_W^{\text{eff}}$ extractions are performed by a variation of M_W about some unphysical value in order to consistently account for the EW corrections encoded within $\Delta\kappa$ in the G_μ scheme. M_W is typically used to perform the variation as it is the least well-measured input parameter. The central value taken is typically around $M_W \sim 79.95$ GeV when one only considers LO in QCD, with the EW corrections in $\Delta\kappa$ reweighting M_W back to the physical value when a parameter scan is performed.

The dominant theoretical uncertainty in most hadron-collider $\sin^2 \theta_W^{\text{eff}}$ extractions comes from the uncertainties in PDF determinations [5, 256, 257]. However, as we have seen previously, Drell-Yan processes can also be used provide strong constraints on PDFs, particularly in the quark sector. This allows the use of PDF profiling and similar techniques [258, 259] to systematically reduce PDF uncertainties when measuring electroweak parameters, particularly when one directly uses the cross sections differential in $\cos \theta^*$ in the fit as opposed to the reconstructed A_{FB} . If we consider the cross section differential in each of m_{ll} , y_{ll} and $\cos \theta^*$ as introduced in (5.2.3), the PDF sensitivity is again enhanced with respect to simply measuring

²Left-right polarisation asymmetry A_{LR} and combined left-right-forward-backward asymmetries have also been constructed for $\sin^2 \theta_W^{\text{eff}}$ extraction in electron-positron colliders, see e.g. [242].

$d\sigma/d\cos\theta^*$, as each of these observables allows us to probe different aspects of the PDF content:

- y_{ll} gives a strong sensitivity to the behaviour of the cross section with Bjorken- x given their strong correlation.
- m_{ll} probes the u/d quark ratio as the relative Z and γ^* contributions vary with m_{ll} through the mass dependence of the Z propagator, giving a considerable variation in relative u and d -type contributions across the m_{ll} spectrum.
- Higher order QCD terms modify the $\cos\theta^*$ decay angle dependence through qg , $\bar{q}g$, and gg initiated channels which open up at NLO and NNLO, and give the measurement sensitivity to gluon and sea quark PDF contributions.

Whilst the systematic uncertainties are larger when using cross-section data rather than A_{FB} , the direct use of such differential data provides substantial sensitivity to PDFs and allows a competitive $\sin^2 \theta_W^{\text{eff}}$ determination to be made. Beyond this, standalone differential cross section data provides important input data for future PDF fits. In the following sections, we will introduce such a triple differential measurement and consider some of the associated theoretical challenges with the goal of producing consistent NNLO QCD corrections for an associated $\sin^2 \theta_W^{\text{eff}}$ fit.

5.3 ATLAS Drell-Yan Triple Differential (Z3D) Measurement

In 2017, the ATLAS collaboration presented a measurement of the inclusive Drell-Yan process at $\sqrt{s} = 8$ TeV [260], collected using 20.2 fb^{-1} of data taken in 2012 using combined electron and muon decay channels³. The results taken were triply

³We will henceforth refer to this measurement as Z3D in order to distinguish this from the complementary angular analysis also performed by ATLAS on 8 TeV data [261].

Central-Central	Central-Forward	
$p_T^l > 20 \text{ GeV}$	$p_{T,F}^l > 20 \text{ GeV}$	$p_{T,C}^l > 25 \text{ GeV}$
$ y^l < 2.4$	$2.5 < y_F^l < 4.9$	$ y_C^l < 2.4$
$46 \text{ GeV} < m_{ll} < 200 \text{ GeV}$	$66 \text{ GeV} < m_{ll} < 150 \text{ GeV}$	

Table 5.1: Selection criteria for the central-central and central-forward fiducial regions in the ATLAS measurement of [260].

Observable	Central-Central	Central-Forward
$m_{ll} \text{ [GeV]}$	[46,66,80,91,102,116,150,200]	[66,80,91,102,116,150]
y_{ll}	[0,0.2,0.4,0.6,0.8,1,1.2, 1.4,1.6,1.8,2,2.2,2.4]	[1.2,1.6,2,2.4,2.8,3.6]
$\cos \theta^*$	[-1,-0.7,-0.4,0,0.4,0.7,1]	[-1,-0.7,-0.4,0,0.4,0.7,1]
Total Bin Count:	504	150

Table 5.2: Binnings for the central-central and central-forward fiducial regions in the ATLAS measurement of [260].

differential in the di-lepton invariant mass m_{ll} , di-lepton rapidity y_{ll} and the scattering angle in the Collins-Soper frame $\cos \theta^*$, with the measurement divided into two regions. These are defined by different selection criteria: a central-central (CC) region where both leptons were observed in the central rapidity region of the ATLAS detector, and a central-forward (CF) region where one lepton is required to be central whilst the other is measured in the forward detector region. The full fiducial cuts and binnings are shown in Tables 5.1 and 5.2.

The original measurement in [260] was presented alongside theoretical results generated at NLO QCD using Powheg-Box [262–265] with Pythia 8 [266] to model parton showering, hadronisation and underlying event effects alongside NLO EW corrections [267]. The distributions were then corrected using a set of NNLO QCD + NLO EW k -factors differential only in the invariant di-lepton mass m_{ll} generated using FEWZ 3.1 [124], which varied from 1.035 for the lowest m_{ll} bin to 1.025 in the highest. A fit of $\sin^2 \theta_W^{\text{eff}}$ to the data by the ATLAS collaboration is underway, with preliminary results presented in [5]. It is in this context that the implementation of Drell-Yan at NNLO in QCD within the NNLOJET framework is used, with the secondary goal of exploiting the data alongside a consistent set of NNLO results for Drell-Yan type processes produced using NNLOJET for PDF fitting purposes.

One must be much more careful in the QCD theory predictions when performing a full $\sin^2 \theta_W^{\text{eff}}$ fit to such highly differential data, as is our aim. The kinematics of the measurement regions mean that there is a complex structure to the higher order corrections across the bins which can be exploited for increased precision, and leads to acceptance effects that must be taken into account. Interfacing the QCD predictions with the appropriate EW corrections for multiple values of $\sin^2 \theta_W^{\text{eff}}$ must also be feasible in order for a scan of $\sin^2 \theta_W^{\text{eff}}$ to be performed, and this requires careful attention to avoid consistency issues between the two theory inputs.

Whilst differential NNLO QCD results for the Drell-Yan process have been known for almost two decades and there are many available codes producing these results (see e.g. [124, 125, 127, 128]), accurate and exclusive results typically require non-trivial CPU resources to compute. This is particularly true when producing multi-differential results, and it is for this reason that generating accurate predictions for the 654 separate bins of the Z3D analysis remains technically challenging today. These issues are multiplied when producing results for a parameter fit, where multiple sets of such results are required for parameter variation, uncertainty estimation and closure tests. As a result, one can consider the numerical demands of producing such predictions to be more comparable to those required for VJ production than in more standard single or double differential inclusive Drell-Yan distributions.

5.4 Kinematics of the Z3D Measurement

We begin our discussion of the production of Drell-Yan theoretical predictions by considering the kinematics of the Z3D measurement. There is a rich kinematic structure within inclusive Drell-Yan production, which becomes increasingly apparent when taking multi-differential measurements. This is particularly true when considering results differential in both $\cos \theta^*$ and y_{ll} , where indirect kinematic constraints severely restrict the available phase space for the leptonic decay of the Z boson. As we will see, these constraints occur naturally as an artefact of rapidity cuts both in

the Born phase space and beyond.

5.4.1 Born Level ($\mathcal{O}(\alpha^2\alpha_s^0)$) Kinematics

At LO, the kinematics are particularly simple, and serve as a good illustration of how phase space constraints can be induced by the fiducial cuts. We begin with the definition of $\cos\theta^*$ from (5.2.1), and use the standard momentum parameterisation of a 4-vector in terms of rapidity and p_T for each of the outgoing leptons:

$$\begin{aligned} p_l^\mu &= (p_l^x, p_l^y, p_l^z, E_l) \\ &= (p_T^l \cos\theta, p_T^l \sin\theta, E_T^l \sinh(y^l), E_T^l \cosh(y^l)). \end{aligned} \quad (5.4.1)$$

From this we can construct the separate component parts of $\cos\theta^*$, noting that for massless leptons, $E_T^l = p_T^l$:

$$\begin{aligned} l_i^\pm &= \frac{1}{\sqrt{2}} p_l^T \exp(\pm y_i) \\ 2l_i^+ l_i^- &= (p_T^i)^2 \\ l_1^+ l_2^- - l_1^- l_2^+ &= p_{T,1}^l p_{T,2}^l \sinh(\Delta y_{ll}) \\ Q^2 &= E_{12}^2 - (p_{12}^z)^2 - Q_T^2 \\ Q^2 + Q_T^2 &= E_{12}^2 - (p_{12}^z)^2 \\ &= 2(l_1^+ + l_2^+)(l_1^- + l_2^-) \\ &= 2l_1^+ l_1^- + 2l_2^+ l_2^- + 2l_1^+ l_2^- + 2l_2^+ l_1^- \\ &= (p_{T,1}^l)^2 + (p_{T,2}^l)^2 + 2p_{T,1}^l p_{T,2}^l \cosh(\Delta y_{ll}) \end{aligned} \quad (5.4.2)$$

At Born level, $p_{T,1}^l = p_{T,2}^l = p_T^l$ and $Q_T = 0$ as there is no recoiling system, and we can directly reconstruct (5.2.1) in order to give our expression for $\cos\theta^*$,

$$\cos\theta^* = \frac{\sinh(\Delta y_{ll})}{1 + \cosh(\Delta y_{ll})} = \tanh\left(\frac{\Delta y_{ll}}{2}\right). \quad (5.4.3)$$

This immediately allows one to derive constraints on $\cos\theta^*$ which are induced through constraints on Δy_{ll} .




Figure 5.5: Phase space constraints in the $(|y_l|, \cos \theta^*)$ plane at Born level (LO) for the central-central Z3D Drell-Yan fiducial region. Overlaid are the measurement bins, integrated over m_{ll} .

Central-Central (CC) Region

For the case of rapidity cuts symmetric between the leptons and about the origin, as is the case in the central-central region of the Z3D measurement, this procedure is particularly simple. If we note that the minimal value of $|\Delta y_l|$ from the cuts is 0, and that (5.4.3) is symmetric in Δy_l , we can see that constraints on $\cos \theta^*$ come from the upper bounds on $|\Delta y_l|$. For a given y_l value with lepton rapidity cut $|y^l| < |y_{\text{max}}^l|$, and

$$y_l = \frac{y_1^l + y_2^l}{2}, \quad (5.4.4)$$

the greatest value of Δy_l permitted by the cuts is $2(y_{\text{max}}^l - |y_l|)$, which leads to the constraint

$$\cos \theta^* \leq \frac{\sinh(2(y_{\text{max}}^l - |y_l|))}{1 + \cosh(2(y_{\text{max}}^l - |y_l|))}. \quad (5.4.5)$$

This defines a region in $(|y_l|, \cos \theta^*)$ space which is forbidden at LO in QCD, as shown in Fig. 5.5 for the CC region of the Z3D measurement, where we see that central di-lepton rapidities occupy almost the entire $\cos \theta^*$ range, with stronger phase




Figure 5.6: Bin classifications at LO for the central-central Z3D Drell-Yan fiducial region in the $(|y_{ll}|, \cos \theta^*)$ plane. Overlaid are the measurement bins, integrated over m_{ll} .

space restrictions as one increases y_{ll} .

One can then use this to classify the measurement bins in $(|y_{ll}|, \cos \theta^*)$ space into three categories, that depend on whether the associated fiducial regions can be fully accessed at LO, partially accessed at LO, or are completely forbidden. The bin classifications for the central-central region are shown in Fig. 5.6 where, as one would expect, the majority of bins in the central region ($y_{ll} \sim 0$) are fully allowed, while beyond $y_{ll} = 1.4$ the restrictions take effect.

One important corollary of this is that the forbidden bins, shown in red, will be described at best at NLO accuracy within the full NNLO calculation. These bins can only be populated starting at $\mathcal{O}(\alpha_s)$ for a full fixed-order NNLO Drell-Yan calculation in analogy to the vector boson transverse momentum distribution (see Chapter 3). In effect, the cuts of the lepton rapidities have induced an indirect transverse momentum cut which becomes exposed when one is simultaneously differential in $\cos \theta^*$ and y_{ll} .

Bound	Region 1 $ y_{ll} < y_{ll} ^{\text{mid}}$	Region 2 $ y_{ll} > y_{ll} ^{\text{mid}}$	Midpoint $ y_{ll} ^{\text{mid}}$
Upper	$\Delta y_{ll}^{\text{max}} = 2(y_{ll} - y_{C,\text{min}}^l)$	$\Delta y_{ll}^{\text{max}} = 2(y_{F,\text{max}}^l - y_{ll})$	$\frac{1}{2}(y_{F,\text{max}}^l + y_{C,\text{min}}^l)$
Lower	$\Delta y_{ll}^{\text{min}} = 2(y_{F,\text{min}}^l - y_{ll})$	$\Delta y_{ll}^{\text{min}} = 2(y_{ll} - y_{C,\text{max}}^l)$	$\frac{1}{2}(y_{C,\text{max}}^l + y_{F,\text{min}}^l)$

Table 5.3: Maximum/minimum values of y_{ll} permitted in different regions of phase space for the CF Z3D selection.

Central-Forward (CF) Region

A similar procedure can be undertaken for the central-forward region of the Z3D measurement, where one lepton is boosted to the forward calorimeter. In order to extract the LO constraints on Δy_{ll} in the asymmetric case, it is easiest to first divide the phase space into two regions, detailed in Table 5.3, for construction of each of the upper and lower bounds. These regions correspond to values of $|y_{ll}|$ for which a particular lepton rapidity cut provides the limiting value of Δy_{ll} , and hence extremal values for $\cos \theta^*$.

The associated phase space and bin classifications are shown in Fig. 5.7, where one can see that (unlike the CC region), there is a bias in the allowed phase space towards non-central values of $|\cos \theta^*|$. In the context of a $\sin \theta_W$ fit this is particularly interesting, as the cuts mean that the distribution of the cross section is biased towards larger values of y_{ll} . The number of LO-forbidden bins is also greatly increased with respect to the CC-case. This means that inclusive NNLO QCD predictions for Drell-Yan production will less efficiently describe the bulk of the data than for the CC region. As the forward-backward asymmetry can be measured more precisely at large values of rapidity because the incoming quark and antiquark are better defined in the Collins-Soper frame, one could consider exploiting this in order to construct an experimental binning in which as many bins are fully accessible at LO as possible (statistical and detector constraints notwithstanding).




Figure 5.7: Bin classifications at LO for the central-forward Z3D Drell-Yan fiducial region in the $(|y_{11}|, \cos \theta^*)$ plane. Overlaid are the measurement bins, integrated over m_{ll} .

5.4.2 Constraints Beyond Born Level

We now consider the rather more involved case where the di-lepton system has a non-zero transverse momentum through recoil against some partonic radiation. Our aim now changes somewhat; we are now looking to evaluate the minimum transverse momentum required to populate the bins forbidden at LO rather than to simply evaluate the bounds of this region.

Taking once again the components of Eqn. (5.4.2), one can generate the general form of $\cos \theta^*$ by considering the p_T dependence of the Z boson, using the identity $Q = \sqrt{(Q^2 + Q_T^2) - Q_T^2}$:

$$\cos \theta^* = \frac{2p_{T,1}^l p_{T,2}^l \sinh(\Delta y_{ll})}{\sqrt{(p_{T,1}^l)^2 + (p_{T,2}^l)^2 + 2p_{T,1}^l p_{T,2}^l \cosh(\Delta y_{ll})}} \times \frac{1}{\sqrt{(p_{T,1}^l)^2 + (p_{T,2}^l)^2 + 2p_{T,1}^l p_{T,2}^l \cosh(\Delta y_{ll}) - Q_T^2}}, \quad (5.4.6)$$

where we again denote the di-lepton transverse momentum as Q_T . Directly deriving constraints from (5.4.6) is considerably less trivial than the Born case, given the




Figure 5.8: Minimum Q_T values required for the different regions of the $(|y_{ll}|, \cos \theta^*)$ plane in the central-central region. Overlaid are the Z3D measurement bins, integrated over m_{ll} .

complex interplay of the different fiducial cuts, but can still be performed for the two fiducial regions of the Z3D measurements under certain approximations.

For the central-central region, the minimum value of Q_T at a given point in the $(|y_{ll}|, \cos \theta^*)$ plane can be obtained simply by rearranging for Q_T^2 , and substituting in the minimum lepton transverse momenta $p_{T,\min}^l$ and maximum rapidity difference Δy_{ll}^{\max} permitted by the cuts:

$$Q_T^2 \geq 2(p_{T,\min}^l)^2 \left(1 + \cosh(\Delta y_{ll}^{\max}) - \frac{\sinh^2(\Delta y_{ll}^{\max})}{(1 + \cosh(\Delta y_{ll}^{\max})) \cos^2 \theta^*} \right). \quad (5.4.7)$$

This constraint means that for a given Q_T value, there is a region of $\cos \theta^*$ space that can not be populated, determined by the accessible lepton p_T and y_l values given by the cuts. These Q_T constraints are shown in Fig. 5.8 for the central-central Z3D region under the assumption that

$$y_{ll} = \frac{y_1^l + y_2^l}{2}, \quad (5.4.8)$$

which is true when the leptons have equal transverse momentum. We see that the constraints, whilst significant, are not strong enough to prevent population of the forbidden bins beyond Born level for centre-of-mass energies one typically observes at the Tevatron or the LHC.

This effective Q_T cut can be utilised to extend the accuracy of the fixed-order theoretical predictions for the Born-forbidden bins by using calculations of ZJ in a similar manner to the p_T^V spectrum in Chapter 3, the requirement for a non-zero Q_T implicitly requiring partonic radiation. Given that fixed-order results for the ZJ process are known to α_s^3 as opposed to α_s^2 for inclusive Drell-Yan production, and the triple virtual contribution which lies in the Born phase space doesn't contribute, this represents an improvement equivalent to extending the predictions to N³LO accuracy in the forbidden bins.

However, the assumption made in (5.4.8) is not strictly the case. When we relax the Born phase space constraints, the leptons have different transverse momenta which alters the relationship between the lepton rapidities y_1 , y_2 and y_{ll} . Taking the definition of rapidity:

$$\begin{aligned} y_{ll} &= \frac{1}{2} \log \left(\frac{p_{T,1}^l (\cosh y_1 + \sinh y_1) + p_{T,2}^l (\cosh y_2 + \sinh y_2)}{p_{T,1}^l (\cosh y_1 - \sinh y_1) + p_{T,2}^l (\cosh y_2 - \sinh y_2)} \right) \\ &= \frac{1}{2} \log \left(\frac{e^{y_1} + p_{T,2}^l / p_{T,1}^l e^{y_2}}{e^{-y_1} + p_{T,2}^l / p_{T,1}^l e^{-y_2}} \right) \\ &= \frac{1}{2} (y_1 + y_2) + \frac{1}{2} \log \left(r \cdot \left[\frac{1 + r \cdot e^{y_2 - y_1}}{1 + 1/r \cdot e^{y_2 - y_1}} \right] \right), \end{aligned} \quad (5.4.9)$$

where we have introduced the ratio of lepton transverse momenta as

$$r = p_{T,2}^l / p_{T,1}^l. \quad (5.4.10)$$

Eqn. (5.4.9) is explicitly dependent on both the p_T ordering of the leptons as well as the rapidity ordering, and we observe that it reduces to the correct form in the Born limit $r \rightarrow 1$. The new term weakens the constraints on Q_T , and allows Δy_{ll}^{\max} to take a larger range of values for a given y_{ll} value by introducing a large transverse momentum imbalance between the leptons (up to the maximum directly permitted by the lepton cuts). To evaluate the full Q_T dependence for a given y_{ll} , $\cos \theta^*$ value requires an iterative numerical solution, as the circularity of the system of equations makes an analytic solution impossible.




Figure 5.9: The upper panel shows the minimum values of y_1 and maximum values of Δy_{ll}^{\max} as a function of the lepton transverse momentum ratio r . The lower panel shows the resultant minimum transverse momentum induced by the momentum ratio r . For this we consider the forbidden phase-space point at $y_{ll} = 2$, $\cos \theta^* = 0.4$ in the Z3D central-central fiducial region, and take $y_2 = y_2^{\max} = 2.4$ in order to maximise the rapidity difference.

Since

$$r \in \left\{ \sim p_{T,\min}^l / \frac{\sqrt{s}}{2}, \sim \frac{\sqrt{s}}{2} / p_{T,\min}^l \right\} \quad (5.4.11)$$

(which for the central-central region in the Z3D measurement is equivalent to the approximate range $[0.005, 200]$), this would allow one to effectively produce any Δy_{ll}^{\max} given the correct conditions. This is shown in Fig. 5.9, where we show the minimum values of y_1^l and maximum values of Δy_{ll}^{\max} as a function of r for the forbidden phase-space point $y_{ll} = 2$, $\cos \theta^* = 0.4$, setting y_2^l to the maximum permitted value in order to maximise Δy_{ll}^{\max} . Here we see that if one produces values of $r \sim 8$, one could imagine that it is possible to generate any lepton rapidity separation and thus

any value of Q_T . However, hidden within r , there is a second restriction on Q_T , governed by the requirement for transverse momentum conservation:

$$\vec{p}_{T,1}^l + \vec{p}_{T,2}^l + \vec{Q}_T = 0. \quad (5.4.12)$$

From here, one can see that the minimal value for Q_T is given when the lepton p_T values are back to back, which allows us to conclude that there is a second competing Q_T bound at

$$p_{T,\min}^l(r-1) = Q_{T,\min}. \quad (5.4.13)$$

which has no effect for $r = 1$, but for $r = 8$ and $p_{T,\min}^l = 20$ GeV gives a minimum Q_T of 140 GeV. Somewhat counter-intuitively, one can in effect *decrease* the required Q_T for a given point in $(|y_{ll}|, \cos\theta^*)$ space by *generating* a lepton p_T imbalance through a non-zero Q_T . The lower panel of Fig. 5.9 shows the variation of this $Q_{T,\min}$ value with r for a given forbidden bin.

Given that we hope to use fixed-order predictions to calculate predictions for these forbidden bins, it is first useful to evaluate the Q_T spectrum to assess the potential impact of large logarithms in m_{ll}/Q_T . As we have seen already, if present such logarithms could in principle be resummed using tools such as RADISH [144] to N³LL accuracy. We show the normalised Q_T spectra at $\mathcal{O}(\alpha_s)$ for each y_{ll} bin in the $0.4 < \cos\theta^* < 0.7$, $46 \text{ GeV} < m_{ll} < 200 \text{ GeV}$, $1.2 < |y_{ll}| < 2.4$ region of the Z3D analysis in Fig. 5.10. This gives the evolution of the p_T spectrum with y_{ll} as one passes from fully allowed bins at $1.2 < |y_{ll}| < 1.4$ (green) through the mixed (yellow) region $1.4 < |y_{ll}| < 2.0$ to the fully forbidden (red) region $|y_{ll}| > 2.0$. For the forbidden bins we observe that logarithmically divergent behaviour is indeed not present at low Q_T , with no Q_T values below 5 GeV⁴. This can be understood as the volume of the phase space in which low Q_T production is permitted decreasing at a faster rate than the matrix element diverges as $Q_T \rightarrow 0$.

Particularly when one considers cross sections integrated over Q_T as in the Z3D

⁴In our original $y_{ll} = \frac{1}{2}(y_1^l + y_2^l)$ approximation, this low Q_T limit was 13.5 GeV.




Figure 5.10: The normalised Q_T spectrum for the $0.4 < \cos \theta^* < 0.7$, $46 < m_{ll} < 200$ region for each rapidity bin of the Z3D central-central measurement region between $y_{ll} = 1.2$ and $y_{ll} = 2.4$. The results are produced to $\mathcal{O}(\alpha_s)$, with the colours labelling as before the allowed (green), mixed (orange) and forbidden (red) bins. The kink observed at $1/\sigma \cdot d\sigma/dQ_T = 10^{-4}$ is a consequence of the linearisation of the axes between $Q_T \pm 10^{-4}$ to allow the negative contribution at $Q_T = 0$ to be shown.

measurement, any residual breakdown in the perturbative series will be largely suppressed by contributions at moderate and high values of Q_T where fixed-order results are demonstrably well-behaved. If we consider the discussion in Section 3.4 where logarithms in Q_T emerge from miscancellations between real and virtual contributions at low Q_T , this suppression is a direct result of the lack of $Q_T = 0$ virtual contributions combined with a kinematic suppression of low Q_T real contributions. This is visible in Fig. 5.10, where the peaks in the Q_T distributions for the two forbidden bins occur at 40 GeV and 160 GeV despite being allowed down to 5 GeV and 80 GeV respectively. This is in stark contrast to the mixed and allowed bins, where such logarithms give a large enhancement to the low Q_T cross section. These distributions are only rendered finite when one includes the $Q_T = 0$ contribution as is the case when one integrates out Q_T to form the inclusive cross section. As a consequence of this kinematic suppression, we can conclude that fixed-order results are indeed reliable in the forbidden regions of phase space.

At this point, we can consider the contributions as being intrinsically ZJ in nature due to the implicit Q_T requirement. As a consequence, it becomes possible in these bins to extend the results to $\mathcal{O}(\alpha_s^3)$ through the use of the ZJ calculation in NNLO-JET, which gives exactly the contributions that would be found in a full calculation of the inclusive Drell-Yan cross section to N³LO. This will have the impact of enhancing the accuracy of these predictions in the high- y_{ll} region, where the asymmetry A_{FB} is largest. The same cannot be said for the partially forbidden, partially allowed ‘mixed’ bins, where one will encounter the divergence at the boundary between the two regions and logarithmically divergent $Q_T = 0$ contributions are present. In these bins, one is restricted to NNLO.

In principle, the above discussion would allow one to adjust the experimental binnings in order to maximise the precision of the available theory in future experimental measurements. Were one to construct bin edges that align with the kinematic boundary, one could consider using the inclusive NNLO calculation in the allowed region, and solely using the NNLO ZJ calculation along with resummation in the

forbidden region, as it would amount to a systematic removal of the mixed bins. The resummation would also be rendered optional if a small margin $y_{ll} \sim \mathcal{O}(0.01)$ between the forbidden and allowed regions were required, approximately equivalent to the difference between $y_{ll} = 2$, $\cos \theta^* = 0.4$ and the LO kinematic bound, in order to remove the forbidden region in which $Q_T^{\text{min}} \sim 0$.

Since the current discussion is related only to the kinematics of the process and not the gauge group, similar observations can be made for the higher order EW corrections to the process. For the associated real contributions, the transverse momentum required can be created through photon emission, such that in the forbidden regions, there are two “LO” contributions, one at $\mathcal{O}(\alpha^2 \alpha_s)$, and one at $\mathcal{O}(\alpha^3)$, where due to the relative coupling constants, the QCD contribution will dominate. One notes also that the dominant contributions to higher order EW corrections come in the form of massive virtual contributions, with radiative corrections largely suppressed, so the magnitude of this effect is considerably smaller than in the QCD case.

Once again, a similar procedure can be repeated for the CF case to generate the Q_T dependence of the phase space constraints. For the region of phase space above $\cos \theta^* \sim 0.9$, limited by the maximum Δy_{ll} permitted by the cuts, one can proceed simply by rearranging (5.4.6), this time retaining the full p_T^l dependence. Substituting the $\Delta y_{ll}^{\text{max}}$ values from Table 5.3 for the appropriate $|y_{ll}|$ region, alongside minimum p_T^l values permitted by the cuts, this gives the minimum Q_T dependence of the upward Born-forbidden region as

$$Q_T^2 \geq (p_{T,C}^{l,\text{min}})^2 + (p_{T,F}^{l,\text{min}})^2 + 2p_{T,C}^{l,\text{min}} p_{T,F}^{l,\text{min}} \cosh(\Delta y_{ll}^{\text{max}}) - \frac{\left(2p_{T,C}^{l,\text{min}} p_{T,F}^{l,\text{min}} \sinh(\Delta y_{ll}^{\text{max}})\right)^2}{\cos^2 \theta^* \left((p_{T,C}^{l,\text{min}})^2 + (p_{T,F}^{l,\text{min}})^2 + 2p_{T,C}^{l,\text{min}} p_{T,F}^{l,\text{min}} \cosh(\Delta y_{ll}^{\text{max}})\right)}. \quad (5.4.14)$$

We can now consider the new case in which the bound is given by $\Delta y_{ll}^{\text{min}}$, in the lower region of the central-forward $(y_{ll}, \cos \theta^*)$ plane. Taking the transverse components from Eqn. (5.4.1), one can write

$$Q_T^2 = (p_{T,F}^l)^2 + (p_{T,C}^l)^2 + 2p_{T,F}^l p_{T,C}^l \cos(\Delta \theta_{FC}), \quad (5.4.15)$$

where the angular separation of the two leptons $\Delta\theta_{FC} = |\theta_F - \theta_C|$. Using this in conjunction with Eqn. (5.4.14), one can then identify

$$\cos(\Delta\theta_{FC})^{\min} \geq \cosh(\Delta y_{ll}^{\max}) - \frac{2p_{T,C}^{l,\min} p_{T,F}^{l,\min} \sinh^2(\Delta y_{ll}^{\max})}{\cos^2 \theta^* \left((p_{T,C}^{l,\min})^2 + (p_{T,F}^{l,\min})^2 + 2p_{T,C}^{l,\min} p_{T,F}^{l,\min} \cosh(\Delta y_{ll}^{\max}) \right)}, \quad (5.4.16)$$

such that when Q_T^2 is minimised, so is $\cos(\Delta\theta_{FC})$. However, for the constraints derived from minimal values Δy_{ll}^{\min} , there is not an equivalent meaningful lower bound on $\cos(\Delta\theta_{FC})$, as it first saturates at $\cos(\Delta\theta_{FC}) = -1$. It is this saturation that complicates the picture when considering the minimum values of Q_T in the region below the Δy_{ll}^{\min} , corresponding to $\cos \theta^* \lesssim 0.9$, as one can not rely on $\cos(\Delta\theta_{FC})$ being minimised to some value > -1 by Δy_{ll} .

In order to minimise Q_T for a given $(\cos \theta^*, y_{ll})$ point in this region, one must find values of $p_{T,F}^l$ and $p_{T,C}^l$ that are consistent with $\cos(\Delta\theta_{FC}) = -1$. This is straightforward if one enforces that $\Delta\theta_{FC} = \pi$, $\cos(\Delta\theta_{FC}) = -1$ in order to minimise Q_T , such that (for arbitrary lepton ordering 1,2)

$$p_{T,1}^l = Q_T + p_{T,2}^l = Q_T + p_T^l; \quad (5.4.17)$$

which ensures that one lepton is parallel to \vec{Q}_T in the \vec{p}_T plane. One can then substitute this directly into (5.4.6) and solve for Q_T , to give:

$$Q_T^2 \geq \frac{2p_T^{l,2} (1 + \cosh(\Delta y_{ll}^{\min}))^2 (\cos^2 \theta^* - \tanh(\Delta y_{ll}^{\min})^2)}{\cos^4 \theta^*} \times \left[\frac{\cos^2 \theta^* \cosh(\Delta y_{ll}^{\min})}{1 + \cosh(\Delta y_{ll}^{\min})} - \tanh(\Delta y_{ll}^{\min})^2 + \left(\frac{\cos^2 \theta^* - \tanh(\Delta y_{ll}^{\min})^2}{(1 + \cosh(\Delta y_{ll}^{\min}))} (\cos^2 \theta^* (\cosh(\Delta y_{ll}^{\min}) - 1) - (1 + \cosh(\Delta y_{ll}^{\min})) \tanh(\Delta y_{ll}^{\min})^2) \right)^{\frac{1}{2}} \right]. \quad (5.4.18)$$

It is straightforward to see that this is minimised for the smallest value of p_T^l accessible to *both* $p_{T,C}^l$ and $p_{T,F}^l$ permitted by the cuts

$$p_T^l = p_{T,\min}^l = \max(p_{T,C}^{l,\min}, p_{T,F}^{l,\min}), \quad (5.4.19)$$




Figure 5.11: Minimum Q_T values required for the different regions of the $(|y_{ll}|, \cos \theta^*)$ plane in the central-forward region. Overlaid are the Z3D measurement bins, integrated over m_{ll} .

which for the Z3D measurement gives $p_{T,\min}^l = 25$ GeV. Now this is in place, one can calculate the minimum values of Q_T required across the $(y_{ll}, \cos \theta^*)$ plane, which is shown in Fig. 5.11 in the $y_{ll} = \frac{1}{2}(y_1^l + y_2^l)$ approximation.

One interesting effect in the CF region is the emergence of an ultra-forbidden region due to the constraints induced by Δy_{ll}^{\min} , which is excluded to all orders in perturbation theory. This is present as $\cos \theta^* \rightarrow 0$, and is defined by the region where $Q_T^{\min} > \frac{\sqrt{s}}{2}$, such that there can never be enough energy present in the event to overcome the minimum Q_T and allow an event to occur.

