Photoproduction Processes as a Probe of the Strong Nuclear Force in Relativistic Heavy Ion Collisions at the LHC

Benjamin Gilbert

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2024

© 2024

Benjamin Gilbert

All Rights Reserved

Abstract

Photoproduction Processes as a Probe of the Strong Nuclear Force in Relativistic Heavy Ion Collisions at the LHC

Benjamin Gilbert

The intense electromagnetic fields surrounding highly charged, ultra-relativistic ions at the LHC provide an intense flux of quasi-real photons. This high-energy photon flux provides a precise probe of the structure of the nucleus in photonuclear interactions and a tool to precisely study the properties of photoproduction in two-photon interactions. This thesis presents recent results from studies of photoproduction processes in $\sqrt{s_{NN}}$ = 5•02 TeV Pb+Pb collisions at the LHC, using data collected in 2015 and 2018. A measurement of the differential cross-section for photonuclear jet production is performed with an integrated luminosity of 1•72 nb ^{1}. This measurement is triple-differential in two different sets of kinematic variables ([$_{T}$, G_{A}, I_{W}] and [$_{T}$, H_{jets}, $<_{jets}$]) which allow for a direct measurement of the kinematics of struck partons in the nuclear target, and it is fully unfolded in three dimensions. After performing detailed studies of the rate for nuclear breakup in these collisions, the photonuclear jet cross-sections agree with leading-order predictions at the level of 10%. The uncertainty on this measurement and full treatment of its correlated uncertainties will allow for it to significantly constrain the nuclear parton distributions over a wide region of parton kinematics. A measurement was also performed of dimuon production via two-photon fusion in Pb+Pb collisions with nuclear overlap. This measurement studied the relative deflection of muons in each pair, and its results indicated that these distributions are substantially modified as a function of collision centrality. This modification is mostly well-described by initial-state models, which improve the modelling of the photon flux by incorporating correlations between the collision impact parameter and photon $?_{T}$. Additional studies were performed to constrain the potential impact of any deflections due to the presence of strong magnetic fields in the hot, dense medium produced in heavy ion collisions.

Table of Contents

Acknowledgments xxxvi

Dedication xxxviii

Chapter 1: Introduction 1

Chapter 2: Background 4

2.1 Quantum Chromodynamics 4

2.1.1 The QCD Lagrangian and Basic Formulation 4

2.1.1.1 The Feynman Rules of Perturbative QCD 8

2.1.1.2 The QCD Beta Function 10

2.1.1.3 Lattice QCD 13

2.1.2 Confinement and Asymptotic Freedom 16

2.1.3 QCD Factorization and the Parton Model 17

2.1.4 Deep Inelastic Scattering 19

2.1.5 Parton Distribution Functions and DGLAP Evolution 22

2.1.6 Diffraction in QCD 28

2.2 Descriptions of Nuclear Structure 28

2.2.1 Nuclear Modification of the Parton Distribution Functions 28

2.2.1.1 The EMC Effect 30

2.2.1.2 Nuclear Shadowing and Anti-Shadowing 31
2.2.1.3 Global Fits to Determine nPDFs 34
2.2.2 Gluon Saturation at Low-G 39
2.3 Hadronic Particle Production 40
2.3.0.1 Fragmentation Functions 40
2.3.1 Parton Showers: Fundamentals and Models 41
2.3.1.1 Initial-State Parton Showers 44
2.3.2 Models of Hadronization 45
2.3.2.1 The Lund String Model 45
2.3.2.2 Cluster Hadronization 47
2.4 Properties of Relativistic Heavy Ion Collisions 47
2.4.1 Collision Centrality and Geometry 48
2.4.2 The Quark-Gluon Plasma 52
2.5 Photoproduction and Ultra-Peripheral Collisions 53
2.5.1 The Equivalent Photon Approximation 55
2.5.2 The Structure of Quasi-Real Photons 58
2.5.3 Particle Production in Ultra-Peripheral Collisions 61
2.5.3.1 Rapidity Gaps 61
2.5.3.2 Electromagnetic Dissociation and Nuclear Break-up 63
2.5.3.3 Photonuclear Scattering 68
2.5.3.4 Two-Photon Processes 71
2.5.3.5 Photo-diffractive Scattering 77
2.6 Monte Carlo Generators and Models 81

2.6.1 Leading Order Generators 81
2.6.1.1 Pythia 81
2.6.1.2 HERWIG 83
2.6.2 Next to Leading Order Generators 84
2.6.2.1 POWHEG 84
2.6.2.2 Sherpa 86
2.6.3 STARlight 86

Chapter 3: Experimental Apparatus and Techniques 87
3.1 Particle Collisions at the Large Hadron Collider 87
3.1.1 The LHC Injection Chain 87
3.1.2 The LHC Ring, Magnets, and Optics 88
3.1.3 Luminosity Measurement and Calibration 89
3.1.4 Luminosity Determination Technique 90
3.2 The ATLAS Detector 95
3.2.1 The ATLAS Coordinate System 95
3.2.2 The Magnet System and Beams 96
3.2.3 The ATLAS Inner Detector 96
3.2.4 Calorimeter Systems 98
3.2.4.1 Fundamentals of Particle Passage Through Matter 99
3.2.4.2 The Liquid Argon Calorimeter 103
3.2.4.3 The Tile Hadronic Calorimeter 108
3.2.4.4 Calorimeter Readout 110

3.2.5 The Muon Spectrometer 111

3.2.6 Forward Detectors 113

3.2.6.1 The Zero-Degree Calorimeter (ZDC) 113

3.2.6.2 The Luminosity Cerenkov Integrating Detector (LUCID) 114

3.2.7 The ATLAS Trigger System 115

3.3 Event and Object Reconstruction 118

3.3.1 Charged Particle Tracks 118

3.3.2 Calorimeter Clusters 120

3.3.3 Combined Muons, Electrons, and Photons 124

3.3.4 Particle-Flow Objects 125

Chapter 4: Jet De nition, Reconstruction, and Calibration 130

4.1 Jet Algorithms in QCD 130

4.1.1 Requirements of Formal Jet De nitions 130

4.1.2 Top-Down and Bottom-up Jet Finders 131

4.1.3 The Anti-k_t Jet Clustering Algorithm 133

4.2 Calorimetric Jet Constituents 134

4.3 Jet Reconstruction and Calibration 136

4.3.1 MC-Based Calibrations 137

4.3.1.1 Eta-JES Calibration 137

4.3.1.2 Global Sequential Calibration 138

4.3.2 In-situ JES Calibration 141

4.3.2.1 Eta-Intercalibration 142

4.3.2.2 Absolute Energy Scale Measurement 147
4.3.3 MC-Derived UPC Correction 151
4.3.4 Jet Energy Resolution 152
4.3.4.1 MC-Based Measurement by Truth Matching 153
4.3.4.2 In-situ Measurement of the JER 153

Chapter 5: Photonuclear Jet Production 161
5.1 Datasets and Reconstruction 161
5.1.1 Datasets and Trigger Strategy 161
5.1.2 MC Sample Generation with PYTHIA 8 162
5.1.3 Vertex, Track, and Topocluster Requirements 164
5.2 Gap Analysis 165
5.3 Template Fits of Gap Distributions 170
5.3.1 0nXnTemplate Fits 171
5.3.2 XnXnTemplate Fits 177
5.4 ZDC Analysis, Calibration, and Trigger Ef ciency 177
5.5 Trigger Ef ciencies 180
5.5.1 Jet Trigger Ef ciencies 180
5.5.2 L1 Trigger Ef ciencies 184
5.5.3 Electromagnetic Pileup Correction 184
5.6 Measuring Nuclear Break-up 187
5.7 Data Distribution Consistency Checks 190
5.7.1 Data Consistency Checks by Detector Side 190

5.7.2 Data-Monte Carlo Comparisons 192

5.8 Studies of Geometric Acceptance 195

5.9 Reconstruction Ef ciencies 201

5.10 Summary of Analysis Before Unfolding 202

5.10.1 Analysis Cut-Flow Overview and Full Sample Requirements 202

5.10.2 Cross-Sections Before Unfolding 205

5.11 Unfolding 208

5.11.1 Resolution Prior to Unfolding 208

5.11.2 Response Matrices and Binning 210

5.11.3 Iteration Dependence 213

5.11.4 Impact of Unfolding on Results 216

5.11.5 Unfolding Closure Tests 216

5.12 Systematic Uncertainties 221

5.12.1 Luminosity 221

5.12.2 EM Pileup Correction 222

5.12.3 Event Selection Ef ciency 222

5.12.4 Sensitivity to the Unfolding Prior 227

5.12.5 Jet Response 231

5.12.6 Full Systematics Breakdowns 242

5.12.7 Correlation in Systematic Uncertainties 249

Chapter 6: Muon Photoproduction in Non-UPC Events. 250

6.0.1 Datasets and Reconstruction 250

6.0.2 Event Selections 251

6.0.2.1 Event and Muon-pair Selections 251

6.0.2.2 Centrality De nition 254

6.0.3 Trigger and Reconstruction Ef ciencies 257

6.0.4 Background Subtraction 259

6.0.4.1 Heavy Flavor Backgrounds 259

6.0.4.2 Comparing Background Subtraction Methods 263

6.0.4.3 Drell-Yan Backgrounds 266

6.0.4.4 Excess observed at large U and $:_?$ 269

6.0.4.5 Distributions after Subtraction 271

6.0.5 Unfolding 276

6.0.6 De nition of Observables 277

6.0.7 Systematic Uncertainties 278

Chapter 7: Results 282

7.1 Photonuclear Jet Production 282

7.1.1 Measuring Nuclear Breakup 282

7.1.2 Sensitivity to Nuclear PDF Effects 284

7.1.3 Photonuclear Jet Cross-Sections 284

7.2 Muon Photoproduction in Non-UPC Events 301

7.2.1 Theoretical calculations 301

7.2.2 Cross-sections and relative yield measurements 301

7.2.3 U and $:_?$ distributions 303

7.2.4 Characterizing the centrality dependence of v_2 distributions 308

7.2.5 Rapidity difference and event-plane dependence of v_2. 310

Chapter 8: Conclusion 313

References 319

Appendix A: Unfolding 343

Appendix B: Non-Physics Backgrounds to UPC Jets 346

B.1 Upstream Splash Events 346

B.2 Beam Halo Muons 348

B.3 Out-of-Time Pile-up 351

Appendix C: Fiducial and Geometric Acceptance for UPC Jets 354

List of Figures

2.1 The $\&$-dependence of the QCD coupling strength μ_B, from Ref. [9]. The 5-loop calculation of the QCD ∇ function, assuming $\mu_B{}^{1}$" / º = 0•1180 0•0009, is compared to world data for independent measurements of μ_B. 14

2.2 A two-dimensional representation of the eld structure of lattice QCD calculations, from Ref. [9]. Quark elds lie on gauge sites, and the gauge elds (lie on the links between sites. 15

2.3 A diagram showing the stretching and breaking of a ux tube between two quarks as they move apart. 17

2.4 A diagram of deep inelastic electron-nucleon scattering for an electron with incoming four-momentum and outgoing four-momentum η^0, momentum transfer @ and incoming nucleon four-momentum $\%$, from Ref. [9] 20

2.5 Feynman diagrams of the different parton splittings which give rise to the four leading-order DGLAP splitting functions 24

2.6 The kinematic coverage of world data (left) to determine the free proton PDFs and world data on DIS measurements of the free proton structure function $\mathcal{F}_2^p(\eta, \&^2)$, both from Ref. [9] 26

2.7 The free proton PDFs determined at NNLO in the MSHT20 t to world data from Ref. [53]. The dependence on G is shown for two scales $\&^2$ = 10 GeV^2 and $\&^2 = 10^4$ GeV^2, the former of which relies more on xed target and HERA data, while the latter is more sensitive to Tevatron and LHC data. 27

2.8 The rough shape of nPDF modi cations as a function of G for xed $\&^2$ from Ref. [69]. The t parameters for the EPPS21 nPDF t are shown, in addition to the nuclear modi cation function. 30

2.9 The correlation between the slope of the structure functions in the EMC region (0•35 G 0•7) and the SRC scaling factor $a_2{}^{1}$ •3º, from Ref. [73] 31

2.10 The different diagrams which contribute to the leading twist calculation of nuclear shadowing, from Ref. [77]. Diagrams (a) is the cross-section from the impulse approximation, while diagrams (b) and (c) yield the shadowing correction due to coherent multiple scattering effects. 33

2.11 A recent comparison of the results from three leading global nPDF ts as a function of G, for &^{2} = 10 GeV^2., from Ref. [79] 38

2.12 Left: The spacetime picture of string fragmentation in 1+1 dimensions in the Lund string model. Right: Formation of hadrons in the string fragmentation model by "popcorn production". Both gures are reproduced from Ref. [98]. 46

2.13 A diagram of a hadronic heavy ion collision from a side view (left) and front view (right). The impact parameter of the nucleus-nucleus collision is (B) is different from the impact parameter of the NN collision with respect to the nucleus A (B) or (nucleus B (B 1). Reproduced from Ref. [116]. 50

2.14 The relationship between different collision geometry and activity properties in Glauber MC from Ref. [116]. 52

2.15 A side-view diagram of two colliding ions in an Ultra-Peripheral Collision from Ref. [1] (left) and a transverse view from Ref. [140] (right). The right diagram demonstrates the directions of electric and magnetic elds, along with the resulting polarization vectors (S). 55

2.16 The photon PDFs from leading NLO ts for the up quark (left) and the gluon (right) from Ref. [148]. While the gluon PDF falls rapidly as G 0, the up quark PDF has a maximum near 0•6 0•7, demonstrating the valence structure of the photon. While the models have some substantial disagreements in scale, the shapes are reasonably consistent. 61

2.17 Feynman diagrams showing the photoproduction of a dijet pair in direct (left) or resolved (right) W A UPCs. In the case of the direct photon, no particles are produced between the photon-going nucleus and the hard-scattering, while the resolved photon remnant partially lls the gap in that case. 62

2.18 A diagram showing an example photonuclear jet production event the gap quantities relevant to inclusive photonuclear measurements. The blue circles are jets, and the black circles are particles produced in the event. The edge gaps are the darker shaded intervals, and both green intervals compose . Figure by A. Angerami. . 63

2.19 A diagram of a photonuclear UPC from a side view (left) and front view (right). The impact parameter of the nucleus-nucleus collision is (B) is different from the impact parameter of the W A collision with respect to the photon-emitting nucleus (B) or the struck nucleus (B). Inspired by Figure 3 from Ref. [116]. 66

2.20 Left: The different probabilities relevant for calculating the nuclear breakup probability as a function impact parameter. Right: The photon uxes which arise from modelling W A ! jets with or without the full calculation of the breakup probability. The sum of the 0nXn and 0n0n contributions is is de ned to sum to $\%_{no\ had}$ as does the sum of $\%_{EM}$ and $\%_{no\ EM}$. Figures by A. Angerami. 67

2.21 Projected reductions in nPDF uncertainties for the gluon PDF resulting from the inclusion of data on photonuclear jet production, from Ref. [158]. The left panel shows the projected uncertainty reduction if all point-to-point errors are 5% of the cross-section, while the right panel shows the reduction for 10% errors. In both cases, no knowledge of the correlation in the uncertainties is assumed. 70

2.22 Feynman diagrams from Refs. [2, 162] showing the different processes where leptons are produced via two-photon fusion in UPCs. The right-most and center diagrams show leading-order production mechanisms for electron and muon pairs, while the latter shows the case of nal-state radiation of a photon by one lepton. . 72

2.23 Left: The core+tail t of dimuon acoplanarity distributions by the CMS collaboration from Ref. [5]. Right: The coherent+dissociative photoproduction t of di-electron acoplanarity distributions by the ATLAS collaboration from Ref. [162] . 73

2.24 Left: The dependence of di-electron production on the angle of a single electron with respect to the pair measured by the STAR collaboration. Right: The STAR measurement of the different $\%_{\overline{q^2}}$ distributions for UPC and peripheral Au+Au collisions Both gures are reproduced from Ref. [165]. 74

2.25 The oscillations in the photon Wigner distribution as a function of transverse momentum for different values of impact parameter from Ref. [170]. 77

2.26 Feynman diagrams of different two-photon fusion processes sensitive to physics beyond the Standard Model from Refs. [175, 178]. 78

2.27 Left: Constraints on Λ_g set by ATLAS measurements of two-photon production of g, g pairs compared to existing limits from world data, from Ref. [175]. Right: World limits on the existence of axion-like particles where ATLAS and CMS measurements of light-by-light scattering contribute to a substantial region of mass vs. coupling, from Ref. [178]. 78

2.28 A Feynman diagram for direct dijet photoproduction (left) and an NLO QCD prediction for the cross-section of photo-diffractive jet production as a functi on of the jet system rapidity (right). Both gures are reproduced from Ref. [196]. 80

2.29 Estimates of the fraction of photonuclear jet events which come from photo-diffractive scattering as a function of G_A from Ref. [200]. The predictions vary substantially with the choice of factorization-breaking correction, but the contribution is limited to the very small G_A region. 81

2.30 The schematic event structure of a ! C event in PYTHIA 8 from Ref. [98]. The generator simulates the process with a hardness-ordering, moving from the hard-scattering matrix element in the center out to fully hadronized outgoing particles. . 82