5.5 Acceptances and k -factors

Not only does the above kinematic discussion allow a systematic improvement in the evaluation of certain bins, it also directly informs the structure of the acceptances and the k -factors across the phase space binning. Here we use the standard definition of the acceptance as the ratio of a cross-section computed with fiducial cuts to the corresponding cross-section without any fiducial cuts. A high acceptance ~ 1 indicates that the measurement is relatively independent of the cuts, and the bulk of the total cross section contribution is contained within the fiducial region.




Figure 5.12: Acceptances for the CC Z3D fiducial region. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos\theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos\theta^*$ values are denoted by the central colour of each point, and the exterior colour and shape label the bin as allowed, partially allowed or forbidden.

This can be seen when considering the structure of all bins together, as shown in Fig. 5.12 for the acceptances in the CC region using NNLO QCD predictions with and without fiducial cuts. To display the full triple differential structure, we define the bin index using the index of each observable \mathcal{O}^{idx} from low to high as

$$\text{Bin No.} = 72 \cdot m_{ll}^{\text{idx}} + 6 \cdot y_{ll}^{\text{idx}} + \cos \theta^{*\text{idx}}, \quad (5.5.1)$$

such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). We use this as our x -axis.

One can see a sharp decline in acceptance within each invariant mass bin as one moves into the partially allowed and then forbidden regions with increasing y_{ll} . This is expected, as the Q_T restriction for a given point in $(y_{ll}, \cos \theta^*)$ space is caused entirely by the lepton rapidity cuts, and the bulk of the fixed-order Drell-Yan cross section lies at low Q_T meaning that a Q_T cut greatly decreases the acceptance. The structure between invariant mass bins, where acceptance increases with m_{ll} independent of the values of $\cos \theta^*$ and y_{ll} , are an artefact of the imposition of lepton p_T cuts. Lower invariant mass events are correlated with lower centre of mass energies, which are less likely to have the lepton transverse momenta required to pass the fiducial cuts. As a result, more of the cross section lies outside of the fiducial region, resulting in a reduced acceptance in the low- m_{ll} phase space regions.

That the variation is so extreme reflects the large impact of the cuts on the cross-section across phase space. It is important when using data for extractions of e.g. $\sin^2 \theta_W^{\text{eff}}$ that fiducial cuts do not introduce any systematic bias into the final results, and thus it is important to be aware of any limitations that low acceptance regions of phase space might have. Indeed, for $\sin^2 \theta_W^{\text{eff}}$ extraction, one can impose an acceptance cut below which multi-differential asymmetry A_{FB} values are constructed. This strategy greatly reduces the dependence of the result on the definition of the fiducial region due to the relative independence of the angular coefficients from the cuts. This occurs at the cost of a decreased PDF sensitivity as one loses the

ability to fit directly at the cross section level.

To remove this dependence on the fiducial cuts as best as possible whilst still maintaining the PDF sensitivity, the procedure followed in the ATLAS Z3D $\sin^2 \theta_W^{\text{eff}}$ extraction is to use an acceptance cut of $\mathcal{O}(95\%)$ whereby for high acceptances cross sections are directly fitted, and below this asymmetries are constructed and then fitted. One can then vary the exact value of this cut to ensure that any final extracted value is independent of the acceptance cut to a desirable level. From the kinematics, we can then understand the majority of below-cut bins as lying within the forbidden and mixed regions of phase space where they suffer kinematic suppression directly as a result of the cuts applied.

The use of an acceptance cut also has a secondary effect of ensuring that the theoretical predictions are robust and relatively insensitive to higher order corrections. The low acceptance regions are strongly correlated with large k -factors in the theory predictions as the phase space restrictions become relaxed at higher orders. This occurs as a result of the partonic radiation, which generates kinematic configurations inaccessible at lower orders, and can be seen in Fig. 5.13, which shows the NNLO/NLO QCD k -factor.

The effect is particularly evident in the forbidden region, where the LO contribution is identically zero, such that the perturbative series effectively begins at $\mathcal{O}(\alpha_S)$. As a result, the NNLO/NLO ($\mathcal{O}(\alpha_S^2)/\mathcal{O}(\alpha_S)$) k -factor in these regions only captures the inclusion of the first additional perturbative order, which typically gives ($\mathcal{O}(20\%)$) corrections for processes in which a single vector boson is produced. In Fig. 5.13 this is the case, where the majority of forbidden bins lie outside of the y -axis range, corresponding to corrections of $> 10\%$.

These large k -factors in the forbidden region can however be remedied by the inclusion of the known $\mathcal{O}(\alpha_S^3)$ corrections through exploitation of the ZJ calculation implemented within NNLOJET, as discussed in Section 5.4.2. This effect is shown




Figure 5.13: NNLO/NLO ($\mathcal{O}(\alpha_S^2)/\mathcal{O}(\alpha_S)$) k -factors for the CC Z3D fiducial region. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos \theta^*$ values are denoted by the central colour of each point, and the exterior colour and shape label the bin as allowed, partially allowed or forbidden. The majority of k -factors for the forbidden region are outside the bounds of the plot.

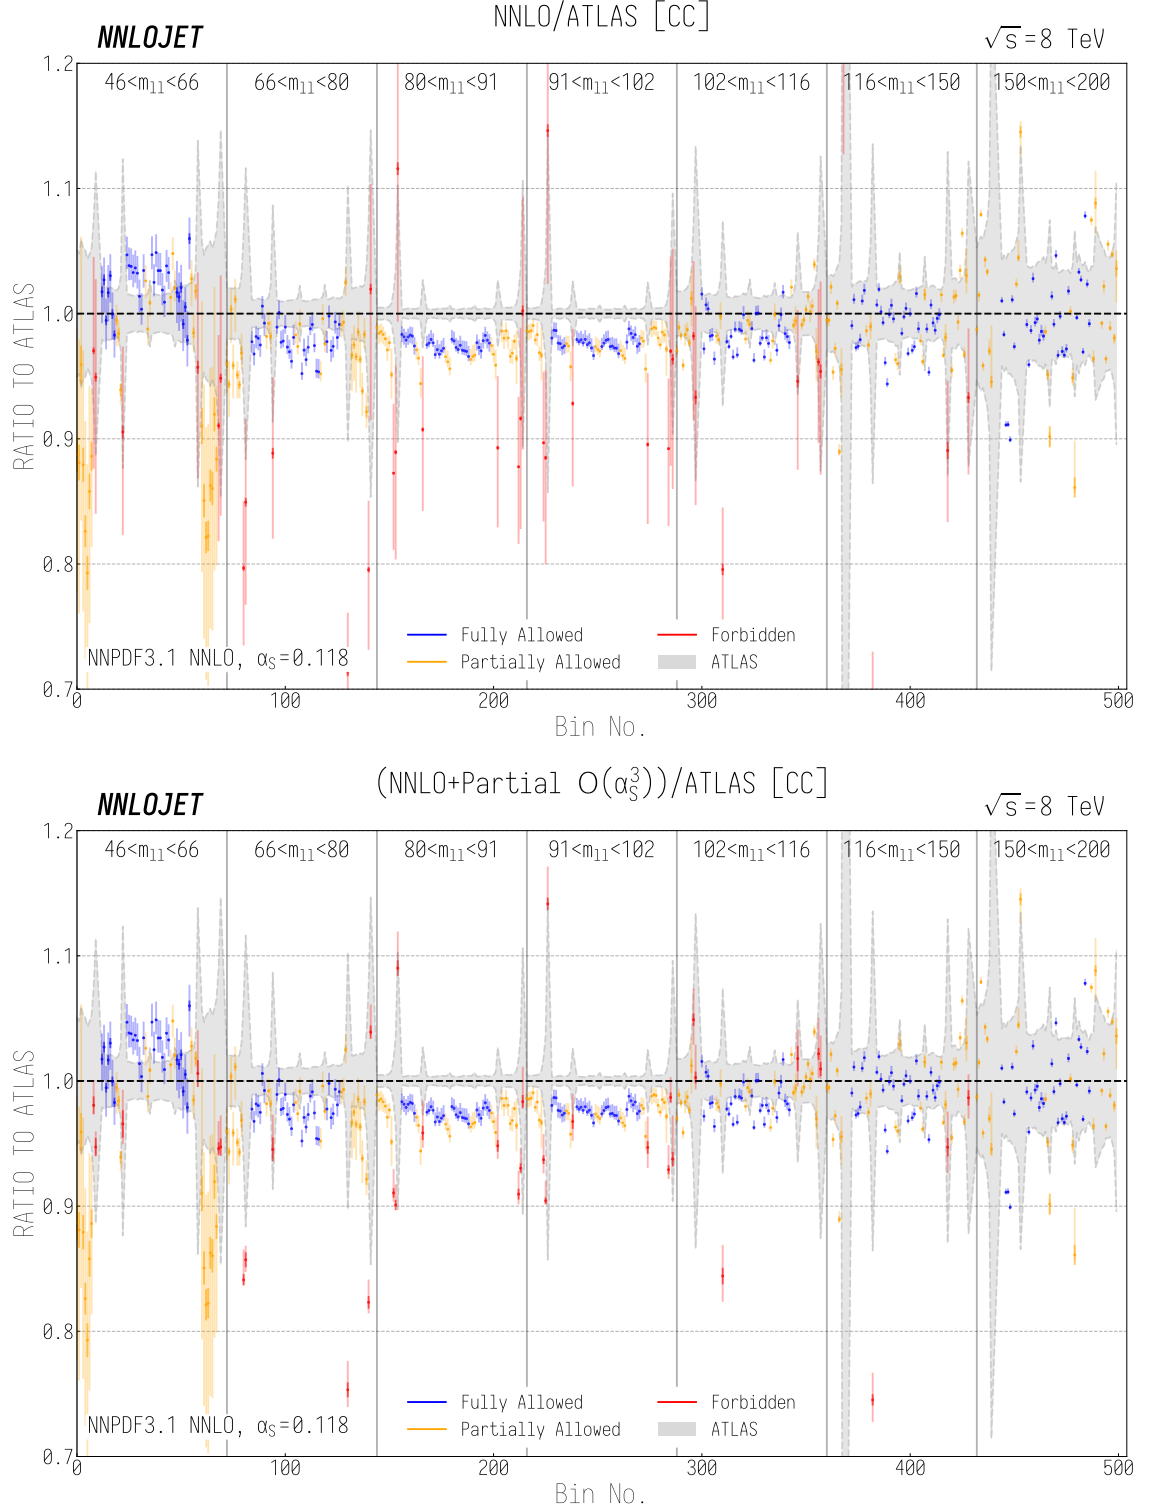


Figure 5.14: The ratio of the NNLO QCD predictions to ATLAS data in the central-central region of the Z3D analysis. The upper plot shows the $\mathcal{O}(\alpha_s^2)$ theory predictions, and the lower plot includes the $\mathcal{O}(\alpha_s^3)$ contributions for the forbidden bins. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right). Luminosity uncertainties of $\sim 1.8\%$ are not included.

between the two panels of Fig. 5.14, where we redefine the bin index as

$$\text{Bin No.} = 72 \cdot m_{ll}^{\text{idx}} + 12 \cdot \cos \theta^{*\text{idx}} + y_{ll}^{\text{idx}}, \quad (5.5.2)$$

and we see that these corrections both stabilise the predictions in the forbidden bins substantially and considerably reduce the scale variation uncertainty in the theory results. This indicates a substantial improvement in the reliability of the fixed-order results. It is noticeable that at low invariant masses, the scale variation uncertainty is much larger than for the high invariant mass counterparts. This can be traced back to the running of α_S having a steeper gradient at low scales, and a larger factorisation scale dependence as the PDF and hard cross section are less easily distinguished. This is particularly evident in the lowest m_{ll} bin, where the scale uncertainties in the partially allowed bins are $\mathcal{O}(10\%)$, even at $\mathcal{O}(\alpha_S^2)$. The inclusion of the $\mathcal{O}(\alpha_S^3)$ terms makes the largest impact here, to the point that the scale variation error in the forbidden bins become smaller than the corresponding fully and partially allowed bins. From this it is reasonable to conclude that the full N³LO Drell-Yan results once available will give the largest improvement to scale uncertainties at low m_{ll} .

If we now consider the agreement of the pure QCD predictions with the data, we see that in the region of the Z pole the theory undershoots the data by 2 – 3%. This discrepancy can be traced back to two primary causes. The first is the absence of NLO EW corrections, which for this measurement give an $\mathcal{O}(1\%)$ increase to the cross section between 80 GeV < m_{ll} < 102 GeV, as shown in Figure 5.15. This is in contrast to the extremal 46 GeV < m_{ll} < 66 GeV and 150 GeV < m_{ll} < 200 GeV bins where the EW corrections are $\mathcal{O}(-1\%)$ and $\mathcal{O}(-1.5\%)$ respectively. Each of these gives shape corrections that considerably improve the agreement with data. The second factor is the luminosity uncertainty (neglected in the figures), which is estimated at $\mathcal{O}(1.8\%)$ for the measurement and is strongly correlated between bins.

It is also notable that the forbidden bins correspond to the regions in which the statistical uncertainty on the data is largest, as a result of the kinematic suppression




Figure 5.15: EW (NLO+partial higher order)/LO k -factors for the CC Z3D fiducial region as produced for $M_W = 79.939$ GeV using Powheg [262–265]. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos \theta^*$ values are denoted by the colour of each point.




Figure 5.16: EW (NLO+partial higher order)/LO k -factors for the CF Z3D fiducial region as produced for $M_W = 79.939$ GeV using Powheg [262–265]. The bin number is as defined in (5.5.1), with appropriate modifications for the CF region, such that the major m_{ll} bins are divided into 5 y_{ll} sub-bins from 0-2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right). The different $\cos \theta^*$ values are denoted by the colour of each point.

discussed above. This uncertainty also increases away from the Z pole region, due to the reduced cross section away from the Breit-Wigner peak.

We show the equivalent QCD-only plots in Figure 5.17 for the central-forward region of the analysis, where we use the equivalent bin index notation adjusted for the central forward bin edges:

$$\text{Bin No.} = 30 \cdot m_{ll}^{\text{idx}} + 6 \cdot \cos \theta^{*\text{idx}} + y_{ll}^{\text{idx}}. \quad (5.5.3)$$

Whilst the uncertainty on the data is considerably larger than in the central region, we again see a reasonable description of the data by the theory, albeit with a slight overshoot in the region of the Z pole. The inclusion of the $\mathcal{O}(\alpha_s^3)$ corrections to the forbidden bins again substantially decreases the scale uncertainty of the associated bins, save for the extremal invariant mass bins. Due to the extreme kinematic suppression in these bins relative to other regions, the statistical uncertainty on the theory predictions is considerable and in some cases the dominant source of uncertainty. However, these bins typically have extremely large experimental uncertainties so for fitting purposes this is sufficient.

5.6 PDF Variation

We now turn to the effects of PDF variation on the results presented. Whilst we have so far neglected the effects of PDF uncertainties, these are a dominant theory uncertainty and must be taken into account. We show in Figure 5.18 the ratios of the central members of the **MMHT14** [38] and **CT14** [238] PDF sets to our benchmark **NNPDF3.1** [117] results in the central-central region, using as an index the definition of (5.5.1). The comparison between different PDF sets is primarily representative of methodological differences between the PDF fitting collaborations, incorporating effects due to fitting procedures, parameterisations, experimental data sets, input theory and so on.

For the central members of both **MMHT14** and **CT14** sets we see a shape difference

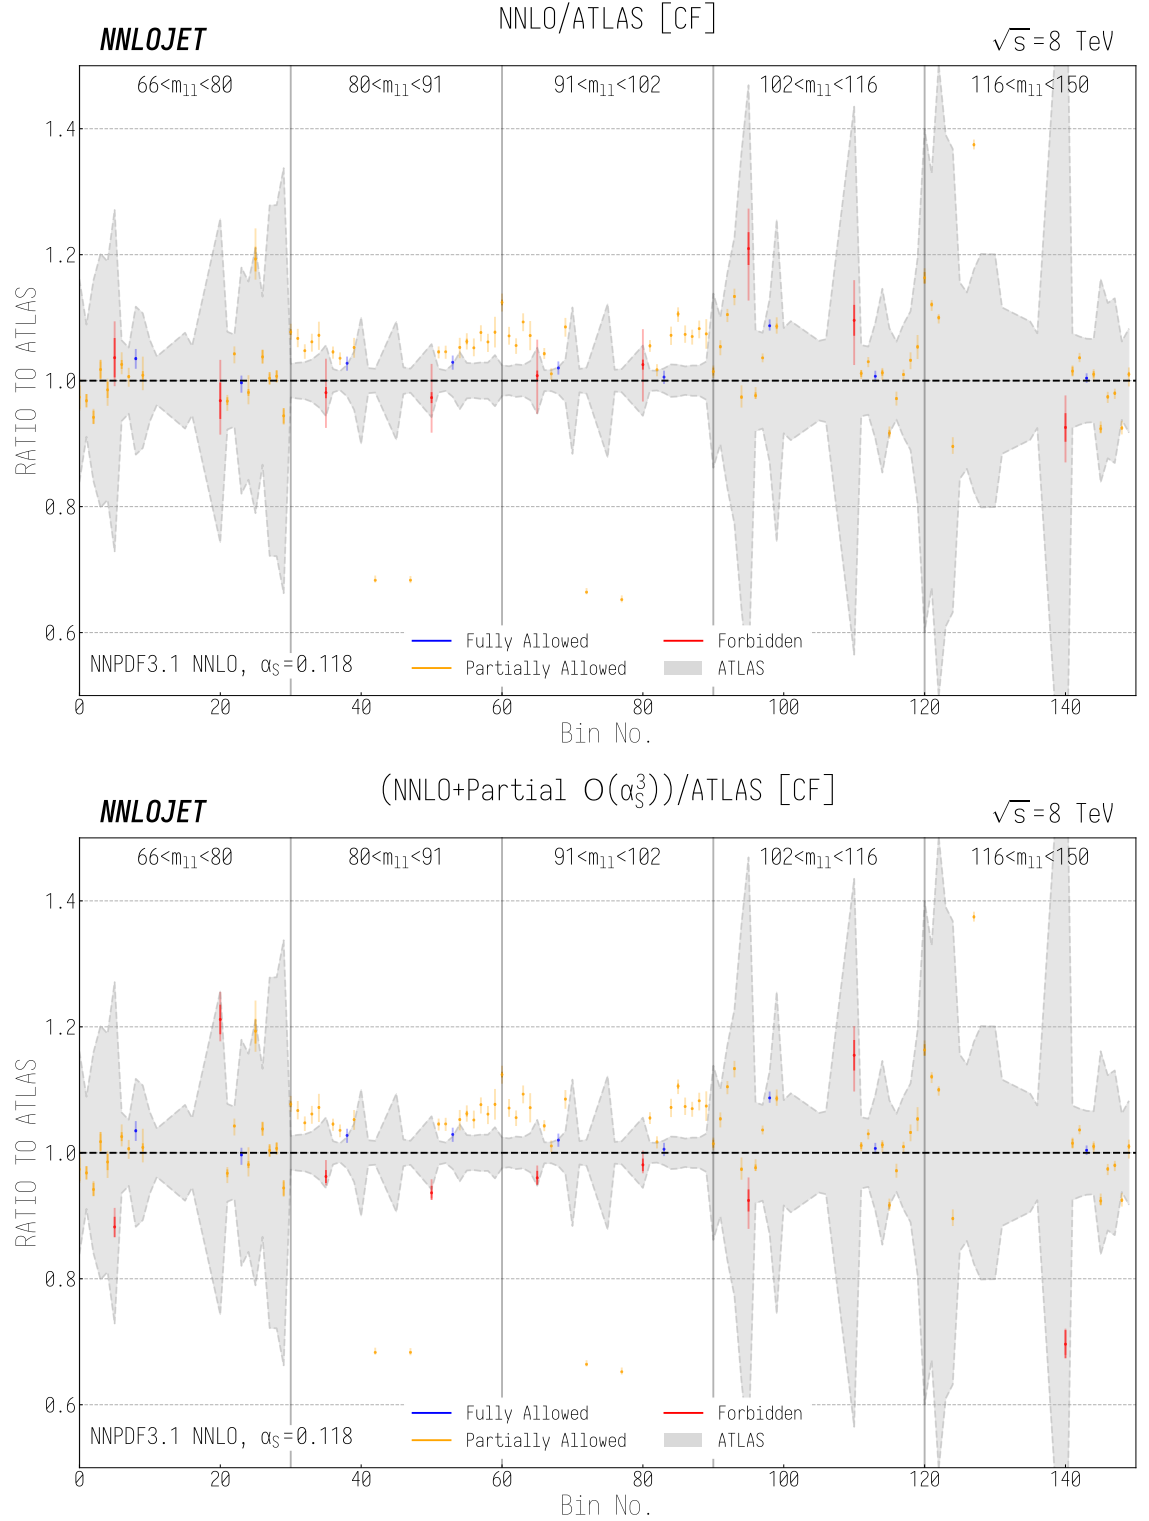


Figure 5.17: The ratio of the NNLO QCD predictions to ATLAS data in the central-forward region of the Z3D analysis. The upper plot shows the $\mathcal{O}(\alpha_s^2)$ theory predictions, and the lower plot includes the $\mathcal{O}(\alpha_s^3)$ contributions for the forbidden bins. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.3), such that the major m_{ll} bins are divided into $6 \cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 6 y_{ll} sub-bins from 0 to 2.4 (left to right). Luminosity uncertainties of $\sim 1.8\%$ are not included.




Figure 5.18: The ratio of the central predictions of the MMHT14 (top panel) and CT14 (bottom panel) PDF sets to the predictions of NNP3.1 data in the central-central region of the Z3D analysis. Both plots show the predictions in all bins to $\mathcal{O}(\alpha_s^2)$. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty. The bin number is as defined in (5.5.1), such that the major m_{ll} bins are divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right) which are in turn divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right).

across the variation in rapidity, with central rapidities showing a $\mathcal{O}(3 - 4\%)$ difference with respect to the NNPDF3.1 results, which decreases to $\mathcal{O}(1 - 2\%)$ in the forward regions of the measurement. This can be interpreted primarily as the impact of different sea and valence quark distributions between the three sets, given the dominant incoming partonic sub-process in Drell-Yan production is quark-antiquark annihilation. These are analysed in more detail in [117], where the primary driver of difference between the sets occurs in the anti-quark distributions at $Q \sim M_Z$, and are visualised in Fig. 5.19. These effects are lessened at high $y_{ll} \leftrightarrow x$, where the valence quark contributions dominate over the sea quark and the central members of the PDF sets exhibit better agreement. The kinematic regions covered by each of the central-central and central-forward regions can be seen in Figure 5.20, where one sees the distinctive “split” kinematic region associated with the central-forward selection.

This comparison between PDF sets does not, however, account for uncertainties *within* each PDF set, which are parameterised through $\mathcal{O}(30 - 100)$ additional Hessian or eigenvector sets. In order to evaluate these using standard NNLO techniques, one must perform a separate NNLO calculation for each set member, which whilst technically possible is prohibitively computationally expensive. At NLO, grid techniques are well-established for dealing with this issue, where the PDF dependence of the (differential) cross section is stored using look-up tables which allow *a posteriori* evaluation of PDF uncertainties [268, 269].

Whilst these are being extended to NNLO for certain processes [270, 271], results are not yet widespread and largely still in development. Standard practice is to reweight NLO results for PDF variation obtained using these look-up tables with NNLO/NLO k -factors, a technique which is also used within the fitting of the PDFs themselves. The closure of this method can be checked using dedicated NNLO runs either for specific members of a single PDF set, or for central members of different sets, where good agreement is generally found. The PDF set uncertainties for the results shown in Figure 5.18 have been evaluated in this manner, and are large




Figure 5.19: The ratio of MMHT14 and CT14 PDF sets to the central member of the NNPDF3.1 set as a function of Bjorken- x . From the top left panel clockwise, the panels correspond to the d valence, u valence, u sea quark and d sea quark contributions to the various PDF sets at $Q^2 = M_Z^2$. The uncertainties given are the PDF uncertainties.




Figure 5.20: The kinematic regions in the (x, Q^2) plane associated with the central-central (yellow) and central-forward (green) fiducial selections of the Z3D measurement. The total kinematic reach of the ATLAS detector at $\sqrt{s} = 8$ TeV is shown in red.

enough to accommodate the differences between PDF sets. As a result, the results produced using the MMHT14, CT14 and NNPDF3.1 central members are not mutually inconsistent, and indeed when they are profiled one observes improved agreement between PDF sets. Further well-understood data can be folded in to this profiling procedure in order to further reduce PDF uncertainties at the differential cross section level.

5.7 $\sin^2 \theta_W^{\text{eff}}$ Variation

Thus far, all results presented have been calculated in the G_μ scheme, using $M_W = 79.939$ GeV in order to account for the effects of $\sin^2 \theta_W^{\text{eff}}$ as is standard practice. At LO EW ($\Delta\kappa = 0$), with $M_Z = 91.1876$ GeV, this corresponds to $\sin^2 \theta_W^{\text{eff}} = \sin^2 \theta_W = 0.23150$, which is in agreement with the central value of the current best measurement. For a $\sin^2 \theta_W^{\text{eff}}$ fit to be performed however, this value must be varied within some range about the expected value such that a χ^2 fit or similar can be

used to determine the best fit value and uncertainties. This variation is primarily an electroweak effect, and can be performed by an EW reweighting about the fixed-order QCD results in order to avoid performing a full NNLO QCD run for each $\sin^2 \theta_W^{\text{eff}}$ point included in the fit. The closure of this technique can be checked by producing dedicated NNLO QCD results for a given set of $\sin^2 \theta_W^{\text{eff}}$ values.

The ratio of two such dedicated runs are shown in Fig. 5.21 for the central-central (upper) and central-forward (lower) fiducial regions. For clarity, we choose extremal values about the central $\sin^2 \theta_W^{\text{eff}}$ value, corresponding to

$$\frac{M_W}{M'_W} = \frac{79.864 \text{ GeV}}{80.014 \text{ GeV}} \longleftrightarrow \frac{\sin^2 \theta_W^{\text{eff}}}{\sin^2 \theta_W^{\text{eff}'}} = \frac{0.23294}{0.23005}. \quad (5.7.1)$$

This allows us to see how the sensitivity of the measurement is distributed kinematically before the inclusion of electroweak corrections, and we adopt the indexing of Eqn. (5.5.2) to make this distribution clear.

Concentrating on the central values ($-0.4 < \cos \theta^* < 0.4$) of the cross section, we can see that the largest variation with $\sin^2 \theta_W^{\text{eff}}$ comes in the low ($46 \text{ GeV} < m_{ll} < 66 \text{ GeV}$) invariant mass region, where cancellation between the negative $|Z/\gamma^*|$ interference $\sin^2 \theta_W^{\text{eff}}$ term and positive pure Z terms in the cross section drives the sensitivity. At the cross section level, the variation with $\sin^2 \theta_W^{\text{eff}}$ decreases with increasing m_{ll} in both central-central and central-forward selections.

The change in asymmetry between the two $\sin^2 \theta_W^{\text{eff}}$ values is governed by the variation of the cross section with $\cos \theta^*$ across each invariant mass bin and behaves differently due to the steep gradient of A_{FB} with m_{ll} in the central m_{ll} region (see Fig. 5.4). Here we see the largest differences between forward and backward contributions in these m_{ll} regions about the Z peak, most prominently in the forward rapidity bins. That these forward bins vary the most is an artefact of the Collins-Soper frame, in which the incoming quark and antiquark are best identified statistically in the forward region, and therefore the asymmetry is more sensitive to $\sin^2 \theta_W^{\text{eff}}$ in these regions. In the CC region, this behaviour also naturally correlates with the mixed and forbidden bins (yellow and red respectively). In the CF selection

this behaviour is exaggerated (albeit with larger scale variation uncertainties), and about the Z peak the $0.7 < \cos \theta^* < 1$ region changes by several percent.

That a measurable change in the cross sections is possible is crucial to the competitiveness of the Z3D measurement as a probe of the effective weak mixing angle. From [5], the preliminary uncertainty on $\sin^2 \theta_W^{\text{eff}}$ extracted from the associated 8 TeV angular coefficient measurement is 36×10^{-5} , with a central value of 0.23140. This has been made using the same ATLAS Run 1 data set, and the Z3D analysis is likely to be competitive with this more standard approach to the $\sin^2 \theta_W^{\text{eff}}$ extraction.

5.8 Running Width Corrections

There is one final technical point to be made with regards to the binning chosen for the Z3D measurement. The existence of a bin edge at 91 GeV gives an exceptional sensitivity of the cross section in the peak region to the precise location of the Z pole, which can be shifted by $\mathcal{O}(35 \text{ MeV})$ through EW effects [272]. If one generates predictions using the standard fixed propagator width Γ_Z and PDG mass M_Z , a step effect is observed across the bin boundary when considering the ratio of theory to data, as shown in Fig. 5.22. As a result, it is crucial to have a consistent description of the EW parameterisation of the cross section when performing a fit, which is largely summarised in the form of the propagator and choice of input values.

Near the Z pole, Dyson resummation gives a Z propagator of the form

$$D_Z = \frac{i}{s - M_Z^2 + \Pi_Z(s)}, \quad (5.8.1)$$

where $\Pi_Z(s)$ is the renormalised Z self energy in some arbitrary renormalisation scheme. For stable particles, $\Im(\Pi_Z(s)) = 0$, giving the propagator a pole at physical mass $s = M_{\text{phys}}^2$, whereas the propagator becomes complex for unstable particles due to the non-vanishing absorptive component in the associated self energies. When one works at LO EW, the explicit choice of renormalisation scheme is made exclusively




Figure 5.21: The ratio of the $\mathcal{O}(\alpha_s^2)$ predictions for the Z3D measurement made at $M_W = 79.864$ GeV and $M_W = 80.014$ GeV, corresponding to $\sin^2 \theta_W^{\text{eff}} = 0.23294$ and $\sin^2 \theta_W^{\text{eff}} = 0.23005$ for the CC (top) and CF (bottom) measurement regions. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty. The CC bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right). The CF bin number is defined analogously.




Figure 5.22: The “step effect” in the ratio of theory to data across the bin edge at $m_{ll} = 91$ GeV in the CC Z3D fiducial region, using the on-shell mass with a fixed propagator width.

through the choice of input parameters to the calculation. As a result, one must take particular care in this choice in order to ensure self-consistency, particularly in the measurement of a scheme-dependent parameter such as $\sin^2 \theta_W^{\text{eff}}$.

5.8.1 On-Shell (OS) scheme

The most common scheme used at LO EW is the the on-shell scheme. In the on-shell mass scheme, $\Pi_Z(s) = \Pi_Z^{\text{OS}}(s)$, and in the vicinity of the the Z pole,

$$D_Z = \frac{i}{s - M_Z^2 + iM_Z^{\text{OS}}\Gamma_Z^{\text{OS}}}, \quad (5.8.2)$$

where the on-shell width

$$\Gamma_Z^{\text{OS}} = \frac{1}{M_Z^{\text{OS}}} \Im(\Pi_Z^{\text{OS}}(M_Z^{\text{OS}2})), \quad (5.8.3)$$

and the on-shell mass is defined by the physical mass $M_Z^{\text{OS}} = M_{\text{phys}}$ if the particle were stable, equivalent to imposing as a renormalisation condition

$$M_Z^{\text{OS}2} = M_Z^2 + \Re(\Pi_Z(M_Z^{\text{OS}2})), \quad (5.8.4)$$

such that M^{OS^2} is the zero of the real part of the denominator in (5.8.1). As shown in [273], this is gauge dependent beyond the one-loop level, i.e. at $\mathcal{O}(\alpha^4)$ if one includes leptonic decays, although it suffices for LO EW calculations. It is the on-shell mass that has been measured to an extraordinary precision at LEP, taking the value $M_Z = 91.1876 \pm 0.0021$ [242], and it is when one naïvely makes this choice that the above step effect is encountered.

This can be simply extended to include an all-orders class of EW corrections, as shown in [272]. About the Z pole, and neglecting the light fermion masses, one can approximate the absorptive part of the Z self energy as

$$\Im(\Pi_Z^{\text{OS}}(s)) = s \frac{\Gamma_Z^{\text{OS}}}{M_Z^{\text{OS}}} \Theta(s), \quad (5.8.5)$$

where the Heaviside Θ function ensures that one doesn't permit a width in t -channel contributions with negative propagator invariant s . This naturally leads to a “running width” description of the Z boson resonance in the s -channel, using the propagator

$$D_Z = \frac{i}{s - M_Z^2 + is\Gamma_Z^{\text{OS}}/M_Z^{\text{OS}}}, \quad (5.8.6)$$

so-called because of the appearance of the energy s in the imaginary width part. This is the scheme used for the on-shell Z mass determinations at LEP and the Tevatron as well as the LEP $\sin^2 \theta_W^{\text{eff}}$ result [242], and correctly accounts for the step effect across the $m_{ll} = 91$ GeV bin edge we observe in Fig. 5.22 by effectively shifting the location of the Z peak.