3.1 The CERN accelerator complex and LHC injection chain from Ref. [217]. The injection chain for heavy ion beams includes the LINAC 3, LEIR, PS, SPS, and LHC. 88

3.2 A vdM scan curve from BCID 502 of the 2015 Pb+Pb run. The two different curves correspond to the C and H scans, respectively, where the beam displacement is held constant at 0 for the other scan. 92

3.3 The ratio of luminosity measurements from different sources for each run of the 2015 Pb+Pb data-taking period. The ratios are xed (anchored) to agree in the vdM scan run, and the assigned normalization uncertainty is shown as the dashed black lines. 93

3.4 Vertical (left) and three-dimensional (right) cut-out images of the ATLAS inner-detector subsystems from Ref. [227] 97

3.5 A cut-out image of the ATLAS calorimeter subsystems from Ref. [227] 98

3.6 The different sources of electron or positron (left) and photon (right) energy loss in matter as a function of particle energy from Ref. [9]. In the right gure, $\hat{}_{nuc}$ and $\hat{}_e$ refer to energy loss in a nuclear or electron eld, respectively. 102

3.7 The values of $\hat{}_0$ and _ in different materials as a function of atomic charge number from Ref. [9] 104

3.8 The ATLAS EM calorimeter depth, in radiation lengths, as a function of $|\eta|$ from Ref. [236] 105

3.9 Cutout diagrams of the LAr EM Barrel calorimeter accordion design (left) and HEC honeycomb layering and readout design (right) from Ref. [236] 106

3.10 The ATLAS calorimeter depth, in hadronic interaction lengths, as a function of $|\eta|$ from Ref. [236] 107

3.11 Cutout diagrams of a Tile hadronic calorimeter module (left) and the LAr Forward Calorimeter (right) from Refs. [237] and [236] 109

3.12 Cutout diagrams of the muon spectrometer, demonstrating the different chambers in the barrel and endcap from Ref. [241] 112

3.13 Diagram of the relative positioning of the ATLAS detector, ZDCs, and LHC beams from Ref. [243] 114

3.14 A schematic description of the ATLAS trigger and Data Acquisition system after Phase-I upgrades between Run 1 and Run 2 from Ref. [247]. 116

3.15 The performance of ATLAS track reconstruction inside R = 0.4 Jets from Ref. [252]. The comparison of the ' distribution between data and MC (left) shows the precision with which MC can model these effects, especially in low p_T jets. The reconstruction ef ciency as a function of jet radius (right) shows that while the performance still degrades in the dense core of high p_T jets, that degradation is very small for those jets discussed in Chapter 4. 120

3.16 The ATLAS calorimeter cell noise levels as a function of pseudorapidity for different subsystems from Ref. [255]. The purely electronic noise (left) is most relevant for measurements presented in this thesis, whereas the noise in high pile-up events (right) is much larger. 121

3.17 The pulse shape in the EM Barrel of the ATLAS Liquid Argon calorimeters from Ref. [255]. The unipolar triangular pulse is the current pulse generated by fast ionizing particles with characteristic drift time 450 ns. The shaped pulse is superimposed, with closed circles marking each 25 ns interval following the original collision. 122

3.18 A comparison of the design relative energy resolutions parameterized using information from test beam measurements. The calorimeter and tracker resolution parameterizations are from the ATLAS Technical Design Report [227]. The resolutions become commensurate near 40 GeV, where the Particle-Flow algorithm turns over from using track-based energy measurements to calorimetric ones. . . . 126

3.19 A schematic description of the ATLAS Particle Flow algorithm from Ref. [262]. . . 127

4.1 The different jets and their areas clustered from an example event for four different algorithms, from Ref. [273]. 134

4.2 Fitted mean energy response of isolated, truth-matched $R = 0.4$ P-Flow jets before calibration as a function of E_T for each bin in η. Points are shown which correspond to p_T ¡ 15 GeV, where each color corresponds to a different η bin. The coverage of each energy bin changes with E_T since all points are only shown of they correspond to p_T ¡ 15 GeV. 138

4.3 Fitted mean energy response of isolated, truth-matched $D=4$ P-Flow jets after applying the EM+JES calibration as a function of η for each bin in η. Points are shown which correspond to p_T ¡ 12 GeV, where each color corresponds to a different η bin. This plot demonstrates closure at the level of 2% everywhere and to within a percent near mid-rapidity. 139

4.4 Inputs into η and flavor dependent components of the JES uncertainty as a function of jet p_T for one representative bin in η. Left: the difference in gluon response of various event generators with respect to the PYTHIA sample, which is used to derive the calibration. Right: difference between quark and gluon response for various event generators. 141

4.5 Systematic uncertainty component due to η and flavor response resulting from the MC-based calibration procedure. Uncertainty is shown as a fraction of the total jet energy. 142

4.6 Binned η-intercalibration relative ratios as a function of η for p_T bins 15-17, 20-25, 30-35, 40-50, 65-90 and 110-185 GeV, with the binned calibration factors in the denominators with a line indicating the smoothed nominal result. 146

4.7 Smoothed, symmetrized eta-intercalibration systematic uncertainties as a function of η evaluated at p_T^{jet} values of 15, 20, 30, 40, 65, and 115 GeV. 147

4.8 Distributions of total cross-section for events passing nominal $Z\to\mu\mu$ Jet Balance selections, comparing data to MC, for Z^0 decays to muons (left) and electrons (right). Since no efficiency is applied for the electron quality requirement, the rate in the electron channel is slightly lower. 149

4.9 Example POWHEG+PYTHIA 8 template fits for the electron channel of the 20 Ÿ p_T^{ℓ} Ÿ 25 GeV bin (left) and 30 Ÿ p_T^{ℓ} Ÿ 35 GeV bin of the muon channel (right). The points show the distribution of $\sigma_T^{jet} \cdot p_T^{\ell}$ in data, while the lines demonstrate the fit using an MC template to derive the jet energy scale factor. The bottom panel shows contributions to the negative log-likelihood chi-squared calculation, which is terminated at the red dashed line. 150

4.10 Left: Results for the binned calibration factors and smoothed results as a function of jet p_T, with the high-η calibration result included for comparison. Right: Breakdown of smoothed statistical and systematic uncertainties on the in-situ determination of the absolute Jet Energy Scale. 151

4.11 Left: The closure of the jet energy scale calibration after the MC-derived UPC correction is applied. Right: The binned jet energy resolution determined in MC by truth matching after jets are fully calibrated. 152

4.12 The Jet Energy Resolution evaluated using dijet balance in the Pythia 8 MC sample is shown as a function of jet p_T and η. These values are used to compute systematic variations to measure the JER systematic uncertainty. The resolution can exhibit substantial transitions between η regions, due to detector geometry. 154

4.13 Example fits of the dijet balance convolution fitting procedure for UPC data (left) and MC (right). The blue distribution is the truth dijet balance, and the black markers are the reconstructed distribution. The magenta lines are fits to the truth, and red lines are convolution fits to the reconstructed distribution. The chi-squared of the distribution for the fit is computed using points extending out to the vertical magenta lines. 155

4.14 A summary of results from random cone studies in data and MC, where the independent results from the positive and negative detector sides are shown for data, to help provide a metric of the procedure's consistency. 157

4.15 Results of the combination fitting procedure in two example η bins for data (black) and MC (red), compared to the results from truth-matched single jets (blue). The middle panel shows the difference in the resolutions from the Pythia 8 dijet balance fits, and the bottom panel shows their ratios to those fits. 158

4.16 The breakdown of systematic uncertainties as a function of p_T for selection of η bins, with the total uncertainty computed in quadrature, excluding the MC non-closure systematic which is shown separately. 160

5.1 This figure shows the official topocluster significance threshold recommendations applied by ATLAS during?? running from Ref. [288], which are used in this analysis to remove background topoclusters from the gap analysis. An additional 1fb significance requirement is applied to cells in the tile calorimeter. 165

5.2 The left figure above shows comparisons of the 1-dimensional edge gap distributions for events passing all other photonuclear signal selections, oriented by detector side. The right figure shows edge gaps with the same selections, instead oriented with respect to the photon-nucleus collision system. 166

5.3 The left figure above shows comparisons of the 1-dimensional sum of gaps distributions for events passing all other photonuclear signal selections, oriented by detector side. The right figure shows sum of gaps with the same selections, instead oriented with respect to the photon-nucleus collision system. 167

5.4 Two-dimensional distributions of W η versus η_W (left) and η versus η (right) for events selected by the 2018 trigger set with an offline ZDC XOR and jet requirement. 167

5.5 The correlation of $_W$ [versus charged particle multiplicty (left) and the charged particle multiplicity distribution for several different candidate gap cuts (right) as well as the clearly hadronic (red) and photonuclear (blue, magenta, orange) contributions, compared to PYTHIA 8 (green). The cut applied in this analysis at 2.5 is shown in magenta on the left plot. 168

5.6 The distribution of the sum of gaps (left) and edge gaps (right) for the photon-going side of each event is shown for data with full photonuclear selections and Monte Carlo. 169

5.7 The distribution of both gap quantities for the nucleus-going side of each event is shown for data with full photonuclear selections and Monte Carlo. The cut on the edge gap at 3 occurs in a region where the data distribution continues but the MC has almost entirely fallen off. 169

5.8 Template ts of the photon-going sum-of-gaps distribution for different intervals of I_W The rst bin is outside the bounds of the ducial acceptance of the measurement. 173

5.9 Template ts of the photon-going sum-of-gaps distribution for different intervals of I_W The red distribution is the photonuclear signal, the pink distribution is photodiffractive background, and the green distribution is the reverse gaps background. . 174

5.10 Left: The background fraction corresponding to several $_W$[selections. Right: The background fraction corresponding to several different selections on $_A$. . . 175

5.11 The direct fraction extracted from the template ts shown in Figure 5.8 as a function of I_W(red) compared to the results from CJKL photon PDFs used in Pythia 8 (black). The magenta line corresponds to the upper acceptance edge for the measurement, after which ts do not always converge. 176

5.12 Template ts of the maximum edge gap distribution for different intervals of I_W with an XnXn requirement applied. The red and green distributions are the photonuclear signal, while the pink distribution is the hadronic background. 178

5.13 The background fraction corresponding to several $_W$[selections. No attempt is made to measure breakup in the three highest bins, where the ts also fail in some cases. 178

5.14 Shown here are the ZDC trigger ef ciencies for an example high-luminosity run, 366754. The left panel shows the A-Side ZDC trigger ef ciency in black, while the right plot shows the C-Side distribution. The ZDC energy spectra (red markers) are super-imposed to demonstrate the location of the turn-on. 179

5.15 Shown here are one-dimensional distributions of the A-Side and C-Side ZDC energies for events selected with an L1 ZDC XOR jet trigger with different linear and log x-axes to see the full spectrum as well as single-neutron and two-neutron peaks. Since some events shown here do not satisfy a 0nXn, the nucleus direction is chosen as the side with more energy in its ZDC. 179

5.16 The positive and negative [dependences of the jet trigger turn-ons used in this analysis. The ef ciency is evaluated as a function of of ine R = 0.4 P-Flow [jet Black points are for the j10 trigger, red points are j15, blue are j20, and green are j30. Black and red points are typically hidden under the blue points. 181

5.17 The ?$_T$-dependences of the jet trigger turn-ons used in this analysis. The ef ciency is evaluated as a function of of ine R = 0.4 P-Flow [jet Black points are for the j10 trigger, red points are j15, blue are j20, and green are j30. The red lines are Fermi function ts used to smooth the j10 and j20 curves. 182

5.18 Maps of the smoothed trigger ef ciencies for the j10 (left) and j20 (right) triggers as a function of ?$_T$ in bins of [, where the blue dashed line denotes the ?$_{jet}$ threshold where the trigger starts being used. 183

5.19 The impact of jet trigger ef ciency corrections on the cross-sections before unfolding, demonstrating that although the single-jet trigger ef ciencies are sometimes large, large corrections are not applied to the cross-section due to dijet requirement. 183

5.20 The one-dimensional distributions of measured L1TE values for events passing all event-level cuts and the j10 0eta490 trigger used for the physics measurement in slices of $_T$ are shown here. The red line on the left plot is the L1TE trigger threshold cut. The red curves are Gaussian ts used to extrapolate the fraction below the threshold, where the extrapolation is shown as a pink, shaded band. . . . 185

5.21 The $_T$-dependence of the L1TE turn-on extrapolated from the high-statistics jet triggers 186

5.22 The left panel shows the event-by-event interactions per bunch crossing computed for all events in this analysis. The right panel shows the different correction factors applied for each of the three physics triggers used for the cross-section measurement. These corrections are computed on a per-lumiblock basis, and the distributions are constructed using the values in all events passing these selections. 187

5.23 Comparisons between jet kinematic distributions measured for the main analysis sample (black) and the ZDC AND sample (red) used to measure nuclear break-up effects. These results are determined as absolute cross-sections. 188

5.24 Here, the measured fraction of photonuclear events which do not break up are shown in bins of $_W$ 189

5.25 Comparisons between neutron multiplicity distributions measured in the photon-going and nucleus-going directions of ZDC AND events (left) and the fraction of break-up events from different neutron multiplicities (right). 190

5.26 This distribution shows the comparison between the leading dijet (left) and H_{jets} (right) for events measured with the photon going on the A-Side (black) or C-Side (red) of the detector. 191

5.27 This distribution shows the comparison between (left) and $<_{jets}$ (right) for events measured with the photon going on the A-Side (black) or C-Side (red) of the detector. 192

5.28 This distribution shows the comparison between (left) and I_W (right) for events measured with the photon going on the A-Side (black) or C-Side (red) of the detector. 192

5.29 The differential cross-sections in between the leading dijet pair (left) and H_{jets} (right) are shown here for data and MC. 193

5.30 The differential cross-sections in $_T$ (left) and $<_{jets}$ (right) are shown here for data and MC. 194

5.31 The differential cross-sections G_A (left) and I_W (right) are shown here for data and MC. The agreement is good except for low and high regions most sensitive to nuclear shadowing, the break-up corrections, and NLO effects. 195

5.32 The jet multiplicity distributions are shown here for data and MC. The data has more events with 3 or 4 jets than the Monte Carlo, where the most important discrepancy is that 2-to-3 jet ratio is 2•5% higher in data than in MC. Since the rates for multi-jet events are so low, small total differences are re ected as large fractional differences in the ratio. 195

5.33 Shown here are the distributions of the ratio $_{jets}$•), which are used to introduce additional ducial requirements on this measurement. The upper and lower selections which are applied are shown as lines on both plots. 197

5.34 This gure demonstrates the impact of applying the 0•9 Ÿ " $_{jets}$•) Ÿ 4•0 selection to the G_A vs. I_W correlation, for several slices of $_T$. The left panels show the distributions prior to applying the selection, while the right plots show the resulting distribution after applying the cut. 198

5.35 The shift in the mean within a bin of G_A and I_W (top) or H_{jets} and $<_{jets}$ for slices in $_T$ due to the geometric acceptance of that bin. The points are the mean values, while the error bars denote the error on the calculated mean. The bin volumes are also shown, for all bins which are included in the nal cross-section measurement. 200

5.36 Shown here are the two-dimensional projections of the low-p_T regions of the reconstruction efficiencies for both sets of variables used in this analysis. The correction is very small outside of the region where jets are near the rapidity range of the gap selection. 203

5.37 Cross-sections before unfolding shown differentially in η and H_{jets} (top left), p_T and $\langle_{jets}$ (top right), p_T and G_A (middle left), p_T and I_W (middle right), $\langle_{jets}$ and H_{jets} (bottom left), or G_A and I_W (bottom right). These distributions are computed with corrections for trigger inefficiency and all event selections applied, normalized by luminosity to provide a raw cross-section. 207

5.38 Shown here are the reconstructed to truth ratios for η, $\langle_{jets}$, G_A, and I_W of all events in MC passing reconstructed and truth fiducial requirements, from the fully re-weighted PYTHIA 8 MC sample. The z-axis is the number of events in that bin normalized by the bin width. Open markers are the peaks of Gaussian fits to resolution slices. Closed markers are the arithmetic means. 209

5.39 The reconstructed to truth ratio in p_T for the full MC sample is shown on the left. The right panel shows the reconstructed to truth ratio of H_{jets}. The z-axis is the number of events in that bin normalized by the bin width. Open markers are the peaks of Gaussian fits to resolution slices. Closed markers are the arithmetic means. 210

5.40 The Monte Carlo response in p_T for binning finer than the actual measurement (left) and the binning used for the final unfolded result (right). 212

5.41 The Monte Carlo response in $\langle_{jets}$ for binning finer than the actual measurement (left) and the binning used for the final unfolded result (right). 212

5.42 The Monte Carlo response in H_{jets} for binning finer than the actual measurement (left) and the binning used for the final unfolded result (right). 213

5.43 The Monte Carlo response in G_A for binning finer than the actual measurement (left) and the binning used for the final unfolded result (right). 213

5.44 The Monte Carlo response in I_W for binning finer than the actual measurement (left) and the binning used for the final unfolded result (right). 214

5.45 On the left, the trace of the covariance matrix for the hard-scattering kinematic unfolding is shown as a function of unfolding iteration. The right panel demonstrates the variation in χ^2 between the default unfolding and a re-weighted sample using TUJU 19 PDFs as the prior. 215

5.46 The left panel demonstrates the difference in χ^2 between the default unfolding and a re-weighted sample using TUJU 19 PDFs as the prior from iteration to iteration. The right panel shows the fractional change. 215