This can be seen by rewriting the propagator in (5.8.6) in terms of fixed energy independent values, using the conversion factor

$$\gamma = \left(\frac{\Gamma_Z^{\text{OS}}}{M_Z^{\text{OS}}} \right) \quad (5.8.7)$$

to define the shifted mass, width and normalisation factor N

$$M'_Z = \frac{M_Z}{\sqrt{1 + \gamma^2}}, \quad \Gamma'_Z = \frac{\Gamma_Z}{\sqrt{1 + \gamma^2}}, \quad N = \frac{1}{1 + i\gamma} \quad (5.8.8)$$

such that [272]

$$\frac{i}{s - M_Z^2 + is\Gamma_Z^{\text{OS}}/M_Z^{\text{OS}}} = \frac{iN}{s - M_Z'^2 + i\Gamma_Z' M_Z'}. \quad (5.8.9)$$

For the PDG values of M_Z^{OS} and Γ_Z^{OS} , we have [8]

$$\begin{aligned} M_Z^{\text{OS}} &= 91.1876 \pm 0.0021 \text{ GeV} & \Gamma_Z^{\text{OS}} &= 2.4952 \pm 0.0023 \text{ GeV} & \gamma &= 0.027363 \\ M_Z' &= 91.1535 \text{ GeV} & \Gamma_Z' &= 2.4943 \text{ GeV}, \end{aligned} \quad (5.8.10)$$

which amounts to an effective reduction in the Z mass of 31.4 MeV and “undoes” the step effect through a redistribution of the cross section between the $80 \text{ GeV} < m_{ll} < 91 \text{ GeV}$ and $91 \text{ GeV} < m_{ll} < 102 \text{ GeV}$ bins. It is this running width scheme we use throughout this chapter unless otherwise stated; this is chosen in accordance with the EW corrections to be used in the Z3D analysis and for consistency with previous measurements of $\sin^2 \theta_W^{\text{eff}}$. That this accounts for the aforementioned step effect can be seen in e.g. Fig. 5.14.

5.8.2 Pole Mass (PM) scheme

A second scheme, the pole mass scheme, is also in widespread use. Here, the Z mass is defined through the complex-valued pole in (5.8.1)

$$\mu^2 - M_Z^2 - \Pi_Z^{\text{PM}}(\mu^2) = 0, \quad \mu^2 \equiv M_Z^{\text{PM}^2} - iM_Z^{\text{PM}}\Gamma_Z^{\text{PM}} \quad (5.8.11)$$

such that μ^2 is the zero of the complex denominator in (5.8.1). As the complex pole is a property of the S-matrix, this definition is manifestly gauge invariant to all orders, meaning that as a theoretical choice, the pole mass scheme is more robust than the on-shell scheme.

At LO EW, the PM scheme Z mass and width can be directly related their on-shell shifted counterparts as

$$M_Z' = M_Z^{\text{PM}}, \quad \Gamma_Z' = \Gamma_Z^{\text{PM}}, \quad (5.8.12)$$

such that this scheme also directly replicates the mass shift required to negate the step effect. However the two schemes differ in one crucial respect, the overall com-

	Fixed Width OS [fb]	Running Width OS [fb]	Fixed Width PM [fb]
$\cos \theta^* > 0$	5984.732 ± 0.029	5983.167 ± 0.029	5992.36 ± 0.13
$\cos \theta^* < 0$	5569.212 ± 0.027	5574.664 ± 0.027	5575.67 ± 0.12
A_{FB}	$0.035962 \pm 3 \times 10^{-6}$	$0.035341 \pm 3 \times 10^{-6}$	$0.036021 \pm 4 \times 10^{-5}$

Table 5.4: Tabulated A_{FB} values for the $80 < m_{ll} < 100$, $1.6 < y_{ll} < 2.5$ bin of [261], performed using different EW input parameter schemes.

plex normalisation contained in N . This complex phase directly affects the $|Z/\gamma^*|$ interference contribution to the cross section, and leads to a substantial modification of the asymmetry in the region of the Z pole where sensitivity to the interference is large (see Fig. 5.3) and the asymmetry itself is small. In this region the normalisation can lead to large differences between the schemes which is important to understand for high precision analyses.

To illustrate this, we construct the forward backward asymmetry in the region $80 \text{ GeV} < m_{ll} < 100 \text{ GeV}$, $1.6 < y_{ll} < 2.5$ used by the ATLAS 8 TeV angular coefficient analysis [261] in the fixed width scheme with OS EW input parameters, running width scheme with OS EW parameters and fixed width scheme with PM parameters. Here we see that whilst in each case the change in the individual forward and backward cross sections is marginal, at the per-mille level, the inclusion of the running width corrections induces a much larger shift of 1.7% to the constructed asymmetry due to the relative effect on the forward and backward cross sections induced by the change in the $|Z/\gamma^*|$ interference contribution. In comparison, the two fixed width schemes with OS and PM input parameters vary much less. Whilst 1.7% may not seem to be a large change, in the context of an asymmetry analysis where many systematic uncertainties cancel it is considerable.

The size of this effect on the asymmetry is largely dependent on the binning chosen for the measurement, as it relies on cross-cancellation between forward and backward cross sections when integrating over m_{ll} . We see this in Fig. 5.23, where we compare fixed width and running width values of $A_4 = \frac{8}{3} A_{FB}$ constructed using a variety of different binnings of the same data in the OS scheme. Small shifts to the forward and backward cross sections induced by inclusion of the running width propagator

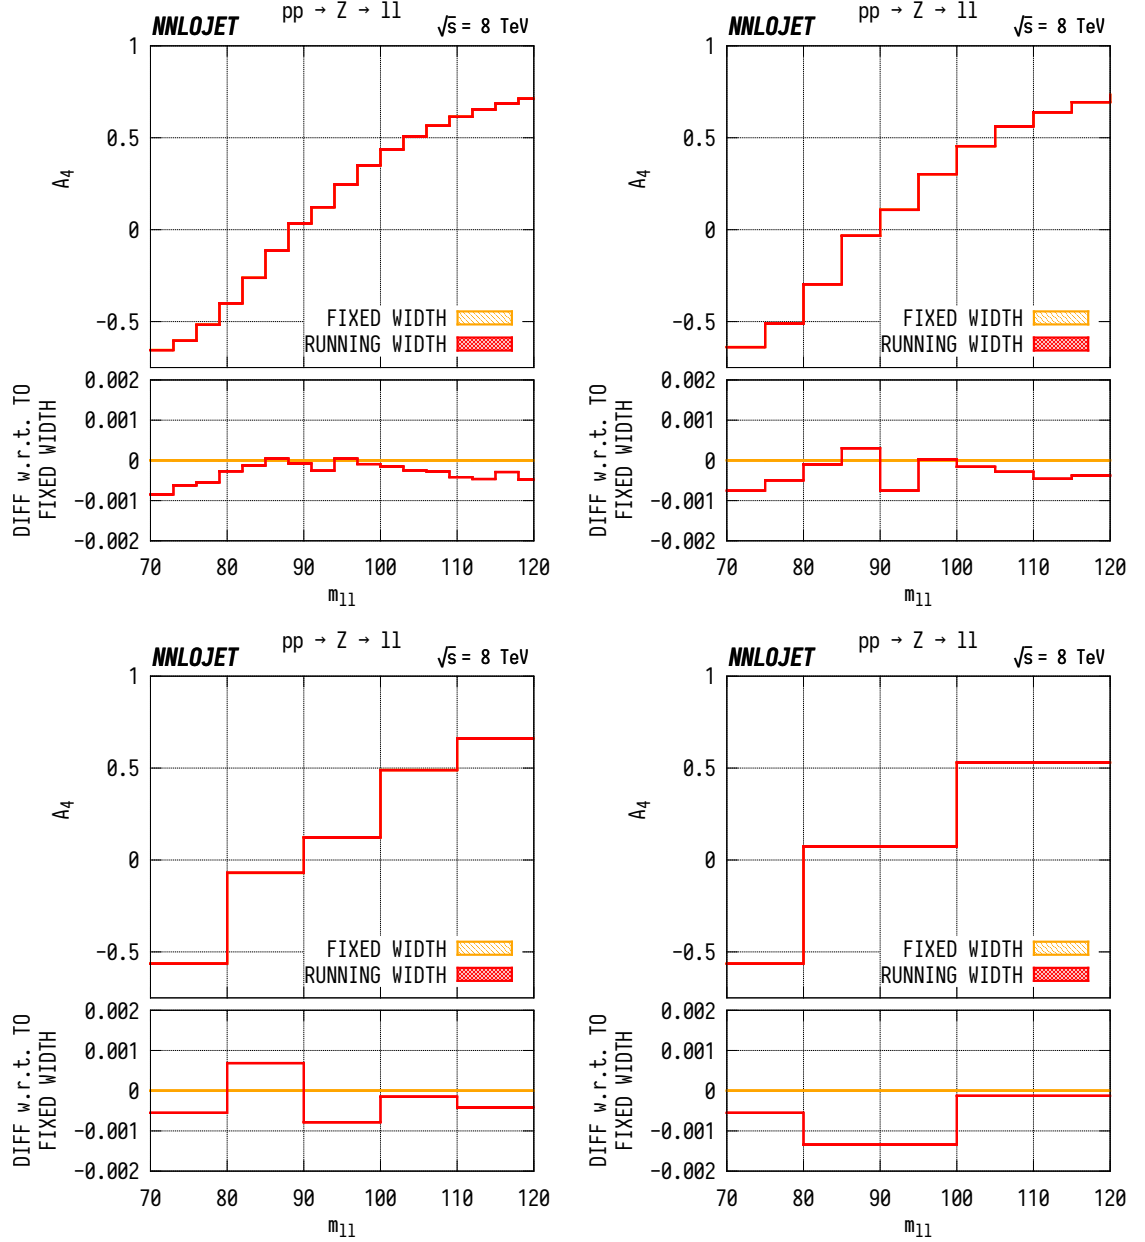


Figure 5.23: LO predictions for A_4 as a function of m_{ll} in the running and fixed width schemes, in both cases using OS input values. The different plots show the same data with a different binning in m_{ll} .

become magnified around the Z peak as the change in sign in A_{FB} causes a large cross cancellation between the two regions

$$\begin{aligned} m_{ll} \lesssim M_Z &\rightarrow \sigma^{(\cos \theta^* > 0)} < \sigma^{(\cos \theta^* < 0)}, \\ m_{ll} \gtrsim M_Z &\rightarrow \sigma^{(\cos \theta^* > 0)} > \sigma^{(\cos \theta^* < 0)}. \end{aligned} \quad (5.8.13)$$

This cross cancellation means that the size of the effect is *highly* dependent on the binning in m_{ll} chosen. If one uses a wide bin of 20 GeV across the m_{ll} peak as in the lower right panel of Fig. 5.23, the correction can reach $\mathcal{O}(2\%)$ which invariably has a large effect on any subsequent $\sin^2 \theta_W$ extraction. This large effect was not seen during the LEP extractions of $\sin^2 \theta_W^{\text{eff}}$, where e^+e^- collisions were performed for a variety of fixed $\sqrt{s} = m_{ll}$ and as such no cross-cancellation was observed. As a result, it is crucial to understand and consistently control for these effects when performing $\sin^2 \theta_W^{\text{eff}}$ extractions.

5.9 Combined EW+QCD Predictions using

Powheg and NNLOJET

We now conclude with the results for the cross section differential in rapidity for each bin in $\cos \theta^*$ and m_{ll} using the nominal value of $\sin^2 \theta_W^{\text{eff}} = 0.23150$. We construct combined QCD×EW results for the central-central region in Figures 5.24 and 5.25, including the $\mathcal{O}(\alpha_S^3)$ contributions to the forbidden bins and Powheg EW corrections. We use the G_μ EW scheme including running width effects, with the following OS input parameters:

$$\begin{aligned} M_Z &= 91.1876 \text{ GeV} & M_W &= 79.939 \text{ GeV} & G_\mu &= 1.663787 \times 10^{-5} \text{ GeV}^{-2} \\ \Gamma_Z &= 2.4952 \text{ GeV} & \Gamma_W &= 2.085 \text{ GeV} & \sin^2 \theta_W &= 0.23150 \\ \alpha(M_Z) &= 0.0110796, \end{aligned} \quad (5.9.1)$$

in combination with the central member of the NNPDF3.1_as_0118 PDF set. The unphysical value of M_W corresponds to that required to set $\sin^2 \theta_W = \sin^2 \theta_W^{\text{eff}}$ in the G_μ scheme at LO EW. The difference between the $\cos \theta^* > 0$ and $\cos \theta^* < 0$ contributions which constitutes the asymmetry is shown with the hatched area, where we see that the A_{FB} is indeed much smaller around the Z peak than in the extremal m_{ll} regions. For each plot, the forbidden and mixed bins lie to the right hand side, with the final four rapidity bins for $0.7 < |\cos \theta^*| < 1$ (red) and final two rapidity bins for $0.4 < |\cos \theta^*| < 0.7$ (blue) being forbidden and therefore supplemented with $\mathcal{O}(\alpha_s^3)$ corrections.

Corresponding results for the central-final region are given in Fig. 5.26, where we observe the same pattern of asymmetry with m_{ll} . Here the $\mathcal{O}(\alpha_s^3)$ enhanced bins are the left-most two rapidity bins for $0 < |\cos \theta^*| < 0.4$ (green) and left-most rapidity bin for $0.4 < |\cos \theta^*| < 0.7$ (blue).

The ratio to data for both the CC and CF regions is shown in Fig. 5.27, where we see for the first time in the CC region that the combination of EW and QCD corrections brings the theory closer to data than the QCD-only predictions of Fig. 5.14. The remaining discrepancy lies well within the remaining luminosity uncertainty which we do not show. The same is not true for the CF region, where we see consistent over-prediction of the data by theory. Here the predictions become very sensitive to the high- x valence quark distribution within the PDFs, and as we saw in Fig. 5.19 as well as in previous chapters, there is a considerable uncertainty in the high- x d -valence contribution within PDF sets. As a result, when PDF uncertainties are also included, one observes better agreement.

The higher order QCD results contained within this chapter form a subset of those provided to the ATLAS collaboration for use in a fit of $\sin^2 \theta_W^{\text{eff}}$ to the Z3D data. Those results not shown include further variations of $\sin^2 \theta_W^{\text{eff}}$ for closure tests, alongside results for different values of α_s and central scale choices. It is also anticipated that advances in NNLO grid technology will allow for a full NNLO evaluation of PDF uncertainties in the near future.

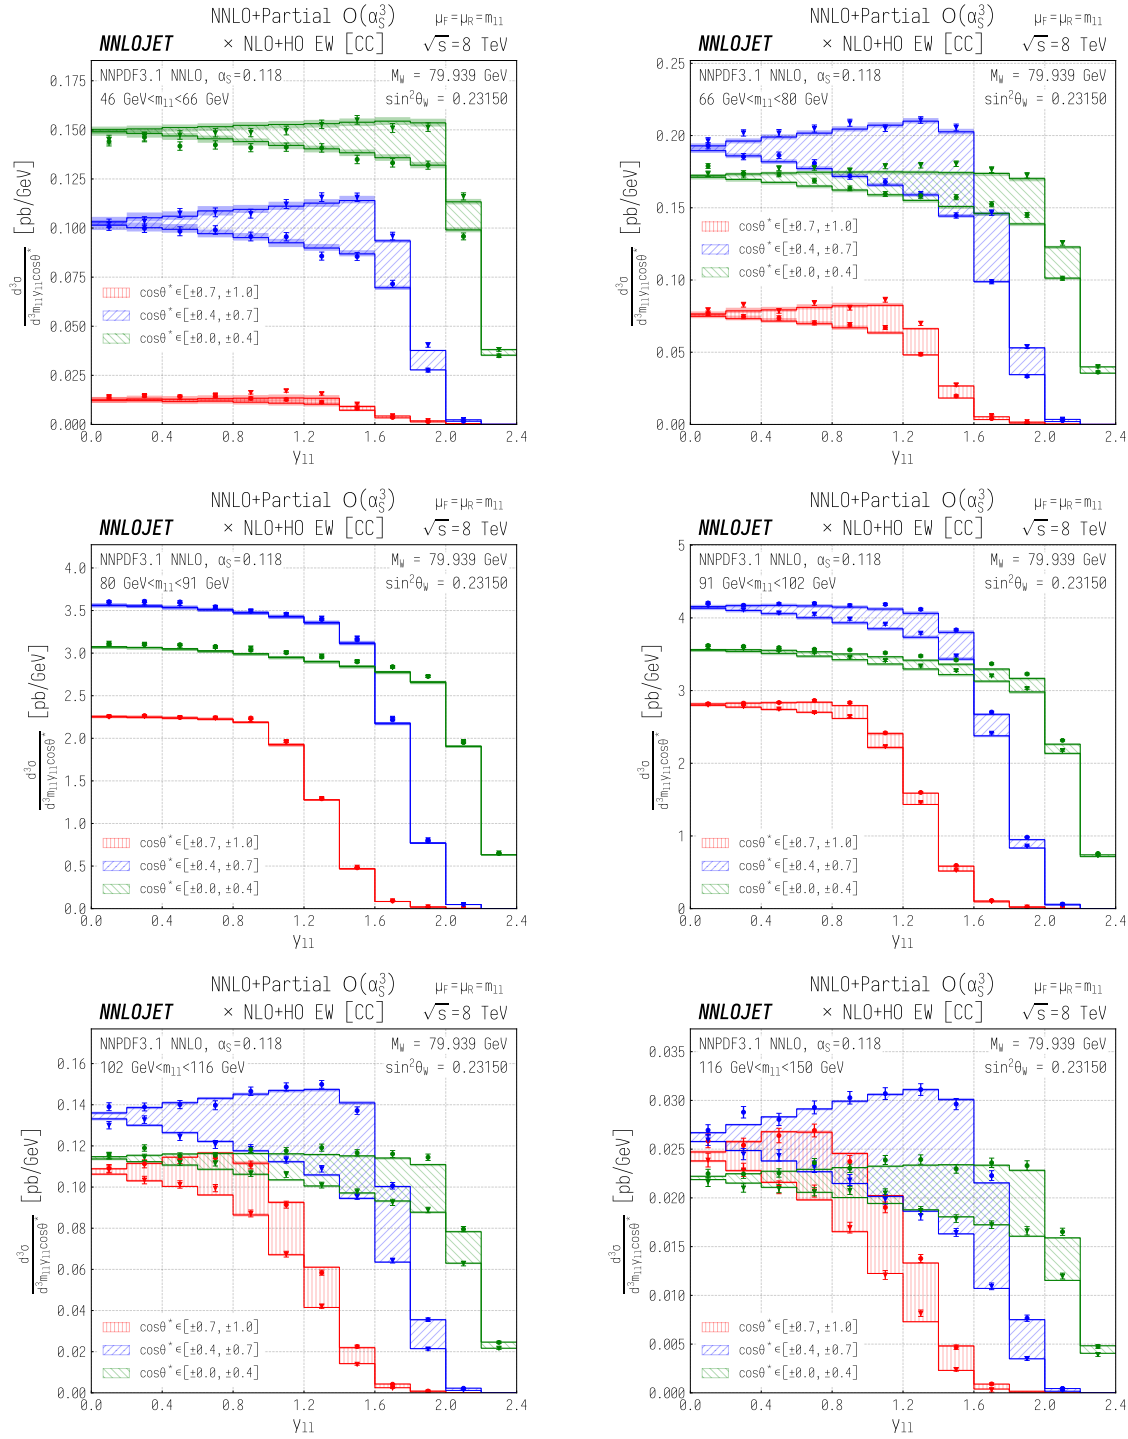


Figure 5.24: (NNLO + partial $\mathcal{O}(\alpha_s^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-central region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos\theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass m_{ll} .




Figure 5.25: (NNLO + partial $\mathcal{O}(\alpha_s^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-central region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos \theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass $m_{\ell\ell}$.

5.10 Summary

In this chapter we have given an overview of the extraction of the effective Weinberg angle $\sin^2 \theta_W^{\text{eff}}$ using $\sqrt{8}$ TeV triple-differential Drell-Yan data from the ATLAS collaboration. Considering the kinematics of this process, we find a rich structure which directly informs the higher-order QCD corrections in terms of acceptances and k -factors and allows the extension of the theory input to $\mathcal{O}(\alpha_s^3)$ in certain regions of phase space.

We have also demonstrated the discriminating power of the triple-differential data in terms of PDFs, potentially allowing a substantial reduction in PDF uncertainties on $\sin^2 \theta_W^{\text{eff}}$ fits using data from hadron colliders. Finally we discussed EW scheme considerations when producing combined QCD+EW results, and showed a comparison of NNLOJET results in combination with higher-order EW corrections from Powheg to the ATLAS data. These constitute a selection of the theoretical inputs to a future ATLAS $\sin^2 \theta_W^{\text{eff}}$ fit.

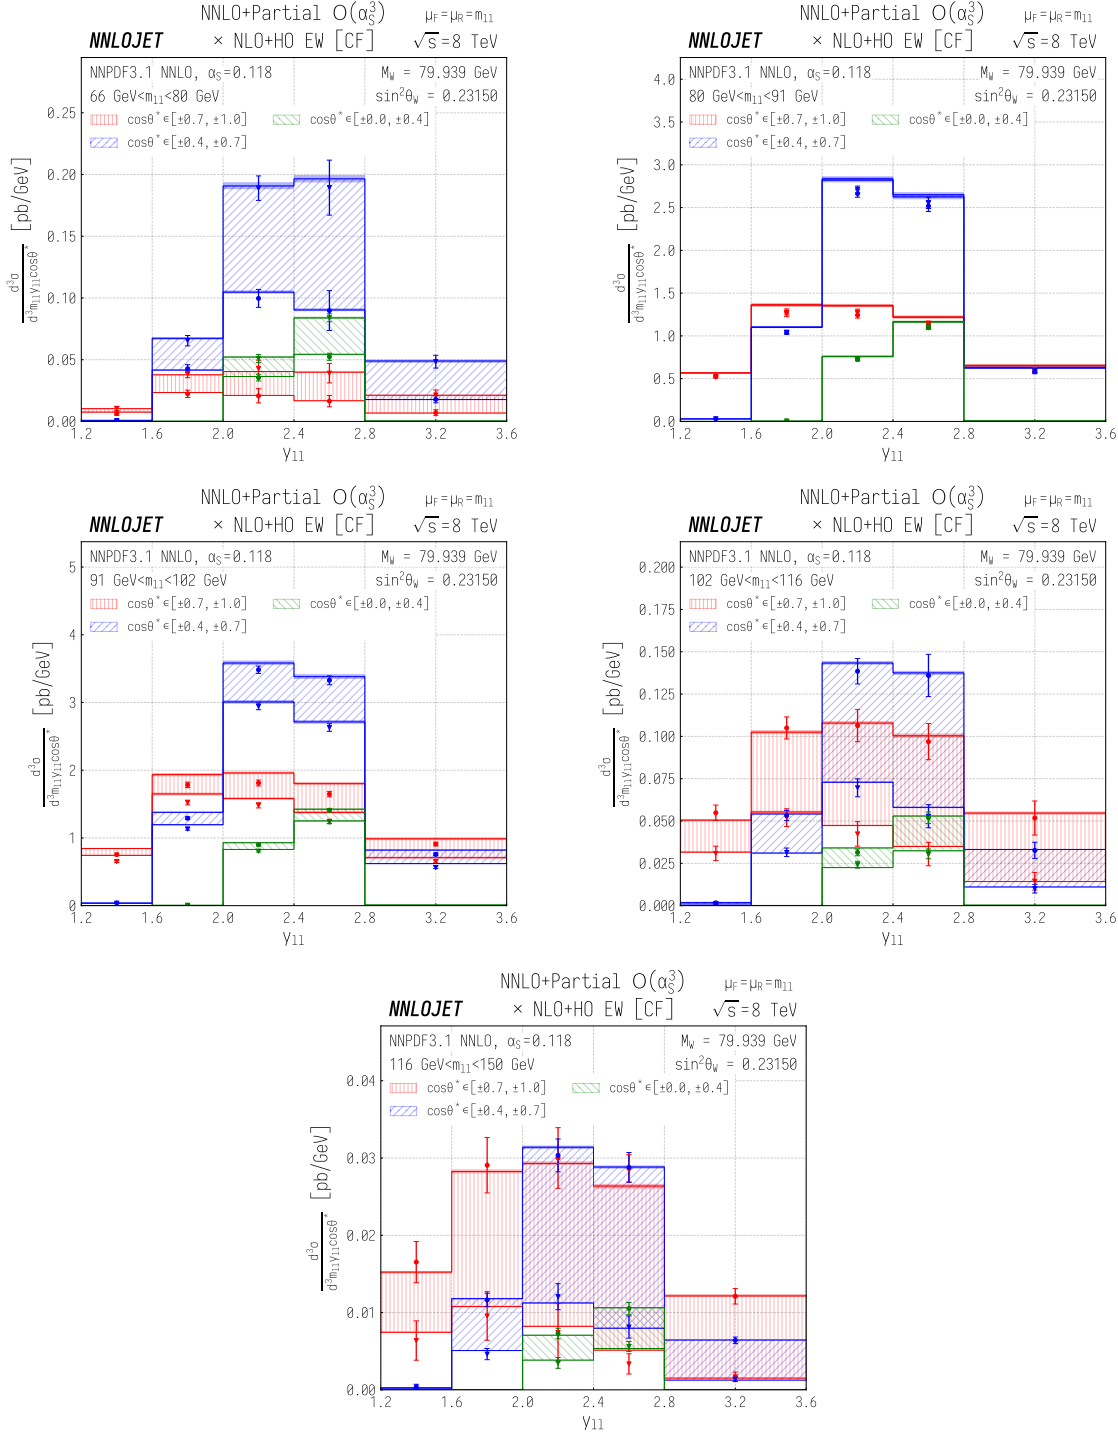


Figure 5.26: (NNLO + partial $\mathcal{O}(\alpha_s^3)$ QCD) \times (NLO+HO EW) predictions for cross sections in the central-forward region of the Z3D analysis in the G_μ scheme with $M_W = 79.939$ GeV, corresponding to $\sin^2 \theta_W = 0.23150$. The solid lines correspond to the theory predictions, about which the shaded band corresponds to the scale uncertainty. The markers correspond to the ATLAS results and associated uncertainty, and the hatched region gives the asymmetry for each of the three regions in $|\cos \theta^*|$. Each panel shows a separate bin in the di-lepton invariant mass m_{ll} .




Figure 5.27: The ratio of NNLO + partial $\mathcal{O}(\alpha_s^3)$ QCD predictions reweighted with NLO + partial higher order EW k -factors to ATLAS data in the central-central (upper) and central-forward (lower) region of the Z3D analysis. The light error bar on the theory predictions corresponds to the scale variation uncertainty and the dark error bar corresponds to the statistical uncertainty, whilst the grey shaded region shows the ATLAS experimental uncertainty. The bin number is as defined in (5.5.2), such that the major m_{ll} bins are divided into 6 $\cos \theta^*$ sub-bins from -1 to 1 (left to right) which are in turn divided into 12 y_{ll} sub-bins from 0 to 2.4 (left to right) for the central-central case. The CF binning is analogously defined. Luminosity uncertainties of $\sim 1.8\%$ are not included.

Chapter 6

Jet Production in

Charged-Current Deep Inelastic

Scattering

Jet production in charged-current (CC) deep inelastic scattering (DIS) provides an important testing ground for both the strong and electroweak sectors of the Standard Model. Inclusive single jet CC DIS enables structure-function measurements at high Bjorken- $x \gtrsim 0.01$ [274–276] as well as precision tests of the chiral structure of the Standard Model in the case that the incoming leptons are polarised, due to the linear dependence of the cross section on the polarisation fraction. Furthermore, one can measure the weak mixing angle $\sin^2 \theta_W$ alongside the W boson mass M_W [277, 278] to complement efforts in hadron-hadron and lepton-lepton collisions.

Requiring the production of two or more jets provides sensitivity to the value of α_s at leading order (LO) in QCD, and allows one to start to probe jet observables in more detail. At the HERA collider, CC events have been observed with final states containing up to four jets, and fully differential results have been presented for production of up to three jets by the ZEUS collaboration [279] and di-jet production by the H1 collaboration [280]. These high-multiplicity events, currently only measured with large statistical uncertainties, become only more relevant at larger

centre-of-mass energies such as at the proposed LHeC [281] due to the improvements in luminosity and kinematic reach.

Both single and multi jet production are crucial processes for our understanding of flavour content in parton distribution functions (PDFs), due to the preferential couplings of the W bosons to quarks dependent on their charge. At leading order, single jet inclusive production is characterised by the basic scattering process $W^\pm q \rightarrow q'$, whereas for di-jet production at LO both initial state gluons and quarks are present for the first time through the production channels $W^\pm g \rightarrow q\bar{q}'$ and $W^\pm q \rightarrow gq'$. As the W^+ (W^-) bosons couple separately to the down(up)-quarks inside the proton, these processes provide strong constraints on the valence quark flavour content of parton distribution functions (PDFs) at high Bjorken- $x \gtrsim 0.01$.

Whilst the most precise measurements of CC DIS have been made at HERA through lepton initiated DIS (lepton-hadron scattering), it is also possible for CC DIS to proceed through neutrino-hadron scattering. While generally taking place at lower energies than at leptonic colliders, neutrino initiated DIS experiments allow complementary measurements to leptonic DIS in different kinematic regimes, useful not only for structure function determinations, measurements of α_s [282–285] and PDF determinations, but also in understanding of e.g. backgrounds for neutrino oscillation experiments [286]. From the theory side, neutrino DIS calculations are related to leptonic DIS by crossing symmetry, such that the translation between the two is essentially trivial.

The exclusive next-to-leading order (NLO) QCD contributions to single and di-jet production in CC DIS have been known for some time [287], and the inclusive CC structure functions have more recently been calculated to next-to-next-to-next-to-leading order (N^3 LO) in QCD [288–294]. These inclusive results give renormalisation and factorisation scale uncertainties smaller than the (statistically dominated) experimental error for the majority of H1 and ZEUS measurements at HERA [279,280]. However, for the LHeC, with an anticipated luminosity a thousand times larger than at the HERA experiment [281], more precise predictions would be required to become

competitive with statistical experimental uncertainties. A centre-of-mass design energy of $\sqrt{s} \approx 1.5$ TeV would also allow such an experiment to examine the content of the proton at a larger range of values of Bjorken- x and gauge boson virtuality Q^2 than previously possible at HERA. To be able to fully exploit the statistical precision that would be possible, the calculation of jet production in CC DIS to higher orders in QCD is essential.

CC DIS is also relevant for current state-of-the-art LHC predictions, in particular the production of Higgs bosons through vector boson fusion (VBF-Higgs). In the structure-function approximation [295] the latter can be described well as “double-DIS”, where each leg is constructed from independent DIS structure functions, with non-factorisable colour exchanges strongly suppressed by both colour factors and kinematics. This provides a strong motivation for improved NC (through Z boson exchange) and CC DIS predictions as they constitute the separate legs of the VBF-Higgs calculation. The N³LO inclusive cross section for single Higgs-boson [296] and double Higgs-boson [297] production were calculated recently using this structure-function approximation.

However, in order for theory predictions to be directly comparable to experimental data, they must be able to account for arbitrary (infrared safe) cuts on the final states produced, a requirement which also allows predictions of multiple-differential exclusive cross sections. This is in contrast to inclusive calculations, which yield results for the full phase space by using analytical techniques to integrate out final-state information; their comparison to experiment then requires ad-hoc extrapolations of data from the measured fiducial regions to the full phase space. As a result, exclusive fixed-order calculations have become the benchmark for comparisons of experiment with theory as no such approximations are required.

We begin this chapter with an introduction to DIS kinematics in Section 6.1 before presenting in Section 6.2 the first theoretical calculation of di-jet production in CC DIS to NNLO in QCD, presented in [6]. We then give an overview of the projection-to-Born (p2B) subtraction method in Section 6.3, before moving to the exclusive




Figure 6.1: Born level kinematics for a single jet leptonic CC DIS event.

calculation of inclusive jet production in CC DIS to N³LO accuracy in QCD [7] in Section 6.4.

6.1 DIS Kinematics

The kinematics of an inclusive leptonic CC DIS event, shown at LO in Figure 6.1, take the generic form

$$p(P) + \ell(k) \rightarrow \nu(k') + X(p_X), \quad (6.1.1)$$

where p is the incoming proton, ℓ the incoming charged lepton, ν the outgoing neutrino and X a generic hadronic final state, with their corresponding four-momenta in brackets. The process is mediated by a W boson of momentum $q = k - k'$ with $Q^2 = -q^2 > 0$, and can be fully described by the standard DIS variables

$$s = (P + k)^2, \quad x = \frac{Q^2}{2P \cdot q}, \quad y = \frac{P \cdot q}{P \cdot k}. \quad (6.1.2)$$

Here x is the usual Bjorken variable, and y the scattering inelasticity (the fraction of the incoming lepton energy transferred to the proton in the proton rest frame).