5.47 Both panels demonstrate the change in χ^2 from the previous iteration as a function of unfolding iteration, with a linear scale on the left and log scale on the right. . . . 216

5.48 Shown here is the impact of unfolding the measured inclusive p_T distribution for detector response for the first 5 iterations compared to the raw data for the final-state kinematic unfolding (left) and hard-scattering kinematic unfolding (right). . . 217

5.49 Shown here is the impact of unfolding the measured inclusive $\langle_{jets}$ (left) and H_{jets} distributions for detector response for the first 5 iterations compared to the raw data. 217

5.50 Shown here is the impact of unfolding the measured inclusive G_A (left) and I_W (right) distributions for detector response for the first 5 iterations compared to the raw data. 218

5.51 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in $_T$. These closure tests are done on the final-state kinematic variable set. 219

5.52 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in H_{jets}. These closure tests are done on the final-state kinematic variable set. 219

5.53 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in $\langle_{jets}$. These closure tests are done on the final-state kinematic variable set. 220

5.54 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in $_T$. These closure tests are done on the hard-scattering kinematic variable set. 220

5.55 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in G_A. These closure tests are done on the hard-scattering kinematic variable set. 221

5.56 Shown here are the tautological (left) and split-sample (right) closure test of the 3D unfolding in I_W These closure tests are done on the hard-scattering kinematic variable set. 221

5.57 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ due to systematic variations of both gap selections. 224

5.58 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to l_W distributions in slices of $_T$ due to systematic variations of both gap selections. 225

5.59 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to H_{jets} distributions in slices of $_T$ due to systematic variations of both gap selections. 226

5.60 Smoothed results are shown for the ratio of data and MC cross-sections as a function of G_A and l_W for different regions in $_T$. The smoothing is only extended to regions within the geometric acceptance of the gap selections, and these results are interpolated to derive re-weighting factors. 229

5.61 Smoothed results are shown for the ratio of data and MC cross-sections as a function of $<_{jets}$ and H_{jets} for different regions in $_T$. The smoothing is only extended to regions within the geometric acceptance of the gap selections, and these results are interpolated to derive re-weighting factors. 230

5.62 The fraction of gluon jets in the photonuclear Pythia 8 MC sample determined by matching outputs of the hard-scattering process from the truth record to truth jets . 233

5.63 The fraction of gluon jets in the photonuclear Pythia 8 MC sample determined by matching outputs of the hard-scattering process from the truth record to truth jets for Direct events (left) and Resolved events (right) 233

5.64 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ due to the Jet Energy Scale. 234

5.65 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to l_W distributions in slices of $_T$ due to the Jet Energy Scale. 235

5.66 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to H_{jets} distributions in slices of $_T$ due to the Jet Energy Scale. 236

5.67 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ due to the Jet Energy Resolution. 239

5.68 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to l_W distributions in slices of $_T$ due to the Jet Energy Resolution. 240

5.69 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to H_{jets} distributions in slices of $_T$ due to the Jet Energy Resolution. . . . 241

5.70 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. This sample is isolated 0•0023 Ÿ I_W Ÿ 0•004. 243

5.71 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. This sample is isolated 0•004 Ÿ I_W Ÿ 0•008. 244

5.72 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. This sample is isolated 0•008 Ÿ I_W Ÿ 0•015. 245

5.73 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. This sample is isolated 0•015 Ÿ I_W Ÿ 0•027. 246

5.74 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to G_A distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. 247

5.75 Shown here is a breakdown of the fractional systematic uncertainties which will be applied to H_{jets} distributions in slices of $_T$ from all combined sources. The total systematic uncertainty is shown in red and the total statistical uncertainty is shown in black. 248

6.1 Distributions of 3_{0pair} (left) and $^1I_0 \sin\backslash^o_{pair}$ (right) for muon pairs passing the preselections (black), and for pairs additionally passing the LFid (Ÿ 0•06, U Ÿ 0•012) selection (red). The error bars in both panels correspond to the statistical uncertainties, which are typically too small to be seen. The pairs passing the Fid-U selection have much smaller 3_{0pair} and $^1I_0 \sin\backslash^o_{pair}$ values than pairs passing the preselections, due to the large HF background in preselected pairs. Figure reproduced from Ref. [6]. 253

6.2 Left: distributions of Σ_T^{FCal} for different muon pair selections: pairs passing preselections only, additionally passing the z_{0pair} Ÿ 0•1 mm and $|z_0 \sin\theta|_{pair}$ Ÿ 0•2 mm requirements, and additionally passing the asymmetry and acoplanarity (Ÿ 0•06, U Ÿ 0•012) requirements (labeled as “Signal candidates”). The Σ_T^{FCal} distribution for minimum-bias events (arbitrarily scaled) is also shown for comparison. Right: the same distributions, excluding the minimum-bias distribution, plus that for the UPC event selection, shown over a restricted range of Σ_T^{FCal} values. The distributions extend to negative values due to electronic noise in the calorimeter. Figure reproduced from Ref. [6]. 255

6.3 Distributions of the dimuon invariant mass (left) and p_T (right) for the dimuons in the signal region before (red) and after (black) the application of muon reconstruction and trigger ef ciency corrections. Figure reproduced from Ref. [6]. 258

6.4 Results of template ts to measured z_{0pair} distributions for pairs passing the muon preselections, the $|z_0 \sin\theta|_{pair}$ requirement and the Fid-U selection (Ÿ 0•06 and U Ÿ 0•012). Each panel represents a different centrality interval. The error bars shown on the data and the templates represent statistical uncertainties only. For many of the points, the error bars are smaller than the size of the marker. Figure reproduced from Ref. [6]. 261

6.5 The signal fractions obtained from the template ts and Equation (6.1) for the Fid-U selection, plotted as a function of centrality for several p_T intervals. The rightmost points represent the >80% centrality interval, which includes the UPC and Unassigned intervals. The error bars indicate statistical uncertainties only. Figure reproduced from Ref. [6]. 262

6.6 The heavy- avor U (left), : ? (middle), and (right) distributions for the 0–5% centrality interval. The U and : ? distributions are obtained by applying ¡ 0•06 to preselected pairs. The distribution is obtained by applying U ¡ 0•012. The error bars correspond to statistical uncertainties. The lines indicate ts to constant functions for U and : ? and to a linear function for . Figure reproduced from Ref. [6]. 263

6.7 Measured U distributions for pairs passing the full set of W ` › ` selections and in the Fid-U region for nine centrality intervals. The error bars on the points indicate statistical uncertainties. Also shown are the background levels estimated from the z_{0pair} template tting (red, ne-dashed lines) and from the asymptotic ts (blue, coarse-dashed lines). The widths of the lines indicates the statistical uncertainty of the background estimates. Figure reproduced from Ref. [6]. 264

6.8 Measured $\Delta\phi$ distributions for pairs passing the full set of WW `›` selections and in the Fid-$\Delta\phi$ region for nine centrality intervals. The error bars on the points indicate statistical uncertainties. Also shown are the background levels estimated from the 3_{0pair} template fitting (red, fine-dashed lines) and from the asymptotic fits (blue, coarse-dashed lines). The widths of the lines indicates the statistical uncertainty of the background estimates. Figure reproduced from Ref. [6]. 265

6.9 Differential cross-sections versus u (left) and $|\Delta\phi|$ (right) for Drell–Yan production of dimuons in $p\bar{p}$ $\sqrt{s_{NN}}$ = 5.02 TeV nucleon–nucleon collisions from POWHEG +PYTHIA 8 for three nuclear PDF sets. The error bars indicate statistical uncertainties. The insets show the differential cross-sections for pairs with $\ddot{y}$ 0.06. They are plotted over ranges of u or $|\Delta\phi|$ that are twice as wide as the corresponding fiducial regions. The statistical uncertainties are correlated between the different nPDF sets because all results derive from the same set of POWHEG+PYTHIA 8 events. Figure reproduced from Ref. [6]. 267

6.10 Top row: comparison of the measured excess background yield for p_T ¡ 4 GeV with the estimated DY contribution obtained using the nCTEQ15 and nNNPDF2.0 nuclear PDF sets, which represent the extreme ranges of POWHEG+PYTHIA 8-predicted rates. The error bars on the excess background indicate statistical uncertainties, and the error bars on the DY calculations represent the quadrature sum of the statistical and systematic uncertainties shown in Table 6.3. The data-points for the DY calculations are shifted slightly along the u axis, for clarity. Bottom row: The differences between the observed background excess and the DY rates obtained using nNNPDF2.0 PDF set, expressed as a fraction of the fiducial WW `›` yields. The left and right panels correspond to the Fid- and Fid-$\Delta\phi$ selections, respectively. Figure reproduced from Ref. [6]. 270

6.11 HF, DY background-subtracted u distributions for pairs satisfying the Fid-u selection and having p_T ¡ 4 GeV in different centrality intervals. In a few of the panels, the distributions are scaled to allow a common u axis range for the plots. Also shown for comparison are the generated and reconstructed distributions obtained from the reconstructed STARLIGHT simulation samples. The STARLIGHT distributions are scaled vertically to allow direct comparison with the other distributions. Figure reproduced from Ref. [6]. 273

6.12 HF, DY background-subtracted $|\Delta\phi|$ distributions for pairs satisfying the Fid-$|\Delta\phi|$ selection and having p_T ¡ 4 GeV in different centrality intervals. In a few of the panels, the distributions are scaled to allow a common $|\Delta\phi|$ axis range for the plots. Also shown for comparison are the generated and reconstructed distributions obtained from the reconstructed STARLIGHT simulation samples. The STARLIGHT distributions are scaled vertically to allow direct comparison with the other distributions. Figure reproduced from Ref. [6]. 274

6.13 HF, DY background-subtracted asymmetry distributions for pairs having p_T > 4 GeV in different centrality intervals. In a few of the panels, the distributions are scaled to allow a common H-axis range for the plots. Also shown for comparison are the generated and reconstructed distributions obtained from the reconstructed STARLIGHT simulation samples. The STARLIGHT distributions are scaled vertically to allow direct comparison with the other distributions. Figure reproduced from Ref. [6]. 275

6.14 The U (top) and ϕ (bottom) distributions of dimuons before and after unfolding. The left and right panels correspond to the 0–10% and UPC intervals, respectively. Figure reproduced from Ref. [6]. 277

7.1 The fraction of photonuclear jet events passing the fiducial requirements which do not break up as a function of $W_{\gamma N}$. The systematic uncertainties are shown as shaded blue bands, and the statistical uncertainties are black lines. A quadratic fit in $W_{\gamma N}$ is shown in red. These results indicate a strong dependence of the break-up rate on $W_{\gamma N}$ and an overall high rate of break-up due to additional Coulomb interactions, reaching 60% at high $W_{\gamma N}$. 283

7.2 The ratio of the photonuclear cross-sections, with or without nuclear modifications, in nCTEQ15 WZ+SIH (squares) and nNNPDF 3.0 (diamonds). Results are integrated over the $W_{\gamma N}$ range $3.7 \times 10^{4} < W_{\gamma N} < 0.027$, and each panel shows a separate range in p_T. The total uncertainty on this measurement in each bin is compared to the size of these modifications, where the red bands show the total scale uncertainty, the grey bands are the quadrature sum of the residual systematic uncertainty, and the yellow bands show the statistical uncertainty. 285

7.3 Triple-differential cross-sections, $\frac{d^3\sigma}{dp_T dH_{jets} d\eta_{jets}}$, as a function of H_{jets}, in four η_{jets} intervals with η_{jets} increasing from the top left to the bottom right. The cross-sections are shown for the selected intervals in η_{jets} and integrated over the p_T acceptance. For each plot systematic uncertainties are shown the upper panel as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a $W_{\gamma N}$-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and the data. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. 286

7.4 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T dH_{jets} d\phi_{jets}}$, as a function of ϕ_{jets}, in several H_{jets} intervals. The cross-sections are integrated over the acceptance. For each plot systematic uncertainties are shown the upper panel as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and the data for a selection of the H_{jets} intervals. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 289

7.5 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T d\xi_A dl_W}$, as a function of ξ_A for different bins of η_T for events with emitted photon energies in the kinematic range $0.0023 < l_W < 0.0043$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 290

7.6 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T d\xi_A dl_W}$, as a function of ξ_A for different bins of η_T for events with emitted photon energies in the kinematic range $0.0043 < l_W < 0.0079$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 291

7.7 Triple-differential cross-sections, $\frac{d^3\sigma}{d\eta_T d\xi_A d\lambda_W}$, as a function of ξ_A for different bins of η_T for events with emitted photon energies in the kinematic range $0.0079 < \lambda_W < 0.015$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a λ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 292

7.8 Triple-differential cross-sections, $\frac{d^3\sigma}{d\eta_T d\xi_A d\lambda_W}$, as a function of ξ_A for different bins of η_T for events with emitted photon energies in the kinematic range $0.015 < \lambda_W < 0.027$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a λ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 293

7.9 Triple-differential cross-sections, $\frac{d^3\sigma}{d\eta_T d\xi_A d\lambda_W}$, as a function of ξ_A for different bins of η_T for events with emitted photon energies in the kinematic range $3.7 \cdot 10^{-4} < \lambda_W < 0.027$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a λ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. 294

7.10 Triple-differential cross-sections $\frac{d^3\sigma}{d\,{}_T\,d\eta_A\,dl_W}$, as a function of ${}_T$ for different bins of η_A for events with emitted photon energies in the kinematic range $0.0023 < l_W < 0.0043$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_A. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 295

7.11 Triple-differential cross-sections $\frac{d^3\sigma}{d\,{}_T\,d\eta_A\,dl_W}$, as a function of ${}_T$ for different bins of η_A for events with emitted photon energies in the kinematic range $0.0043 < l_W < 0.079$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_A. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 296

7.12 Triple-differential cross-sections $\frac{d^3\sigma}{d\,{}_T\,d\eta_A\,dl_W}$, as a function of ${}_T$ for different bins of η_A for events with emitted photon energies in the kinematic range $0.0079 < l_W < 0.015$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a l_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_A. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 297

7.13 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T d\eta_A d\xi_W}$, as a function of η_T for different bins of η_A for events with emitted photon energies in the kinematic range $0.015 < \xi_W < 0.027$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a ξ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_A. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. . . . 298

7.14 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T d\eta_A d\xi_W}$, as a function of η_T for different bins of η_A for events with emitted photon energies in the kinematic range $3.7 \times 10^{-4} < \xi_W < 0.027$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a ξ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for a selection of the bins of η_A. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. 299

7.15 Triple-differential cross-sections $\frac{d^3\sigma}{d\eta_T d\eta_A d\xi_W}$, as a function of ξ_W for different bins of η_T for events with struck parton energies in the kinematic range $0.010 < \eta_A < 0.166$. In the upper panel, systematic uncertainties are shown as shaded boxes, while statistical uncertainties shown as vertical lines are usually smaller than the size of the markers. A theoretical comparison is shown to cross-sections computed using PYTHIA 8 with nCTEQ15 WZ+SIH PDFs, a photon flux from STARLIGHT, and a ξ_W-dependent breakup fraction, as shown in Section 7.1.1. The bottom panels show the ratio between the theory prediction and data for several bins of η_T. The red bands in the ratio panels are the quadrature sum of scale uncertainties on the cross-section, while the gray band shows the remaining systematic uncertainty. The yellow band shows the point-to-point statistical uncertainty. 300

7.16 The cross-sections for $WW \to \ell\nu\ell\nu$ production in non-UPC Pb+Pb collisions as a function of centrality (left) and p_T (right). The error bars indicate statistical (lines) and systematic uncertainties (shaded boxes). The solid lines in the left panel represent the results of STARLIGHT calculations (see text) for the given p_T interval. The data points in the right panel for the 0–20% and 60–80% centrality intervals are staggered, for clarity. 301

7.17 Measured normalized yields of $W \to \ell\nu$ pairs passing the Fid-η selection as a function of centrality for four p_T intervals. The error bars indicate statistical (lines) and systematic (shaded boxes) uncertainties. 303

7.18 Differential cross-sections as a function of U for $W \to \ell\nu$ pairs passing the Fid-U selection. Each panel represents a different centrality or E_T^{FCal} interval, with the last panel representing the UPC interval. The error bars indicate combined statistical and systematic uncertainties, excluding the background subtraction uncertainties, which are indicated by a shaded band at $3 \times U = 0$, and overall normalization uncertainties, which are quoted on each panel as "Scale". 304

7.19 Differential cross-sections as a function of η for $W \to \ell\nu$ pairs passing the Fid-η selection. Each panel represents a different centrality or E_T^{FCal} interval, with the last panel representing the UPC interval. The error bars indicate combined statistical and systematic uncertainties, excluding the background subtraction uncertainties, which are indicated by shaded bands at $3 \times \eta = 0$, and overall normalization uncertainties, which are quoted on each panel as "Scale". 305

7.20 Background-subtracted U (left column) and η (right column) distributions measured in different p_T intervals for three centrality intervals: 0–20%, 20–40%, and 60–80%. The distributions are normalized to have unit integral. The error bars indicate statistical and systematic uncertainties added in quadrature. 307

7.21 Moments of the $W \to \ell\nu$ η distributions as a function of centrality compared with the PWF predictions. The error bars indicate combined statistical and systematic uncertainties. In some bins, the statistical uncertainties, determined both by the number of counts and by their η distribution, dominate. 308

7.22 Moments of the $W \to \ell\nu$ U distributions as a function of centrality compared with the QED and PWF predictions. The error bars on the data points indicate combined statistical and systematic uncertainties. In some bins the statistical uncertainties, determined both by the number of counts and by their U distribution, dominate. 309

7.23 Dimuon η distributions in three intervals of H in the 20–40% (left) and 40–80% (right) centrality intervals, respectively. The distributions are self-normalized to allow the distributions to be directly compared. The error bars show combined statistical and systematic uncertainties. 310

7.24 Differential cross-sections for $W \to \ell\nu$ production as a function of η for two intervals of q: $0 \leq q \leq c \times 2$ and $c \times 2 \leq q \leq c$ in the 20–40% (left) and 40–80% (right) centrality intervals, respectively. The error bars show combined statistical and systematic uncertainties. 311

7.25 The average (left) and the standard deviation (right) of the `, ` : ? distributions in the 20–40% and 40–80% centrality intervals as a function of H (top) and q (bottom). Results are shown for H intervals of width 0.5 for H Ÿ 2•0, and two additional intervals covering 2•0 H Ÿ 3•0 and 3•0 H Ÿ 4•8. The q results are shown for four equal intervals covering c ¼ The error bars show combined statistical and systematic uncertainties. The points are staggered for clarity of presentation. 312

8.1 Left: Kinematic coverage of different data sets which are used in the EPPS 21 [69] nPDF set, compared to the coverage of measurements presented in this thesis. Projected EIC coverage bands for two possible band energies are also shown from Ref. [93]. Right: Comparisons between photonuclear jet production data presented in this thesis and several leading nPDF models. 315

B.1 Distributions of the cluster multiplicity (left) and the vertex positions (right) for photonuclear events passing the trigger, gap, and of ine ZDC selections in blue. Events sequentially passing the different splash cuts are shown for the vertex cuts (red), and the cluster-track correlation cut with track rejection cut (blue). 348

B.2 The left plot shows the correlation between tracks and inner detector clusters (2•5). The red line drawn is a cut applied to remove splash events. The right plot is the 2D correlation between raw tracks and tracks, for events passing the cut on the left side and vertex cuts. The red line denotes the cut applied to remove this background. 348

B.3 The left plot shown here is the jet distribution for jets in j[j Ÿ 2•5 without any tracks in the event. The right plot is the jet time distribution for jets with the same selection. All events included here pass ZDC, gap, and trigger requirements, as well as splash cuts. 349

B.4 The left plot is the 2D correlation between un-calibrated, matched EM Topo jet times and the respective jet after applying the -dependent correction to the mean jet time. The right plot is the same distribution with calibrated jet times. The residual bands far from = 0 indicate several residual backgrounds, including beam halo muons, noisy clusters, and a small amount of residual splash events. . . 350

B.5 Here, the corrected P-Flow (red) and matched EM Topo (black) jet times are shown, demonstrating the substantially improved timing resolution of EM Topo jets. The Gaussian t to the EM Topo jet time distribution shows the extracted width of the main band of the distribution, motivating a cut at 4 ns (magenta lines). 351

B.6 The left plot shows the energy measured in the photon-going FCal for events passing physics, splash, and beam-halo muon cuts. The red line denotes a cut applied in the data to remove pile-up. The right plot demonstrates the correlation between negative-energy cluster multiplicity and total negative energy in the event. 352

B.7 Here, the negative cluster multiplicity distribution (left) and total negative $?_{T}$ in the event (right) are shown for events passing all other event-level requirements. Several candidate cuts are shown along with a Gaussian t to the main peak. The red cut is 25 negative clusters and 3 GeV of negative $?_{T}$. The blue cut is 30 negative clusters and 5 GeV of negative $?_{T}$. The magenta cut is 35 negative clusters and 10 GeV of negative $?_{T}$. 353

B.8 These plots show the relationship between track multiplicity and cluster multiplicity within the acceptance of the inner detector for events passing the OOT pile-up cut (left) and failing it (right). 353

C.1 This gure demonstrates the fraction of bin volumes within the ducial region of this measurement which are within the region selected by the geometric volume criteria. These corrections are applied to the bin-width normalization of the nal measurement, and any bin with less than 50% acceptance is omitted from the nal results. 355

C.2 This gure demonstrates the fraction of bin volumes within the ducial region of this measurement which are within the region selected by the geometric volume criteria. These corrections are applied to the bin-width normalization of the nal measurement, and any bin with less than 50% acceptance is omitted from the nal results. 356

C.3 This gure demonstrates the fraction of the ducial cross-section for that bin contained within the selected geometric volume region for that bin. A selection is applied to only retain bins with a fraction greater than 0.975. 357

C.4 This gure demonstrates the fraction of the ducial cross-section for that bin contained within the selected geometric volume region for that bin. A selection is applied to only retain bins with a fraction greater than 0.975. 358

C.5 This gure demonstrates the fraction of the total cross-section for that bin which passes the ducial requirements, for the region contained within the selected geometric volume. A selection is applied to only retain bins with a fraction greater than 0.975. 359

C.6 This gure demonstrates the fraction of the total cross-section for that bin which passes the ducial requirements, for the region contained within the selected geometric volume. A selection is applied to only retain bins with a fraction greater than 0•975. 360

List of Tables

2.1 The Feynman rules for perturbative QCD from Ref. [11], with a general choice of gauge de ned by two parameters ξ and $\tilde{\xi}$, using gauge- xing terms from Equations 2.12 and 2.13. In the case of a physical gauge, the ghost propagator and ghost vertex are set to 0. 11

2.2 Summary of world data included in recent ts to nPDF data, adapted from Ref. [79]. 35

2.3 Parameters of the coherent photons emitted in UPCs for different existing collider facilities. Adapted from Refs. [1, 2]. For asymmetric collisions, the photon is assumed to be emitted from the higher-Z nucleus. Energies are quoted in the target rest frame, by convention. 54

3.1 Relative systematic uncertainties on the determination of the visible cross-section f_{vis} and luminosity measurements for the 2015 and 2018 5.02 TeV Pb+Pb data. . . 94

3.2 Summary of ATLAS calorimeter sub-system coverages and granularities, adapted from the design parameters in Ref. [233]. All listed components are based on the LAr sampling calorimeter technology besides the Tile calorimeter. The segmentation in $|\eta|$ of the Forward Calorimeter is listed as approximate, since its actual segmentation is in longitudinal distance. 100

4.1 The full set of triggers used for deriving the eta-intercalibration for this measurement. Min-bias triggers from 2017 were necessary in order to cover the η-low-region. The L1 chains either require lled bunch crossing (RD0_FILLED) or a hardware-based jet trigger (JXX). 143

5.1 List of triggers used in this measurement for the 2018 data sample, the run number at which the trigger was enabled and disabled, the average combined (level-1 times HLT) prescale and total luminosity sampled by the trigger. Some jet triggers sample the same jet p_T threshold, but they are not active during the same part of the 2018 running period. 163

5.2 Information about each Monte Carlo sample used to simulate the different jet kinematics relevant for this analysis. The samples were ltered on the initial-state with ?) and the nal-state with $_T$. 163

5.3 The full cut- ow of event selections used in this analysis is shown here. All event counts require two of ine 15 GeV P-Flow jets, and the fraction removed is calculated from the count immediately preceding it. 202

6.1 The $_T^{FCal}$ intervals used for this measurement, along with the corresponding centrality ranges and $h_{AA}i$ values. No centrality calibration is available for the intervals labeled ET0–ET3. 254

6.2 Numbers of muon pairs in each centrality interval for different ranges and the two ducial selections. Also listed are the total W ` ` yields. The interval labeled as Unassigned includes pairs that contribute to the inclusive yield but do not appear in any of the listed centrality intervals, as they do not satisfy the ZDC coincidence requirement imposed on the non-UPC centrality intervals, and also do not satisfy the exclusive UPC requirement. 257

6.3 Effective nucleon–nucleon cross-sections obtained from POWHEG+PYTHIA 8 for the production of Drell–Yan muon pairs in p_{NN} = 5•02 TeV collisions using different nuclear PDF sets. The systematic uncertainties of the ducial cross-sections obtained by propagating PDF uncertainties through POWHEG+PYTHIA 8 are also shown. A separate 15% uncertainty in the cross-sections, due to factorization and renormalization scale uncertainties, is not included in the shown uncertainties. . . 267

Acknowledgements

I would like to begin by thanking my advisor, Brian Cole, for more contributions to my development as a physicist than I can count. Your relentless enthusiasm for understanding every aspect of our work made each question more interesting and each challenge more surmountable. I distinctly remember a number of times during my PhD when it seemed like you made the impossible possible, by knowing how to break a problem down to the essential pieces and having faith in our ability to understand each one. I am deeply grateful to you for challenging me to do the kind of work that I am now so proud to have been a part of.

This work was only made possible by the incredible collaborators I have met through ATLAS. Aaron Angerami has been a mentor to me as I have developed my own skills and understanding in this eld. Your belief in my ability to solve problems even when I was unsure of it myself was instrumental for so much of the progress toward results in this thesis. I am also always grateful for the refreshingly honest perspective you bring to our discussions about physics. Martin Rybar helped me get started in ATLAS while we still shared an of ce at Columbia. You have so generously shared your knowledge and time with me, and I do not know where I would be without your guidance and friendship. Peter Steinberg, Riccardo Longo, Witold Kozanecki, Dennis Perepelitsa, and Mateusz Dyndal have all taught me so much throughout my PhD as well.

The Columbia group has been essential in navigating this process, and I am so glad to have been part of this team. Bill Zajc has has always been so welcoming to me and the other ATLAS students, sharing stories, advice, and humor. Soumya Mohapatra and Blair Seidlitz have been amazing friends and collaborators, and I've so appreciated every “group drink” with them. I am especially grateful for the other graduate students who went through this process with me: Pengqi Yin, Wenkai Zou, and Luke Krauth. From group lunches to the occasional gripe session, the dif cult times of my PhD were so much more bearable because of your friendship. I have been lucky to befriend the younger students joining our group, Yuhan, Emma, Hanpu, and Shuhang. Your presence has helped make our of ce feel like a community these past few years.

I have had so many great friends while in New York, from Columbia and beyond. The late-

night of ce hangouts with Michael, Roman, and Noah made me feel welcome here my rst year. I will be forever indebted to all past members of the Columbia physics bagel club: Claire, Niccolò, Mick, Rocco, Josh, and Daniel. I would also like to thank Ashley, Simon, and Karly for helping me remember to nd the humor in every insane occurrence of these past few years. I have been so lucky to have Mario Sukkar so close by this past year, and I am grateful to Emily Cox and Margaret Haltom for their consistent friendship, support, and spirited conversation. I would also like to thank Walter Kunkle for being a caring, steadfast best friend for half my life.

Thank you to my brother, Zach, the most resilient, hard-working, and rambunctious person I know. I am so proud to that you are my brother, and it means the world to me that you take so much pride in me. I owe so much of who I am to my dad, Mitch, whose love and support both got me through my PhD and made me into the person I am today. Through your kindness, deep love for your family, and steady work ethic, you have been such an incredible role model for me my entire life. My girlfriend, Maddy, has been the person who got me through the dif cult times near the end of my PhD to reach this point. Thank you for always indulging my desire to talk about physics and taking so much pride in my achievements, even if I sometimes have trouble with it myself. This achievement means so much more to me because I can share it with you.

The last two years have been exceptionally dif cult, and I have more gratitude that I can put into words for everyone who helped me navigate my mother's illness and passing, which inevitably became intertwined with the end of my PhD. I will always be grateful to Brian for the exibility he gave me to spend time with my mom when she was sick, while still helping me to make progress towards completing my analyses. I want to thank Aaron for making sure I knew the UPC jet analysis was in good hands, even when I had too much on my plate to keep up with every detail. My dad, Mitch, was basically a superhero when we navigated those months of illness as a family, and without his strength to manage that impossible situation, I would not have been able to keep working towards this goal for all this time. Lastly, I want to thank Maddy, who has helped keep me focused on what really matters, even after being thrown into such an impossible situation herself. I am so grateful for you.

Dedication

For my mother, Irene

“If I knew the way, I would take you home.”

Chapter 1: Introduction

The purpose of this thesis is to describe a set of measurements using photoproduction processes in 5.02 TeV Pb+Pb collisions as a probe of Quantum Chromodynamics (QCD) at the Large Hadron Collider (LHC). QCD is the theory of the strong interaction between quarks and gluons, which features a complex structure arising from asymptotic freedom, the surprising property that the interaction strength of QCD is small at high energies and large at low energies. The structure arising from this property has led to a substantial number of open questions about QCD which have persisted in the 50 years since the theory's development. Photoproduction processes, interactions where one or both participants are real photons, however, are comparatively quite simple, since they are mostly described by Quantum Electrodynamics (QED). This thesis will present to measurements which take advantage of these properties, by using photoproduction processes as a probe of QCD systems which do not introduce much of their own additional complexity. These results include precision measurement of two such photoproduction processes, photonuclear jet production and photoproduction of muon pairs, as means of studying the structure of QCD phenomena.

Measurements of photoproduction in heavy ion collisions are enabled by the intense electromagnetic elds surrounding highly charged relativistic nuclei, which act as an intense source of high-energy, nearly-real [1, 2] photons. This photon ux is typically studied in Ultra-Peripheral Collisions (UPCs), where the two colliding nuclei do not overlap, causing photon-mediated interactions to be the only source of particle production. In UPCs, both photonuclear (γA) and two-photon ($\gamma\gamma$) collisions may occur, where each type of process allows for this study of different aspects of both QED and QCD. While both types of collisions produce a variety of nal states, photonuclear interactions are broadly useful for characterizing properties of the struck nucleus, while two-photon processes allow for high-precision studies of QED and the physics of photoproduction. In the case of photonuclear interactions, exclusive nal states have been studied

extensively in UPCs, typically through the production of vector mesons. This thesis will present results which take a less common approach, the study of inclusive photonuclear processes, where the production of high-energy nal states allows them to directly probe the structure of the nuclear target.

Nuclei are complex objects, with a structure governed by the high-coupling regime of QCD which is quite poorly understood from rst principles. Early measurements of nuclear structure [3] demonstrated this complexity, with the unexpected result that the shapes of quark and gluon (parton) distributions measured inside a nucleus in high-energy scattering processes differ from those of free protons. These nuclear Parton Distribution Functions (nPDFs) cannot be calculated from rst-principles in QCD, and they are essential for understanding both the basic structure of matter which makes up our world and the complex properties of QCD at high densities. Since the discovery of nPDF modi cations, they have become a subject of intense study, involving a wide range of measurements across different collision systems, species of nuclei, and nal-states. This thesis will present a new observable sensitive to the nPDFs in a region of phase space with little other available data [4], jet production in inclusive γA collisions. Through the application of unfolding methods and novel jet calibration techniques, these results can be compared directly to theoretical calculations in order to extract nPDF effects.

Jets are energetic, highly collimated “sprays” of hadrons produced in high-energy particle collisions. They are the result of a high energy parton produced in a scattering process, which then emits a shower of collimated QCD radiation, resulting in a collection of energetic hadrons. These hadrons share the total energy of the initiating parton, so by measuring jets, one may reconstruct the energy of that parton, probing the kinematics of the system which produced it. In the case of photonuclear jet production, the kinematics of the hard-scattering system approximately map onto three variables de ned in terms of the jet kinematics: ξ_T, ξ_W and ξ_A. The variable ξ_T provides a metric of the energy scale of the process, while ξ_W and ξ_A measure the fraction of either nucleus's energy which is carried by the photon (ξ_W) or struck parton (ξ_A) in the photonuclear interaction. Thus, by measuring the jet kinematics, one may directly probe the distribution of partons within

the nuclear target at a well-de ned energy scale Q^2, accessing the nPDFs. Jets are typically challenging to measure, though, and since these studies necessitate high precision to constrain nPDFs, the work in this thesis will present a new, dedicated jet calibration for UPCs, which employs novel techniques to measure jet energies with high precision in a UPC environment.

While photoproduction processes may be used to study the structure of nuclei, they are also a laboratory for precision QED in two-photon interactions. In particular, the study of W^+W^- collisions producing muons [5] has precisely characterized the photon spectra in UPCs through the relative angular de ection of muons in a dimuon pair. Results presented in this thesis of dimuon photoproduction in typical (non-UPC) Pb+Pb collisions [6] extend the study of these processes considerably. These results improve on previous studies [7] using a much larger data sample available from the full LHC Run 2 data set. The improved statistical precision from this large sample allows for detailed studies of the dimuon relative de ection, measured by $\Delta\phi = 1 - \phi_{\mu\mu}/\pi$ and $\Delta\phi^* = \tan(\phi^*_{\eta}/2)\sqrt{1-\tanh(\Delta\eta/2)}$, and its dependence on collision centrality. These measurements probe the dependence of photoproduction processes on impact parameter, allowing for detailed comparisons to new theoretical calculations. They also provide results which help to constrain the strength of magnetic elds in the hot, dense nuclear matter formed in heavy ion collisions, the Quark-Gluon Plasma.