In the laboratory frame, the Born level kinematics of a single-jet CC DIS event can be fully reconstructed from the incoming beam energies and outgoing neutrino

momentum, using momentum conservation:

$$p_{\text{in},B} = xP, \quad p_{\text{out},B} = xP + q. \quad (6.1.3)$$

This remains true regardless of the presence of QCD radiation, and as such inclusive DIS structure functions are generally written as functions of x and Q^2 . However, these variables do not provide sensitivity to the jet kinematics of a given event which are present at $\mathcal{O}(\alpha_s)$ and beyond. In order to probe these it is common to study jet variables similar to those usually considered in hadron-hadron collisions, namely rapidities, transverse momenta and invariant mass distributions which in turn require exclusive theory calculations.

6.2 Di-Jet Production in CC DIS at NNLO

In [6], we presented the first calculation of di-jet production in CC DIS to (α_s^3) accuracy. The calculation was performed within the NNLOJET framework using antenna subtraction, using the same matrix elements as in the calculation of WJ production in Section 2.6 (up to an initial-final crossing), and we summarise the results here.

The ZEUS collaboration measured jet distributions in the collision of 920 GeV protons with polarised 27.6 GeV electrons/positrons corresponding to a centre-of-mass energy of $\sqrt{s} = 318.7$ GeV [279]. These measurements were taken as functions of x , Q^2 , leading jet transverse energy E_j^T and leading jet pseudorapidity η_j for inclusive jet production, and Q^2 , transverse energy E_{12}^T , average pseudorapidity η_{12} and invariant mass M_{12} of the two leading jets for di-jet production. In the experimental analysis, the jets are p_T ordered and clustered in the laboratory frame, applying the k_T -clustering algorithm with $R = 1$. The experimental results are presented for both $e^+ - P$ and $e^- - P$ scatterings, and are corrected for polarisation effects to give unpolarised cross sections.

In our calculation, electroweak parameters are defined in the G_μ -scheme, with W boson mass, $M_W = 80.398$ GeV, width $\Gamma_W = 2.1054$ GeV, and Z boson mass




Figure 6.2: Predictions at LO (green right-hatched), NLO (orange left-hatched), and NNLO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_{12} , E_{12}^T and M_{12} distributions for inclusive di-jet production in $e^- - P$ collisions. The bands correspond to scale uncertainties as described in the main text.

$M_Z = 91.876$ GeV, using the NNPDF31_nnlo PDF set with $\alpha_s(M_Z) = 0.118$ [117]. The leading jet is required to have transverse energy $E_1^T > 14$ GeV, and the sub-leading jet $E_2^T > 5$ GeV, and both must lie within the rapidity region $-1 < \eta_j < 2.5$. We also correct the theory distributions for hadronisation and QED radiative effects using the multiplicative factors provided in [279].

A comparison between NNLOJET results and ZEUS data for cross sections differential in η_{12} , E_{12}^T , M_{12} and Q^2 for inclusive di-jet production in unpolarised $e^- - P$ collisions is shown in Fig. 6.2. Corresponding results for unpolarised $e^+ - P$ collisions are shown in Fig. 6.3. For both $e^- - P$ and $e^+ - P$ collisions, theory and data show reasonable agreement, and we observe overlapping NLO and NNLO scale

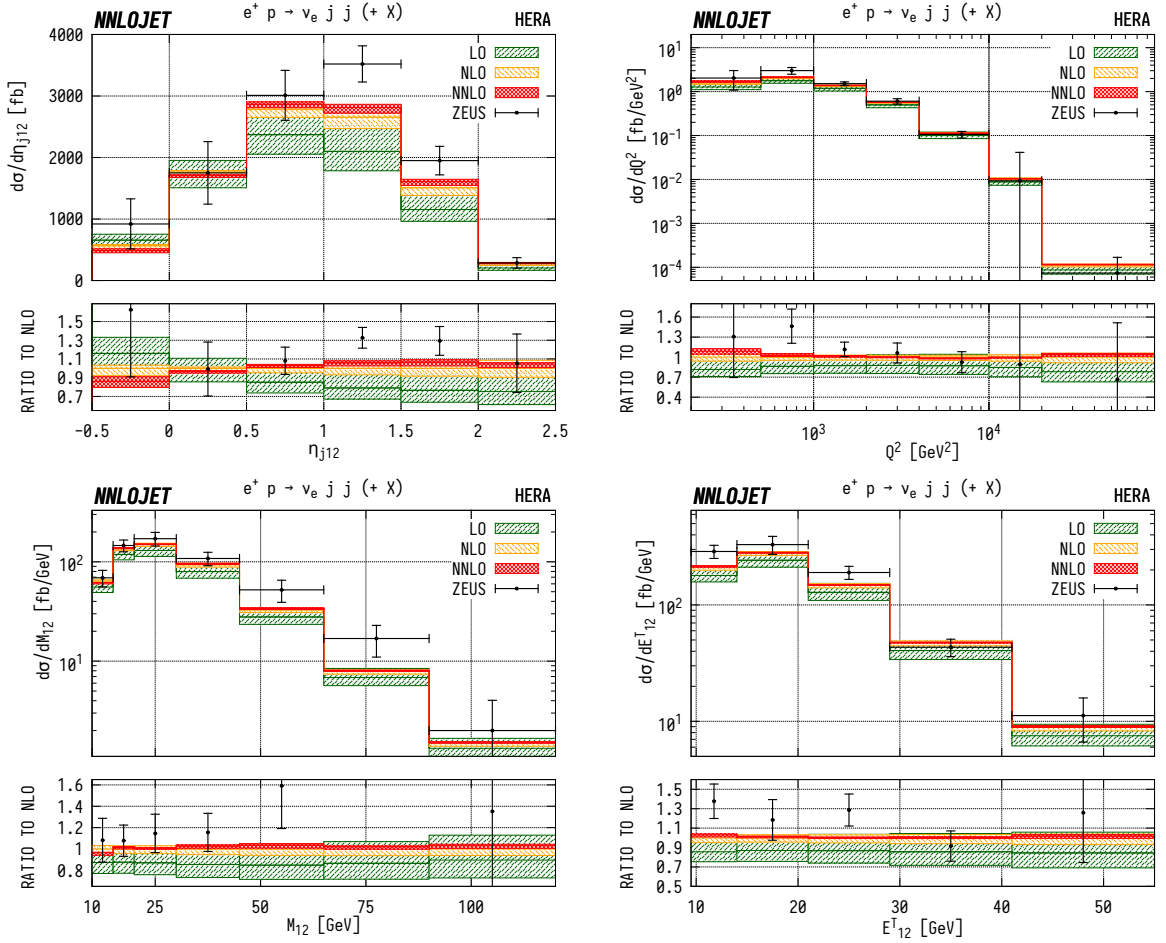


Figure 6.3: Predictions at LO (green right-hatched), NLO (orange left-hatched), and NNLO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_{12} , E_{12}^T and M_{12} distributions for inclusive di-jet production in $e^+ - P$ collisions. The bands correspond to scale uncertainties as described in the main text.

uncertainty bands with a reduction of scale variation uncertainties by typically a factor of two or better from NLO to NNLO. In general, the inclusion of the NNLO corrections do not generally improve the agreement with data with respect to NLO, primarily due to the dominant statistical uncertainty on the experimental data.

For the η_{12} distributions, moderately large and negative higher-order QCD corrections in the lowest bins are observed where NNLO scale variation uncertainties are in some cases larger than at NLO. These uncertainties can be explained by the observation that at NLO, the scale band that lies at the top in the first bin switches to the bottom in the fourth bin and the scale band at the bottom moves up to top at the same time. This turnover of scale bands results in artificially small scale

variation uncertainties at NLO, underestimating the uncertainty from truncation of the perturbative series. This is no longer the case at NNLO, where the scale errors provide a more realistic estimation of the uncertainty and the shape of the NNLO distribution better matches the data than at NLO. We note that the shape corrections at low η_{12} correspond to the low- x region of phase space, in which the centre-of-mass energy of the collisions are smallest and the QCD calculation is least perturbatively stable.

It was reported in [279] for the NLO calculation that the M_{12} distributions beyond $M_{12} \sim 70$ GeV are substantially below the data. We also observe this effect, albeit not to the same extent as in the MEPJET calculation [287], indicating that the reported issue regarding the agreement of MEPJET with other programs at NLO [298] was not the primary source of tension with the data. Using these di-jet results, we are now in a position to combine them with inclusive structure function results to perform exclusive single jet CC DIS calculations to N³LO in QCD.

6.3 Projection to Born Subtraction

The method of projection to Born (p2B) forms a subtraction scheme that exploits the fact that for certain processes, the inclusive kinematics can be fully defined by the non-QCD content in the event at Born level. In DIS, this information is contained within the Bjorken- x of the incoming parton alongside the vector boson momentum q , assuming that the beam energies are known. One is then able to construct a well-defined mapping for any multiplicity of final state partons on to a single set of Born kinematics, as x and q are defined independently of the partonic kinematics:

$$\begin{aligned} \{p\}_N &\xrightarrow{P2B} \{p_{\text{in}}; p_{\text{out}}\} \\ d\mathcal{O} &\xrightarrow{P2B} d\mathcal{O}_B, \end{aligned} \tag{6.3.1}$$

where we define $p_{\text{in}}, p_{\text{out}}$ by (6.1.3). This can then be used to construct an order independent subtraction scheme for process X where such a mapping is present, using as ingredients:

- the inclusive calculation of X at $N^k\text{LO}$,
- the exclusive calculation of X in association with a jet at $N^{k-1}\text{LO}$,

in order to produce fully exclusive predictions at $N^k\text{LO}$. Such processes where p2B could be applied include inclusive Drell-Yan and Higgs production in hadron-hadron collisions, single jet inclusive DIS and di-jet production in lepton-lepton collisions. In practice, it has been performed for VBF-Higgs and di-Higgs production at NNLO in the DIS approximation [49, 50], Higgs decay to b quarks at $N^3\text{LO}$ [299], as well as for photonic DIS at $N^3\text{LO}$ within the NNLOJET framework in [300] at the time of writing.

The master formula for an exclusive production of X (multiply-)differential in observable(s) \mathcal{O} using p2B is given as

$$\frac{d\sigma_X^{N^k\text{LO}}}{d\mathcal{O}} = \frac{d\sigma_{X+j}^{N^{k-1}\text{LO}}}{d\mathcal{O}} - \frac{d\sigma_{X+j}^{N^{k-1}\text{LO}}}{d\mathcal{O}_B} + \frac{d\sigma_X^{N^k\text{LO, incl}}}{d\mathcal{O}_B}, \quad (6.3.2)$$

where the B subscript indicates that the observable has been mapped to and evaluated in the Born level phase space. The first two terms on the right hand side can be evaluated point-by-point in a $N^{k-1}\text{LO}$ calculation of $X+j$ production, such that the term projected into the Born phase space becomes a local counterterm. These counter terms have the same weight as the original event, and generate an exact cancellation in the limit that the QCD radiation in the original event becomes fully unresolved and it lies in the Born phase space. This is in effect using the $X+j$ matrix element/subtraction terms themselves as the counterterm to keep the integrand finite in all singular limits not already subtracted in the $N^{k-1}\text{LO}$ $X+j$ calculation.

In order for a subtraction scheme to be well defined, one must reintroduce the integrated counterterms such that the original integrand is not changed in any way.

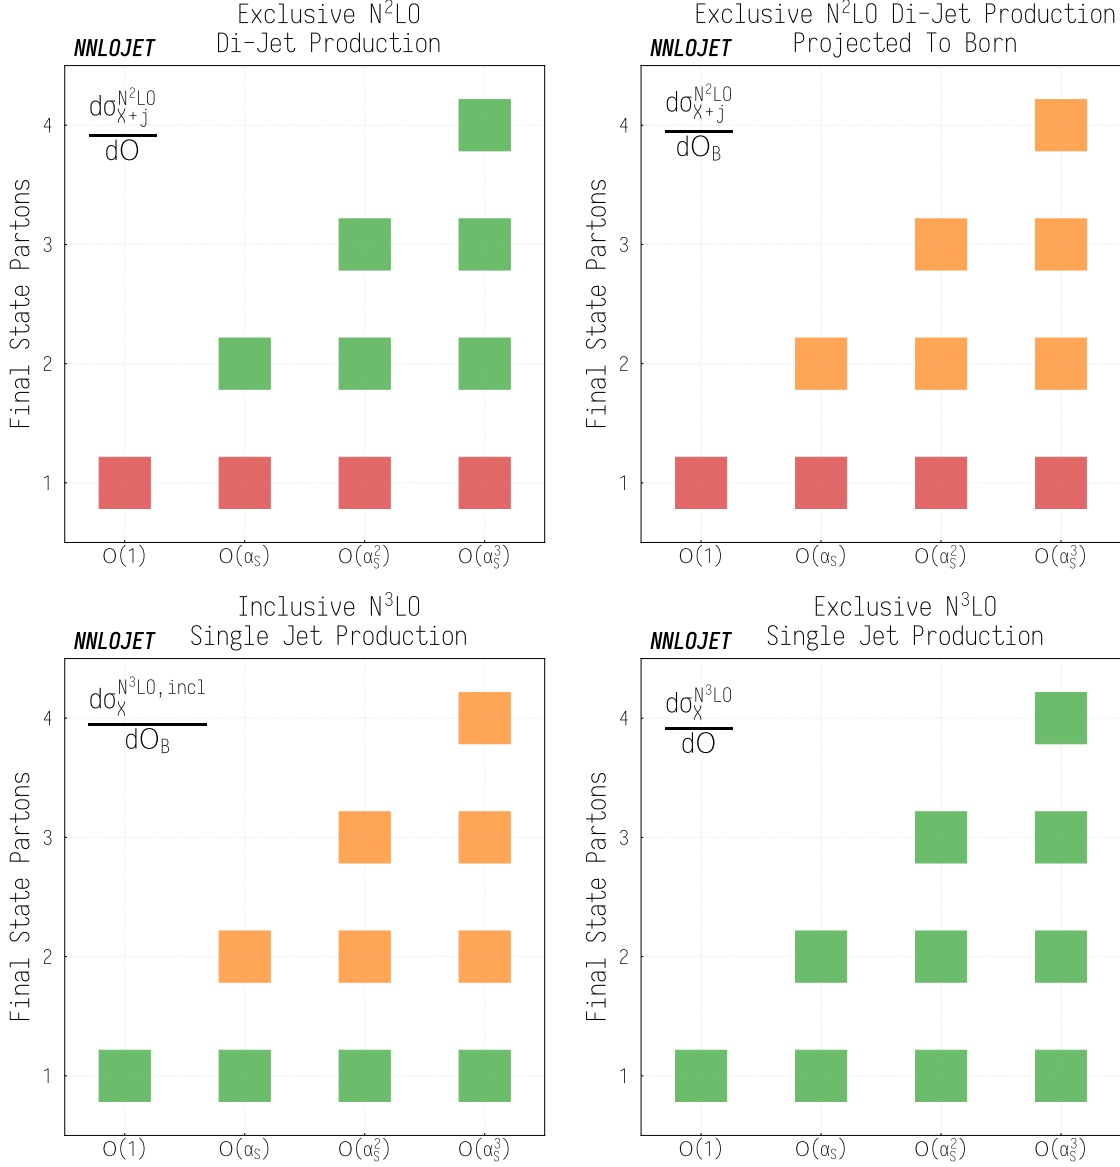


Figure 6.4: Schematic diagrams of the separate contributions to inclusive single jet production in CC DIS at N³LO using the method of p2B. The green points correspond to multiplicities at which the contribution fully describes the event kinematics. The orange points give the integrated out contributions, which act as a counterterm for the di-jet calculation in the divergent phase space regions where it approaches Born kinematics, and are naturally present in the inclusive structure function calculation. The red points correspond to the regions with no access to the associated event kinematics.

Channel	Structure Function	Antenna	Ratio
e^+q	-53.110 ± 0.059	-53.060 ± 0.096	1.00094 ± 0.00213
e^+g	-322.211 ± 0.083	-322.092 ± 0.095	1.00037 ± 0.00039
$e^+\bar{q}$	-87.537 ± 0.044	-87.628 ± 0.074	0.99896 ± 0.00098
Total	-462.842 ± 0.079	-462.784 ± 0.149	1.00013 ± 0.00036

Table 6.1: Comparison of inclusive cross sections calculated independently through antenna subtraction and structure functions for the CC DIS e^+p $\mathcal{O}(\alpha_S^2)$ term, broken down by incoming parton channel.

This is where p2B displays remarkable economy, as the counterterm integrated over phase space is exactly equivalent to the radiative contribution to the inclusive cross section. This radiative part is mapped to the Born phase space during the analytic integration, such that when the three terms in (6.3.2) are combined, the fully exclusive cross section at $N^k\text{LO}$ is recovered. The separate contributions by multiplicity are shown schematically for each of the individual terms in Fig. 6.4.

Writing out the $N^3\text{LO}$ term in the perturbative series explicitly, this gives

$$\begin{aligned}
\frac{d\sigma_X^{N^3\text{LO}}}{d\mathcal{O}} = & \int_{\Phi_{n+3}} \left(\frac{d\sigma_X^{RRR}}{d\mathcal{O}} - \frac{d\sigma_X^S}{d\mathcal{O}} \right) - \int_{\Phi_{n+3}} \left[\left(\frac{d\sigma_X^{RRR}}{d\mathcal{O}} - \frac{d\sigma_X^S}{d\mathcal{O}} \right) \right]_{\mathcal{O} \rightarrow \mathcal{O}_B} \\
& + \int_{\Phi_{n+2}} \left(\frac{d\sigma_X^{RRV}}{d\mathcal{O}} - \frac{d\sigma_X^T}{d\mathcal{O}} \right) - \int_{\Phi_{n+2}} \left[\left(\frac{d\sigma_X^{RRV}}{d\mathcal{O}} - \frac{d\sigma_X^T}{d\mathcal{O}} \right) \right]_{\mathcal{O} \rightarrow \mathcal{O}_B} \\
& + \int_{\Phi_{n+1}} \left(\frac{d\sigma_X^{RVV}}{d\mathcal{O}} - \frac{d\sigma_X^U}{d\mathcal{O}} \right) - \int_{\Phi_{n+1}} \left[\left(\frac{d\sigma_X^{RVV}}{d\mathcal{O}} - \frac{d\sigma_X^U}{d\mathcal{O}} \right) \right]_{\mathcal{O} \rightarrow \mathcal{O}_B} \\
& + \frac{d\sigma_X^{N^k\text{LO, incl}}}{d\mathcal{O}_B}
\end{aligned} \tag{6.3.3}$$

as the differential structure of antenna subtraction calculation used in combination with an inclusive calculation through the method of p2B.

6.3.1 Implementation in the NNLOjet Framework

Following the implementation of di-jet production CC DIS in NNLOJET to NNLO accuracy, the only remaining ingredient required for evaluation of the $N^3\text{LO}$ corrections to single jet production using p2B is the known inclusive $N^3\text{LO}$ calculation term $\frac{d\sigma_X^{N^k\text{LO, incl}}}{d\mathcal{O}_B}$. These results are publicly available in the form of FORTRAN codes for the individual Wilson coefficient functions for fixed scales $\mu_R = \mu_F = Q$, and

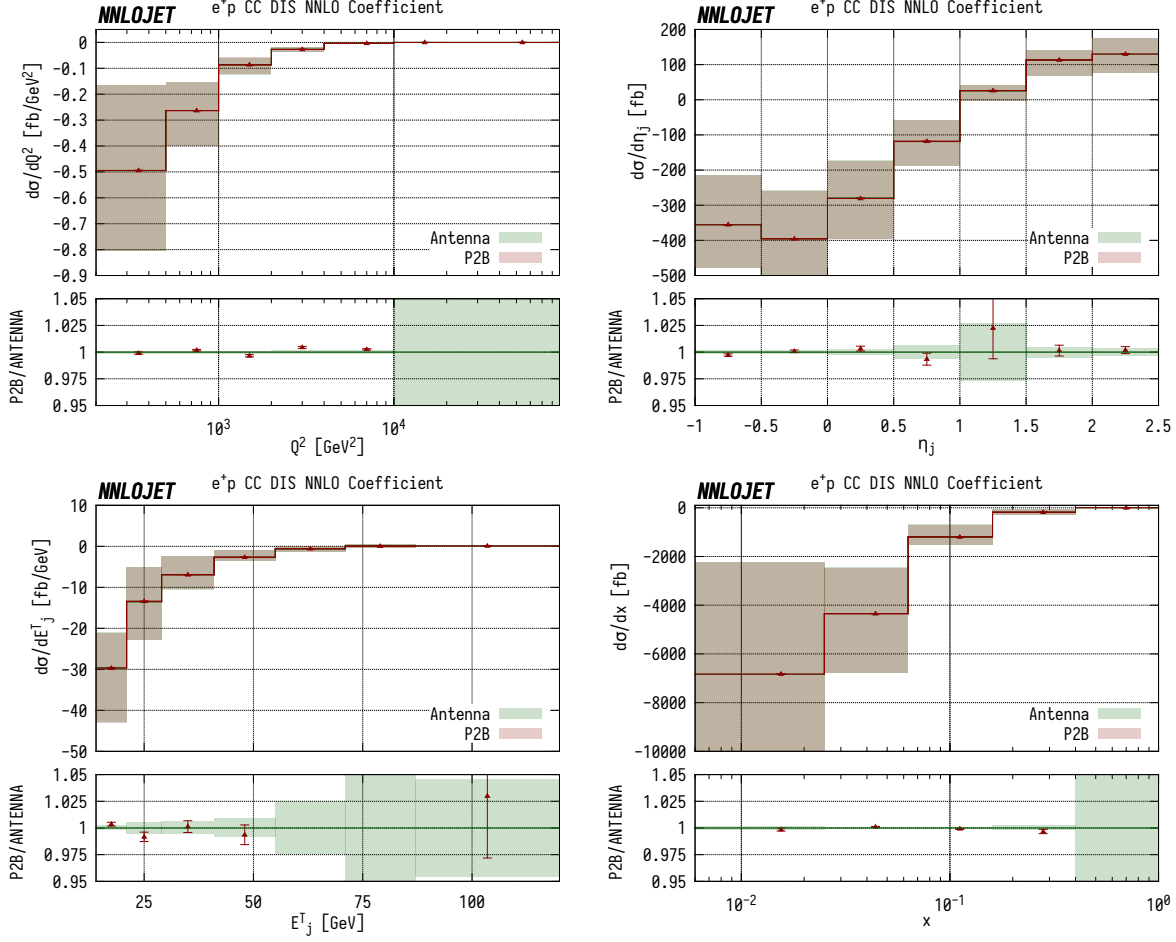


Figure 6.5: Comparison of the $\mathcal{O}(\alpha_s^2)$ term in single-jet inclusive e^+p CC DIS calculated via antenna subtraction and p2B differential in Q^2 , η_j , E_j^T , and x . The shaded regions correspond to the 7 point scale variation band calculated about central scale $\mu_F = \mu_R = Q$, and the vertical error bands correspond to the statistical error. The upper panel of each plot is the absolute correction to the cross section from the $\mathcal{O}(\alpha_s^2)$ terms, and the lower panel is the ratio to the central scale of the antenna calculation.

have been implemented in the NNLOJET framework with the inclusion of full scale dependence. Further details on the form of the inclusive cross section are given in Appendix B, and a breakdown of the $\mathcal{O}(\alpha_s^2)$ term by incoming partonic channel is given in Table 6.1 alongside a comparison with the results from the equivalent calculation performed using antenna subtraction from [6]. For the F_L and F_3 difference structure functions, we use the approximate parameterised forms given in [293], as the full results are not publicly available; these terms contribute less than 1% to the N³LO corrections, and so any difference with the full result is expected to be negligible.

Having validated the inclusive cross section, we can now turn our attention to the validation of the p2B framework in NNLOJET, using again the exclusive results from the equivalent antenna subtraction calculation as a benchmark up to $\mathcal{O}(\alpha_s^2)$. In this case it is instructive to not only compare the standard distributions Q^2 and x , now calculable with fiducial cuts, but also further distributions that were not accessible in the structure function approach beyond LO, in order to ensure that everything is implemented correctly. For these distributions, we use the leading jet transverse energy E_j^T and pseudorapidity η_j , alongside the fiducial cuts applied to single jet inclusive production in [6].

Results, including scale variation bands, for both calculations are shown in Fig. 6.5 where we observe good differential agreement between the two processes when considering the $\mathcal{O}(\alpha_s^2)$ term only. It should be noted that in the E_j^T distribution there are some slight discrepancies in bins where the $\mathcal{O}(\alpha_s^2)$ approaches zero; this occurs due to the instability of Monte-Carlo numerical integration about zero.

The adaptation of the VEGAS algorithm to the radiative contributions in the p2B integrand (lines 1–3 in (6.3.3)) is particularly challenging, as the p2B counterterm has the same weight as the original event if both pass the jet cuts. In this case, a zero-weight event is returned, which consequently means that meaningful adaptation to the total cross section is impossible, even when the event and the projected counter-event have different differential kinematics. Whilst this is most problematic for the case of very inclusive experimental cuts where a large proportion of all events pass the cuts, it in general means that a reweighting of the integrand by e.g. jet rate observables can increase the efficiency of VEGAS adaptation.

6.4 Single Jet Production in CC DIS at N³LO

We are now in a position to give the results of the fully exclusive single-jet inclusive CC DIS to N³LO in QCD, as first presented in [7]. Alongside CC di-jet results, CC single jet inclusive results at the same centre-of-mass energy, $\sqrt{s} = 318.7$ GeV, were

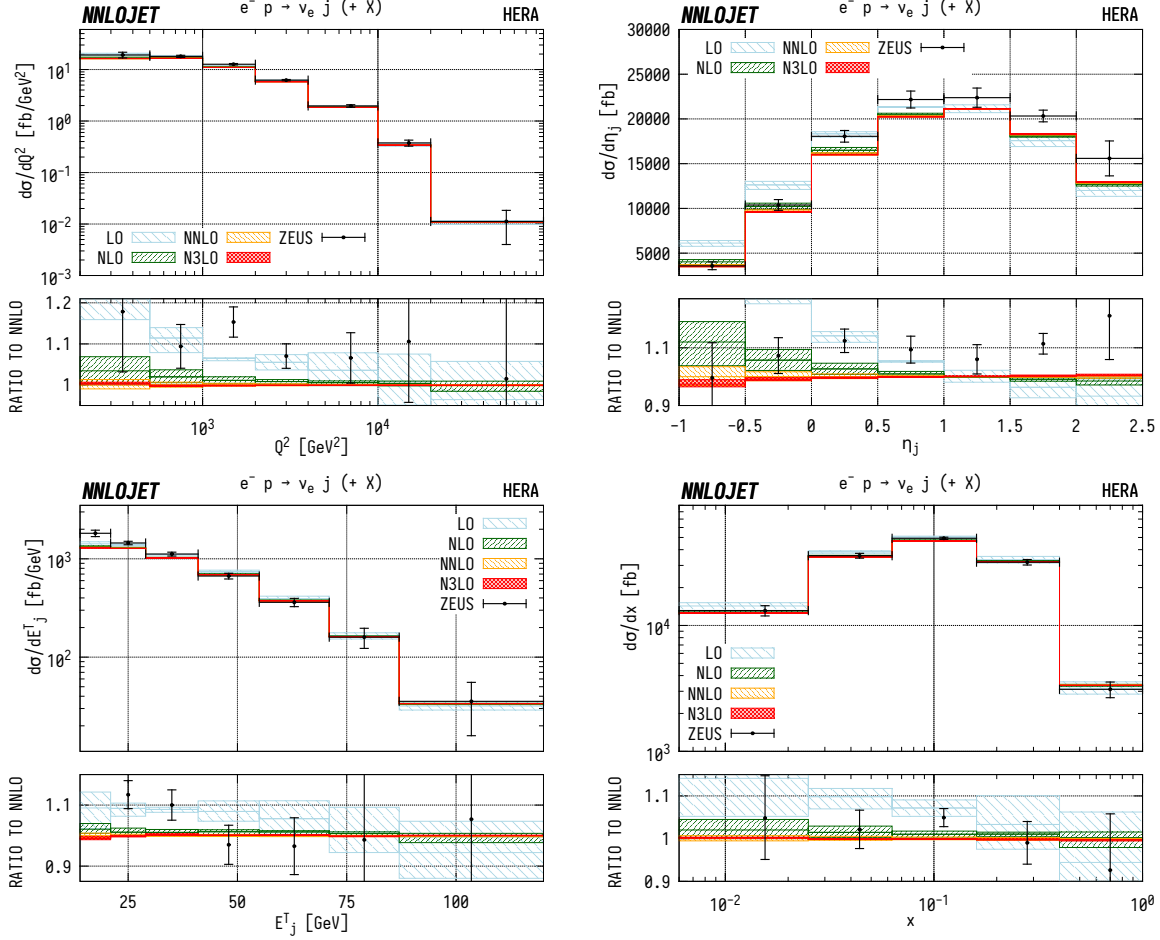


Figure 6.6: Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^-p collisions.

also presented in [279] by the ZEUS collaboration. These measurements were taken for both e^+p and e^-p collisions as functions of x , Q^2 , leading-jet transverse energy E_j^T and pseudorapidity η_j for inclusive jet production, and as before, the jets are E_T ordered and clustered in the laboratory frame using the k_T -clustering algorithm. Data are presented for both e^+p and e^-p collisions, and in the first instance are corrected for polarisation effects to give unpolarised cross sections. The electroweak parameters used for the theory predictions are the same as for di-jet production in Section 6.2, and we use the central renormalisation (μ_R) and factorisation (μ_F) scales $\mu_F^2 = \mu_R^2 = Q^2$. Scale variation uncertainties are estimated by varying μ_R and μ_F independently by factors of 0.5 and 2, but restricted to $0.5 \leq \mu_R/\mu_F \leq 2$.

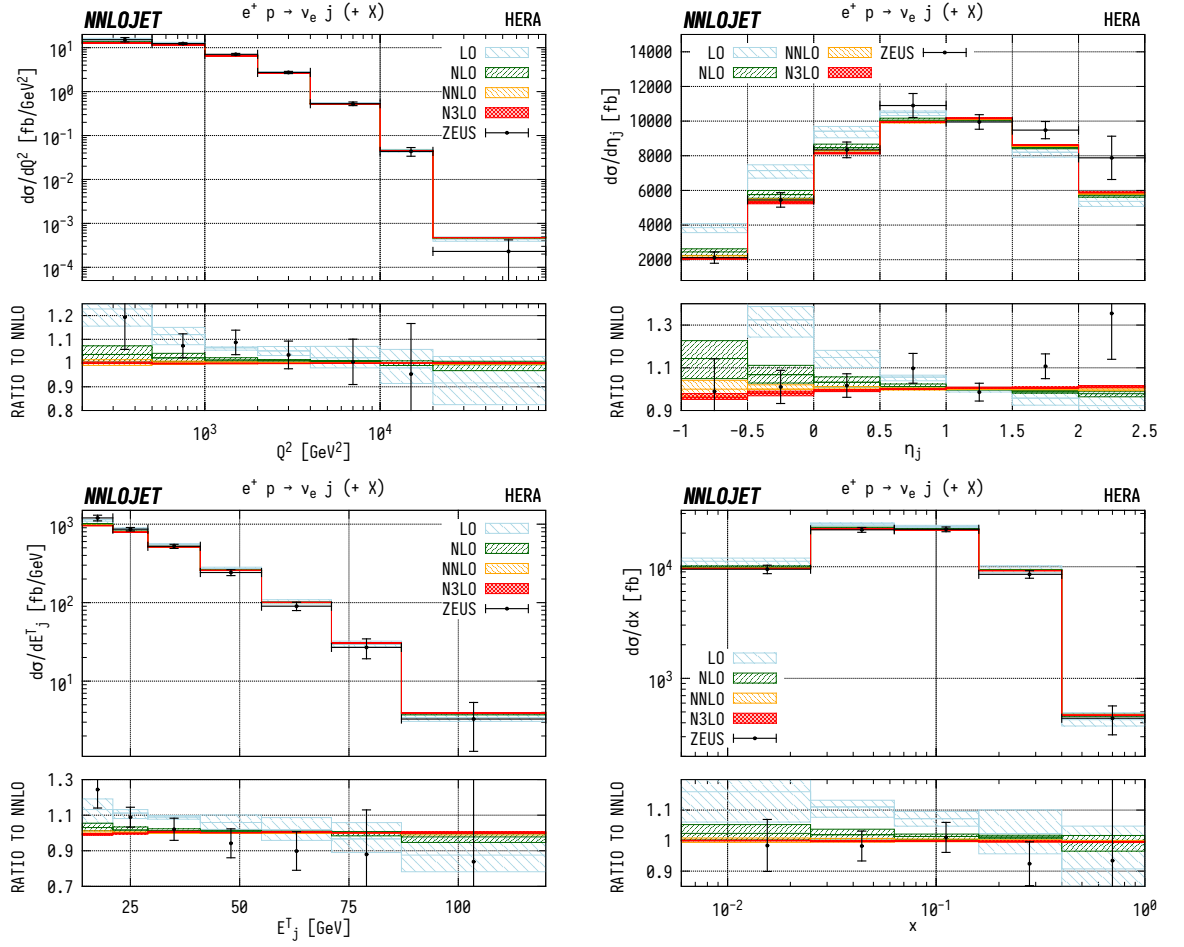


Figure 6.7: Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^+p collisions.