This thesis is organized as follows. Chapter 2 discusses the relevant theoretical background of QCD, photoproduction, and heavy ion physics which provide the context and motivation for the measurements presented here. Chapter 3 surveys relevant aspects of the LHC, ATLAS detector, and experimental methods employed to perform these measurements. Chapter 4 describes a novel jet calibration which was derived in order to perform the measurements of this thesis, along with some theoretical background in jet measurements. Chapters 5 and 6 discuss the data analysis for the photonuclear jet and dimuon photoproduction measurements, respectively. Chapter 7 then discusses the results of this thesis, differential cross-sections for photonuclear jet production and dimuon photoproduction. Chapter 8 then highlights the key results of this thesis and discusses their implications.

Chapter 2: Background

2.1 Quantum Chromodynamics

2.1.1 The QCD Lagrangian and Basic Formulation

In Quantum Field Theory [8], the eld content and interactions of a theory are entirely described by the Lagrangian density, where the dynamics of the theory are then given by applying the quantum mechanical path integral. For Quantum Chromodynamics [9], the theory describing the strong nuclear force, the Lagrangian density is described in Equation 2.1. QCD belongs to a class of Quantum Field Theories known as Yang-Mills [10], where each theory corresponds to a speci c compact Lie group, under which it exhibits local transformation symmetry (the gauge symmetry). In the case of QCD, it exhibits SU(3) gauge symmetry, distinguishing it from other Yang-Mills theories, such as QED or the theory of the weak nuclear interaction, which are the Yang-Mills theories corresponding to U(1) and SU(2), respectively. The QCD Lagrangian can be written as:

$$L_{QCD} = \sum_{@=1}^{\tilde{O}^5} k_{@} \left(i\gamma^{\mu} D_{\mu} - m_{@} \right) k_{@} - \frac{1}{4} G^{a}_{\mu\nu} G^{\mu\nu}_{a} \qquad (2.1)$$

Here, $k_{@}$ is the spin $\frac{1}{2}$ quark eld of avor @ and U_{μ} is the gluon four-potential. The strength of QCD interactions is set by the coupling $6 = \sqrt{4cU_B}$ in natural units. The group structure of SU(3) is encoded in the QCD Lagrangian through both $\grave{a}^{a}_{0}$ and the generators of SU(3), $\grave{a}^{1}_{0} \ldots \grave{a}^{o}_{0}$, which are themselves $\#^2_2 - 1$ distinct $\#_2 \times \#_2$ traceless, Hermitian matrices, where $\#_2 = 3$ is the number of color charges associated with the gauge group. The $\#_2$ notation denotes the fundamental representation of the SU(3) generators. $G^{a}_{\mu\nu}$ is the gluon eld strength tensor de ned in terms of

the gluon vector potential, strong coupling, and structure constants of SU(3), as:

$$\text{[illegible]} \quad (2.2)$$

It is particularly useful to express the Lagrangian density in terms of covariant derivative, which is de ned to be manifestly Lorentz covariant, [illegible] In addition to its manifest Lorentz covariance, it also combines all components of the theory which encode the structure of the SU(3) gauge group in a single object. The indices 8 and 9 each represent one of $\#_2$ colors:

$$\text{[illegible]} \quad (2.3)$$

Using the covariant derivative, the QCD Lagrangian density can then be re-written (with sums over repeated indices written explicitly in order to more clearly show their de nitions and ranges) as:

$$\text{[illegible]} \quad (2.4)$$

The 8 and 9 indices both run over $\#_2 = 3$ color charges, the spinor indices 0 and 1 run over the four Dirac matrices, the Lorentz indices $\grave{a}$ and a run over 3 ¸ 1 space-time dimensions, and @ runs over all quark avors, $=_5 = 6$. Additionally, the index 0 runs over the $\#_2^2 \quad 1$ different combinations of color charges which may be carried by the gauge elds, which have both a color charge and anti-charge in one of eight linearly independent color eigenstates which exclude the forbidden color singlet eigenstates.

Since this Lagrangian is the Yang-Mills theory for the group SU(3), it includes all terms with spin-1•2 fermion elds or spin 1 boson elds and their couplings which are renormalizable. The 3 color charges in QCD are the conserved charges corresponding to the SU(3) gauge symmetry. The key property distinguishing QCD from simpler theories such as QED is that SU(3) is a non-Abelian group, meaning that its elements do not necessarily commute with one another. The implications of this property are manifest in a perturbative picture through the gluon self-coupling,

and many important properties of QCD emerge as a direct consequence of the non-Abelian structure of the SU(3) group. This property is visible in the Lie algebra for SU(3), which is de ned by the commutation relation of the generator matrices:

$$\left[\frac{\lambda_a}{2}, \frac{\lambda_b}{2} \right] = i f_{abc} \frac{\lambda_c}{2} \tag{2.5}$$

In this commutation relation, the non-Abelian structure of QCD is manifest, since non-Abelian theories are de ned by having non-zero structure constants f_{abc}. In order to make full use of the SU(3) gauge symmetry in the QCD Lagrangian density, it is useful to write the transformations of the eld content under local gauge transformations $U(x)$, with color indices a and b, some set of transformation angles θ_a, and the generators of SU(3), described by:

$$U(x) = \exp\left[i \sum_{a=1}^{8} \theta_a(x) \frac{\lambda_a}{2} \right]_{ab} \tag{2.6}$$

The transformation properties of the fermionic elds are fairly straightforward:

$$\psi_a(x) \to U_{ab}(x) \psi_b(x) \tag{2.7}$$

Given this transformation of the fermionic eld, in order for the QCD Lagrangian density to exhibit gauge invariance, the covariant derivative also must have a straightforward transformation property (with assumed color indices):

$$D_\mu \psi_a \to U_{ab} D_\mu \psi_b \tag{2.8}$$

This necessary transformation of the covariant derivative then implies that the gauge elds transform as:

$$D_\mu \psi_a \to U_{ab} D_\mu \psi_b + \frac{i}{g} \partial_\mu U_{ab} \psi_b \tag{2.9}$$

After fully de ning the transformation properties of the components in the Yang-Mills QCD Lagrangian, these properties can be used in order to x the gauge of the theory. Since SU(3) is

non-Abelian, though, the gauge- xing requires the introduction of a gauge- xing function, as a new part of the effective Lagrangian. The introduction of these gauge- xing terms also introduces additional complex scalar “ghost” elds to the theory, yielding an effective QCD Lagrangian (after making choices regarding the particular gauge) of:

$$L^{eff}_{QCD} = L_{QCD} \, , \, L_{gf} \, , \, L_{ghost} \tag{2.10}$$

Here, L_{QCD} is the manifestly gauge-invariant Yang-Mills QCD Lagrangian, L_{gf} is the gauge- xing function, and L_{ghost} is the Lagrangian for ghost elds, which arise due to the Jacobian in the path integral of the gauge- xing function being non-constant for non-Abelian gauge theories, where they enforce the unitarity of the theory after gauge- xing. The ghost elds $2_0{}^1C^a$ and $2_0{}^1C^a$, are complex Grassmann elds, which means they anti-commute in spite of being scalar elds and is important when the Feynman rules for ghost diagrams are derived in Section 2.1.1.1. The ghost Lagrangian includes both a kinetic and interacting term:

$$L_{ghost} = \quad m\grave{}\, 2^0 m\, 2^0 \, , \; 6\, 5^{012} \; {}^{2}\hat{}\, m\grave{}\, 2^0 2^1 \tag{2.11}$$

From this Lagrangian, it is evident that the ghost elds interact with gauge elds through the second term where the 5^{012} constants appear. Thus, in an Abelian gauge theory where the constants 5^{012} are always zero, the ghosts have no interactions with physical elds of the theory. While it is not immediately obvious from the ghost Lagrangian, there are several reasons why ghost elds are not actually physical and therefore any amplitude to produce them in a scattering process must be 0. They do not appear in the QCD Lagrangian and are purely a construct of the gauge- xing procedure, making them clearly un-physical. Additionally, they would violate the spin-statistics theorem [8] if they did exist as physical elds. The un-physical nature of ghost elds can be seen quite obviously when one actually chooses a gauge by writing down a gauge- xing function. These

functions can be chosen almost freely, but two popular classes of choices are [11]:

$$L_{gf} = \frac{\#\text{Ö}^{3}}{2} \sum_{0=1}^{1} {}^{1}m \; \grave{}_{0}{}^{o2} - 1 \; \ddot{Y} \_ \ddot{Y} \; 1 \tag{2.12}$$

$$L_{gf} = \frac{\#\text{Ö}^{3}}{2} \sum_{0=1}^{1} {}^{1}= \; \grave{}_{0}{}^{o2} - \_ ! \; 1 \tag{2.13}$$

The set of gauges described by Equation 2.12 are called the “covariant” gauges, the most common of which is _ = 1, referred to as Feynman gauge. The set of gauges described in Equation 2.13 are referred to as either the “axial” or “physical” gauges. Here, ` is some arbitrary four-vector which de nes the choice of gauge, and a common example of a useful physical gauge is choosing for it to be light-like ($=^2 = 0$), referred to as light-cone gauge. The latter set of gauges eliminate the need for ghost elds by taking _ ! 1 . Often, though, the explicit Lorentz covariance of the former set of gauges is more useful than the elimination of ghost elds.

2.1.1.1 The Feynman Rules of Perturbative QCD

Perturbation theory provides a framework in which predictions for observables may be calculated from QCD as an expansion in the coupling constant of the strong force, $\text{º}_B{}^{1}\grave{}$, where ` is the renormalization scale of the calculation. This coupling constant is expressed as a function of some dimensionful energy scale `, since when the theory is renormalized, the coupling varies with the energy scale. The renormalization of QCD and the implications for the running of its coupling are described in further detail in Section 2.1.1.2. This variation of the coupling strength with renormalization scale is important, because the underlying assumption of perturbation theory is that the coupling is suf ciently small that one may neglect higher-order terms in the perturbative expansion. Since º_B varies with energy scale, perturbative QCD can only make meaningful predictions for processes at a certain range of energy scales, beyond which other methods are required, and precise calculations of observables are much more challenging.

In order to express the Feynman rules of perturbative QCD [11], the Lagrangian density must

be quantized by adding gauge- xing (L_{gf}) and ghost (L_{ghost}) Lagrangian densities, as shown in equation 2.10. The gauge- xing Lagrangian densities in Equations 2.12 and 2.13 demonstrate that the only impact of the choice of gauge on the Feynman rules, though, is in the rule for the gluon propagator. In order to derive the Feynman rules, it is useful to separate the effective QCD Lagrangian density into its free- eld and interacting parts:

$$L^{eff}_{QCD} = L^{0}_{quark}, L^{0}_{gluon}, L^{0}_{ghost}, L_{int} \tag{2.14}$$

Of the three free- eld parts of the Lagrangian density, the quark part is gauge-independent, while the gluon and ghost terms depend on choice of gauge. The technical separation between these terms is achieved by taking the coupling constant to 0, where the remaining parts of the Lagrangian density are the free- eld terms. Since these terms contain no interactions between different elds, the free- eld Lagrangian density can be expressed as three parts, one corresponding to each kind of eld:

$$L^{0}_{quark} = \sum_{8=1}^{\#2} \sum_{@=1}^{\#5} k_{@-8} \, 8 X_{8} W m < @ X_{8} k_{@-9} \tag{2.15}$$

$$L^{0}_{gluon} = \frac{1}{2} \sum_{0=1}^{\#2} {}^{1}m_{a-0} m^{a}_{0} \, m_{a} {}_{-0} m^{ao}_{0}, L_{gf} \tag{2.16}$$

$$L^{0}_{ghost} = m^{2}_{0} m^{2}_{0} \tag{2.17}$$

The remaining terms provide the interacting Lagrangian density, which dictates the vertex Feynman rules. This Lagrangian density can be expressed (with implicit sums over repeated indices) as:

$$L_{int} = 6 \, {}_{-0} k_{@} W)^{1}_{0} {}^{o} k_{@} \, 6 \, 5^{012} {}^{1}m_{a-0} {}^{o}_{1} {}^{a}_{2} \\ \frac{1}{4} 5^{012} {}_{-1} {}_{a-2} 5^{034} {}_{-3} {}_{a-4}, 6 \, 5^{012} {}^{2}m^{2}_{0} {}^{1} \tag{2.18}$$

From these different Lagrangians, the Feynman rules for perturbative QCD can be derived (for some general choice of gauge parameterized by and =). With the notation of a slash through a variable denoting a Lorentz contraction of it with a Dirac matrix, these Feynman rules are

written in Table 2.1. It should be noted that any closed loop of ghost elds receives an extra factor of 1 on its contribution to the amplitude, since they are complex Grassmann elds. The Feynman rules provide, in principle, all the necessary ingredients for calculations in perturbative QCD, although these calculations often become quite complicated due to the assortment of spinor and color indices as well as the non-Abelian structure of the theory. In order to derive useful results from xed-order perturbative calculations, though, the rst step is to determine when the coupling, $6 = \sqrt{4cU_B}$, is small. In order to do so, the Feynman rules can be employed to renormalize QCD and determine the variation in the coupling with energy scale, as is discussed in Section 2.1.1.2.

2.1.1.2 The QCD Beta Function

The coupling strength of QCD famously varies with energy scale, leading to properties such as asymptotic freedom. This modi cation to the coupling arises when the theory is renormalized, which requires the introduction of a renormalization scale, $\grave{}$, which regulates the theory. This scale, however, is not a parameter of the theory – it does not appear anywhere in the QCD Lagrangian density, and it is introduced in order to remedy problems with calculating scattering amplitudes. Since $\grave{}$ is not a physical parameter, the choice of its value cannot impact physical observables of the theory. This basic requirement, applied through the formalism of the Wilsonian renormalization group [12], provides a mathematical approach to determine the variation of U_B with energy scale.

While a full derivation of the running of the coupling using perturbative QCD is beyond the scope of this thesis, thorough calculations to one loop appear in References [13] and [14]. These calculations begin with the requirement that any generic physical cross-section, $f\,{}^{1}?–\grave{}\,{}^{o}$, as a function of some set of particle momenta, $?$, and the renormalization scale, $\grave{}$, can be written as a perturbative expansion in the renormalized coupling:

$$f\,{}^{1}?–\grave{}\,{}^{o} = \sum_{=1}^{\infty} 0_{=}\,{}^{1}?–\grave{}\,{}^{o}U_B^{=} \tag{2.19}$$

Feynman Rule	Labelled Diagram	Contribution to amplitude
Quark Propagator		[illegible]
Gluon Propagator		[illegible] (Covariant Gauge) [illegible] (Physical Gauge)
Ghost Propagator		[illegible]
Quark-Gluon Vertex		[illegible]
Ghost-Gluon Vertex		[illegible]
3-Gluon Vertex		[illegible]
4-Gluon Vertex		[illegible]

Table 2.1: The Feynman rules for perturbative QCD from Ref. [11], with a general choice of gauge de ned by two parameters, and = ` , using gauge- xing terms from Equations 2.12 and 2.13. In the case of a physical gauge, the ghost propagator and ghost vertex are set to 0.

Since μ is not a physical parameter, the cross-section must be independent of it:

$$\mu^2 \frac{m}{m\mu^2} f\left(\frac{Q}{m}, \alpha_s(\mu)\right) = \mu^2 \frac{m}{m\mu^2} \sum_{n=1}^{\infty} \alpha_s^n \left(\frac{mQ}{m\mu^2}, \frac{m\alpha_s}{m\mu^2}\right) = 0 \tag{2.20}$$

This requirement can be re-formulated in order to define the β function for QCD, which describes the variation of α_s with μ in a way which enforces the condition that physical observables do not depend on μ. It also demonstrates that the β function itself can be expanded in α_s, lending itself to perturbative calculation. This definition in terms of the β function by imposing this condition yields the Renormalization Group Equation:

$$\mu^2 \frac{3\alpha_s}{3\mu^2} = \beta(\alpha_s) = \sum_{n=0}^{\infty} \beta_n \alpha_s^{n+2} \tag{2.21}$$

Perturbative calculation of the QCD β-function shows that the term proportional to α_s vanishes, so the expansion begins with the term proportional to α_s^2. The minus sign in Equation 2.21 is included by convention, so positive values for the β_n parameters then indicate that α_s decreases with increasing μ. The values of the leading two terms of the β function are [9]:

$$\beta_0 = \frac{1}{12\pi}\left(11 - \frac{2}{3} n_f\right) = \frac{7}{4\pi} \quad \text{(3 colors, 6 flavors)} \tag{2.22}$$

$$\beta_1 = \frac{1}{24\pi^2}\left(51 - \frac{19}{3} n_f\right) = \frac{39}{24\pi^2} \quad \text{(3 colors, 6 flavors)} \tag{2.23}$$

The explicit dependence on the number of colors and number of quarks is instructive, and it is useful in the case that $\beta_0 < 0$ for some number of quarks, which would exclude them from the calculation. Regardless, though, the leading β_0 and β_1 terms for QCD are positive for any theory with 16 or fewer quark flavors, which the Standard Model certainly satisfies. This property also holds for higher order calculations of the β function [15, 16]. The sign of the QCD β function gives it the property of asymptotic freedom [17, 18], where the coupling grows stronger for lower-energy processes and weaker for higher-energy processes. For processes involving some physical hard scale (often the momentum transfer Q), it is useful to set the renormalization scale μ to Q,

then to leading order, the variation of the strong coupling with $\&$ can be written as:

$$U_B^1\&^{2o} = \frac{U_B^1 `_0^{2o}}{1 ¸ V_0 U_B^1 `_0^{2o} \ln \&^2 • `_0^2} \qquad (2.24)$$

This simple solution to the differential equation allows the value of the coupling at any $\&$ scale, to be computed from the value of the coupling at some reference scale for which it is known, $`_0$. This equation appears to depend both on the choice of the reference scale and the value $U_B$ at that scale, but these separate dependencies obfuscate that it is independent of the choice $`_0$. In order to make this property more obvious, we introduce an extra scale $_{QCD}$, which provides a dimensionful scale at which the coupling would diverge in perturbation theory:

$$_{QCD} = `_0 4^{1•1/2V_0 U_B^1 `_0^{2oo}} ! \quad U_B^1\&^{2o} = \frac{1}{V_0 \ln \&^2 • {}_{QCD}^2} \qquad (2.25)$$

Since $_{QCD}$ exhibits some renormalization scheme dependence (due to its dependence at higher orders on V_2¸ terms which are themselves scheme-dependent), it is difficult to describe without ambiguity for physical processes. Instead, it is often more robust to measure the value of U_B at a given scale, from which values of $_{QCD}$ can be extracted to a given order. Existing world data for $U_B^1\&^{2o}$ is shown in Figure 2.1, compared to the predicted trend from the renormalization group approach, demonstrating the high precision and robustness of this approach.

2.1.1.3 Lattice QCD

Because of the behavior of $U_B^1\&^{2o}$ described in Section 2.1.1.2, the QCD coupling is large in the kinematic region $\&$ (. $_{QCD}$) in which most QCD systems exist. For describing QCD in this kinematic region, though, fixed-order perturbative methods fail due to the large coupling. Other, non-perturbative methods have thus been important for understanding the low-energy behavior of these systems. One of the most successful such approaches is Lattice QCD (LQCD), which was originally proposed by Wilson in 1974 [19] as lattice gauge theory. In this approach, the

Figure 2.1: The $\&$-dependence of the QCD coupling strength α_s, from Ref. [9]. The 5-loop calculation of the QCD β function, assuming $\alpha_s(M_Z) = 0.1180 \pm 0.0009$, is compared to world data for independent measurements of α_s.

theory is formulated on discretized Euclidean space-time, typically as a hypercubic lattice with spacing a. In this lattice description of space-time, the quark fields are placed on lattice sites, and the gauge fields are placed on the links between those sites, as represented in Figure 2.2. This configuration allows for the lattice spacing a, to serve as the regulator for UV divergences in the theory, rendering calculations finite. The discretization of the action can be performed in several different ways, but all valid discrete forms of the action must give consistent results in the continuum limit, $a \to 0$. Matrix elements can then be computed through direct evaluation of the path integral, using a discretized form of the action.

Discretization of the gauge action in the simplest way [19] works well, but it leads to discretization errors which grow proportional to a^2. Naively discretizing the fermion action leads to problems, however, due to an effect known as "fermion doubling" [20, 21] where the process introduces $2^3 - 1$ spurious fermion fields for 3 space-time dimensions. This problem is solved via the use of Wilson fermions [19] at the cost of chiral symmetry or staggered fermions [20] which retain a smaller, more manage-able number of doublers. Discretization errors from both the gauge

Figure 2.2: A two-dimensional representation of the eld structure of lattice QCD calculations, from Ref. [9]. Quark elds lie on gauge sites, and the gauge elds (lie on the links between sites.

and fermion actions are substantially improved through the modi cation of the action to better account for the short-wavelength behavior of the theory which is poorly approximated by the lattice [22, 23]. Analytic expressions for these corrections can be computed perturbatively since, due to the asymptotically free nature of QCD, the short-wavelength behavior relies on the low-coupling region of the theory.

Lattice QCD calculations proceed by directly evaluating the path integral for a discretized action on the lattice. Due to the high dimensionality of this problem, it requires Monte Carlo integration methods, meaning that lattice calculations suffer from statistical uncertainties from MC sampling of the integral. The computing cost to reduce these errors grows as [23] $L^4 m_\pi^2 a^6$, where L is the length of the simulated hypercube, m_π is the pion mass, and a is the lattice spacing. Because of the rapid growth of computing time with shrinking lattice spacing, large improvements in minimizing discretization errors [24] have been necessary to greatly improve the precision of lattice calculations. These improvements have allowed for precise lattice calculations for a wide range of QCD observables [9] including quark masses, form factors, and decay constants, as a few examples.

2.1.2 Con nement and Asymptotic Freedom

The key consequence of the renormalization of QCD, the scale-dependence of its coupling constant (Equation 2.24), implies that the theory is asymptotically free at high energy, but also that its coupling sharply rises at low energy. In this low energy (long-distance) limit, the theory becomes strongly coupled, preventing any strongly-interacting particles to exist outside of a color singlet state at low energy. This property, known as con nement [19], has not yet been proven rigorously [25], but a simple qualitative explanation in terms of the potential is instructive. The potential in QCD can be determined on the lattice to behave as approximately Coulomb-like on short distance scales but to have a long-range behavior of growing linearly with separation,

$$+ \frac{1}{A} + fA \qquad (2.26)$$

Here, is the strength of the short-distance Coulomb term and is the string tension, which has similar values across other models as well [26]. The picture represented by this linear term in the potential is that of a string-like QCD ux tube composed of gluon elds stretching between the quark and anti-quark. The potential term represents the tension in this string which, when stretched far enough, provides enough energy that it is favorable for it to break it and produce a new pair from the vacuum. A simple illustration of this process for pair moving away from each other is shown in Figure 2.3. This “string model” for describing the phenomenology of interactions between quarks and gluons as they transition from high to low energy is useful in modelling high energy hadronic interactions, and it is described in much greater detail in Section 2.3.2.1.

This simpli ed model describes a picture through which con nement occurs, where no particle with color charge can exist far from all other such particles, since it will always be energetically favorable to produce a color singlet state by pulling pair from the vacuum. This property of con nement at low energies, combined with asymptotic freedom at high energies, allows QCD to reconcile many early independent theories explaining the physics of hadrons. In particular,

Figure 2.3: A diagram showing the stretching and breaking of a ux tube between two quarks as they move apart.

the quark model [27, 28, 29], which predicted the measured spectrum of hadrons [30] could be reconciled with the lack of any free quarks in nature through con nement. Similarly, the “quarks” described in this model can behave like the point-like, weakly-interacting constituents of hadrons [31, 32] observed in deep-inelastic scattering (DIS) experiments [33], since they are asymptotically free at high energies. Thus, these two properties of QCD at high and low energies allow this theory to provide a uni ed explanation for several models of hadronic physics which, at the time of their development, seemed completely disparate.

2.1.3 QCD Factorization and the Parton Model

While the separation of scales allows for the application of perturbative methods to calculate high-energy matrix elements at some xed order, real-world scattering processes do not exhibit this clean separation. For QCD interactions occurring on short distance scales, while the hard-scattering matrix element is perturbatively calculable, outgoing free quarks and gluons from such a process must propagate for long distances prior to measurement. Thus, any measurable cross-section becomes sensitive to the long-distance (non-perturbative) behavior of QCD. Further

complicating this scenario is that no sources of free quarks and gluons exist in nature, so QCD scattering processes almost always involve incoming objects with some hadronic structure governed by non-perturbative QCD.

The tool by which these different regimes are combined in order to predict high-energy cross-sections in QCD is known as factorization [34]. Factorization is a framework in which the short-distance and long-distance behavior is systematically separated, allowing the cross-section for a given process to be written as the convolution of different pieces, each of which describes one regime of QCD, but not both. By introducing a factorization scale, in addition to the renormalization scale, , the calculation directly cuts off the regions relying on long-range ( ) and short-range ( ) descriptions. Schematically, the cross-section for the collision of two hadrons, and , to produce a hadron, can be written as:

$$3f \quad , \quad ! \quad , - {}^{1}? \; - ? \, {}^{o} = \tilde{O}_{0-1-2-3}^{\;1} \; 3b_0 3b_1 3l \; 5_0 \cdot {}^{1}b_0 - ` \, {}^{o} \, 5_1 \cdot {}^{1}b_1 - ` \, {}^{o} \; f^{\wedge}_{01!\;2,\;3} \, {}^{1}?_0 - ?_1 - U_B^{1} ` \, {}}^{o} - \&^{2o} \; {}_2 \, {}^{1}l - ` \, {}^{o} \tag{2.27}$$

Here, the integration variables b_0 and b_1 represent the momentum fraction of hadron or . The fragmentation functions, $_2$ encode the non-perturbative process by which outgoing partons fragment into hadrons, and they are discussed in greater detail in Section 2.3.0.1. The functions $5_0 \cdot {}^{1}b - ` \, {}^{o}$ are the Parton Distribution Functions (PDFs) for a parton in hadron . The PDFs are distributions representing the number density of a given parton (quark or gluon) within the hadron, carrying a fraction of the hadron's momentum. They are not perturbatively calculable in QCD, and the experimental determination of them is a central topic of this thesis described in detail in Section 2.1.5 for free nucleons and Section 2.2.1 for nuclei. The PDFs provide essential information for studying high-energy hadronic collisions, since they encode the non-perturbative structure of hadrons in terms of high-energy partons which are themselves amenable to perturbative methods.

The formalism of factorization is essential for the different distributions entering into these

cross-section calculations to be well-de ned. The prescription itself relies on the separation in scales allowing one to neglect quantum mechanical interference between the different components of the calculation. Through this feature of factorization, the different pieces exhibit universality, where they are completely independent of the particular process they are used to calculate. Since the PDFs are universal, they are amenable to determination from a huge variety of observables, without any loss of precision from reconciling the results of different physical processes. This aspect of the PDFs is exploited in global PDF ts, as described in Section 2.1.5.

Unfortunately, factorization methods do not necessarily hold for all cross-sections in QCD, and proofs of their validity only exist for certain kinds of collisions and processes. These proofs rely on basic arguments of causality and unitarity [35], where the processes causing the formation of the initial and nal states are too far separated in time from the hard-scattering process to in uence them. While complete proofs have existed for a long time in DIS and collisions [34], the only process in hadron-hadron collisions for which factorization has been fully proven is Drell-Yan [36, 37] production of di-leptons. In diffractive collisions, factorization has been rigorously proven [38], but it is also broken in the case of dijet photoproduction [39], which is described further in Section 2.5.3.5. For many other QCD systems, factorization is still applied, but it is only approximate in these cases.

2.1.4 Deep Inelastic Scattering

This observation of point-like constituents of hadrons [31, 32] arose from measurements of deep inelastic electron-proton scattering [33]. In these experiments, the very high energy probe causes the highly virtual photon exchange to effectively shatter the proton, producing some number of outgoing hadrons. The basic Feynman diagram picture of DIS, along with particle kinematics, is shown in Figure 2.4. In order to describe this process, it is useful to introduce the kinematic quantities:

$$G = \frac{\&^2}{2@\,\%} = \frac{\&^2}{2"\ ^1\ \ ^0} - H = \frac{@\,\%}{:\ \ \%} = \frac{\ \ \ \ \ \ \ \ ^0}{\ \ \ \ \ \ \ \

Figure 2.4: A diagram of deep inelastic electron-nucleon scattering for an electron with incoming four-momentum and outgoing four-momentum, momentum transfer and incoming nucleon four-momentum, from Ref. [9]

Where: and are the incoming and outgoing four-momentum of the electron; : is the four-momentum of the virtual photon, and are the incoming and outgoing electron energies in the target rest frame, is the mass of the struck nucleon, and = . is the four-momentum of the incoming nucleon, and is its mass. The variable, referred to as Bjorken, is the scaling variable, which can, in leading order pQCD, be shown to be the fraction of the nucleon's momentum carried by the struck quark in the in nite momentum frame. This simply the fraction of the electron's energy which is transferred to the nucleon in the target rest frame. It is also useful to de ne the mass of the recoil system, as:

$$, \; ^2 = {}^1\%, \; @^2 = " \; ^2, \; 2" \; {}^1 \qquad {}^0 \; \; \&^2 \tag{2.29}$$

Additionally, the center-of-mass energy squared of the lepton-nucleus system is:

$$B = {}^1 : \; , \; \%^2 = \frac{\&^2}{GH} \; " \; ^2, \; <^2 \tag{2.30}$$

In terms of these variables, deep inelastic scattering is formally de ned [9] as satisfying two conditions: deep ($\&^2$ ¡¡ " ^{2}) and inelastic (^{2} ¡¡ " ^{2}). When these conditions are met, the

differential cross-section for DIS can be expressed in terms of these quantities as:

$$\frac{3^2 f}{3G\,3H} = \frac{4cU_{EM}^2}{GH\&^2} \quad 1 \quad H \quad \frac{G^2H^2 \,''\, {}^2}{\&^2} \quad {}_2{}^1G-\&^{2o} \;,\; H^2G \;{}_1{}^1G-\&^{2o} \tag{2.31}$$

Here, ${}_1{}^1G-\&^{2o}$ and ${}_2{}^1G-\&^{2o}$ are the nucleon structure functions, and U_{EM} is the ne-structure constant. While this expression holds for unpolarized DIS, the inclusion of parity-violating processes, such as for neutrino-nucleon DIS, would require a third structure function. For charged spin-$\frac{1}{2}$ elds, the structure functions are related by the Callan-Gross relations [40], $2G\,{}_1 = {}_2$. This relationship differs from the case for spin-0 elds, where ${}_1 = 0$. These structure functions encode the interaction between the photon and the nucleon, and they are connected to the eld theory description of the photon-nucleon interaction via the hadronic tensor, $,^{\grave{a}}$, de ned as [34]:

$$,^{\grave{a}} = \frac{1}{4c} \quad 3^4H4^{8@H} \quad \tilde{O}_{-} \quad h\;j\;9^{\grave{}\,1H\,o}\,j-\,i\,h-\,j\;9^{a\,1\,0o}\,j\;\;i \tag{2.32}$$

Here, represents the state of the incoming nucleon, the sum runs over hadronic nal states, and the electromagnetic current density $9^{\grave{}}$ is This tensor can be directly related to the structure functions ${}_1$ and ${}_2$ as:

$$,^{\grave{a}} = \;{}_1{}^1G-\&^{2o} \quad 6^{\grave{a}} \;,\; \frac{@\grave{@}^{a}}{@^{a}} \;,\; {}_2{}^1G-\&^{2o} \;\frac{{}^{1}\%^{\grave{}} \quad @\grave{\%} \;@\grave{@}^{a\,o\,1}\%^{a} \quad @^{a}\%\;@\grave{@}^{a\,o}}{\%\;@} \tag{2.33}$$

Through Equations 2.32 and 2.33, the structure functions provide a direct connection between the hard-scattering matrix elements and the structure functions. The structure functions then serve as a connection to the structure of hadrons in QCD via the parton model [31, 41], where both structure functions can be written as [34]:

$${}_1{}^1G-\&^{2o} = \tilde{O}_{0\;G}^{\;1} \;{}^1\frac{3b}{b}5_{0}\cdot\;{}^1b-{}^{o}\;{}_{10}\;\frac{G\;\&}{b\;\grave{}}-U_{B}{}^{1\grave{}\,o} \;,\; O^{1}\,1\bullet\&^{2o} \tag{2.34}$$

$${}_2{}^1G-\&^{2o} = G\;\tilde{O}_{0\;G}^{\;1} \;{}^1\frac{3b}{b}5_{0}\cdot\;{}^1b-{}^{o}\frac{b}{G}\;{}_{20}\;\frac{G\;\&}{b\;\grave{}}-U_{B}{}^{1\grave{}\,o} \;,\; O^{1}\,1\bullet\&^{2o} \tag{2.35}$$

The $\mathcal{O}(\alpha_s)$ terms are calculable in perturbation theory [42] to some fixed order in α_s and separate the process-dependent part of the structure functions from the universal PDFs. The first terms on the RHS for Equations 2.34 and 2.35 are referred to as the leading-twist contribution, since additional terms are suppressed by higher orders of Q^2. Included in this description is that, since Q^2 only serves as a regulator for the $\mathcal{O}(\alpha_s)$ the leading-twist structure functions are independent of Q^2 for sufficiently high Q^2. This property of very weak dependence on Q^2 at fixed x, known as Bjorken scaling [31, 32], demonstrates that the scattering in DIS occurs from point-like constituents in the struck nucleon. These point-like constituents were named partons by Feynman, and they were later identified as spin $\frac{1}{2}$ quarks via the experimental confirmation of the Callan-Gross relation [40].

2.1.5 Parton Distribution Functions and DGLAP Evolution

While the PDFs arise naturally in discussions of the structure functions, they are fundamental objects themselves which encode the collinear structure of hadrons. They can be expressed directly in terms of matrix elements of the hadronic state $|P\rangle$. These expressions are simplest when written in light-cone gauge, as [43, 44]:

$$q(x) = \frac{1}{4\pi} \int dz^- e^{ixP^+z^-} \langle P|\bar{\psi}(0^-, z^-, 0_\perp)\gamma^+\psi(0^-, 0^-, 0_\perp)|P\rangle \tag{2.36}$$

$$g(x) = \frac{1}{4\pi} \int dz^- e^{ixP^+z^-} \langle P|F^{+a}(0^-, z^-, 0_\perp)F^{+a}(0^-, 0^-, 0_\perp)|P\rangle \tag{2.37}$$

Here, the indices 0 and 1 correspond to the two color indices of the gluon field. P is the hadron momentum, P^+ is the plus component of the hadron momentum, and P^- is the minus component of the hadron momentum where $P^- = \frac{M^2}{2P^+}$ and the transverse components of both are zero. These definitions are explicitly the matrix elements of operators which count the number of quarks or gluons at a given momentum fraction x. From these definitions, it is clear that the PDFs are not simply a means of expressing the structure functions but rather the fundamental description of collinear hadronic structure in terms of partonic consituents. These definitions are not amenable to direct analytical calculation, though, so our understanding of the PDFs relies on interpreting

experimental results.