The calculations are performed using the NNPDF31_nnlo PDF set with $\alpha_s(M_Z) = 0.118$ [117] at NNLO with $\alpha_s(M_Z) = 0.118$. It should be noted that the splitting functions for the PDFs are fully known only at NNLO¹ [288,289], so for the purposes of this calculation we have used NNLO PDFs. We do not expect that this will have any impact on the final results due to the small size of the overall correction at N³LO.

As previously, each event must pass the DIS cuts

$$Q^2 > 200 \text{ GeV}^2, \quad y < 0.9, \quad (6.4.1)$$

¹For the status of the N³LO calculations see [301–303]

and the leading jet pseudorapidity must lie in the range $-1 < \eta_j < 2.5$ with minimum transverse energy $E_j^T > 14$ GeV. The theory distributions are corrected for hadronisation and QED radiative effects using the multiplicative factors provided in [279].

A comparison of NNLOJET predictions to ZEUS data for full cross sections differential in Q^2 , η_j , E_j^T and x in single jet inclusive production in unpolarised e^-p collisions is shown in Fig. 6.6, with corresponding results for unpolarised e^+p collisions shown in Fig. 6.7. In general, we find good agreement between theory and data, with overlapping scale uncertainty bands for NNLO and N³LO predictions and a typical reduction in scale variation uncertainties going from NNLO to N³LO by a factor of two or better. Stabilisation of the perturbative QCD prediction can be observed for the first time below $\eta_j = 0$ at this order. In the Q^2 distribution, the convergence of the prediction can now be seen in all bins, with the N³LO predictions contained fully within the NNLO scale variation bands. For low values of x and Q^2 , the predictions for e^-p and e^+p collisions begin to coincide as contributions from sea quarks and gluons inside the proton become dominant and differences between W^+ and W^- exchange diminish. At larger values of x , valence-type quark distributions of the different charges determine the behaviour of the distributions and we see a difference between the two charges.

We can also use the polarised distributions from [279] to perform a qualitative test of the chirality of the Standard Model against our N³LO results. As SM W bosons couple only to left handed leptons, we expect a linear dependence of the cross section on the incoming lepton polarisation fraction P :

$$\sigma_{CC}^{e^\pm p}(P) = (1 + P) \cdot \sigma_{CC}^{e^\pm p}(P = 0). \quad (6.4.2)$$

This can be used to set constraints on possible right-handed W -boson content in BSM models, and has previously been used to exclude the existence of right handed weak bosons with masses below 214 GeV in e^-p scattering [304]. Using (6.4.2), we are able to straightforwardly rescale our unpolarised results in order to directly

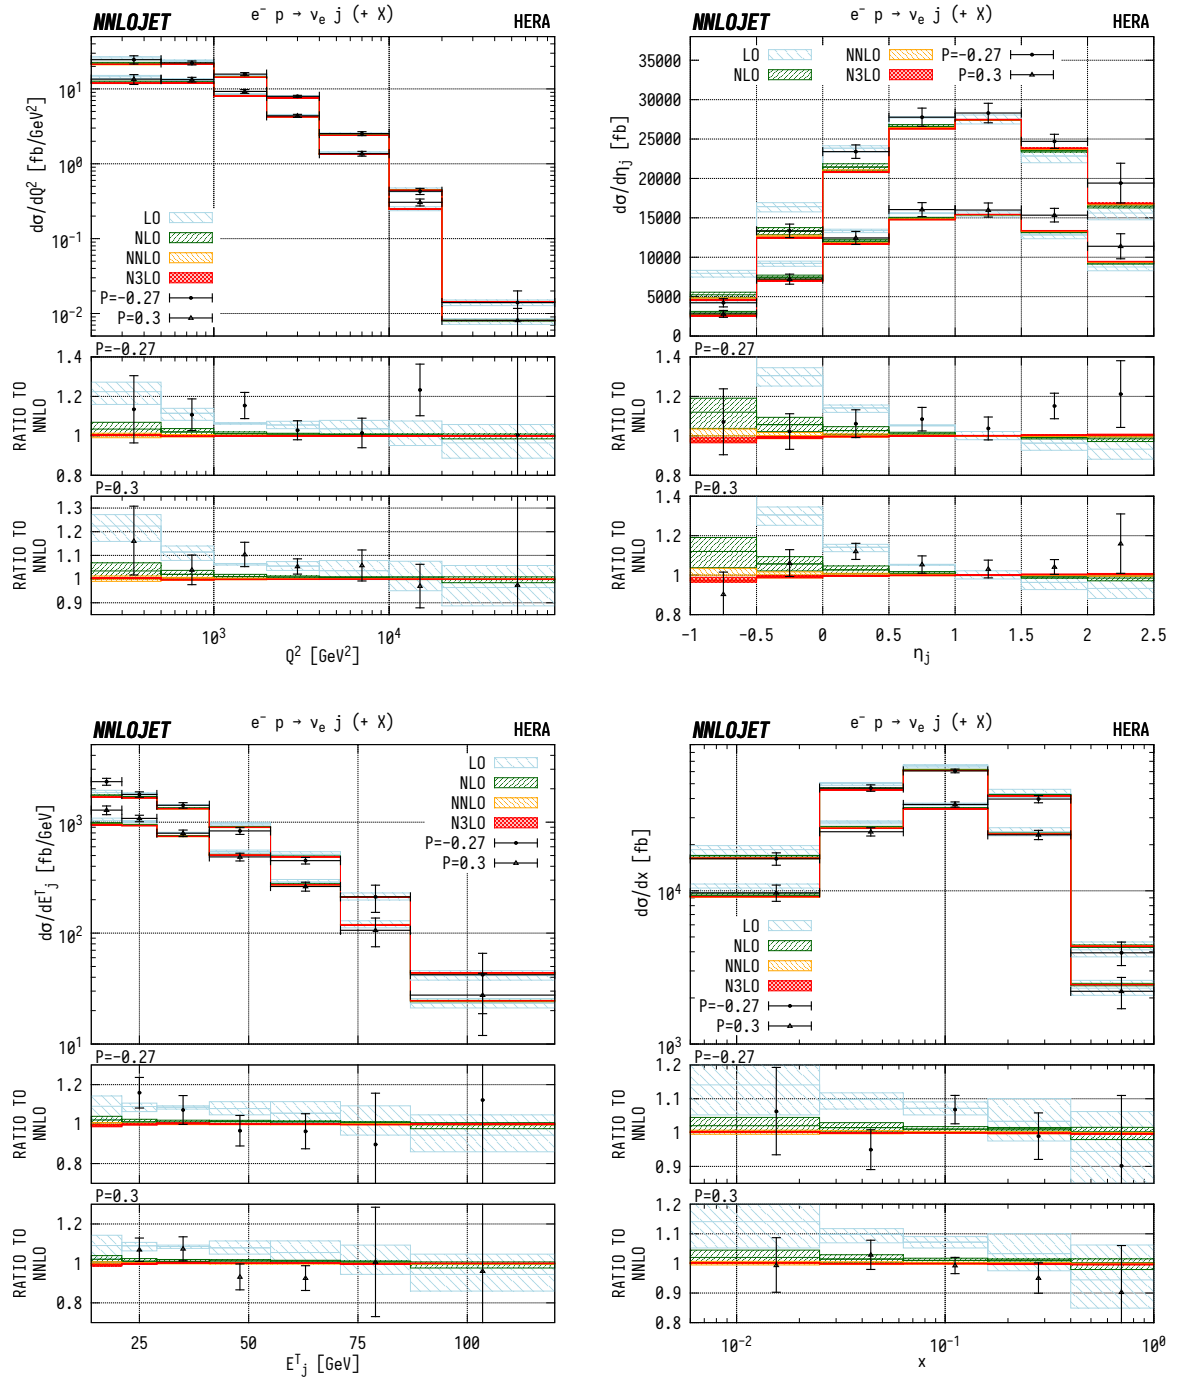


Figure 6.8: Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^-p collisions. The upper panel shows the differential cross sections for $P = -0.27$ and $P = 0.30$, and the lower panels show the individual ratios of data and experiment to the NNLO result.

	Pol.	LO	NLO	NNLO	N ³ LO	ZEUS
e^+p	-0.37	$18.33^{+0.08}_{-0.19}$	$17.06^{+0.22}_{-0.19}$	$16.78^{+0.10}_{-0.07}$	$16.76^{+0.03}_{-0.01}$	$17.55^{+0.83}_{-0.83}$
	$+0.32$	$38.40^{+0.17}_{-0.41}$	$35.75^{+0.46}_{-0.4}$	$35.17^{+0.21}_{-0.15}$	$35.11^{+0.07}_{-0.01}$	$34.51^{+1.29}_{-1.29}$
	0	$29.09^{+0.13}_{-0.31}$	$27.08^{+0.35}_{-0.30}$	$26.64^{+0.16}_{-0.12}$	$26.60^{+0.05}_{-0.01}$	$26.88^{+0.98}_{-0.98}$
e^-p	$+0.29$	$73.09^{+0.44}_{-0.60}$	$68.88^{+0.56}_{-0.62}$	$68.04^{+0.28}_{-0.21}$	$67.96^{+0.08}_{-0.01}$	$70.54^{+1.21}_{-1.20}$
	-0.27	$41.36^{+0.25}_{-0.34}$	$38.98^{+0.32}_{-0.35}$	$38.50^{+0.16}_{-0.12}$	$38.46^{+0.05}_{-0.00}$	$40.53^{+1.02}_{-1.01}$
	0	$56.66^{+0.34}_{-0.46}$	$53.40^{+0.44}_{-0.48}$	$52.75^{+0.22}_{-0.17}$	$52.69^{+0.06}_{-0.00}$	$56.18^{+0.93}_{-0.92}$

Table 6.2: Total cross sections at each perturbative order compared to ZEUS data from Ref. [279] for each polarisations in the experimental data. The uncertainties quoted for the theoretical predictions are from scale variations only, and statistical uncertainties are negligible.

compare to polarised distributions to check the above linear dependence against experiment.

This is performed for two polarisations for each of e^-p ($P = -0.27, 0.3$) and e^+p ($P = -0.37, 0.32$) CC DIS, the differential results of which are shown in Figures 6.8 and 6.9 respectively and summarised for the total cross section in Fig. 6.10². In general, we see similar patterns in the agreement of data to theory as in the inclusive case, with broad agreement across the different polarisations for both the e^+p and e^-p cases. As was the case beforehand, the dominant uncertainty is experimental. From this we can conclude, qualitatively at least, that linear dependence expected in (6.4.2) is not substantially violated. The total cross sections for each polarisation and incoming lepton charge are shown in Table 6.2, where we see good agreement for e^+p initiated CC DIS and slightly worse for e^-p CC DIS. The agreement is relatively consistent across the polarisations, and the discrepancy between the $e^\pm p$ initiated channels in HERA II data is well known, with a considerable impact on PDF set determination [305].

²The experimental data point at $P = 0$ is constructed using linear combinations of the polarised data, and is therefore not independent of the other two data points.

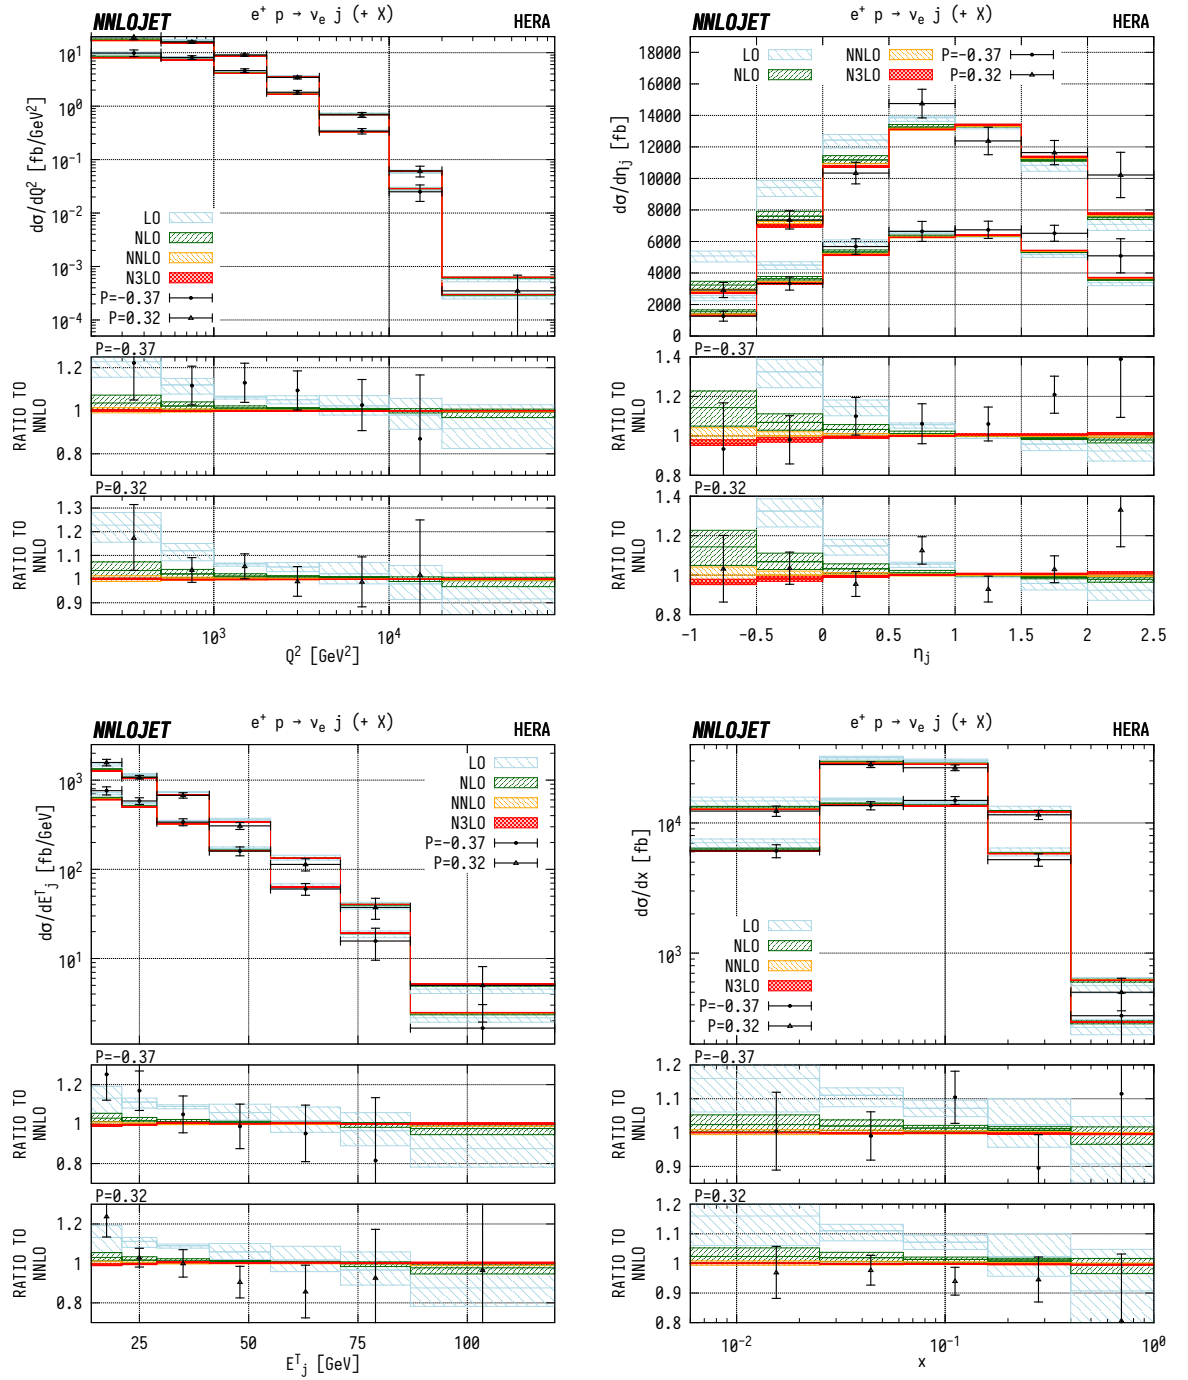


Figure 6.9: Predictions at LO (blue left-hatched), NLO (green right-hatched), NNLO (orange left-hatched) and N³LO (red cross-hatched) are compared to ZEUS data from Ref. [279] for Q^2 , η_j , E_j^T and Bjorken- x for single jet production in e^+p collisions. The upper panel shows the differential cross sections for $P = -0.37$ and $P = 0.32$, and the lower panels show the individual ratios of data and experiment to the NNLO result.

6.5 Summary

In this chapter we have derived the NNLO QCD corrections to di-jet production in charged-current DIS using the method of antenna subtraction. The results of this calculation are compared to experimental data taken by the ZEUS collaboration, where we observe reasonable agreement for both electron and positron scattering.

We then introduce the projection-to-Born subtraction scheme which allows one to combine inclusive calculations with an exclusive differential calculation at one order lower with one additional parton emission in order to generate fully exclusive calculations. This method is then applied to inclusive jet production at N³LO in charged-current DIS. The results are again compared to ZEUS data, both polarised and unpolarised, where we observe a substantial reduction in the theory error and reasonable agreement with data. The results of these calculations are readily applicable to Higgs boson production through vector boson fusion in the DIS approximation, and in principle allow one to derive the first N³LO results for a $2 \rightarrow 3$ process in hadron-hadron collisions.




Figure 6.10: Cross sections as a function of polarisation in single jet leptonic CC DIS at N 3 LO in QCD. Uncertainty bands from scale variation, albeit small, are included as a shaded band around the central prediction. Results are compared to ZEUS data from Ref. [279].

Chapter 7

Conclusions

In this thesis we have considered the application of antenna subtraction to charged current processes in both DIS and hadron-hadron collisions to perform calculations to NNLO accuracy and beyond within the NNLOJET framework. Having completed these calculations, we have also detailed some of the phenomenological implications of the new results, as well as the utilisation of the new-found precision by experiments in order to perform PDF and $\sin^2 \theta_W$ determinations.

In Chapter 3 we presented the first calculations of the W transverse momentum spectrum to $\mathcal{O}(\alpha_s^3)$ alongside the first calculation of the W/Z ratio at $\mathcal{O}(\alpha_s^3)$ including N³LL transverse momentum resummation. This will prove invaluable for the upcoming ATLAS determination of the W boson mass. It also provides the radiative contributions for inclusive W boson production at N³LO, which if using Q_T subtraction as has been performed in the inclusive Higgs case, requires only the calculation of the Q_T beam function to N³LO in combination with the work presented here.

In Chapter 4 we detailed the phenomenological consequences of NNLO vector boson plus jet calculations, in particular the impact on PDF determinations from the enhanced precision in forward regions of phase space. Alongside this, we also presented more esoteric distributions that can be used to probe the strengths and weaknesses of fixed-order calculations.

In Chapter 5 we performed a phenomenological study of the triple differential Drell-Yan analysis at $\sqrt{s} = 8$ TeV by the ATLAS collaboration [260], and detailed some of the results that will be used in an upcoming $\sin^2 \theta_W$ determination as well as a PDF fit primarily using ATLAS data. In particular, this saw an in-depth discussion of the kinematics of the process, with the implication that the effective N³LO contributions could be produced and used in certain areas of phase space using the ZJ calculation in the NNLOJET framework of [87–91].

We have also presented in Chapter 6 the first calculations of di-jet production in charged current leptonic DIS to NNLO accuracy using antenna subtraction and single jet inclusive production to N³LO accuracy using a combined antenna subtraction and projection to Born approach. Alongside providing an important and substantial reduction in theory uncertainty for current and possible future DIS experiments, they encompass the subtraction structure and inclusive structure functions required to perform an exclusive calculation of Higgs production in vector boson fusion to N³LO accuracy in the double-DIS approximation.

It is anticipated that the NNLOJET collaboration will move from strength to strength using both the results of this thesis and others, with a large number of highly important precision predictions possible in the near future. Beyond QCD corrections, the addition of higher order contributions in EW theory would allow for a substantial increase in precision, and could be enabled by the publicly available programs that allow automated evaluation of EW matrix elements. This would provide an opportunity to apply the antenna subtraction method to the Abelian QED sector which would be interesting in its own right.

The advantage of such a flexible and automated framework alongside the substantial testing infrastructure is that development becomes considerably simplified, allowing a wide variety of calculations to be performed in a straightforward manner. At this point, the main challenge of producing such predictions is the considerable computational cost caused primarily by the inability of the current generation of numerical integration techniques to handle complex multidimensional integrands.

Appendix A

Validation and Testing Procedures in NNLOJET

In order to ensure correct results when performing calculations such as those presented in this thesis, it is crucial to test all machinery thoroughly both at the time of implementation and on an ongoing basis. To this end, an extensive suite of tests has been implemented within NNLOJET to check for the correctness and self-consistency of results, which has proven invaluable for both development and the long term stability of the code. These tests can be divided into four relatively broad categories: matrix element tests, subtraction methodology tests, integrated cross section tests and regression tests.

The first three categories are primarily to validate that a given calculation is correct, whilst regression testing is aimed at maintaining the stability and integrity of the codebase in the longer term. All tests are written under the principle that they can be easily used for all processes contained within NNLOJET (where applicable) in order to both ensure test accuracy and minimisation of effort. In this appendix, we consider each of these four classes of test in turn, considering matrix element testing in Section A.1, subtraction term validation in Section A.2, integrated cross section tests in Section A.3 and finally regression testing in Section A.4.

A.1 Point-wise Matrix Element Testing

The first set of tests we consider are those concerning the correctness of the real and 1-loop matrix elements (MEs). These MEs contain the majority of the process-specific source code that cannot currently be auto-generated through associated MAPLE and FORM scripting, and so any initial implementation is very likely to introduce typographical errors. There are a number of external libraries, such as MadGraph [306] and OpenLoops [307], which provide generic implementations of MEs calculated procedurally and are thus largely free from these issues, however a typical ME evaluation using these tools is generally too slow for widespread use in a performant NNLO code. Despite this, they can be used for validation of any independent NNLOJET ME implementation by comparing results for a range of selected phase space points and ensuring that they agree to numerical precision.

This is typically done in two complementary ways depending on the process being tested and the availability of the ME in the external packages. The first is through separate programs which are compiled against both the NNLOJET MEs and results from the external tool, and then check for agreement across a set of phase space points. Care needs to be given to any difference in conventions and input parameters between the programs, particularly for loop matrix elements where there can arise differences in the pole structure from ϵ -dependent global factors.

The second is through an OpenLoops interface which has been recently implemented in NNLOJET, and allows direct validation without the construction of a separate program. The advantage of this is to remove uncertainty associated with the infrastructure surrounding the matrix element, such as final state symmetry factors. However, any final state symmetrisations applied to MEs for convenience in the antenna subtraction formalism need to be removed to obtain direct agreement so certain source code alterations are required. Both point testing methods allow for partonic and flavour breakdowns of the matrix elements, which gives fine-grained resolution to help locate possible issues in the NNLOJET implementation.

Unfortunately at the time of writing there are no equivalent tools to MadGraph and OpenLoops providing MEs at the two-loop level for validation of the double virtual matrix elements present in NNLOJET, although procedures such as numerical unitarity may allow this in future. However, in general we use adapted versions of the original results provided by the authors in order to minimise the chance that errors occur.

A.2 Subtraction Term Tests

Once the MEs have been implemented and validated, the next step is to test the cancellation of infrared (IR) divergences between the ME and the associated subtraction counterterms. There are multiple such tests which operate both numerically and analytically dependent on the number of loops and radiated partons in the ME.

A.2.1 Unresolved Limits Of Real Matrix Elements

The first set of tests apply to all real emission subtraction terms in soft and collinear limits where the matrix element diverges. In these regions of phase space, the subtraction terms should approach the value of the matrix element to cancel the singularity and thus regulate the integrand. In order to test this behaviour numerically for a given divergence, one can consider the ratio

$$R(x) = \frac{\hat{\sigma}^{ME}(x)}{\hat{\sigma}^{SUB}(x)} \quad (\text{A.2.1})$$

where the parameter x regulates the proximity to the divergence at $x = x_0$ in a controlled manner¹.

One expects that as $x \rightarrow x_0$, $R(x) \rightarrow 1$ such that the matrix element is exactly cancelled by the subtraction term in the divergent limit, with the cancellation im-

¹There is a large degree of freedom in the choice of x , as long as it allows one to smoothly approach the divergent limit at some value of $x = x_0$. In the NNLOJET test suite, x is generally aligned with a Mandelstam invariant $s_{i\dots k}$ which either vanishes (for collinear limits) or approaches the total event energy (for soft limits).




Figure A.1: An example spikeplot for a sub-leading colour $qg \rightarrow gg\bar{q} W^+$ double real matrix element in a limit where one final state gluon goes collinear with the incoming quark and another goes soft. As the x parameter approaches 0, the phase space points are driven further into the divergent region, and the subtraction term more closely matches the divergence of the matrix element.

proving as $x \rightarrow x_0$. By generating sets of phase space points² for several values of x about x_0 and generating a histogram of the resultant R values, this behaviour can be explicitly checked, as is demonstrated for an example soft-collinear divergence in Fig. A.1. For this example, the sub-leading colour $qg \rightarrow gg\bar{q} W^+$ double real matrix element, the limit where one final state gluon goes collinear with the incoming quark and another goes soft is approached as $x \rightarrow 0$. One can see that as x decreases from 10^{-7} to 10^{-9} the distribution of R becomes more sharply peaked about 1, indicating good regulation of the matrix element by the subtraction term in the divergent limit. In order to verify that the entirety of the subtraction term structure is correct, one generates figures such as Fig. A.1 for each possible unresolved limit in each of the R, RR and RV terms in the NNLO calculation.

²In NNLOJET this generation is implemented using a modified version of RAMBO [308].

A.2.2 Pole Cancellation Of Virtual Matrix Elements

Unlike the implicit singularities contained in the real MEs, the virtual MEs contain explicit poles in the dimensional regulator ε , which must also cancel exactly with the pole structure of the associated subtraction terms. This cancellation can be checked directly both analytically during subtraction term generation (in the two-loop case) and numerically at runtime (for both the one- and two-loop cases) within the NNLOJET infrastructure, and ensures that the correct contributions have been included (up to finite flavour changing terms).

Numerical 1- and 2-Loop Poles

The full pole structure has been implemented for all of the integrated antenna functions within NNLOJET, which allows generation of the full subtraction term pole structure for arbitrary phase space points. Coupled with the explicit divergences in the virtual MEs in the codebase, this means that the cancellation of the ε poles between subtraction term and ME can be assessed at each order in ε , broken down by both colour order and incoming partonic channel. At one-loop level, one only observes poles up to order $1/\varepsilon^2$ in the regulator, whereas at two-loop level, poles up to order $1/\varepsilon^4$ are present, both of which provide a stringent check on the construction of the virtual subtraction terms. As the poles in ε are present across the entirety of phase space, the phase space points do not have to be generated in specific limits, meaning that this test can be performed directly in NNLOJET at runtime where it is dynamically controlled by a input flag in the runcard.

Analytic 2-Loop Poles

For the case of the two loop matrix elements/subtraction terms, the pole cancellation in NNLOJET can also be tested analytically when the code for the subtraction terms is generated. This is made possible as the double virtual subtraction terms in NNLOJET are auto-generated using a MAPLE [309] based scripting language via an

intermediate stage written in FORM [310]. When called, the FORM code is passed the pole structure in the form of Catani operators [74] of both the relevant two-loop matrix elements and integrated antenna functions, allowing it to analytically verify pole cancellation for all negative powers of ε when the production code is produced.

A.2.3 Finite Term Consistency Checks

Having asserted that the unresolved divergences and the explicit virtual poles in the MEs are cancelled exactly by the subtraction terms, there are still a number of possibilities for inconsistency within the subtraction structure, arising from two primary locations:

1. mass factorisation counterterms, against which initial-state collinear singularities are cancelled order by order into a scheme dependent redefinition of the bare PDFs,
2. degeneracy in the choice of antenna functions used to construct the subtraction terms.

The former means that there may still be explicit poles that are not cancelled fully in the virtual integration, as they are not included in the checks of pole cancellations in the virtual MEs, and the latter means that there are multiple ways of constructing subtraction terms which contain different finite contributions yet still cancel the singularities in the integrands. Any inconsistencies between the finite contributions in the real and virtual counterterms will lead give a non-cancelling contribution to and hence a non-trivial change in the total cross section, equivalent at NNLO to a violation of the identity

$$\int_{\Phi_{N+2}} d\sigma^S + \int_{\Phi_{N+1}} d\sigma^T + \int_{\Phi_N} d\sigma^U = 0 \quad (\text{A.2.2})$$

which is required for subtraction to hold.

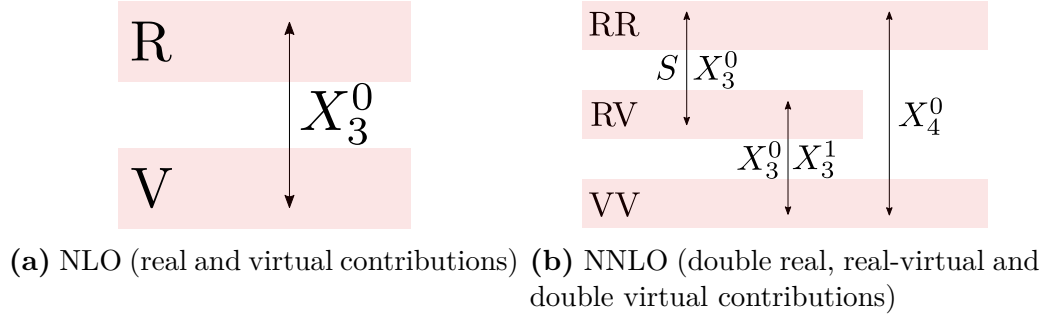


Figure A.2: Cancellations of separate terms between different phase space multiplicities in the antenna subtraction method at NLO and NNLO in QCD.

In order to ascertain that errors are not being introduced through the above, it is important to check directly the consistency of the real and virtual subtraction terms. To this end, a suite of automated, process independent tests have been introduced into the NNLOJET codebase which ensure that for each unintegrated antenna function present in the real contributions, a corresponding integrated antenna is present in the virtual contributions with the appropriate symmetry factors.

This correspondence holds within each colour factor and set of incoming partons, and can be further checked independently for each different class³ of subtraction terms. At NLO only one antenna class is present, $X_3^0|_V^R$, whereas at NNLO there are a total of 5 antenna classes: $X_3^0|_{RV}^{RR}$, $X_3^0|_{VV}^{RV}$, $X_3^1|_{VV}^{RV}$, $X_4^0|_{VV}^{RR}$, and $S|_{RV}^{RR}$. These are shown in Fig. A.2.

The tests themselves are constructed in a manner which exploits the automation of large parts of the NNLOJET codebase. This allows the contents of the subtraction terms to be extracted directly from the MAPLE input files, whilst the colour, flavour and partonic decomposition is given straightforwardly from the driver auto-generation. A set of Unix scripts parse the non-automated driver source code to reconstruct the subtraction term \leftrightarrow ME correspondence along with any further symmetry factors applied for convenience, which can then be used to decompose the appropriate real and virtual subtraction terms by colour, flavour and incoming

³Here we define a subtraction class as containing all integrated and corresponding unintegrated antenna functions of a given type (e.g. X_3^0 and X_3^0) between two phase space multiplicities (e.g. RR and RV), which we label as $X_3^0|_{RV}^{RR}$.

partonic configuration for each class of antenna function. At this point all that remains is to check that an appropriate unintegrated antenna function exists for each integrated antenna function, and therefore that Eqn.A.2.2 is respected. If symmetries of the ME or antenna functions have been exploited, functionality to provide symmetry mappings is also present.

A.3 Integrated Cross Section Tests

Having validated the implementation of both the ME and the subtraction scheme, it still remains to check the behaviour of the results after integration. The following tests assume that the implementation of the phase space generation is sound; this is a relatively generic portion of the calculation that is normally shared between processes with equivalent outgoing mass topologies.

A.3.1 Technical Cut Dependence

In order to prevent numerical precision issues causing a miscancellation between subtraction terms and matrix elements in unresolved limits of phase space, a technical cut is implemented in NNLOJET as

$$s_{\min} < y_0 \cdot \hat{s}, \quad (\text{A.3.1})$$

where \hat{s} is the centre of mass energy, and y_0 is the technical cut parameter, such that any phase space point with an invariant $s \leq s_{\min}$ is discarded. The value of y_0 (typically $\mathcal{O}(10^{-7})$) is chosen such that the cross section is independent of y_0 to within the target statistical uncertainties. If y_0 is chosen to be too large, non-negligible regions of phase space can be cut away, giving a material impact on the total cross section, and if it is too small, large numerical precision issues will be encountered, greatly slowing the convergence of integration. Integrating a given cross section to infinitesimal statistical uncertainty will always eventually resolve the technical cut dependence.

In order to check that this behaves as expected, one can vary y_0 by several orders of magnitude and measure any associated changes of the cross total section. Whilst this is vital to find an appropriate value of y_0 for any physical results produced which give stable and fast results up to an acceptable uncertainty, it also provides an extremely strong secondary check that the unresolved limits of the real matrix elements are finite. If they are not correctly cancelled, a divergent contribution to the integrand will be regulated by y_0 , the variation of which will lead to a strong variation in the cross section.