The parton distribution functions are subject to several sum rules [34, 43] which follow from the conservation laws for several quantities: charge, avor, and momentum. The total charge and momentum sum rules for the proton PDFs require a sum over all parton avors:

$$\tilde{O}_{8}^{1}\int_{0}^{1} dG\, 5_{8}(?)\,{}^{1}G-^{2o} = 1-\tilde{O}_{8}^{1}\int_{0}^{1} dG\, @_{8}\, 5_{8}(?)\,{}^{1}G-^{2o} = @_{p} \qquad (2.38)$$

Where $@_{8}$ is the charge of each parton avor and $@_{p}$ is the charge of the proton. In addition to preserving the total charge and momentum of the proton, the PDFs must also re ect the conservation of avor in QCD, where the total conserved avor comes from the valence structure (two up quarks and one down quark). The avor sum rules for the light quarks thus are:

$$\int_{0}^{1} dG\, 5_{D}(?) - 5_{D}(?) = 2-\int_{0}^{1} dG\, 5_{3}(?) - 5_{3}(?) = 1-\int_{0}^{1} dG\, 5_{B}(?) - 5_{B}(?) = 0 \qquad (2.39)$$

Similar sum rules to the strangeness sum rule exist for the heavy quark PDFs, since there cannot be any net production of those avors either in QCD. The sum rules provide important constraints on the PDFs which aid in the process of constraining them from ts to world data, as described later in this section and in section 2.2.1.

While the PDFs themselves are not directly calculable in perturbative QCD, they provide distributions of partons which can be described by perturbative methods. Thus, while perturbative methods cannot directly calculate the PDFs, they can provide information about how the PDFs evolve with the renormalization and factorization scales, so long as for all scales along that evolution, $U_{B}{}^{1`2o}$ is small enough for perturbative methods to apply. The equations which govern that evolution are known as the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [45, 46, 47, 48] which, to leading order, are of the form:

$$`^{2}\frac{m\, 5_{@}{}^{1}b-^{2o}}{m`^{2}} = \frac{U^{1`2o}}{2c}\int_{b}^{1}\frac{3H}{H}\,\tilde{O}_{9=1}^{5}\, 5_{@9}{}^{1}H-^{2o}\%_{@@9}\,\frac{b}{H}\,,\; 5_{6}{}^{1}H-^{2o}\%_{@6}\,\frac{b}{H} \qquad (2.40)$$

Figure 2.5: Feynman diagrams of the different parton splittings which give rise to the four leading-order DGLAP splitting functions

$$\frac{m\,5^{1}b-^{2o}}{m^{2}} = \frac{U^{1}\,^{2o}}{2c}\;{}_{b}\;\frac{3H}{H}\;\sum_{9=1}^{2=5} 5_{@^{9}}{}^{1}H-^{2o}\%_{6@^{9}}\;\frac{b}{H}\;,\;5_{6}{}^{1}H-^{2o}\%_{66}\;\frac{b}{H} \qquad (2.41)$$

Where $@^{9}$ represents a quark of avor 9, so the sum runs over all quark and anti-quark avors, for a total of $2=_{5}$ terms. If one notes that he DGLAP equations are of the same form as Equation 2.21, it is clear that they are essentially the renormalization group equations of the quark and gluon PDFs, arising from RG ow in QCD. The functions of the form $\%_{01}{}^{1}G^{9}$ are the splitting functions, which provide the probability of a parton 1 splitting to some pair of partons, one of which is parton 0, with a fraction of the parent parton's momentum G. These are motivated by the leading-order Feynman diagrams for parton splitting processes shown in Figure 2.5, and certain rules of these diagrams and assumptions allow the number of splitting functions to be simpli ed signi cantly. First, since quark avor must be conserved in the $\%_{@@^{9}}$ function, it must be diagonal, implying:

$$\%_{@@^{9}}{}^{1}G^{9} = X_{89}\,\%_{@@}{}^{1}G^{9} \qquad (2.42)$$

Where $\%_{@@}$ is the splitting function for a quark to split to a quark of the same avor. Additionally, for simplicity, all quarks are assumed to be massless, allowing the gluonic splitting functions to be

independent of quark avor. After applying these simpli cations, the splitting functions are [47]:

$$
\begin{aligned}
\%_{@@} &= \frac{4}{3}\,\frac{1 \,\text{¸}\, G^2}{1 \quad G}\,\text{¸} = \frac{4}{3}\,\frac{1 \,\text{¸}\, G^2}{{}^{1}1 \quad G^{9}\,\text{¸}}\,\text{¸}\,\frac{3}{2}X^{1}G \quad 1^{o} \\
\%_{@6} &= \frac{1}{2}\,G^2 \,\text{¸}\, {}^{1}1 \quad G^{9\,2} - \%_{6@} = \frac{4}{3}\,\frac{1 \,\text{¸}\, {}^{1}1 \quad G^{9\,2}}{G} \\
\%_{66} &= 6\,\frac{1 \quad G}{G}\,\text{¸}\, G^{-}1 \quad G^{9}\,\text{¸}\,\frac{G}{{}^{1}1 \quad G^{9}\,\text{¸}}\,\text{¸}\,\frac{11}{2} \quad \frac{=_5}{3}\,X^{1}1 \quad G^{9}
\end{aligned}
\tag{2.43}
$$

Here, $=_5$ is the number of quark avors, and the X functions are introduced to assist in regularizing the $\%_{@@}$ and $\%_{66}$ for $G! \quad 1$. The second aspect of this regularization involves the notation $x^{1}G^{9}¼_{¸}$, which denotes the replacement of $^{1}G^{9}$ by a distribution (for some suitably regular test function, $5^{1}G^{9}$):

$$
\int_0^1 3G^{-}5G^{9}\,\text{»}\,{}^{1}G^{9}¼_{¸} = \int_0^1 3G^{-}5^{1}G^{9} \quad 5^{1}1^{oo}\,{}^{1}G^{9} \tag{2.44}
$$

This prescription for DGLAP evolution is LO in $\mathcal{L}_B$, but the NLO [49, 50] and NNLO [51, 52] splitting functions and evolution equations have been calculated. These prescriptions have applications beyond the evolution of the PDFs, and they are of use in additional contexts such as the description of parton showers or hadronic photon structure.

In order for the DGLAP evolution equations to be useful, one must start from existing knowledge of the PDFs. This knowledge comes from a wide range of experimental measurements which have been conducted over the last 50 years, where the increased center-of-mass energy available at new colliders has greatly improved the G and $\&^2$ reach of world data over time. Figure 2.6 demonstrates the G and $\&^2$ coverage of different sources for world data, where about half of the total constraint on PDF ts now comes from the LHC [53]. Due to the universality of the PDFs, data from a wide variety of measurements in hadron-hadron collisions at the LHC and the Tevatron can be included alongside DIS data from both the Hadron Electron Ring Accelerator (HERA) and xed target experiments. These data are combined through global tting procedures from various groups, most notably MSHT [53], NNPDF [54], CTEQ [55], ATLAS [56], ABMP [57], and CJ [58]. Additionally, dedicated ts exist for exclusively HERA data, chie y HERAPDF [59].

Figure 2.6: The kinematic coverage of world data (left) to determine the free proton PDFs and world data on DIS measurements of the free proton structure function $F_2(x, Q^2)$ (right), both from Ref. [9]

These PDF fitting procedures typically they consist of parameterizations in terms of high-order polynomials in x at some initial scale Q_0^2. The PDFs are then evolved using DGLAP evolution (typically at NNLO), to the scales at which different data is available. Then, the PDFs are used as inputs to fixed-order perturbative calculations of the cross-sections measured for different observables. Predictions can then be made for the input experimental data (typically at NNLO or NLO, depending on the experimental precision and availability of calculations), and a chi-squared is computed via comparisons to data. Accounting for quark masses is important for the requisite level of precision, and the most commonly applied method is the General Mass Variable Flavor Number Scheme (GM-VFNS) [60], which interpolates between schemes with or without considering heavy quark flavors, using the relationship between the hard scale and heavy quark masses.

The chi-squared is minimized, and the resulting PDFs then have uncertainties determined using the Hessian method [61, 62], where the fit parameters are expanded around their best fit values.

Figure 2.7: The free proton PDFs determined at NNLO in the MSHT20 fit to world data from Ref. [53]. The dependence on x is shown for two scales $Q^2 = 10\,\mathrm{GeV}^2$ and $Q^2 = 10^4\,\mathrm{GeV}^2$, the former of which relies more on fixed target and HERA data, while the latter is more sensitive to Tevatron and LHC data.

The PDF uncertainties can then be decomposed into eigenvalue and eigenvector variations on the parameters, allowing for the quadrature sum of independent variations on the PDFs to represent the total uncertainty. The notable exception to this approach is the NNPDF collaboration [54], who use a neural network to compute the PDF and fit via machine learning techniques. Since this method does not allow for expanding the parameters around a local minimum, they determine uncertainties via fitting Monte Carlo replicas [63], where the data points are allowed to fluctuate by their uncertainties to produce variations which are re-fitted. The uncertainty on the PDF is the standard deviation of resulting replicas.

An example of the free proton PDFs from the MSHT20 fit is shown in Figure 2.7, where the behavior of the different PDFs clearly varies substantially with Q^2. The gluon PDF dominates over others at low x due to its very rapid growth. The valence quark densities are surpassed by gluon densities by $x \approx 0.1$, where the crossover varies with Q^2. At lower Q^2, the uncertainties become larger at very low x, limiting the potential to observe the continuation of the trend. Sea quark PDFs also rise rapidly at lower x, but they are each at least an order of magnitude smaller than the gluon density. These different PDF sets, along with nuclear PDFs discussed in Section 2.2.1, are

available in the LHAPDF database [64], which is used for producing theoretical comparisons to different nPDF sets shown in Chapter 5.

2.1.6 Diffraction in QCD

The study of diffractive processes in QCD has a long and complex history, as evidenced by the dif culty in nding a consistent de nition for them between experiment and theory. The typical experimental de nition for a diffractive collision [9] is one where large regions of the detector have no particle production (large rapidity gaps) and/or incoming particle(s) remain intact but are de ected slightly. Theoretically, diffractive processes are those resulting from the scattering of an incoming plane-wave in some region of impact parameter. High-energy diffractive processes are typically described within the context of Regge theory [65], which is based on studying the singularities of amplitudes in the complex angular momentum plane in quantum scattering theory.

Regge theory was rst developed as a method of explaining the dependence of the elastic cross-section for hadronic collisions on the momentum transfer, $\&^2$ [9]. By studying the amplitude of a general plane wave scattering from a target in relativistic quantum mechanics, Regge theory proceeds by applying unitarity, analyticity, and symmetry constraints on the scattering amplitude [66]. Applying these constraints in addition to the optical theorem allows one to derive the inelastic, elastic, and total scattering amplitudes through the treatment of “Regge trajectories”, relationships between $\&$ and poles in the complex angular momentum plane, $U^1\&$. The Regge trajectories can be represented in a quasi-particle picture as a set of quasi-particles referred to as Reggeons. The pole with the largest intercept, referred to as the pomeron, has vacuum quantum numbers and dominates diffractive processes at high energies.

2.2 Descriptions of Nuclear Structure

2.2.1 Nuclear Modi cation of the Parton Distribution Functions

While descriptions of nuclear structure were being developed using the parton model and studies of deep-inelastic scattering, it was initially thought that nuclear effects would be largely neg-

ligible [67]. This argument, that the energy scale of nuclear binding effects (10 MeV) was much smaller than the scale of DIS momentum-transfers ($Q^2 \gtrsim 1$ GeV), was proven wrong when the EMC collaboration published first results on the modification of the structure functions in iron [3]. Since then, the field of nuclear modifications to the PDFs has rapidly grown into a significant source of insight into nuclear physics, perturbative QCD, non-perturbative QCD, and the inter-play between all three.

Analogous to the parton distributions of a free proton, $f_i(x, Q^2)$, discussed up to this point, one can define parton distributions of larger nuclei as:

$$f_i^{A/Z}(x, Q^2) = Z f_i^{p/A}(x, Q^2) + (A - Z) f_i^{n/A}(x, Q^2) \tag{2.45}$$

Where $f_i^{A/Z}$ is the nuclear PDF for a nucleus with atomic number Z and mass A, $f_i^{p/A}$ is the modified proton PDF for a nucleus with mass number A, and $f_i^{n/A}$ is the modified neutron PDF. The modified neutron PDF is calculated by applying isospin inversion ($u \leftrightarrow d$) to the modified proton PDF. The ratio between $f_i^{p/A}$ and f_i^{p} is referred to as the nuclear modification of the PDF, an effect where, even accounting for isospin composition and the different number of nucleons, the PDFs of a proton bound within a nucleus differ from those of a free proton.

Differences between the nucleon PDFs in nuclei and the free nucleon PDFs can be broken down into four regions due to the shape of the modification, which is shown in Figure 2.8. These regions, from lowest to highest x, are: shadowing ($x \lesssim 0.02$), anti-shadowing ($0.02 \lesssim x \lesssim 0.3$), EMC ($0.3 \lesssim x \lesssim 0.7$), and Fermi motion ($x \gtrsim 0.7$). The following section will describe theoretical and experimental progress on understanding these modifications, with the exception of Fermi motion. These corrections, which arise from the motion of nucleons near the Fermi surface of the nucleus, have been expected since the 1970s [68], but they are not relevant to the results presented in this thesis.

Figure 2.8: The rough shape of nPDF modi cations as a function of xed &^{2} from Ref. [69]. The t parameters for the EPPS21 nPDF t are shown, in addition to the nuclear modi cation function.

2.2.1.1 The EMC Effect

The EMC effect, referring to the unexpected suppression of the nuclear structure functions in the region 0•25 . G . 0•7, is named for the European Muon Collaboration, who rst discovered this effect in 1982 [3]. Since that discovery, understanding the source of the EMC effect has been the subject of a substantial theoretical effort [70, 68], for which many explanations were proposed, but two have emerged as the remaining plausible causes in recent decades. The rst explanation involves an -dependent modi cation of mean- eld effects in nuclei [71], relying on the mean eld model which successfully predicts many aspects of low-energy nuclear structure. The other explanation, which has become increasingly favored recently in light of experimental results from Jefferson Lab [72, 73], is that the EMC effect arises as a result of Short-Range Correlations (SRCs) between nucleons [74].

The explanation of the EMC effect in terms of SRCs [74, 72] begins with the idea that pairs of nucleons within a nucleus temporarily uctuate to short-lived correlated pairs. These pairs

Figure 2.9: The correlation between the slope of the structure functions in the EMC region (G 0•7) and the SRC scaling factor a_2 0•35 1 • 3º, from Ref. [73]

of nucleons are highly virtual and have substantial spatial overlap in their parton distributions, which leads to the modi cations of those distributions. Jefferson Lab measurements [73] have demonstrated a suggestive correlation (shown in Figure 2.9) between the strength of the EMC effect and the "SRC scaling factor a_2 1 • 3º, which is related to the relative abundance of SRC pairs in a given nucleus. While this correlation does not directly demonstrate that SRCs are the cause of the suppression, it does suggest that SRCs are important for understanding the source of the EMC modi cation.