A.3.2 Renormalisation Scale Variation

For a cross section at a given fixed scale μ_0 and perturbative order which at LO contains terms of $\mathcal{O}(\alpha_s^n)$, one can analytically evaluate the value of the cross section at a second fixed scale using (B.2.6) if each term in the perturbative series

$$\sigma(\alpha_s(\mu_0)) = \sum_{i=0}^{\infty} \left(\frac{\alpha_s(\mu_0)}{2\pi} \right)^{n+i} \sigma^{(i)} \quad (\text{A.3.2})$$

up to the desired order is known separately. Comparing the analytic results of μ_R variation to the numerical output from NNLOJET provides a stringent check on the implementation of the subtraction terms and in certain cases scale dependent terms in MEs/structure functions, which can be decomposed by partonic channel and colour level⁴. As this requires some rescaling of the integrated cross section by values of α_s at known values of μ_R , this test cannot be performed for dynamical scales. An example of the results of this test applied to inclusive single jet production in CC DIS can be seen in Fig. A.3.

A.3.3 3rd Party Validation

The final set of validation checks for the implementation of a calculation in NNLOJET is perhaps the most crucial; the comparison of the fixed-order results to equi-

⁴Decomposition by colour level in some cases requires modification to the β functions in (B.2.2).




Figure A.3: A comparison of the scale variation of the $\mathcal{O}(\alpha_s^2)$ term in single-jet inclusive CC DIS between NNLOJET and the result obtained analytically through running of the renormalisation scale μ_R using the renormalisation group equations. We take $\mu_R = 80\text{GeV}$ as a reference scale, and vary in the range $20 - 400\text{GeV}$.

valent calculations performed external to NNLOJET. This is of course only possible for those cases in which alternative predictions are available so it is not universally applicable, but if agreement is found with an independent calculation one can have good confidence in the results. Even when full results are not available for comparison this can still be useful as fixed-order predictions contain contributions from lower order calculations that can be independently checked against external tools. As NLO QCD predictions are readily available in multiple automated packages, the large parts of an NNLO calculation can be validated in this way even when full NNLO results are not available.

Taking to NNLO the production of arbitrary final state X in proton-proton collisions

$$p + p \rightarrow X \tag{A.3.3}$$

this will contain the independently verifiable sub-processes (labelling jets as j)

$$p + p \rightarrow X; \quad p + p \rightarrow X + j; \quad p + p \rightarrow X + jj \tag{A.3.4}$$

at LO in QCD, and

$$p + p \rightarrow X; \quad p + p \rightarrow X + j \quad (\text{A.3.5})$$

at NLO in QCD. This allows all of the 0- and 1-loop matrix elements and NLO subtraction structure to be checked in detail, and by running with PDF sets with various partonic contributions removed one can also access a partonic breakdown in these tests if needed. It is also crucial in terms of understanding the NNLOJET methodology to compare differentially against results acquired using alternative techniques, including Q_T -slicing, N -jettiness slicing, sector decomposition and inclusive calculations, as this can give one a practical understanding of areas in phase space where methodological issues become prevalent.

It should be noted for completeness that this test is not always useful; any 3rd party calculation can equally contain errors/issues which in practice impact agreement between codes (see e.g. [49, 311]).

A.4 Regression Tests

If at this point the implementation of a process within NNLOJET is found to pass all of the above tests, one can have meaningful confidence in its correctness at that given time. However, for codebases that are changing on a regular basis, particularly with input from multiple people, it is crucial to ensure that no errors are accidentally introduced on an ongoing basis. This has been a widespread problem in software engineering for a long time, to which a standard solution is the implementation of suites of regression tests (see e.g. [312]).

In their simplest form, these are relatively naïve comparisons of program output before and after an update to the codebase in order to isolate any unexpected changes in results, where output is produced for a wide range of input contexts in order to test as much of the codebase as possible. For NNLOJET, such a suite has been introduced which contains a large number of runcards which each individually take




Figure A.4: A summary of the NNLOJET regression tests between December 2017 and April 2019. The upper panel shows the test runtime excluding compilation time, and the lower panel shows the total number of tests alongside the number of passing tests as a function of time. The tests are run on a multi-use desktop containing a 4 core, 8 thread Intel i7-4790 CPU with clock speed of 3.60GHz, meaning that large fluctuations in the test runtimes are generally caused by external (non-regression test) load on the system.

a very short time to run (usually $\mathcal{O}(< 10)$ seconds each), and produce both direct ME/subtraction term evaluations for a number of phase space points and integrated results. It should be stressed that these outputs are not phenomenologically relevant, rather they serve only to monitor any changes that may occur, and any errors which are considered correct by the authors will not be found. Results are emailed out on completion in order to notify users in case of test failures and prompt timely resolution of issues. Regression tests also give a speed-up in certain aspects of software development by flagging issues when they arise and allowing users to make changes with some confidence that unforeseen consequences will be understood.

At the time of writing, a total of 481 tests are implemented in NNLOJET of which

227 are ME/subtraction term tests and 254 are integrated output tests. The tests are designed to be easily extendable, and further tests are anticipated as development continues in the future. Including compilation time, the suite takes $\mathcal{O}(1)$ hour to run on a desktop containing a 4 core, 8 thread Intel i7-4790 CPU with clock speed of 3.60GHz. A summary of test history including the number of tests which pass and fail the regression tests between December 2017 and April 2019 is shown in Fig. A.4.

Appendix B

DIS Structure Functions

The structure functions for CC DIS at N³LO in α_s were presented in [288–294], and have been made publicly available in the form of FORTRAN codes for the individual Wilson coefficient functions for scales $\mu_F = \mu_R = Q$. In this appendix we will briefly cover the form of the inclusive cross section in terms of these structure functions and subsequent Wilson coefficients in Section B.1, before a short discussion in Section B.2 of the implementation of the renormalisation and factorisation scale dependences required for cross section scale variation uncertainties.

B.1 Inclusive Cross Section Decomposition

The CC DIS cross section differential in Bjorken- x and inelasticity y can be expressed in terms of the charged current structure functions F_i .

$$\frac{d^2\sigma^{W^\pm}}{dx dy} = \frac{A\pi\alpha^2}{xyQ^2} \times \left[Y_+ F_2 \mp Y_- x F_3 - y^2 F_L \right], \quad (\text{B.1.1})$$

where we have:

$$A = \left(\frac{G_F M_W^2}{4\pi\alpha} \right)^2 \frac{Q^2}{(Q - M_W^2)^2 + \Gamma_W^2 M_W^2} \quad (\text{B.1.2})$$

$$Y_\pm = 1 \pm (1 - y)^2, \quad (\text{B.1.3})$$

The structure functions F_i can be decomposed into a convolution of PDFs with the Wilson coefficient functions C_i

$$F_i^V = \sum_{a=q,g} C_{i,a} \otimes f_a, \quad (\text{B.1.4})$$

where $i = 2, 3, L$, $V = W^\pm$ and the longitudinal structure function $F_L^i = F_2^i - 2xF_1^i$ through the Callan-Gross identity [313]. The inclusive cross section, differential in x and y is then given by [8]

$$\begin{aligned} \frac{d^2\sigma^{W^\pm}}{dx dy} = & \frac{A\pi\alpha^2}{xyQ^2} \int_0^1 dz \int_0^1 d\bar{y} \delta(x - \bar{y}z) \times \\ & \left[Y_+ \frac{x}{N_F} \sum_{i=1}^{N_F} \left[\mp \delta q_{NS}^-(\bar{y}, \mu_F) C_{2,NS}^-(z, Q, \mu_F, \mu_R) \right. \right. \\ & \quad + q_{PS}(\bar{y}, \mu_F) C_{2,q}(z, Q, \mu_F, \mu_R) \\ & \quad \left. \left. + g(\bar{y}, \mu_F) C_{2,g}(z, Q, \mu_F, \mu_R) \right] \right. \\ & \mp Y_- x \sum_{i=1}^{N_F} \left[\mp \delta q_{NS}^+(\bar{y}, \mu_F) C_{3,NS}^+(z, Q, \mu_F, \mu_R) \right. \\ & \quad \left. + q_{NS}^V(\bar{y}, \mu_F) C_3^V(z, Q, \mu_F, \mu_R) \right] \\ & \left. - \frac{y^2 x}{N_F} \sum_{i=1}^{N_F} \left[\mp \delta q_{NS}^-(\bar{y}, \mu_F) C_{L,NS}^-(z, Q, \mu_F, \mu_R) \right. \right. \\ & \quad + q_{PS}(\bar{y}, \mu_F) C_{L,q}(z, Q, \mu_F, \mu_R) \\ & \quad \left. \left. + g(\bar{y}, \mu_F) C_{L,g}(z, Q, \mu_F, \mu_R) \right] \right], \quad (\text{B.1.5}) \end{aligned}$$

with coefficient function combinations (dropping arguments for convenience)

$$\begin{aligned} C_{L,q} &= C_{L,NS}^+ + C_{L,PS} & \delta C_L &= C_{L,NS}^+ - C_{L,NS}^- \\ C_{2,q} &= C_{2,NS}^+ + C_{2,PS} & \delta C_2 &= C_{2,NS}^+ - C_{2,NS}^- \\ C_3^V &= C_{3,NS}^- + C_{3,s} & \delta C_3 &= C_{3,NS}^+ - C_{3,NS}^-, \end{aligned} \quad (\text{B.1.6})$$

where NS , PS and s denote non-singlet, pure-singlet and sea quark contributions respectively, and corresponding PDF combinations

$$q_{PS} = \sum_{i=1}^{N_F} (q_i + \bar{q}_i) \quad (\text{B.1.7})$$

	$C_{2,q}$	$C_{2,g}$	$C_{2,NS}^+$	$C_{2,NS}^-$	δC_2	$C_{2,PS}$
$\mathcal{O}(1)$	$c_{2,NS}^{(0),EM}$	0	$c_{2,NS}^{(0),EM}$	$c_{2,NS}^{(0),EM}$	0	0
$\mathcal{O}(\alpha_s)$	$c_{2,NS}^{(1),EM}$	$c_{2,g}^{(1),EM}$	$c_{2,NS}^{(1),EM}$	$c_{2,NS}^{(1),EM}$	0	0
$\mathcal{O}(\alpha_s^2)$	$c_{2,NS}^{(2),EM} + c_{2,PS}^{(2),EM}$	$c_{2,g}^{(2),EM}$	$c_{2,NS}^{(2),EM}$	$c_{2,NS}^{(2),EM} - \delta c_2^{(2)}$	$\delta c_2^{(2)}$	$c_{2,PS}^{(2),EM}$
$\mathcal{O}(\alpha_s^3)$	$c_{2,NS}^{(3),EM} + c_{2,PS}^{(3),EM}$	$c_{2,g}^{(3),EM}$	$c_{2,NS}^{(3),EM}$	$c_{2,NS}^{(3),EM} - \delta c_2^{(3)}$	$\delta c_2^{(3)}$	$c_{2,PS}^{(3),EM}$

Table B.1: C_2 coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case.

$$q_{NS}^V = \sum_{i=1}^{N_F} (q_i - \bar{q}_i) \quad (\text{B.1.8})$$

$$\delta q_{NS}^\pm = \sum_{i \in \text{u-type}} (q_i \pm \bar{q}_i) - \sum_{i \in \text{d-type}} (q_i \pm \bar{q}_i), \quad (\text{B.1.9})$$

where q_{NS}^V is the valence quark contribution and δq_{NS}^\pm are the flavour asymmetries. g denotes the gluon PDF. The terms in the outer brackets of (B.1.5) correspond to structure functions F_2 , F_3 and F_L , and the differences of the CC coefficient functions δC_i are defined to always be ‘even - odd’ in the Mellin moments of the operator product expansion, with c_i^\pm corresponding to the even and odd moments in Mellin- N space respectively [314]. These differences can be related to the coefficient functions for linear combinations of coefficient functions $C^{\nu p \pm \bar{\nu} p}$ in neutrino DIS¹ as [292]

$$\delta C_{2,L} \equiv C_{2,L}^{\nu p + \bar{\nu} p} - C_{2,L}^{\nu p - \bar{\nu} p}, \quad \delta C_3 \equiv C_3^{\nu p - \bar{\nu} p} - C_3^{\nu p + \bar{\nu} p}. \quad (\text{B.1.10})$$

At lower orders, the structure functions and Wilson coefficients are closely related to those present in electromagnetic DIS (C_i^{EM} , $i \in [2, L]$), which proceeds via photon exchange and does not contain the parity violating F_3 . Defining

$$C_a = \sum_{n=0}^{\infty} \alpha_s^n c_a^{(n)}, \quad (\text{B.1.11})$$

these relationships for the C_i are tabulated in Tables B.1-B.3 order by order in α_s .

It should be noted that there are certain flavour structures² appearing in C_3^\pm which

¹Note this definition means that $C_{2,L}^\pm = C_{2,L}^{\nu p \pm \bar{\nu} p}$, but $C_3^\pm = C_3^{\nu p \mp \bar{\nu} p}$.

²In the literature, these diagrams are labelled as flavour class fl_{02} [291].

	$C_{L,q}$	$C_{L,g}$	$C_{L,NS}^+$	$C_{L,NS}^-$	δC_L	$C_{L,PS}$
$\mathcal{O}(1)$	0	0	0	0	0	0
$\mathcal{O}(\alpha_s)$	$c_{L,NS}^{(1),EM}$	$c_{L,g}^{(1),EM}$	$c_{L,NS}^{(1),EM}$	$c_{L,NS}^{(1),EM}$	0	0
$\mathcal{O}(\alpha_s^2)$	$c_{L,NS}^{(2),EM} + c_{L,PS}^{(2),EM}$	$c_{L,g}^{(2),EM}$	$c_{L,NS}^{(2),EM}$	$c_{L,NS}^{(2),EM} - \delta c_L^{(2)}$	$\delta c_L^{(2)}$	$c_{L,PS}^{(2),EM}$
$\mathcal{O}(\alpha_s^3)$	$c_{L,NS}^{(3),EM} + c_{L,PS}^{(3),EM}$	$c_{L,g}^{(3),EM}$	$c_{L,NS}^{(3),EM}$	$c_{L,NS}^{(3),EM} - \delta c_L^{(3)}$	$\delta c_L^{(3)}$	$c_{L,PS}^{(3),EM}$

Table B.2: C_L coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case.

	C_3^V	$C_{3,NS}^+$	$C_{3,NS}^-$	δC_3	$C_{3,s}$
$\mathcal{O}(1)$	$c_{2,NS}^{(0),EM}$	$c_{2,NS}^{(0),EM}$	$c_{2,NS}^{(0),EM}$	0	0
$\mathcal{O}(\alpha_s)$	$c_{2,NS}^{(1),EM} - A$	$c_{2,NS}^{(1),EM} - A$	$c_{2,NS}^{(1),EM} - A$	0	0
$\mathcal{O}(\alpha_s^2)$	$c_{3,NS}^{(2),+} - \delta c_3^{(2)}$	$c_{3,NS}^{(2),+}$	$c_{3,NS}^{(2),+} - \delta c_3^{(2)}$	$\delta c_3^{(2)}$	0
$\mathcal{O}(\alpha_s^3)$	$c_{3,NS}^{(3),+} - \delta c_3^{(3)} + c_{3,s}^{(3)}$	$c_{3,NS}^{(3),+}$	$c_{3,NS}^{(3),+} - \delta c_3^{(3)}$	$\delta c_3^{(3)}$	$c_{3,s}^{(3)}$

Table B.3: C_3 coefficients required for CC DIS and their relationship to those in EM DIS at each order in α_s . Those highlighted in green are the new components required for the CC case. Here, $A = 2C_F(1+x)$, with colour factor C_F normalised such that $C_F = (N^2 - 1)/2N$.




Figure B.1: Representative Feynman diagram of the fl_{02} flavour class, which is forbidden for $V = W^\pm$ due to charge conservation in the quark loop.

are proportional to $d_{abc}d^{abc}$. Whilst these terms contribute to the cross section for NC DIS, they do not occur in the CC case as the weak boson couples to closed quark loops which is forbidden through charge conservation. These contributions can be straightforwardly isolated and removed in the FORTRAN code accompanying [293], and an example diagram is shown in Fig. B.1.

B.2 Scale Variation

The structure functions presented in the previous section neglect all dependence on the renormalisation and factorisation scales, implicitly taking $\mu_R = \mu_F = Q$, where Q is the momentum transfer of the boson. The dependence of the cross section on these scales is known analytically [315, 316], and can be reconstructed in general using the appropriate evolution equations. These properties are not limited to the construction of cross sections through structure functions and are directly applicable to all results presented in this thesis when truncated to the appropriate order.

B.2.1 Renormalisation Scale Dependence

In order to reconstruct the renormalisation scale dependence of the cross section in (B.1.5), one can solve the renormalisation group equation for QCD order by order in the coupling constant, and then substitute the results into the expression for the fixed-order cross section. For the strong coupling constant α_s , this renormalisation group equation reads:

$$\mu_R^2 \frac{\alpha_s(\mu_R)}{\mu_R^2} = -\alpha_s(\mu_R) \left[\beta_0 \left(\frac{\alpha_s(\mu_R)}{2\pi} \right) + \beta_1 \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^2 + \beta_2 \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^3 + \mathcal{O}(\alpha_s^4) \right], \quad (\text{B.2.1})$$

with the $\overline{\text{MS}}$ -scheme coefficients

$$\begin{aligned} \beta_0 &= \frac{11C_A - 4T_R N_F}{6}, \\ \beta_1 &= \frac{17C_A^2 - 10C_A T_R N_F - 6C_F T_R N_F}{6}, \\ \beta_2 &= \frac{1}{432} (2857C_A^3 + 108C_F^2 T_R N_F - 1230C_F C_A T_R N_F - 2830C_A^2 T_R N_F \\ &\quad + 264C_F T_R^2 N_F^2 + 316C_A T_R^2 N_F^2). \end{aligned} \quad (\text{B.2.2})$$

One can then express the coupling at fixed scale μ'_R in terms of the coupling at μ_R by introducing the logarithm

$$L_R = \log \frac{\mu_R^2}{\mu'^2_R} \quad (\text{B.2.3})$$

and solving iteratively in order to give

$$\begin{aligned} \alpha_s(\mu'_R) = \alpha_s(\mu_R) & \left\{ 1 + \beta_0 L_R \frac{\alpha_s(\mu_R)}{2\pi} + [\beta_0^2 L_R^2 + \beta_1 L_R] \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^2 \right. \\ & \left. + \left[\beta_0^3 L_R^3 + \frac{5}{2} \beta_0 \beta_1 L_R^2 + \beta_2 L_R \right] \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^3 + \mathcal{O}(\alpha_s^4) \right\}. \end{aligned} \quad (\text{B.2.4})$$

In general, the perturbative expansion of a given cross section starts at order α_s^n , where for the case of single-jet inclusive DIS we have $n = 0$. In this expansion, the renormalisation scale is fixed to some value μ'_R for each of the expansion coefficients $\sigma^{(n)} = \sigma^{(n)}(\mu'_R)$ which to N³LO, gives:

$$\begin{aligned} \sigma(\mu'_R, \alpha_s(\mu'_R)) = & \left(\frac{\alpha_s(\mu'_R)}{2\pi} \right)^n \sigma^{(0)} + \left(\frac{\alpha_s(\mu'_R)}{2\pi} \right)^{n+1} \sigma^{(1)} + \left(\frac{\alpha_s(\mu'_R)}{2\pi} \right)^{n+2} \sigma^{(2)} \\ & + \left(\frac{\alpha_s(\mu'_R)}{2\pi} \right)^{n+3} \sigma^{(3)} + \mathcal{O}(\alpha_s^{n+4}). \end{aligned} \quad (\text{B.2.5})$$

At this point, the renormalisation scale dependence of the cross section to N³LO can then be fully reconstructed through substitution of (B.2.4):

$$\begin{aligned} \sigma(\mu_R, \alpha_s(\mu_R), L_R) &= \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^n \sigma^{(0)} + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+1} \left\{ \sigma^{(1)} + n\beta_0 L_R \sigma^{(0)} \right\} \\ &+ \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+2} \left\{ \sigma^{(2)} + (n+1)\beta_0 L_R \sigma^{(1)} + \left[n\beta_1 L_R + \frac{n(n+1)}{2} \beta_0^2 L_R^2 \right] \sigma^{(0)} \right\} \\ &+ \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+3} \left\{ \sigma^{(3)} + (n+2)\beta_0 L_R \sigma^{(2)} + \left[(n+1)\beta_1 L_R + \frac{(n+1)(n+2)}{2} \beta_0^2 L_R^2 \right] \sigma^{(1)} \right. \\ &\quad \left. + \left[n\beta_2 L_R + \frac{n(6n+9)}{6} \beta_0 \beta_1 L_R^2 + \frac{n(n^2+3n+2)}{6} \beta_0^3 L_R^3 \right] \sigma^{(0)} \right\} \\ &+ \mathcal{O}(\alpha_s^{n+4}). \end{aligned} \quad (\text{B.2.6})$$

Having constructed this dependence, it is then relatively straightforward to implement in the expression for the cross section (B.1.5). This can be performed independently order by order within each Wilson coefficient function, and automated through scripting in the FORM symbolic manipulation language.

B.2.2 Factorisation Scale Dependence

In a similar manner to the renormalisation scale dependence, the factorisation scale dependence can be reconstructed using the DGLAP evolution equations [317–320]:

$$\mu_F^2 \frac{d}{d\mu_F^2} f_i(\mu_F, \mu_R) = \sum_j P_{ij}(\alpha_s(\mu_R), \mu_F, \mu_R) \otimes f_j(\mu_F, \mu_R), \quad (\text{B.2.7})$$

where f_i is the PDF for parton flavour i , and P_{ij} is a matrix of splitting functions in flavour space. To make clear the analogy with the μ_R case, the expression for the splitting functions can be expanded order by order in the strong coupling constant [321] to NNLO³

$$\begin{aligned} P_{ij}(\alpha_s(\mu_R), \mu_F, \mu_R) &= \frac{\alpha_s(\mu_R)}{2\pi} P_{ij}^{(0)} + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^2 \left[P_{ij}^{(1)} + \beta_0 L_F^R P_{ij}^{(0)} \right] \\ &+ \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^3 \left[P_{ij}^{(2)} + \left(\beta_1 P_{ij}^{(0)} + 2\beta_0 P_{ij}^{(1)} \right) L_F^R + \beta_0^2 L_F^{R^2} P_{ij}^{(0)} \right] \\ &+ \mathcal{O}(\alpha_s^4), \end{aligned} \quad (\text{B.2.8})$$

where $L_F^R = \log(\mu_R^2/\mu_F^2)$. Understanding that the PDFs as non-perturbative objects only contain factorisation scale dependence and are independent of μ_R , we can conclude that

$$f_i(\mu_F, \mu_R) = f_i(\mu_F), \quad (\text{B.2.9})$$

which allows us to solve (B.2.7) iteratively to obtain an expression for the the PDF at one scale μ_F' in terms of another, μ_F :

$$\begin{aligned} f_i(\mu_F') &= f_i(\mu_F) - \frac{\alpha_s(\mu_R)}{2\pi} P_{ij}^{(0)} \otimes f_j(\mu_F) L_F \\ &- \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^2 \left[P_{ij}^{(1)} \otimes f_j(\mu_F) L_F - \frac{1}{2} P_{ij}^{(0)} \otimes P_{jk}^{(0)} \otimes f_k(\mu_F) L_F^2 \right. \\ &\quad \left. + P_{ij}^{(0)} \otimes f_j(\mu_F) \beta_0 L_F \left(l + \frac{1}{2} L_F \right) \right] \\ &+ \mathcal{O}(\alpha_s^3), \end{aligned} \quad (\text{B.2.10})$$

³As the N³LO splitting functions are not currently known, we neglect terms of this order (α_s^4). The status of current N³LO calculations of the splitting functions to this accuracy can be found in [301–303].

where we take

$$L_F = \log \frac{\mu_F^2}{\mu_F'^2}. \quad (\text{B.2.11})$$

We can now explicitly include the PDF and μ_F dependence in (B.2.6), using the decomposition

$$\sigma(\mu_0, \mu_0, \alpha_s(\mu_0)) = \sum_{i=0}^{\infty} \left(\frac{\alpha_s(\mu_0)}{2\pi} \right)^{n+i} \hat{\sigma}_{ij}^{(i)} \otimes f_i(\mu_0) \otimes f_j(\mu_0), \quad (\text{B.2.12})$$

to give

$$\begin{aligned} \sigma(\mu_R, \mu_F, \alpha_s(\mu_R), L_R, L_F) = & \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^n \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes f_j(\mu_F) \\ & + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+1} \left\{ \hat{\sigma}_{ij}^{(1)} \otimes f_i(\mu_F) \otimes f_j(\mu_F) \right. \\ & \quad + L_R n \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes f_j(\mu_F) \\ & \quad - L_F \left[\hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(0)} \otimes f_k(\mu_F)) \right. \\ & \quad \quad \left. \left. + \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \otimes f_j(\mu_F) \right] \right\} \\ & + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+2} \left\{ \hat{\sigma}_{ij}^{(2)} \otimes f_i(\mu_F) \otimes f_j(\mu_F) \right. \\ & \quad + L_R \left((n+1) \beta_0 \hat{\sigma}_{ij}^{(1)} + n \beta_1 \hat{\sigma}_{ij}^{(0)} \right) \otimes f_i(\mu_F) \otimes f_j(\mu_F) \\ & \quad + L_R^2 \frac{n(n+1)}{2} \beta_0^2 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes f_j(\mu_F) \\ & \quad + L_F \left[- \hat{\sigma}_{ij}^{(1)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(0)} \otimes f_k(\mu_F)) \right. \\ & \quad \quad - \hat{\sigma}_{ij}^{(1)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \otimes f_j(\mu_F) \\ & \quad \quad - \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(1)} \otimes f_k(\mu_F)) \\ & \quad \quad \left. - \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(1)} \otimes f_k(\mu_F)) \otimes f_j(\mu_F) \right] \\ & \quad + L_F^2 \left[\hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \otimes (P_{jl}^{(0)} \otimes f_l(\mu_F)) \right. \\ & \quad \quad + \frac{1}{2} \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(0)} \otimes P_{kl}^{(0)} \otimes f_l(\mu_F)) \\ & \quad \quad + \frac{1}{2} \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes P_{kl}^{(0)} \otimes f_l(\mu_F)) \otimes f_j(\mu_F) \\ & \quad \quad + \frac{1}{2} \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(0)} \otimes f_k(\mu_F)) \\ & \quad \quad \left. \left. + \frac{1}{2} \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \otimes f_j(\mu_F) \right] \right\} \end{aligned}$$

$$\begin{aligned}
& +L_F L_R \left[-(n+1) \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \otimes (P_{jk}^{(0)} \otimes f_k(\mu_F)) \right. \\
& \quad \left. -(n+1) \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \otimes f_j(\mu_F) \right] \Big\} \\
& +\mathcal{O}(\alpha_s^{n+3}).
\end{aligned} \tag{B.2.13}$$

For the DIS case, which we consider here, we have taken $f_j(\mu_0)$ to be the identity operator under convolution as the incoming lepton is not considered to have a QCD PDF, which simplifies (B.2.13) substantially and gives

$$\begin{aligned}
\sigma(\mu_R, \mu_F, \alpha_s(\mu_R), L_R, L_F) = & \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^n \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \\
& + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+1} \left\{ \hat{\sigma}_{ij}^{(1)} \otimes f_i(\mu_F) + L_R n \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \right. \\
& \quad \left. - L_F \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \right\} \\
& + \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{n+2} \left\{ \hat{\sigma}_{ij}^{(2)} \otimes f_i(\mu_F) + L_R \left((n+1) \beta_0 \hat{\sigma}_{ij}^{(1)} + n \beta_1 \hat{\sigma}_{ij}^{(0)} \right) \otimes f_i(\mu_F) \right. \\
& \quad + L_R^2 \frac{n(n+1)}{2} \beta_0^2 \hat{\sigma}_{ij}^{(0)} \otimes f_i(\mu_F) \\
& \quad - L_F \left[\hat{\sigma}_{ij}^{(1)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) + \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(1)} \otimes f_k(\mu_F)) \right] \\
& \quad + L_F^2 \left[\frac{1}{2} \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes P_{kl}^{(0)} \otimes f_l(\mu_F)) \right. \\
& \quad \quad \left. + \frac{1}{2} \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \right] \\
& \quad \left. - L_F L_R \left[(n+1) \beta_0 \hat{\sigma}_{ij}^{(0)} \otimes (P_{ik}^{(0)} \otimes f_k(\mu_F)) \right] \right\} \\
& +\mathcal{O}(\alpha_s^{n+3}).
\end{aligned} \tag{B.2.14}$$

At this point, the Mellin convolution over the Bjorken x variable

$$[a \otimes b](x) = \int_x^1 \frac{dy}{y} a(y) b\left(\frac{x}{y}\right) \tag{B.2.15}$$

provides a further level of complexity with respect to the renormalisation scale, as it can not be solved simply without recourse to integral transformations. However, the summation structure in (B.2.13) can be substantially simplified by taking the linear combinations of PDFs previously considered in (B.1.7)-(B.1.9), which separates the

evolution of all non-singlet quark PDF terms f^{NS} such that they each individually obey

$$\mu_F^2 \frac{d}{d\mu_F^2} f_{q_i}^{NS}(\mu_F, \mu_R) = P_{qq}(\alpha_s(\mu_R), \mu_F, \mu_R) \otimes f_{q_i}^{NS}(\mu_F, \mu_R). \quad (\text{B.2.16})$$

The remaining gluon and pure-singlet terms then evolve in a coupled manner, giving (dropping arguments for simplicity)

$$\mu_F^2 \frac{d}{d\mu_F^2} \begin{pmatrix} q_{PS} \\ g \end{pmatrix} = \begin{pmatrix} P_{qq} & P_{qg} \\ P_{gq} & P_{gg} \end{pmatrix} \otimes \begin{pmatrix} q_{PS} \\ g \end{pmatrix}. \quad (\text{B.2.17})$$

At this point, the standard method of obtaining the factorisation scale dependence is to transform (B.2.14) into Mellin N -space using the Mellin transform:

$$\tilde{f}(N) = \int_0^1 dx \, x^{N-1} f(x). \quad (\text{B.2.18})$$

In Mellin space, the convolution in (B.2.7) becomes a simple multiplication, allowing the differential equation in μ_F^2 to be solved directly and giving the full μ_F evolution in N -space:

$$\mu_F^2 \frac{d\tilde{f}_i(N, \mu_F^2)}{d\mu_F^2} = \frac{\alpha_S(\mu_F)}{2\pi} \tilde{P}_{ij}(N, \alpha_S(\mu_F^2)) \tilde{f}_j(N, \mu_F^2) \quad (\text{B.2.19})$$

The solutions to this differential equation can then be transformed back into x -space using the inverse Mellin transform

$$f(x) = \frac{1}{2\pi i} \oint dN \, x^{-N} \tilde{f}(N), \quad (\text{B.2.20})$$

at which point the convolution can be numerically evaluated, which is the topic of the following section.

B.3 Numerical Implementation

There are some subtleties involved in this numerical implementation related to evaluation of $+$ -functions, which we will briefly discuss. Whilst for convenience we will

consider Wilson coefficient functions, the same principles also apply to the integrated antenna functions implemented in NNLOJET.

In general, and neglecting factorisation scale dependence, each of the Wilson coefficient functions $C_i(x)$ takes the form

$$C_i(x) = \left(A \delta(1-x) + B(x) + \sum_{m \geq 0} \ln^m(1-x) \frac{C_m(x)}{(1-x)} \right) \Big|_+ M(\{x\bar{y}p\}, \mu_R, \mu_F) \quad (\text{B.3.1})$$

where M is some infrared finite function, and we have the usual definition of the $+$ -distribution:

$$\int_0^1 \frac{f(x)}{1-x} dx = \int_0^1 \frac{f(x) - f(1)}{1-x} dx \quad (\text{B.3.2})$$

regulating the divergence at the endpoint $x = 1$. Thus the convolution of the coefficient functions with the PDF requires two separate regions in DIS⁴ - one for the general region $0 \leq x < 1$ and one for the endpoint $x = 1$.

When expressed as a convolution with some PDF f , and including the mass factorisation convolution over x , we obtain

$$\begin{aligned} d\sigma \sim & \int_0^1 \frac{d\bar{y}}{\bar{y}} f(\bar{y}, \mu_F) \int_0^1 dx \left[A(1) - \frac{1}{1-x} \sum_{m \geq 0} \ln^m(1-x) C_m(1) \right] \\ & \times M(\{\bar{y}p\}, \mu_R, \mu_F) d\Phi_n \\ & + \int_0^1 \frac{d\bar{y}}{\bar{y}} f(\bar{y}, \mu_F) \int_0^1 \frac{dx}{x} \left[B(x) + \frac{1}{1-x} \sum_{m \geq 0} \ln^m(1-x) C_m(x) \right] \\ & \times M(\{x\bar{y}p\}, \mu_R, \mu_F) d\Phi_n. \end{aligned} \quad (\text{B.3.3})$$

In the current form, this requires the structure function to be evaluated twice for each phase space point due to the different scalings of the incoming parton momenta $\{xp\}$ and $\{\bar{y}xp\}$. This can be rewritten to require only a single evaluation [45] by

⁴This as opposed to four regions required in hadron-hadron collisions where equivalent convolutions for e.g. RV antenna subtraction terms contain one endpoint for each PDF convolution with $x_{1,2}$. Evaluation of this requires each of the regions $0 \leq x_1 < 1, 0 \leq x_2 < 1$; $0 \leq x_1 < 1, x_2 = 1$; $x_1 = 1, 0 \leq x_2 < 1$ and $x_1 = x_2 = 1$. In e^+e^- collisions, no convolution is needed as there are no PDFs, and the matrix elements contain only terms proportional to $A\delta(1-x)$.

transforming the x integrals using a factor of $d\xi\delta(\xi - \bar{y}x)$, which gives:

$$\begin{aligned} d\sigma \sim \int_0^1 \frac{d\xi}{\xi} \left\{ \int_\xi^1 f(\xi, \mu_F) dx \left[\frac{A(1)}{1-\xi} + \sum_{m \geq 0} \frac{\ln^m(1-\xi)}{(m+1)(\xi+1)} C_m(1) \right. \right. \\ \left. \left. - \frac{1}{1-x} \sum_{m \geq 0} \frac{\ln^m(1-x) C_m(1)}{1-x} \right] \right. \\ \left. + f\left(\frac{\xi}{x}, \mu_F\right) \frac{dx}{x} \left[B(x) + \sum_{m \geq 0} \frac{\ln^m(1-x)}{(1-x)} C_m(x) \right] \right\} \\ \times M(\{\xi p\}, \mu_R, \mu_F) d\Phi_n, \end{aligned} \quad (\text{B.3.4})$$

using the identities:

$$\int_0^1 A(x) \delta(1-x) dx = \int_\xi^1 \frac{A(1)}{1-\xi} dx, \quad (\text{B.3.5})$$

and

$$\begin{aligned} \int_\xi^1 C_m(x) \left(\frac{\ln^m(1-x)}{1-x} dx \right)_+ = \int_\xi^1 \frac{(C_m(x) - C_m(1)) \ln^m(1-x)}{1-x} dx \\ - \int_0^\xi \frac{C_m(1) \ln^m(1-x)}{1-x} dx \end{aligned} \quad (\text{B.3.6})$$

where

$$- \int_0^\xi \frac{C_m(1) \ln^m(1-x)}{1-x} dx = - \int_\xi^1 \frac{C_m(1) \ln^{m+1}(1-\xi)}{(1-\xi)(m+1)} dx. \quad (\text{B.3.7})$$

This is particularly advantageous as in general the coefficient function is computationally far more expensive than the PDF to evaluate for a given phase space point, meaning that trading evaluations of the C_i for evaluations of f results in a substantial increase in efficiency.

Bibliography

- [1] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *Next-to-Next-to-Leading-Order QCD Corrections to the Transverse Momentum Distribution of Weak Gauge Bosons*, *Phys. Rev. Lett.* **120** (2018) 122001, [1712.07543].
- [2] W. Bizon, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, P. F. Monni et al., *The transverse momentum spectrum of weak gauge bosons at $N^3LL+NNLO$* , [1905.05171].
- [3] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *NNLO QCD Corrections to W +jet Production in NNLOJET*, *PoS LL2018* (2018) 041, [1807.09113].
- [4] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and D. M. Walker, *Vector Boson Production in Association with a Jet at Forward Rapidities*, *Eur. Phys. J.* **C79** (2019) 526, [1901.11041].
- [5] ATLAS collaboration, *Measurement of the effective leptonic weak mixing angle using electron and muon pairs from Z -boson decay in the ATLAS experiment at $\sqrt{s} = 8$ TeV*, [ATLAS-CONF-2018-037].
- [6] J. Niehues and D. M. Walker, *NNLO QCD Corrections to Jet Production in Charged Current Deep Inelastic Scattering*, *Phys. Lett.* **B788** (2019) 243–248, [1807.02529].

-
- [7] T. Gehrmann, A. Huss, J. Niehues, A. Vogt and D. M. Walker, *Jet production in charged-current deep-inelastic scattering to third order in QCD*, *Phys. Lett.* **B792** (2019) 182–186, [1812.06104].
- [8] PARTICLE DATA GROUP collaboration, M. Tanabashi et al., *Review of Particle Physics*, *Phys. Rev.* **D98** (2018) 030001.
- [9] A. Salam, *Weak and Electromagnetic Interactions*, *Conf. Proc.* **C680519** (1968) 367–377.
- [10] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (1967) 1264–1266.
- [11] J. Goldstone, A. Salam and S. Weinberg, *Broken Symmetries*, *Phys. Rev.* **127** (1962) 965–970.
- [12] S. L. Glashow, *Partial Symmetries of Weak Interactions*, *Nucl. Phys.* **22** (1961) 579–588.
- [13] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321–323.
- [14] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (1964) 508–509.
- [15] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (1964) 585–587.
- [16] CMS collaboration, S. Chatrchyan et al., *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett.* **B716** (2012) 30–61, [1207.7235].
- [17] ATLAS collaboration, G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett.* **B716** (2012) 1–29, [1207.7214].

- [18] CMS collaboration, A. M. Sirunyan et al., *Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 421, [1809.10733].
- [19] ATLAS collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, [ATLAS-CONF-2019-005].
- [20] ATLAS Collaboration, *Combination of searches for Higgs boson pairs in pp collisions at 13 TeV with the ATLAS experiment.*, [ATLAS-CONF-2018-043].
- [21] CMS collaboration, A. M. Sirunyan et al., *Combination of searches for Higgs boson pair production in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **122** (2019) 121803, [1811.09689].
- [22] ATLAS Collaboration, *Measurement prospects of the pair production and self-coupling of the Higgs boson with the ATLAS experiment at the HL-LHC*, [ATL-PHYS-PUB-2018-053].
- [23] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652–657.
- [24] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, *Phys. Rev. Lett.* **10** (1963) 531–533.
- [25] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, *Phys. Rev. Lett.* **13** (1964) 138–140.
- [26] F. J. Dyson, *Divergence of perturbation theory in quantum electrodynamics*, *Phys. Rev.* **85** (1952) 631–632.
- [27] G. 't Hooft and M. J. G. Veltman, *Regularization and Renormalization of Gauge Fields*, *Nucl. Phys.* **B44** (1972) 189–213.

-
- [28] P. Marquard, A. V. Smirnov, V. A. Smirnov and M. Steinhauser, *Quark Mass Relations to Four-Loop Order in Perturbative QCD*, *Phys. Rev. Lett.* **114** (2015) 142002, [1502.01030].
- [29] S. Catani and M. H. Seymour, *A General algorithm for calculating jet cross-sections in NLO QCD*, *Nucl. Phys.* **B485** (1997) 291–419, [hep-ph/9605323].
- [30] D. A. Kosower, *Antenna factorization of gauge theory amplitudes*, *Phys. Rev.* **D57** (1998) 5410–5416, [hep-ph/9710213].
- [31] F. Bloch and A. Nordsieck, *Note on the Radiation Field of the electron*, *Phys. Rev.* **52** (1937) 54–59.
- [32] T. D. Lee and M. Nauenberg, *Degenerate Systems and Mass Singularities*, *Phys. Rev.* **133** (1964) B1549–B1562.
- [33] T. Kinoshita, *Mass singularities of Feynman amplitudes*, *J. Math. Phys.* **3** (1962) 650–677.
- [34] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*, *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* **8** (1996) 1–435.
- [35] J. C. Collins, D. E. Soper and G. F. Sterman, *Factorization for Short Distance Hadron - Hadron Scattering*, *Nucl. Phys.* **B261** (1985) 104–142.
- [36] J. C. Collins, D. E. Soper and G. F. Sterman, *Soft Gluons and Factorization*, *Nucl. Phys.* **B308** (1988) 833–856.
- [37] J. R. Forshaw, M. H. Seymour and A. Siodmok, *On the Breaking of Collinear Factorization in QCD*, *JHEP* **11** (2012) 066, [1206.6363].
- [38] L. A. Harland-Lang, A. D. Martin, P. Motylinski and R. S. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J.* **C75** (2015) 204, [1412.3989].

- [39] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J. C* **75** (2015) 132, [1412.7420].
- [40] W. T. Giele and E. W. N. Glover, *Higher order corrections to jet cross-sections in e^+e^- annihilation*, *Phys. Rev.* **D46** (1992) 1980–2010.
- [41] W. T. Giele, E. W. N. Glover and D. A. Kosower, *Higher order corrections to jet cross-sections in hadron colliders*, *Nucl. Phys.* **B403** (1993) 633–670, [hep-ph/9302225].
- [42] S. Catani and M. Grazzini, *An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC*, *Phys. Rev. Lett.* **98** (2007) 222002, [hep-ph/0703012].
- [43] J. Gaunt, M. Stahlhofen, F. J. Tackmann and J. R. Walsh, *N -jettiness Subtractions for NNLO QCD Calculations*, *JHEP* **09** (2015) 058, [1505.04794].
- [44] R. Frederix, S. Frixione, F. Maltoni and T. Stelzer, *Automation of next-to-leading order computations in QCD: The FKS subtraction*, *JHEP* **10** (2009) 003, [0908.4272].
- [45] T. Gleisberg and F. Krauss, *Automating dipole subtraction for QCD NLO calculations*, *Eur. Phys. J. C* **53** (2008) 501–523, [0709.2881].
- [46] A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *Antenna subtraction at NNLO*, *JHEP* **09** (2005) 056, [hep-ph/0505111].
- [47] A. Daleo, T. Gehrmann and D. Maitre, *Antenna subtraction with hadronic initial states*, *JHEP* **04** (2007) 016, [hep-ph/0612257].
- [48] J. Currie, E. W. N. Glover and S. Wells, *Infrared Structure at NNLO Using Antenna Subtraction*, *JHEP* **04** (2013) 066, [1301.4693].

- [49] M. Cacciari, F. A. Dreyer, A. Karlberg, G. P. Salam and G. Zanderighi, *Fully Differential Vector-Boson-Fusion Higgs Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **115** (2015) 082002, [1506.02660].
- [50] F. A. Dreyer and A. Karlberg, *Fully differential Vector-Boson Fusion Higgs Pair Production at Next-to-Next-to-Leading Order*, [1811.07918].
- [51] G. Somogyi and Z. Trocsanyi, *A Subtraction scheme for computing QCD jet cross sections at NNLO: Regularization of real-virtual emission*, *JHEP* **01** (2007) 052, [hep-ph/0609043].
- [52] G. Somogyi, Z. Trocsanyi and V. Del Duca, *A Subtraction scheme for computing QCD jet cross sections at NNLO: Regularization of doubly-real emissions*, *JHEP* **01** (2007) 070, [hep-ph/0609042].
- [53] F. Caola, K. Melnikov and R. Röntsch, *Nested soft-collinear subtractions in NNLO QCD computations*, *Eur. Phys. J.* **C77** (2017) 248, [1702.01352].
- [54] F. Caola, K. Melnikov and R. Röntsch, *Analytic results for color-singlet production at NNLO QCD with the nested soft-collinear subtraction scheme*, [1902.02081].
- [55] L. Magnea, E. Maina, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, *Local Analytic Sector Subtraction at NNLO*, *JHEP* **12** (2018) 107, [1806.09570].
- [56] C. Anastasiou, K. Melnikov and F. Petriello, *A new method for real radiation at NNLO*, *Phys. Rev.* **D69** (2004) 076010, [hep-ph/0311311].
- [57] T. Binoth and G. Heinrich, *Numerical evaluation of phase space integrals by sector decomposition*, *Nucl. Phys.* **B693** (2004) 134–148, [hep-ph/0402265].
- [58] M. Czakon, *A novel subtraction scheme for double-real radiation at NNLO*, *Phys. Lett.* **B693** (2010) 259–268, [1005.0274].

- [59] M. Czakon, *Double-real radiation in hadronic top quark pair production as a proof of a certain concept*, *Nucl. Phys.* **B849** (2011) 250–295, [1101.0642].
- [60] M. Czakon and D. Heymes, *Four-dimensional formulation of the sector-improved residue subtraction scheme*, *Nucl. Phys.* **B890** (2014) 152–227, [1408.2500].
- [61] ATLAS collaboration, *Standard Model Summary Plots Spring 2019*, [ATL-PHYS-PUB-2019-010].
- [62] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, [0802.1189].
- [63] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*, pp. 270–279, 1998, [hep-ph/9907280].
- [64] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001, [hep-ph/9707323].
- [65] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock and B. R. Webber, *New clustering algorithm for multi - jet cross-sections in $e^+ e^-$ annihilation*, *Phys. Lett.* **B269** (1991) 432–438.
- [66] S. Catani, Y. L. Dokshitzer, M. H. Seymour and B. R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys.* **B406** (1993) 187–224.
- [67] F. A. Berends and W. Giele, *The Six Gluon Process as an Example of Weyl-Van Der Waerden Spinor Calculus*, *Nucl. Phys.* **B294** (1987) 700–732.
- [68] M. L. Mangano and S. J. Parke, *Quark - Gluon Amplitudes in the Dual Expansion*, *Nucl. Phys.* **B299** (1988) 673–692.

-
- [69] M. L. Mangano, S. J. Parke and Z. Xu, *Duality and Multi - Gluon Scattering*, *Nucl. Phys.* **B298** (1988) 653–672.
- [70] M. L. Mangano, *The Color Structure of Gluon Emission*, *Nucl. Phys.* **B309** (1988) 461–475.
- [71] F. A. Berends and W. T. Giele, *Recursive Calculations for Processes with n Gluons*, *Nucl. Phys.* **B306** (1988) 759–808.
- [72] D. A. Kosower, *Color Factorization for Fermionic Amplitudes*, *Nucl. Phys.* **B315** (1989) 391–418.
- [73] L. J. Dixon, *Calculating scattering amplitudes efficiently*, in *QCD and beyond. Theoretical Advanced Study Institute in Elementary Particle Physics, TASI-95*, pp. 539–584, 1996, [[hep-ph/9601359](#)].
- [74] S. Catani, *The Singular behavior of QCD amplitudes at two loop order*, *Phys. Lett.* **B427** (1998) 161–171, [[hep-ph/9802439](#)].
- [75] F. A. Berends and W. T. Giele, *Multiple Soft Gluon Radiation in Parton Processes*, *Nucl. Phys.* **B313** (1989) 595–633.
- [76] J. M. Campbell and E. W. N. Glover, *Double unresolved approximations to multiparton scattering amplitudes*, *Nucl. Phys.* **B527** (1998) 264–288, [[hep-ph/9710255](#)].
- [77] A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *Infrared structure of $e^+e^- \rightarrow 2$ jets at NNLO*, *Nucl. Phys.* **B691** (2004) 195–222, [[hep-ph/0403057](#)].
- [78] A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *Quark-gluon antenna functions from neutralino decay*, *Phys. Lett.* **B612** (2005) 36–48, [[hep-ph/0501291](#)].

- [79] A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *Gluon-gluon antenna functions from Higgs boson decay*, *Phys. Lett.* **B612** (2005) 49–60, [hep-ph/0502110].
- [80] A. Gehrmann-De Ridder, T. Gehrmann and G. Heinrich, *Four particle phase space integrals in massless QCD*, *Nucl. Phys.* **B682** (2004) 265–288, [hep-ph/0311276].
- [81] A. Daleo, A. Gehrmann-De Ridder, T. Gehrmann and G. Luisoni, *Antenna subtraction at NNLO with hadronic initial states: initial-final configurations*, *JHEP* **01** (2010) 118, [0912.0374].
- [82] T. Gehrmann and P. F. Monni, *Antenna subtraction at NNLO with hadronic initial states: real-virtual initial-initial configurations*, *JHEP* **12** (2011) 049, [1107.4037].
- [83] A. Gehrmann-De Ridder, T. Gehrmann and M. Ritzmann, *Antenna subtraction at NNLO with hadronic initial states: double real initial-initial configurations*, *JHEP* **10** (2012) 047, [1207.5779].
- [84] A. Gehrmann-De Ridder, E. W. N. Glover and J. Pires, *Real-Virtual corrections for gluon scattering at NNLO*, *JHEP* **02** (2012) 141, [1112.3613].
- [85] S. Weinzierl, *NNLO corrections to 3-jet observables in electron-positron annihilation*, *Phys. Rev. Lett.* **101** (2008) 162001, [0807.3241].
- [86] S. Weinzierl, *The infrared structure of $e^+e^- \rightarrow 3$ jets at NNLO reloaded*, *JHEP* **07** (2009) 009, [0904.1145].
- [87] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and T. A. Morgan, *Precise QCD predictions for the production of a Z boson in association with a hadronic jet*, *Phys. Rev. Lett.* **117** (2016) 022001, [1507.02850].

- [88] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and T. A. Morgan, *The NNLO QCD corrections to Z boson production at large transverse momentum*, *JHEP* **07** (2016) 133, [1605.04295].
- [89] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and T. A. Morgan, *NNLO QCD corrections for Drell-Yan p_T^Z and ϕ^* observables at the LHC*, *JHEP* **11** (2016) 094, [1610.01843].
- [90] A. Gehrmann-De Ridder, A. Huss, T. Gehrmann, E. W. N. Glover and T. A. Morgan, *QCD Predictions For The Production Of A Z Boson In Association With A Hadronic Jet*, in *Proceedings, 51st Rencontres de Moriond on QCD and High Energy Interactions: La Thuile, Italy, March 19-26, 2016*, pp. 119–122, ARISF, ARISF, 2016.
- [91] A. Gehrmann-De Ridder, T. Gehrmann, N. Glover, A. Huss and T. A. Morgan, *NNLO QCD corrections for Z boson plus jet production*, *PoS RADCOR2015* (2016) 075, [1601.04569].
- [92] J. Currie, T. Gehrmann and J. Niehues, *Precise QCD predictions for the production of dijet final states in deep inelastic scattering*, *Phys. Rev. Lett.* **117** (2016) 042001, [1606.03991].
- [93] J. Niehues, J. Currie and T. Gehrmann, *NNLO dijet production in DIS with the antenna subtraction method*, *PoS LL2016* (2016) 055.
- [94] J. Currie, T. Gehrmann, A. Huss and J. Niehues, *NNLO QCD corrections to jet production in deep inelastic scattering*, *JHEP* **07** (2017) 018, [1703.05977].
- [95] L. W. Garland, T. Gehrmann, E. W. N. Glover, A. Koukoutsakis and E. Remiddi, *The Two loop QCD matrix element for $e^+ e^- \rightarrow 3$ jets*, *Nucl. Phys.* **B627** (2002) 107–188, [hep-ph/0112081].

- [96] L. W. Garland, T. Gehrmann, E. W. N. Glover, A. Koukoutsakis and E. Remiddi, *Two loop QCD helicity amplitudes for $e^+ e^- \rightarrow$ three jets*, *Nucl. Phys.* **B642** (2002) 227–262, [[hep-ph/0206067](#)].
- [97] E. W. N. Glover and D. J. Miller, *The one loop QCD corrections for $\gamma^* \rightarrow Q\bar{Q}q\bar{q}$* , *Phys. Lett.* **B396** (1997) 257–263, [[hep-ph/9609474](#)].
- [98] Z. Bern, L. J. Dixon, D. A. Kosower and S. Weinzierl, *One loop amplitudes for $e^+ e^- \rightarrow \bar{q}q\bar{Q}Q$* , *Nucl. Phys.* **B489** (1997) 3–23, [[hep-ph/9610370](#)].
- [99] J. M. Campbell, E. W. N. Glover and D. J. Miller, *The One loop QCD corrections for $\gamma^* \rightarrow q\bar{q}gg$* , *Phys. Lett.* **B409** (1997) 503–508, [[hep-ph/9706297](#)].
- [100] Z. Bern, L. J. Dixon and D. A. Kosower, *One loop amplitudes for $e^+ e^-$ to four partons*, *Nucl. Phys.* **B513** (1998) 3–86, [[hep-ph/9708239](#)].
- [101] K. Hagiwara and D. Zeppenfeld, *Amplitudes for Multiparton Processes Involving a Current at $e^+ e^-$, $e^\pm p$, and Hadron Colliders*, *Nucl. Phys.* **B313** (1989) 560–594.
- [102] F. A. Berends, W. T. Giele and H. Kuijf, *Exact Expressions for Processes Involving a Vector Boson and Up to Five Partons*, *Nucl. Phys.* **B321** (1989) 39–82.
- [103] G. P. Lepage, *VEGAS: AN ADAPTIVE MULTIDIMENSIONAL INTEGRATION PROGRAM*, .
- [104] S. D. Drell and T.-M. Yan, *Massive Lepton Pair Production in Hadron-Hadron Collisions at High-Energies*, *Phys. Rev. Lett.* **25** (1970) 316–320.
- [105] UA2 collaboration, P. Bagnaia et al., *Evidence for $Z^0 \rightarrow e^+ e^-$ at the CERN $\bar{p}p$ Collider*, *Phys. Lett.* **B129** (1983) 130–140.

-
- [106] UA1 collaboration, G. Arnison et al., *Experimental Observation of Lepton Pairs of Invariant Mass Around $95 \text{ GeV}/c^2$ at the CERN SPS Collider*, *Phys. Lett.* **B126** (1983) 398–410.
- [107] UA1 collaboration, G. Arnison et al., *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540 \text{ GeV}$* , *Phys. Lett.* **B122** (1983) 103–116.
- [108] UA2 collaboration, M. Banner et al., *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider*, *Phys. Lett.* **B122** (1983) 476–485.
- [109] CDF collaboration, T. Aaltonen et al., *Precise measurement of the W-boson mass with the CDF II detector*, *Phys. Rev. Lett.* **108** (2012) 151803, [1203.0275].
- [110] DØ collaboration, V. M. Abazov et al., *Measurement of the W Boson Mass with the DØ Detector*, *Phys. Rev. Lett.* **108** (2012) 151804, [1203.0293].
- [111] CDF, DØ collaboration, T. A. Aaltonen et al., *Combination of CDF and DØ W-Boson Mass Measurements*, *Phys. Rev.* **D88** (2013) 052018, [1307.7627].
- [112] ATLAS collaboration, M. Aaboud et al., *Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7 \text{ TeV}$ with the ATLAS detector*, [1701.07240].
- [113] TEVATRON ELECTROWEAK WORKING GROUP collaboration, *Combination of CDF and DØ Results on the Width of the W boson*, [1003.2826].
- [114] CMS collaboration, S. Chatrchyan et al., *Measurement of the weak mixing angle with the Drell-Yan process in proton-proton collisions at the LHC*, *Phys. Rev.* **D84** (2011) 112002, [1110.2682].

- [115] LHCb collaboration, R. Aaij et al., *Measurement of the forward-backward asymmetry in $Z/\gamma^* \rightarrow \mu^+\mu^-$ decays and determination of the effective weak mixing angle*, *JHEP* **11** (2015) 190, [1509.07645].
- [116] ATLAS collaboration, G. Aad et al., *Measurement of the forward-backward asymmetry of electron and muon pair-production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **09** (2015) 049, [1503.03709].
- [117] NNPDF collaboration, R. D. Ball et al., *Parton distributions from high-precision collider data*, *Eur. Phys. J.* **C77** (2017) 663, [1706.00428].
- [118] xFITTER DEVELOPERS' TEAM collaboration, F. Giuli et al., *The photon PDF from high-mass Drell–Yan data at the LHC*, *Eur. Phys. J.* **C77** (2017) 400, [1701.08553].
- [119] J. Gao, L. Harland-Lang and J. Rojo, *The Structure of the Proton in the LHC Precision Era*, *Phys. Rept.* **742** (2018) 1–121, [1709.04922].
- [120] U. Blumenschein et al., *Pushing the precision frontier at the LHC with V +jets*, 2018, [1802.02100].
- [121] R. Hamberg, W. L. van Neerven and T. Matsuura, *A complete calculation of the order α_s^2 correction to the Drell-Yan K factor*, *Nucl. Phys.* **B359** (1991) 343–405.
- [122] C. Anastasiou, L. J. Dixon, K. Melnikov and F. Petriello, *High precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at NNLO*, *Phys. Rev.* **D69** (2004) 094008, [hep-ph/0312266].
- [123] K. Melnikov and F. Petriello, *Electroweak gauge boson production at hadron colliders through $\mathcal{O}(\alpha_s^2)$* , *Phys. Rev.* **D74** (2006) 114017, [hep-ph/0609070].
- [124] Y. Li and F. Petriello, *Combining QCD and electroweak corrections to dilepton production in FEWZ*, *Phys. Rev.* **D86** (2012) 094034, [1208.5967].

- [125] S. Catani, L. Cieri, G. Ferrera, D. de Florian and M. Grazzini, *Vector boson production at hadron colliders: a fully exclusive QCD calculation at NNLO*, *Phys. Rev. Lett.* **103** (2009) 082001, [0903.2120].
- [126] S. Höche, Y. Li and S. Prestel, *Drell-Yan lepton pair production at NNLO QCD with parton showers*, *Phys. Rev.* **D91** (2015) 074015, [1405.3607].
- [127] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. Giele, X. Liu et al., *Color singlet production at NNLO in MCFM*, *Eur. Phys. J.* **C77** (2017) 7, [1605.08011].
- [128] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, *Eur. Phys. J.* **C78** (2018) 537, [1711.06631].
- [129] G. Bozzi, S. Catani, G. Ferrera, D. de Florian and M. Grazzini, *Production of Drell-Yan lepton pairs in hadron collisions: Transverse-momentum resummation at next-to-next-to-leading logarithmic accuracy*, *Phys. Lett.* **B696** (2011) 207–213, [1007.2351].
- [130] T. Becher and M. Neubert, *Drell-Yan Production at Small q_T , Transverse Parton Distributions and the Collinear Anomaly*, *Eur. Phys. J.* **C71** (2011) 1665, [1007.4005].
- [131] A. Banfi, M. Dasgupta, S. Marzani and L. Tomlinson, *Predictions for Drell-Yan ϕ^* and Q_T observables at the LHC*, *Phys. Lett.* **B715** (2012) 152–156, [1205.4760].
- [132] S. Catani, D. de Florian, G. Ferrera and M. Grazzini, *Vector boson production at hadron colliders: transverse-momentum resummation and leptonic decay*, *JHEP* **12** (2015) 047, [1507.06937].
- [133] P. Banerjee, G. Das, P. K. Dhani and V. Ravindran, *Threshold resummation of the rapidity distribution for Drell-Yan production at NNLO+NNLL*, *Phys. Rev.* **D98** (2018) 054018, [1805.01186].

- [134] T. Gehrmann, T. Huber and D. Maitre, *Two-loop quark and gluon form-factors in dimensional regularisation*, *Phys. Lett.* **B622** (2005) 295–302, [hep-ph/0507061].
- [135] T. Gehrmann, E. W. N. Glover, T. Huber, N. Ikizlerli and C. Studerus, *Calculation of the quark and gluon form factors to three loops in QCD*, *JHEP* **06** (2010) 094, [1004.3653].
- [136] T. Ahmed, M. Mahakhud, N. Rana and V. Ravindran, *Drell-Yan Production at Threshold to Third Order in QCD*, *Phys. Rev. Lett.* **113** (2014) 112002, [1404.0366].
- [137] K. Melnikov, R. Rietkerk, L. Tancredi and C. Wever, *Double-real contribution to the quark beam function at N^3LO QCD*, *JHEP* **02** (2019) 159, [1809.06300].
- [138] K. Melnikov, R. Rietkerk, L. Tancredi and C. Wever, *Triple-real contribution to the quark beam function in QCD at next-to-next-to-next-to-leading order*, [1904.02433].
- [139] R. Boughezal, C. Focke, X. Liu and F. Petriello, *W-boson production in association with a jet at next-to-next-to-leading order in perturbative QCD*, *Phys. Rev. Lett.* **115** (2015) 062002, [1504.02131].
- [140] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. T. Giele, X. Liu et al., *Z-boson production in association with a jet at next-to-next-to-leading order in perturbative QCD*, *Phys. Rev. Lett.* **116** (2016) 152001, [1512.01291].
- [141] R. Boughezal, X. Liu and F. Petriello, *W-boson plus jet differential distributions at NNLO in QCD*, *Phys. Rev.* **D94** (2016) 113009, [1602.06965].

- [142] R. Boughezal, X. Liu and F. Petriello, *Phenomenology of the Z-boson plus jet process at NNLO*, *Phys. Rev.* **D94** (2016) 074015, [1602.08140].
- [143] J. M. Campbell, R. K. Ellis and C. Williams, *Driving missing data at the LHC: NNLO predictions for the ratio of $\gamma + j$ and $Z + j$* , *Phys. Rev.* **D96** (2017) 014037, [1703.10109].
- [144] W. Bizoń, X. Chen, A. Gehrmann-De Ridder, T. Gehrmann, N. Glover, A. Huss et al., *Fiducial distributions in Higgs and Drell-Yan production at $N^3LL+NNLO$* , *JHEP* **12** (2018) 132, [1805.05916].
- [145] P. Sun, B. Yan, C. P. Yuan and F. Yuan, *Resummation of High Order Corrections in Z Boson Plus Jet Production at the LHC*, [1810.03804].
- [146] H. Ita, Z. Bern, L. J. Dixon, F. Febres Cordero, D. A. Kosower and D. Maitre, *Precise Predictions for $Z + 4$ Jets at Hadron Colliders*, *Phys. Rev.* **D85** (2012) 031501, [1108.2229].
- [147] Z. Bern, L. J. Dixon, F. Febres Cordero, S. Höche, H. Ita, D. A. Kosower et al., *Next-to-Leading Order $W + 5$ -Jet Production at the LHC*, *Phys. Rev.* **D88** (2013) 014025, [1304.1253].
- [148] F. R. Anger, F. Febres Cordero, S. Höche and D. Maître, *Weak vector boson production with many jets at the LHC $\sqrt{s} = 13$ TeV*, *Phys. Rev.* **D97** (2018) 096010, [1712.08621].
- [149] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, *JHEP* **12** (2012) 061, [1209.6215].
- [150] L. Lönnblad and S. Prestel, *Merging Multi-leg NLO Matrix Elements with Parton Showers*, *JHEP* **03** (2013) 166, [1211.7278].
- [151] S. Hoeche, F. Krauss, M. Schonherr and F. Siegert, *QCD matrix elements + parton showers: The NLO case*, *JHEP* **04** (2013) 027, [1207.5030].

- [152] S. Höche, F. Krauss, M. Schönherr and F. Siegert, *W+ n-Jet predictions at the Large Hadron Collider at next-to-leading order matched with a parton shower*, *Phys. Rev. Lett.* **110** (2013) 052001, [1201.5882].
- [153] J. R. Andersen, T. Hapola and J. M. Smillie, *W Plus Multiple Jets at the LHC with High Energy Jets*, *JHEP* **09** (2012) 047, [1206.6763].
- [154] J. R. Andersen, J. J. Medley and J. M. Smillie, *Z/ γ^* plus multiple hard jets in high energy collisions*, *JHEP* **05** (2016) 136, [1603.05460].
- [155] C. M. Carloni Calame, G. Montagna, O. Nicrosini and M. Treccani, *Higher order QED corrections to W boson mass determination at hadron colliders*, *Phys. Rev.* **D69** (2004) 037301, [hep-ph/0303102].
- [156] C. M. Carloni Calame, G. Montagna, O. Nicrosini and M. Treccani, *Multiple photon corrections to the neutral-current Drell-Yan process*, *JHEP* **05** (2005) 019, [hep-ph/0502218].
- [157] C. M. Carloni Calame, G. Montagna, O. Nicrosini and A. Vicini, *Precision electroweak calculation of the charged current Drell-Yan process*, *JHEP* **12** (2006) 016, [hep-ph/0609170].
- [158] C. M. Carloni Calame, G. Montagna, O. Nicrosini and A. Vicini, *Precision electroweak calculation of the production of a high transverse-momentum lepton pair at hadron colliders*, *JHEP* **10** (2007) 109, [0710.1722].
- [159] P. Golonka and Z. Was, *PHOTOS Monte Carlo: A Precision tool for QED corrections in Z and W decays*, *Eur. Phys. J.* **C45** (2006) 97–107, [hep-ph/0506026].
- [160] L. Barze, G. Montagna, P. Nason, O. Nicrosini and F. Piccinini, *Implementation of electroweak corrections in the POWHEG BOX: single W production*, *JHEP* **04** (2012) 037, [1202.0465].

- [161] L. Barze, G. Montagna, P. Nason, O. Nicrosini, F. Piccinini and A. Vicini, *Neutral current Drell-Yan with combined QCD and electroweak corrections in the POWHEG BOX*, *Eur. Phys. J.* **C73** (2013) 2474, [1302.4606].
- [162] C. Bernaciak and D. Wackeroth, *Combining NLO QCD and Electroweak Radiative Corrections to W boson Production at Hadron Colliders in the POWHEG Framework*, *Phys. Rev.* **D85** (2012) 093003, [1201.4804].
- [163] S. Dittmaier and M. Krämer, *Electroweak radiative corrections to W boson production at hadron colliders*, *Phys. Rev.* **D65** (2002) 073007, [hep-ph/0109062].
- [164] S. Dittmaier and M. Huber, *Radiative corrections to the neutral-current Drell-Yan process in the Standard Model and its minimal supersymmetric extension*, *JHEP* **01** (2010) 060, [0911.2329].
- [165] A. Arbuzov, D. Bardin, S. Bondarenko, P. Christova, L. Kalinovskaya, G. Nanava et al., *One-loop corrections to the Drell-Yan process in SANC. I. The Charged current case*, *Eur. Phys. J.* **C46** (2006) 407–412, [hep-ph/0506110].
- [166] A. Arbuzov, D. Bardin, S. Bondarenko, P. Christova, L. Kalinovskaya, G. Nanava et al., *One-loop corrections to the Drell-Yan process in SANC. (II). The Neutral current case*, *Eur. Phys. J.* **C54** (2008) 451–460, [0711.0625].
- [167] W. Placzek and S. Jadach, *Multiphoton radiation in leptonic W boson decays*, *Eur. Phys. J.* **C29** (2003) 325–339, [hep-ph/0302065].
- [168] W. Płaczek, S. Jadach and M. W. Krasny, *Drell-Yan processes with WINHAC*, *Acta Phys. Polon.* **B44** (2013) 2171–2178, [1310.5994].

- [169] U. Baur, S. Keller and D. Wackeroth, *Electroweak radiative corrections to W boson production in hadronic collisions*, *Phys. Rev.* **D59** (1999) 013002, [hep-ph/9807417].
- [170] U. Baur, O. Brein, W. Hollik, C. Schappacher and D. Wackeroth, *Electroweak radiative corrections to neutral current Drell-Yan processes at hadron colliders*, *Phys. Rev.* **D65** (2002) 033007, [hep-ph/0108274].
- [171] U. Baur and D. Wackeroth, *Electroweak radiative corrections to $p\bar{p} \rightarrow W^\pm \rightarrow \ell^\pm \nu$ beyond the pole approximation*, *Phys. Rev.* **D70** (2004) 073015, [hep-ph/0405191].
- [172] S. Alioli et al., *Precision studies of observables in $pp \rightarrow W \rightarrow l\nu_l$ and $pp \rightarrow \gamma, Z \rightarrow l^+l^-$ processes at the LHC*, *Eur. Phys. J.* **C77** (2017) 280, [1606.02330].
- [173] S. Kallweit, J. M. Lindert, P. Maierhöfer, S. Pozzorini and M. Schönherr, *NLO electroweak automation and precise predictions for W+multijet production at the LHC*, *JHEP* **04** (2015) 012, [1412.5157].
- [174] B. Biedermann, S. Bräuer, A. Denner, M. Pellen, S. Schumann and J. M. Thompson, *Automation of NLO QCD and EW corrections with Sherpa and Recola*, *Eur. Phys. J.* **C77** (2017) 492, [1704.05783].
- [175] R. Frederix, S. Frixione, V. Hirschi, D. Pagani, H. S. Shao and M. Zaro, *The automation of next-to-leading order electroweak calculations*, *JHEP* **07** (2018) 185, [1804.10017].
- [176] S. Dittmaier, A. Huss and C. Schwinn, *Mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell-Yan processes in the resonance region: pole approximation and non-factorizable corrections*, *Nucl. Phys.* **B885** (2014) 318–372, [1403.3216].

- [177] S. Dittmaier, A. Huss and C. Schwinn, *Dominant mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell–Yan processes in the resonance region*, *Nucl. Phys. B* **904** (2016) 216–252, [[1511.08016](#)].
- [178] R. Boughezal, A. Guffanti, F. Petriello and M. Ubiali, *The impact of the LHC Z-boson transverse momentum data on PDF determinations*, *JHEP* **07** (2017) 130, [[1705.00343](#)].
- [179] CDF collaboration, T. Aaltonen et al., *Transverse momentum cross section of e^+e^- pairs in the Z-boson region from $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev. D* **86** (2012) 052010, [[1207.7138](#)].
- [180] DØ collaboration, B. Abbott et al., *Measurement of the shape of the transverse momentum distribution of W bosons produced in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev. Lett.* **80** (1998) 5498–5503, [[hep-ex/9803003](#)].
- [181] DØ collaboration, B. Abbott et al., *Measurement of the inclusive differential cross section for Z bosons as a function of transverse momentum in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev. D* **61** (2000) 032004, [[hep-ex/9907009](#)].
- [182] DØ collaboration, B. Abbott et al., *Differential production cross section of Z bosons as a function of transverse momentum at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev. Lett.* **84** (2000) 2792–2797, [[hep-ex/9909020](#)].
- [183] DØ collaboration, B. Abbott et al., *Differential cross section for W boson production as a function of transverse momentum in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Lett. B* **513** (2001) 292–300, [[hep-ex/0010026](#)].
- [184] DØ collaboration, V. M. Abazov et al., *Measurement of the ratio of differential cross sections for W and Z boson production as a function of transverse momentum in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Lett. B* **517** (2001) 299–308, [[hep-ex/0107012](#)].

- [185] DØ collaboration, V. M. Abazov et al., *Measurement of the shape of the boson transverse momentum distribution in $p\bar{p} \rightarrow Z/\gamma^* \rightarrow e^+e^- + X$ events produced at $\sqrt{s}=1.96$ -TeV*, *Phys. Rev. Lett.* **100** (2008) 102002, [0712.0803].
- [186] DØ collaboration, V. M. Abazov et al., *Measurement of the normalized $Z/\gamma^* \rightarrow \mu^+\mu^-$ transverse momentum distribution in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Lett.* **B693** (2010) 522–530, [1006.0618].
- [187] ATLAS collaboration, G. Aad et al., *Measurement of the Transverse Momentum Distribution of W Bosons in pp Collisions at $\sqrt{s} = 7$ TeV with the ATLAS Detector*, *Phys. Rev.* **D85** (2012) 012005, [1108.6308].
- [188] ATLAS collaboration, G. Aad et al., *Measurement of the transverse momentum and ϕ_η^* distributions of Drell–Yan lepton pairs in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Eur. Phys. J.* **C76** (2016) 291, [1512.02192].
- [189] CMS collaboration, S. Chatrchyan et al., *Measurement of the Rapidity and Transverse Momentum Distributions of Z Bosons in pp Collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev.* **D85** (2012) 032002, [1110.4973].
- [190] CMS collaboration, V. Khachatryan et al., *Measurement of the Z boson differential cross section in transverse momentum and rapidity in proton–proton collisions at 8 TeV*, *Phys. Lett.* **B749** (2015) 187–209, [1504.03511].
- [191] CMS collaboration, V. Khachatryan et al., *Measurement of the transverse momentum spectra of weak vector bosons produced in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **02** (2017) 096, [1606.05864].
- [192] LHCb collaboration, R. Aaij et al., *Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **08** (2015) 039, [1505.07024].

- [193] LHCb collaboration, R. Aaij et al., *Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **09** (2016) 136, [1607.06495].
- [194] R. Gauld, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover and A. Huss, *Precise predictions for the angular coefficients in Z-boson production at the LHC*, *JHEP* **11** (2017) 003, [1708.00008].
- [195] T. Becher, C. Lorentzen and M. D. Schwartz, *Resummation for W and Z production at large pT*, *Phys. Rev. Lett.* **108** (2012) 012001, [1106.4310].
- [196] J. H. Kühn, A. Kulesza, S. Pozzorini and M. Schulze, *One-loop weak corrections to hadronic production of Z bosons at large transverse momenta*, *Nucl. Phys.* **B727** (2005) 368–394, [hep-ph/0507178].
- [197] J. H. Kühn, A. Kulesza, S. Pozzorini and M. Schulze, *Electroweak corrections to large transverse momentum production of W bosons at the LHC*, *Phys. Lett.* **B651** (2007) 160–165, [hep-ph/0703283].
- [198] A. Denner, S. Dittmaier, T. Kasprzik and A. Mück, *Electroweak corrections to W + jet hadroproduction including leptonic W-boson decays*, *JHEP* **08** (2009) 075, [0906.1656].
- [199] A. Denner, S. Dittmaier, T. Kasprzik and A. Mück, *Electroweak corrections to dilepton + jet production at hadron colliders*, *JHEP* **06** (2011) 069, [1103.0914].
- [200] G. Parisi and R. Petronzio, *Small transverse momentum distributions in hard processes*, *Nucl. Phys.* **B154** (1979) 427–440.
- [201] J. C. Collins, D. E. Soper and G. F. Sterman, *Transverse Momentum Distribution in Drell-Yan Pair and W and Z Boson Production*, *Nucl. Phys.* **B250** (1985) 199–224.

- [202] C. Balazs and C. P. Yuan, *Soft gluon effects on lepton pairs at hadron colliders*, *Phys. Rev.* **D56** (1997) 5558–5583, [[hep-ph/9704258](#)].
- [203] F. Landry, R. Brock, P. M. Nadolsky and C. P. Yuan, *Tevatron Run-1 Z boson data and Collins-Soper-Sterman resummation formalism*, *Phys. Rev.* **D67** (2003) 073016, [[hep-ph/0212159](#)].
- [204] M. G. Echevarria, A. Idilbi and I. Scimemi, *Factorization theorem for Drell-Yan at low q_T and transverse momentum distributions On-The-Light-Cone*, *JHEP* **07** (2012) 002, [[1111.4996](#)].
- [205] P. F. Monni, E. Re and P. Torrielli, *Higgs transverse-momentum resummation in direct space*, *Phys. Rev. Lett.* **116** (2016) 242001, [[1604.02191](#)].
- [206] M. A. Ebert and F. J. Tackmann, *Resummation of transverse momentum distributions in distribution space*, *JHEP* **02** (2017) 110, [[1611.08610](#)].
- [207] W. Bizon, P. F. Monni, E. Re, L. Rottoli and P. Torrielli, *Momentum-space resummation for transverse observables and the Higgs p_\perp at $N^3LL+NNLO$* , *JHEP* **02** (2018) 108, [[1705.09127](#)].
- [208] X. Chen, T. Gehrmann, E. W. N. Glover, A. Huss, Y. Li, D. Neill et al., *Precise QCD Description of the Higgs Boson Transverse Momentum Spectrum*, *Phys. Lett.* **B788** (2019) 425–430, [[1805.00736](#)].
- [209] T. Becher and M. Hager, *Event-Based Transverse Momentum Resummation*, [[1904.08325](#)].
- [210] F. Caola, J. M. Lindert, K. Melnikov, P. F. Monni, L. Tancredi and C. Wever, *Bottom-quark effects in Higgs production at intermediate transverse momentum*, *JHEP* **09** (2018) 035, [[1804.07632](#)].
- [211] W. L. van Neerven and E. B. Zijlstra, *The $O(\alpha_s^2)$ corrected Drell-Yan K factor in the DIS and \overline{MS} scheme*, *Nucl. Phys.* **B382** (1992) 11–62.

- [212] C. Anastasiou, L. J. Dixon, K. Melnikov and F. Petriello, *Dilepton rapidity distribution in the Drell-Yan process at NNLO in QCD*, *Phys. Rev. Lett.* **91** (2003) 182002, [[hep-ph/0306192](#)].
- [213] K. Melnikov and F. Petriello, *The W boson production cross section at the LHC through $O(\alpha_s^2)$* , *Phys. Rev. Lett.* **96** (2006) 231803, [[hep-ph/0603182](#)].
- [214] S. Catani, G. Ferrera and M. Grazzini, *W boson production at hadron colliders: The lepton charge asymmetry in NNLO QCD*, *JHEP* **05** (2010) 006, [[1002.3115](#)].
- [215] R. Gavin, Y. Li, F. Petriello and S. Quackenbush, *FEWZ 2.0: A code for hadronic Z production at next-to-next-to-leading order*, *Comput. Phys. Commun.* **182** (2011) 2388–2403, [[1011.3540](#)].
- [216] G. P. Salam and J. Rojo, *A Higher Order Perturbative Parton Evolution Toolkit (HOPPET)*, *Comput. Phys. Commun.* **180** (2009) 120–156, [[0804.3755](#)].
- [217] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159–177, [[1410.3012](#)].
- [218] ATLAS collaboration, G. Aad et al., *Measurement of the Z/ γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **09** (2014) 145, [[1406.3660](#)].
- [219] LHCb collaboration, R. Aaij et al., *Measurement of forward W and Z boson production in association with jets in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **05** (2016) 131, [[1605.00951](#)].
- [220] ATLAS collaboration, G. Aad et al., *A measurement of the ratio of the production cross sections for W and Z bosons in association with jets with the ATLAS detector*, *Eur. Phys. J.* **C74** (2014) 3168, [[1408.6510](#)].

- [221] CMS collaboration, V. Khachatryan et al., *Measurements of differential cross sections for associated production of a W boson and jets in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev.* **D95** (2017) 052002, [1610.04222].
- [222] CDF collaboration, T. A. Aaltonen et al., *Measurement of the differential cross sections for W -boson production in association with jets in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D98** (2018) 112005, [1808.02335].
- [223] CDF collaboration, T. Aaltonen et al., *Measurement of inclusive jet cross-sections in $Z/\gamma^* \rightarrow e^+e^- +$ jets production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ -TeV*, *Phys. Rev. Lett.* **100** (2008) 102001, [0711.3717].
- [224] DØ collaboration, V. M. Abazov et al., *Studies of W boson plus jets production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D88** (2013) 092001, [1302.6508].
- [225] DØ collaboration, V. M. Abazov et al., *Measurement of differential $Z/\gamma^* +$ jet + X cross sections in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ -TeV*, *Phys. Lett.* **B669** (2008) 278–286, [0808.1296].
- [226] ATLAS collaboration, G. Aad et al., *Measurements of the W production cross sections in association with jets with the ATLAS detector*, *Eur. Phys. J.* **C75** (2015) 82, [1409.8639].
- [227] ATLAS collaboration, M. Aaboud et al., *Measurement of differential cross sections and W^+/W^- cross-section ratios for W boson production in association with jets at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **05** (2018) 077, [1711.03296].
- [228] ATLAS collaboration, G. Aad et al., *Measurement of the production cross section of jets in association with a Z boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **07** (2013) 032, [1304.7098].

- [229] ATLAS collaboration, M. Aaboud et al., *Measurements of the production cross section of a Z boson in association with jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J.* **C77** (2017) 361, [1702.05725].
- [230] CMS collaboration, V. Khachatryan et al., *Differential cross section measurements for the production of a W boson in association with jets in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Phys. Lett.* **B741** (2015) 12–37, [1406.7533].
- [231] CMS collaboration, A. M. Sirunyan et al., *Measurement of the differential cross sections for the associated production of a W boson and jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev.* **D96** (2017) 072005, [1707.05979].
- [232] CMS collaboration, V. Khachatryan et al., *Measurements of jet multiplicity and differential production cross sections of Z+ jets events in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev.* **D91** (2015) 052008, [1408.3104].
- [233] CMS collaboration, V. Khachatryan et al., *Measurements of differential production cross sections for a Z boson in association with jets in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **04** (2017) 022, [1611.03844].
- [234] CMS collaboration, A. M. Sirunyan et al., *Measurement of differential cross sections for Z boson production in association with jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J.* **C78** (2018) 965, [1804.05252].
- [235] J. M. Lindert et al., *Precise predictions for V+ jets dark matter backgrounds*, *Eur. Phys. J.* **C77** (2017) 829, [1705.04664].
- [236] LHCb collaboration, R. Aaij et al., *Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **01** (2016) 155, [1511.08039].

- [237] DØ collaboration, V. M. Abazov et al., *Measurement of the electron charge asymmetry in $p\bar{p} \rightarrow W + X \rightarrow e\nu + X$ decays in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev.* **D91** (2015) 032007, [1412.2862].
- [238] S. Dulat, T.-J. Hou, J. Gao, M. Guzzi, J. Huston, P. Nadolsky et al., *New parton distribution functions from a global analysis of quantum chromodynamics*, *Phys. Rev.* **D93** (2016) 033006, [1506.07443].
- [239] S. Farry and R. Gauld, *Leptonic W^\pm boson asymmetry in association with jets at LHCb and parton distribution function constraints at large x* , *Phys. Rev.* **D93** (2016) 014008, [1505.01399].
- [240] S. Kallweit, J. M. Lindert, P. Maierhofer, S. Pozzorini and M. Schönherr, *NLO QCD+EW predictions for $V +$ jets including off-shell vector-boson decays and multijet merging*, *JHEP* **04** (2016) 021, [1511.08692].
- [241] M. Rubin, G. P. Salam and S. Sapeta, *Giant QCD K -factors beyond NLO*, *JHEP* **09** (2010) 084, [1006.2144].
- [242] ALEPH, DELPHI, L3, OPAL, SLD, LEP ELECTROWEAK WORKING GROUP, SLD ELECTROWEAK GROUP, SLD HEAVY FLAVOUR GROUP collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, *Phys. Rept.* **427** (2006) 257–454, [hep-ex/0509008].
- [243] CDF collaboration, T. A. Aaltonen et al., *Measurement of $\sin^2 \theta_{\text{eff}}^{\text{lept}}$ using e^+e^- pairs from γ^*/Z bosons produced in $p\bar{p}$ collisions at a center-of-momentum energy of 1.96 TeV*, *Phys. Rev.* **D93** (2016) 112016, [1605.02719].
- [244] DØ collaboration, V. M. Abazov et al., *Measurement of the Effective Weak Mixing Angle in $p\bar{p} \rightarrow Z/\gamma^* \rightarrow \ell^+\ell^-$ Events*, *Phys. Rev. Lett.* **120** (2018) 241802, [1710.03951].

- [245] CDF, DØ collaboration, T. A. Aaltonen et al., *Tevatron Run II combination of the effective leptonic electroweak mixing angle*, *Phys. Rev. D* **97** (2018) 112007, [1801.06283].
- [246] CMS collaboration, A. M. Sirunyan et al., *Measurement of the weak mixing angle using the forward-backward asymmetry of Drell-Yan events in pp collisions at 8 TeV*, *Eur. Phys. J. C* **78** (2018) 701, [1806.00863].
- [247] A. Sirlin, *Radiative Corrections in the $SU(2)_L \times U(1)$ Theory: A Simple Renormalization Framework*, *Phys. Rev. D* **22** (1980) 971–981.
- [248] P. Gambino and A. Sirlin, *Relation between $\sin^2 \theta_W(M_Z)$ and $\sin^2 \theta_W^{\text{eff}}$ (leptonic)*, *Phys. Rev. D* **49** (1994) 1160–1162, [hep-ph/9309326].
- [249] GFITTER collaboration, M. Baak, J. Cúth, J. Haller, A. Hoecker, R. Kogler, K. Mönig et al., *The global electroweak fit at NNLO and prospects for the LHC and ILC*, *Eur. Phys. J. C* **74** (2014) 3046, [1407.3792].
- [250] ATLAS collaboration, *Prospect for a measurement of the Weak Mixing Angle in $pp \rightarrow Z/\gamma^* \rightarrow e^+e^-$ events with the ATLAS detector at the High Luminosity Large Hadron Collider*, [ATL-PHYS-PUB-2018-037].
- [251] J. C. Collins and D. E. Soper, *Angular Distribution of Dileptons in High-Energy Hadron Collisions*, *Phys. Rev. D* **16** (1977) 2219.
- [252] CDF collaboration, T. Aaltonen et al., *Indirect Measurement of $\sin^2 \theta_W$ (M_W) Using e^+e^- Pairs in the Z-Boson Region with $p\bar{p}$ Collisions at a Center-of-Momentum Energy of 1.96 TeV*, *Phys. Rev. D* **88** (2013) 072002, [1307.0770].
- [253] CDF collaboration, T. A. Aaltonen et al., *Indirect measurement of $\sin^2 \theta_W$ (or M_W) using $\mu^+\mu^-$ pairs from γ^*/Z bosons produced in $p\bar{p}$ collisions at a center-of-momentum energy of 1.96 TeV*, *Phys. Rev. D* **89** (2014) 072005, [1402.2239].

- [254] DØ collaboration, V. M. Abazov et al., *Measurement of $\sin^2 \theta_{\text{eff}}^\ell$ and Z -light quark couplings using the forward-backward charge asymmetry in $p\bar{p} \rightarrow Z/\gamma^* \rightarrow e^+e^-$ events with $\mathcal{L} = 5.0 \text{ fb}^{-1}$ at $\sqrt{s} = 1.96 \text{ TeV}$* , *Phys. Rev. D* **84** (2011) 012007, [1104.4590].
- [255] DØ collaboration, V. M. Abazov et al., *Measurement of the Effective Weak Mixing Angle in $p\bar{p} \rightarrow Z/\gamma^* \rightarrow e^+e^-$ Events*, *Phys. Rev. Lett.* **115** (2015) 041801, [1408.5016].
- [256] E. Accomando, J. Fiaschi, F. Hautmann and S. Moretti, *Constraining Parton Distribution Functions from Neutral Current Drell-Yan Measurements*, *Phys. Rev. D* **98** (2018) 013003, [1712.06318].
- [257] E. Accomando, J. Fiaschi, F. Hautmann and S. Moretti, *Neutral current forward-backward asymmetry: from θ_W to PDF determinations*, *Eur. Phys. J. C* **78** (2018) 663, [1805.09239].
- [258] HERAFITTER collaboration, S. Camarda et al., *QCD analysis of W - and Z -boson production at Tevatron*, *Eur. Phys. J. C* **75** (2015) 458, [1503.05221].
- [259] A. Bodek, J. Han, A. Khukhunaishvili and W. Sakumoto, *Using Drell-Yan forward-backward asymmetry to reduce PDF uncertainties in the measurement of electroweak parameters*, *Eur. Phys. J. C* **76** (2016) 115, [1507.02470].
- [260] ATLAS collaboration, M. Aaboud et al., *Measurement of the Drell-Yan triple-differential cross section in pp collisions at $\sqrt{s} = 8 \text{ TeV}$* , *JHEP* **12** (2017) 059, [1710.05167].
- [261] ATLAS collaboration, G. Aad et al., *Measurement of the angular coefficients in Z -boson events using electron and muon pairs from data taken at $\sqrt{s} = 8 \text{ TeV}$ with the ATLAS detector*, *JHEP* **08** (2016) 159, [1606.00689].

- [262] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040, [[hep-ph/0409146](#)].
- [263] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070, [[0709.2092](#)].
- [264] S. Alioli, P. Nason, C. Oleari and E. Re, *NLO vector-boson production matched with shower in POWHEG*, *JHEP* **07** (2008) 060, [[0805.4802](#)].
- [265] 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, [[1002.2581](#)].
- [266] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [[0710.3820](#)].
- [267] W. F. L. Hollik, *Radiative Corrections in the Standard Model and their Role for Precision Tests of the Electroweak Theory*, *Fortsch. Phys.* **38** (1990) 165–260.
- [268] T. Carli, D. Clements, A. Cooper-Sarkar, C. Gwenlan, G. P. Salam, F. Siegert et al., *A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: The APPLGRID Project*, *Eur. Phys. J.* **C66** (2010) 503–524, [[0911.2985](#)].
- [269] FASTNLO collaboration, D. Britzger, K. Rabbertz, F. Stober and M. Wobisch, *New features in version 2 of the fastNLO project*, in *Proceedings, 20th International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS 2012)*, pp. 217–221, 2012, [[1208.3641](#)].
- [270] M. Czakon, D. Heymes and A. Mitov, *fastNLO tables for NNLO top-quark pair differential distributions*, [[1704.08551](#)].

- [271] D. Britzger et al., *Calculations for deep inelastic scattering using fast interpolation grid techniques at NNLO in QCD and the extraction of α_s from HERA data*, [1906.05303].
- [272] D. Yu. Bardin, A. Leike, T. Riemann and M. Sachwitz, *Energy Dependent Width Effects in e^+e^- Annihilation Near the Z Boson Pole*, *Phys. Lett.* **B206** (1988) 539–542.
- [273] A. Sirlin, *Theoretical considerations concerning the Z^0 mass*, *Phys. Rev. Lett.* **67** (1991) 2127–2130.
- [274] H1 collaboration, F. D. Aaron et al., *Inclusive Deep Inelastic Scattering at High Q^2 with Longitudinally Polarised Lepton Beams at HERA*, *JHEP* **09** (2012) 061, [1206.7007].
- [275] ZEUS collaboration, S. Chekanov et al., *Measurement of charged current deep inelastic scattering cross sections with a longitudinally polarised electron beam at HERA*, *Eur. Phys. J.* **C61** (2009) 223–235, [0812.4620].
- [276] ZEUS collaboration, H. Abramowicz et al., *Measurement of high- Q^2 charged current deep inelastic scattering cross sections with a longitudinally polarised positron beam at HERA*, *Eur. Phys. J.* **C70** (2010) 945–963, [1008.3493].
- [277] ZEUS collaboration, H. Abramowicz et al., *Combined QCD and electroweak analysis of HERA data*, *Phys. Rev.* **D93** (2016) 092002, [1603.09628].
- [278] H1 collaboration, V. Andreev et al., *Determination of electroweak parameters in polarised deep-inelastic scattering at HERA*, *Eur. Phys. J.* **C78** (2018) 777, [1806.01176].
- [279] ZEUS collaboration, S. Chekanov et al., *Multi-jet cross-sections in charged current $e^\pm p$ scattering at HERA*, *Phys. Rev.* **D78** (2008) 032004, [0802.3955].

- [280] H1 collaboration, C. Adloff et al., *Dijet production in charged and neutral current $e^+ p$ interactions at high Q^2* , *Eur. Phys. J.* **C19** (2001) 429–440, [hep-ex/0010016].
- [281] LHeC Study Group, *A Large Hadron Electron Collider at CERN: Report on the Physics and Design Concepts for Machine and Detector*, *J. Phys.* **G39** (2012) 075001, [1206.2913].
- [282] NuTeV collaboration, M. Tzanov et al., *Precise measurement of neutrino and anti-neutrino differential cross sections*, *Phys. Rev.* **D74** (2006) 012008, [hep-ex/0509010].
- [283] CHORUS collaboration, G. Onengut et al., *Measurement of nucleon structure functions in neutrino scattering*, *Phys. Lett.* **B632** (2006) 65–75.
- [284] W. G. Seligman et al., *Improved determination of $\alpha(s)$ from neutrino nucleon scattering*, *Phys. Rev. Lett.* **79** (1997) 1213–1216, [hep-ex/9701017].
- [285] J. P. Berge et al., *A Measurement of Differential Cross-Sections and Nucleon Structure Functions in Charged Current Neutrino Interactions on Iron*, *Z. Phys.* **C49** (1991) 187–224.
- [286] DUNE collaboration, R. Acciarri et al., *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE)*, [1512.06148].
- [287] E. Mirkes and D. Zeppenfeld, *Dijet production at HERA in next-to-leading order*, *Phys. Lett.* **B380** (1996) 205–212, [hep-ph/9511448].
- [288] S. Moch, J. A. M. Vermaseren and A. Vogt, *The Three loop splitting functions in QCD: The Nonsinglet case*, *Nucl. Phys.* **B688** (2004) 101–134, [hep-ph/0403192].

- [289] A. Vogt, S. Moch and J. A. M. Vermaseren, *The Three-loop splitting functions in QCD: The Singlet case*, *Nucl. Phys.* **B691** (2004) 129–181, [hep-ph/0404111].
- [290] S. Moch, J. A. M. Vermaseren and A. Vogt, *The Longitudinal structure function at the third order*, *Phys. Lett.* **B606** (2005) 123–129, [hep-ph/0411112].
- [291] J. A. M. Vermaseren, A. Vogt and S. Moch, *The Third-order QCD corrections to deep-inelastic scattering by photon exchange*, *Nucl. Phys.* **B724** (2005) 3–182, [hep-ph/0504242].
- [292] S. Moch, M. Rogal and A. Vogt, *Differences between charged-current coefficient functions*, *Nucl. Phys.* **B790** (2008) 317–335, [0708.3731].
- [293] S. Moch, J. A. M. Vermaseren and A. Vogt, *Third-order QCD corrections to the charged-current structure function F_3* , *Nucl. Phys.* **B813** (2009) 220–258, [0812.4168].
- [294] J. Davies, A. Vogt, S. Moch and J. A. M. Vermaseren, *Non-singlet coefficient functions for charged-current deep-inelastic scattering to the third order in QCD*, *PoS DIS2016* (2016) 059, [1606.08907].
- [295] T. Han, G. Valencia and S. Willenbrock, *Structure function approach to vector boson scattering in $p p$ collisions*, *Phys. Rev. Lett.* **69** (1992) 3274–3277, [hep-ph/9206246].
- [296] F. A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Production at Three Loops in QCD*, *Phys. Rev. Lett.* **117** (2016) 072001, [1606.00840].
- [297] F. A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Pair Production at N^3LO* , *Phys. Rev.* **D98** (2018) 114016, [1811.07906].
- [298] C. Duprel, T. Hadig, N. Kauer and M. Wobisch, *Comparison of next-to-leading order calculations for jet cross-sections in deep inelastic*

- scattering*, in *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*, pp. 142–150, 1999, [hep-ph/9910448].
- [299] R. Mondini, M. Schiavi and C. Williams, *N^3LO predictions for the decay of the Higgs boson to bottom quarks*, [1904.08960].
- [300] J. Currie, T. Gehrmann, E. W. N. Glover, A. Huss, J. Niehues and A. Vogt, *N^3LO Corrections to Jet Production in Deep Inelastic Scattering using the Projection-to-Born Method*, [1803.09973].
- [301] J. Davies, A. Vogt, B. Ruijl, T. Ueda and J. A. M. Vermaseren, *Large- n_f contributions to the four-loop splitting functions in QCD*, *Nucl. Phys.* **B915** (2017) 335–362, [1610.07477].
- [302] S. Moch, B. Ruijl, T. Ueda, J. A. M. Vermaseren and A. Vogt, *Four-Loop Non-Singlet Splitting Functions in the Planar Limit and Beyond*, *JHEP* **10** (2017) 041, [1707.08315].
- [303] A. Vogt, F. Herzog, S. Moch, B. Ruijl, T. Ueda and J. A. M. Vermaseren, *Anomalous dimensions and splitting functions beyond the next-to-next-to-leading order*, *PoS* **LL2018** (2018) 050, [1808.08981].
- [304] H1, ZEUS collaboration, A. F. Zarnecki, *Electroweak Results from HERA*, in *Proceedings, 48th Rencontres de Moriond on Electroweak Interactions and Unified Theories: La Thuile, Italy, March 2-9, 2013*, pp. 489–496, 2013, [1305.4918].
- [305] T.-J. Hou, S. Dulat, J. Gao, M. Guzzi, J. Huston, P. Nadolsky et al., *CTEQ-TEA parton distribution functions and HERA Run I and II combined data*, *Phys. Rev.* **D95** (2017) 034003, [1609.07968].
- [306] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer 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, [1405.0301].
- [307] F. Cascioli, P. Maierhofer and S. Pozzorini, *Scattering Amplitudes with Open Loops*, *Phys. Rev. Lett.* **108** (2012) 111601, [1111.5206].
- [308] R. Kleiss, W. J. Stirling and S. D. Ellis, *A New Monte Carlo Treatment of Multiparticle Phase Space at High-energies*, *Comput. Phys. Commun.* **40** (1986) 359.
- [309] Maplesoft, *Maple 2016*. Waterloo Maple Inc., Waterloo ON, Canada, 2016.
- [310] J. A. M. Vermaseren, *New features of FORM*, [math-ph/0010025].
- [311] J. M. Campbell, R. K. Ellis and S. Seth, *H+1 jet production revisited*, [1906.01020].
- [312] L. White and H. K. N., *Regression Testing*. Encyclopedia of Software Engineering, 2002. 10.1002/0471028959.sof279.
- [313] A. Karlberg, *At the Frontier of Precision QCD in the LHC Era*, Ph.D. thesis, Merton Coll., Oxford, 2016. 1610.06226.
- [314] A. J. Buras, *Asymptotic Freedom in Deep Inelastic Processes in the Leading Order and Beyond*, *Rev. Mod. Phys.* **52** (1980) 199.
- [315] W. L. van Neerven and A. Vogt, *NNLO evolution of deep inelastic structure functions: The Singlet case*, *Nucl. Phys.* **B588** (2000) 345–373, [hep-ph/0006154].
- [316] W. L. van Neerven and A. Vogt, *NNLO evolution of deep inelastic structure functions: The Nonsinglet case*, *Nucl. Phys.* **B568** (2000) 263–286, [hep-ph/9907472].
- [317] V. N. Gribov and L. N. Lipatov, *Deep inelastic e p scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438–450.

-
- [318] L. N. Lipatov, *The parton model and perturbation theory*, *Sov. J. Nucl. Phys.* **20** (1975) 94–102.
- [319] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, *Nucl. Phys.* **B126** (1977) 298–318.
- [320] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.*, *Sov. Phys. JETP* **46** (1977) 641–653.
- [321] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and J. Pires, *Infrared sensitivity of single jet inclusive production at hadron colliders*, *JHEP* **10** (2018) 155, [1807.03692].