2.2.1.2 Nuclear Shadowing and Anti-Shadowing

While the EMC effect was a surprising phenomenon with few available explanations at the time of its discovery, effects which would suggest the presence of nuclear shadowing were rst proposed by Gribov in the 1960s [75]. At the time, Gribov described nuclear shadowing as a multiple-scattering effect in pion-deuteron collisions. More recent descriptions of nuclear shadowing [76, 77] describe similar effects, which manifest through several different approaches to scattering

from several nearby sites. A simplified example of how multiple scatterings lead to a shadowing suppression from Ref. [76] begins with writing down the amplitudes for one-scattering and two-scattering interactions between a parton and a nuclear target, for the case of scalar QCD at high energy. In light-cone coordinates $\mathbf{?}_{\text{¸}}$ (= $?_0$ ¸ $?_l$, ? = $?_0$ $?_l$), the one-scattering amplitude is:

$$
\begin{aligned}
8\mathbf{2}?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o}M_1{}^1@^0 &= 2^1?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o}8\mathbf{C}_{\text{forw}} \int_1^1 3^4G\, d\,{}^1G_{\text{¸}} - G_{\text{¸}}{}^{o}4^{\,8G\cdot ?^0\ \ ?^o} \\
&= 2^1?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o}8\mathbf{C}_{\text{forw}} \int_1^1 3^2G_{\text{¸}})\ {}^1G_{\text{¸}}{}^{o}4^{\,8G_{\text{¸}}\,{}^1?_{\text{¸}}^{0}\ \ ?_{\text{¸}}{}^{o}}
\end{aligned}
\tag{2.46}
$$

Where d is the nuclear spatial density, ${}^1G_{\text{¸}}{}^{o}$ is the nuclear thickness at some impact parameter $jG_{\text{¸}}j$, $8\mathbf{C}_{\text{forw}}$ = f_{NN} for the nucleon-nucleon scattering cross-section f_{NN}, ? is the incoming projectile momentum, and $?^0$ is the outgoing momentum $2^1?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o} = 4c?_{\text{¸}}X^1?_{\text{¸}}^{0}\ \ ?_{\text{¸}}{}^{o}$ is a normalization factor. Then, the two-scattering contribution (for scattering sites G_1 and G_2) can be written down, and it simplifies to:

$$
\begin{aligned}
8\mathbf{2}?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o}M_2{}^1@^0 = 2^1?_{\text{¸}} - ?_{\text{¸}}^{0}{}^{o}\ {}^1\ \ 1^o{}^18\mathbf{C}_{\text{forw}}{}^{o2} \int_1^1 \frac{3^{2:)}}{{}^12c^{o2}} 3G_{1\text{¸}}\ 3G_{2\text{¸}}\ 3^2G_{1)}\ 3^2G_{2)}\ 4^{\,8;^2\,{}^1G_{2\text{¸}}\ \ G_{1\text{¸}}{}^{o}\bullet{}^12?_{\text{¸}}{}^{o}} \\
4^{\,8\times G_{1)}\ {}^1:)\ \ ?_{)}{}^{o}\text{¸}\ G_{2)}\ {}^1?_{)}^{0}\ \ :_{)}{}^{o}1/4} \\
d\ {}^1G_{1\text{¸}} - G_{1)}{}^{o} d\ {}^1G_{2\text{¸}} - G_{2)}{}^{o}\ {}^1G_{2\text{¸}}\ \ G_{1\text{¸}}{}^{o}
\end{aligned}
\tag{2.47}
$$

In particular, it is helpful to note that the first exponential in the RHS of Equation 2.47 encodes the coherence effects of the multiple-scattering, with coherence length $?_{\text{¸}}$ = $2?_{\text{¸}}\bullet :_{)}$. Hence, in the limit $?_{\text{¸}}$ ŸŸ $:_{)}$, coherence effects vanish and the multiple-scattering contributions to the amplitude go to zero [76], causing the cross-section to only result from the superposition of single-scatterings, referred to as the incoherent limit. For this reason, shadowing effects vanish for large values of $G_{\text{¸}}$ which correspond to small values of $?_{\text{¸}}$. Meanwhile, the completely coherent $?_{\text{¸}}$ ¡¡ $:_{)}$ limit simplifies the two-scattering amplitude to:

$$
8M_2{}^1@^0 = \frac{{}^1\ \ 1^o}{2}{}^18\mathbf{C}_{\text{forw}}{}^{o2} \int_1^1 3^2G_{\text{¸}}\,4^{\,8G_{\text{¸}}\,{}^1?_{\text{¸}}^{0}\ \ ?_{\text{¸}}{}^{o}})\ {}^21G_{\text{¸}}{}^{o}
\tag{2.48}
$$

Figure 2.10: The different diagrams which contribute to the leading twist calculation of nuclear shadowing, from Ref. [77]. Diagrams (a) is the cross-section from the impulse approximation, while diagrams (b) and (c) yield the shadowing correction due to coherent multiple scattering effects.

In the forward limit, the one-scattering and two-scattering cross-sections (f^{11o}, f^{12o}) then are:

$$f^{11o} = f_{NN} - f^{12o} = \frac{1}{2}\left(1^{o}\right)^{1} \quad 3^{2}G_{j}(x) \quad {}^{1}G_{j}^{o}(f) \tag{2.49}$$

The fact that f^{12o} is negative demonstrates that the multiple-scattering effects cause suppression, allowing it to explain observed nuclear shadowing effects. It also demonstrates the predicted scaling of the shadowing effects with and a dependence on the collision centrality encoded in the integral over impact parameters. While this toy model is less sophisticated than full models of shadowing, it helps describe how different shadowing models re-produce these dependencies in the data and how the details of each model can in uence the speci c predictions for those trends.

One such leading theory describing nuclear shadowing phenomena is the Leading Twist Approximation (LTA) [77, 78], which differs from most other models since it is consistent with DGLAP evolution at large Q^2 and accounts for differences in the modi cations of quark vs. gluon nPDFs. This theory relies on three components: Gribov's original theory of shadowing corrections [75], QCD factorization, and experimental inputs from studies of diffraction in ep collisions at HERA. Gribov describes the scattering of a pion from a deuteron receiving additional corrections through pomeron exchanges with both nucleons of the deuteron. This situation is analogous to a virtual photon which uctuates hadronically as it interacts with a nucleus and scatters multiple

times from adjacent scattering sites, as shown diagramatically in Figure 2.10. This multiple scattering yields a suppression like that of the toy model (Equation 2.47), but the LTA then relates the cross-section for two-scattering to the diffractive structure functions. Through this relationship with the diffractive structure functions and QCD factorization, the expected shadowing modi cations can be extracted from diffractive data from HERA.

2.2.1.3 Global Fits to Determine nPDFs

Separate from the theoretical interpretations of different aspects of nuclear PDF modi cations, they are quanti ed through a global tting procedure similar to the one described in Section 2.1.5. Several different ts which use a variety of different approaches and subsets of the available data exist, and a recent summary is provided in Ref. [79]. The most recent results from the main collaborations performing these ts are the nCTEQ15HQ [80], EPPS21 [69], nNNPDF3.0 [81], and TUJU21 [82] ts. This section will attempt to provide some details about the key differences between these approaches, brie y discussing each methodology, and then describe the extent of the differences between results and uncertainties within each result for these state-of-the-art ts to the nPDFs. Apart from methodology, differences in the data incorporated in each t are a key factor which yields differences in the t results, and the datasets included in each are presented in Table 2.2. This table is not an exhaustive list of available data which could be incorporated in nPDF ts, but rather an overview of the different inputs to each model.

The general tting procedure applied is similar to the one described in Section 2.1.5, except for certain accommodations which must be made for the nuclear target. The tted nuclear PDFs must be corrected for isospin effects and combined according to Equation 2.45, and then the resulting PDFs for a given nucleus can be used to compute t observables. Three of the procedures discussed here (nCTEQ15, EPPS21, and TUJU21) use the Hessian method [61, 62] to minimize the χ^2 and determine uncertainties, while nNNPDF uses gradient descent and machine learning optimization methods, with Monte Carlo replicas to compute uncertainties [63].

The nCTEQ15 t is an NLO analysis of nPDF data with successive iterations including the

Dataset Description	nCTEQ15HQ	EPPS21	nNNPDF3.0	TUJU21
SLAC/FNAL/CERN/Hermes DIS	X	X	X	X
JLab CLAS+Hall C4 DIS	X	X		
CHORUS/CDHSW a DIS	X	X	X	X
NuTeV/CCFR a DIS	X		X	
FNAL ? Drell-Yan	X	X	X	
FNAL/NA3/NA10 c Drell-Yan		X		
STAR+PHENIX Inclusive c ^{0}	X	X		
STAR+PHENIX [, c	X			
PHENIX	X			
LHC / ^{0} Bosons	X	X	X	X
CMS+ALICE , Bosons	X	X	X	X
ALICE c ^{0}, [, c ,	X			
CMS ?Pb Dijets		X	X	
ATLAS ?Pb Prompt Photons			X	
LHCb ?Pb Prompt ^{0}	X	X	X	
ALICE ?Pb Prompt ^{0}	X			
LHC ?Pb •k , k ¹2(º, ¹1(º	X			

Table 2.2: Summary of world data included in recent ts to nPDF data, adapted from Ref. [79].

original using DIS, DY, and RHIC c ^{0} data [83], followed by the subsequent inclusion of JLab DIS data [84], LHC, and / ^{0} boson data [85], light hadron production data [86], and quarkonium production data [80]. It relies on parameterizing nuclear PDFs as [87]:

$$G\,5_8^{2\bullet}\ {}^1G-\&_0^{3o} = 2_{08}G^{2}{}_{181}1\quad G^{o2}{}_{28}4^{2}{}_{38}G_{1}1\,¸\ 4^{2}{}_{48}G^{o2}{}_{58} \tag{2.50}$$

Where the starting scale $\&_{,0}$, is 1•3 GeV, 8 indexes the nPDF avor for the avor basis D_E, 3_E, 6, $D_¸$ 3, and, $B_¸$ B and the $2_{=8}$ parameters are the =th free parameter for the PDF of avor 8 The 2_{08} parameters are constrained in the t via the sum rules for valence quarks and total momentum. The nuclear -dependence is incorporated directly into the $2_{=8}$ parameters, where each parameter is replaced by:

$$2_{=8}!\quad 2_{=8}^{1}\ {}^{o} = ?_{=8}¸\ 0_{=8}^{1}1\qquad {}^{1:8o}-\qquad\qquad : = 1-\bullet\bullet 5 \tag{2.51}$$

The EPPS21 [69] analysis ts the initial modi cation of each nPDF directly and xes the proton

PDF baseline to the result from CT18A [55]. With this approach, the modi ed proton PDFs are written as:

$$5_{8}^{?\bullet}\ {}^{1}G-\&^{2o} = '\ {}_{8}^{?\bullet}\ {}^{1}G-\&^{2o}\ 5_{8}^{?}{}^{1}G-\&^{2o} \tag{2.52}$$

The nuclear modi cation for a PDF of avor 8 is then $'\ {}_{8}^{?\bullet}\ {}^{1}G-\&^{2o}$, for parton avors $8 = D_E, 3_E, 6, D, 3$, and B, assuming $B = B$. Note that since the t prescriptions in each case are for the initial PDF at the reference scale, so there is no explicit $\&^2$ dependence in the functional form. All variation of the nPDFs with $\&^2$ arises due to DGLAP evolution from the tted result at the reference scale. The nuclear modi cations are then tted with a piecewise smooth form meant to explicitly model the different nPDF regions:

$$'\ {}_{8}^{?\bullet}\ {}^{1}G-\&^{2o} = \begin{cases} 0_{08}\ ,\ 0_{18}{}^{1}G\ \ G_{08}{}^{o}\ 4^{\ G0_{28}\bullet G_{08}}\ \ 4^{\ 0_{28}} & G\ \ddot{Y}\ G_{08} \\ 1_{08}G^{1_{18}}1\ \ G^{o\,1_{28}}4^{G\,1_{38}} & G_{08}\ \ddot{Y}\ G\ \ddot{Y}\ G_{48} \\ 2_{08}\ ,\ 2_{18}{}^{1}2_{28}\ \ G^{o\,1}1\ \ G^{o\ V_8} & G_{48}\ \ddot{Y}\ G\ \ddot{Y}1 \end{cases} \tag{2.53}$$

Here, the free parameters G_{08} and G_{48} are the G values of the anti-shadowing maximum and EMC effect minimum, respectively. The variables $0_{=8}$, $1_{=8}$, $2_{=8}$, and V_8 are other free parameters in the t, but when the function is tted, they are expressed in terms of the height parameters H_{08}, H_{08}, and H_{48} which are the extrema of the modi cation G at $= 0 - G_0 - G_4$, respectively. The EPPS t function parameters are visualized in terms of the modi cation on Figure 2.8. The normalization parameters $H_{(08)}$ are constrained via the sum rules for valence quarks and total momentum. Then, the -dependence of the modi cations is encoded by replacing each height parameter with:

$$H_{=8}! \quad H_{=8}{}^{1}\ {}_{ref}{}^{o}\ \ 1\ \ \frac{}{{}_{ref}}{}^{W_8} \tag{2.54}$$

This parameterization uses $_{ref} = 12$ and assumes a power-law scaling with the tted parameter W_8 0. This scaling behavior holds well for larger nuclei but the monotonicity assumption breaks down for certain small nuclei, in which case, small adjustments are made to this behavior. The

modi cation is only tted for nuclei with 3, and smaller nuclei are assumed to have no modi cation.

The nNNPDF 3.0 [81] analysis takes a different approach to the other models, parameterizing the nuclear PDFs using a neural network to describe the dependence. Taking $Q_0 = 1$ GeV, the nPDFs are parameterized as:

$$G_{8}^{5}\left(\frac{2}{8}, {}^{1}G, \frac{3}{0}\right) = {}_{8}G^{U_{81}}\left(1 - G^{V_{8}}\right) NN_{8}\left({}^{1}G, {}^{o}\right) \tag{2.55}$$

In this description, $NN_8(G, {}^{o})$ represents the output of a neuron in the output layer of the neural network for the PDF of avor 8. The avor basis used in nNNPDF is a specialized linear combination of parton avors called the "evolution basis". This parameterization uses the parameters to enforce the valence and momentum sum rules, and the "pre-processing exponents", U_8, V_8, are free parameters needed to control the behavior at small or large values of G, which are tted simultaneously with the neural network parameters.

The TUJU21 [82] nPDF analysis is the only fully NNLO t for nPDFs presented here, where both the calculations to compare to data when computing χ^2 contributions and the DGLAP evolution use NNLO results. They nd that the χ^2 per degree of freedom for the t signi cantly improves for the NNLO analysis, compared to their analogous procedure using NLO calculations. This requirement for NNLO calculations restricts the data which this t is able to include, as shown in Table 2.2. They parameterize the modi ed proton PDFs using a straightforward functional form and starting scale $Q_0 = 1.3$ GeV:

$$G_{8}^{5}\left(\frac{2}{8}, {}^{1}G, \frac{3}{0}\right) = {}_{08}G^{2_{181}}\left(1 - G^{\rho_{2281}}\right) \cdot {}_{238}G \cdot {}_{248}G^{2_{o}} \tag{2.56}$$

The avor basis used for this parameterization is $8 = D_E, 3_E, 6, D, 3, B$. They had to assume, however, that $D = 3 = B = B$ due to the smaller amount of data incorporated in this t. The -dependence of the nPDFs is parameterized similarly to the method used for nCTEQ15.

Since each of these procedures takes quite a different approach to the t parameterization,

Figure 2.11: A recent comparison of the results from three leading global nPDF ts as a function of G for $\&^2 = 10\,\mathrm{GeV}^2$., from Ref. [79]

datasets included, and optimization method, it is unsurprising that the resulting nPDFs differ substantially. These differences are apparent in Figure 2.11 shows comparisons of the different nPDF modi cations for the three largest ts as a function of G for $\&^2 = 10\,\mathrm{GeV}^2$. The uncertainties on each of the light quark nPDFs are large in EPPS21 and nNNPDF 3.0, causing them to agree within the stated uncertainties, in spite of substantial differences in some regions. The gluon nPDF typically has much better-controlled uncertainties, but in this case, the different nPDF models disagree within their stated uncertainties in certain regions. In particular, the shadowing effects in nNNPDF and EPPS are much stronger than those observed by nCTEQ. Then, nCTEQ and EPPS have a similar size and position of the anti-shadowing maximum, while nNNPDF differs substantially, accompanied by a larger uncertainty. The data presented in this measurement, described in Chapter 7, bring new information to bear on these discrepancies.

2.2.2 Gluon Saturation at Low x

Gluon saturation [88] is a theoretical limit on the gluon densities which must arise due to unitarity bounds on the scattering cross-section. These constraints arise directly from the quantum mechanical black-disk limit, which states:

$$f_{tot} \quad 2c' \; ^{2} \tag{2.57}$$

This requirement on unitarity translates to the Froissart-Martin bound [89] in QCD, which requires that the hadronic cross-section not grow faster than $\ln^2 s$. To observe how this directly restricts the growth of gluon distributions at low x, it is useful to introduce the Balitsky-Fadin-Kuraev-Lipatov (BFKL) evolution equations. Analogous to the DGLAP equations which dictate the evolution of the PDFs with Q^2, the BFKL equations describe the evolution of the PDFs with low Q^2 [90, 91]. The evolution described by these equations is restricted to low x, and it predicts a growth in the number density of gluons proportional to $x^{-\lambda}$, where $\lambda > 0$ [90]. This BFKL growth in the gluon density would yield a growth in the cross section where $\sigma_{tot} \sim s^{\lambda}$, which violates the Froissart bound.

The modi cation of gluon PDF evolution at low x required to prevent this violation of the Froissart bound is referred to as gluon saturation, where the gluon densities “saturate”, rapidly reducing their growth rate. The mechanism by which this happens is similar to that of nuclear shadowing, since it has been noted [92] that the multiple-scattering effects which are responsible for nuclear shadowing are the same as recombination effects in the in nite-momentum frame. In the case of saturation, suf ciently high gluon densities would induce the non-linear regime where these recombination effects substantially compete with gluon splitting. This new regime of QCD is of great theoretical interest, and it will be explored in much greater depth at the planned Electron-Ion Collider (EIC) [93].

2.3 Hadronic Particle Production

2.3.0.1 Fragmentation Functions

The fragmentation functions, $D_a(z, Q^2)$ are universal functions which represent the probability that an outgoing parton a produces a hadron with a fraction z of the parton's momentum. The scale Q is the factorization scale where the perturbative aspects of the cross-section are cut off. Conceptually, the fragmentation functions are analogous to the PDFs: they are non-perturbative objects which can be measured experimentally and fitted but not determined from first principles in QCD. They are directly accessible through e^+e^- annihilation, due to the simplicity of the initial-state, where the differential cross-section with respect to Bjorken z for a hadron can be expressed as [9]:

$$\frac{d\sigma}{dz} = \sigma_0 \, D(z, Q^2) = \int_z^1 \frac{dy}{y} \, C_a(y, \alpha_s(Q^2)) \, D_a\left(\frac{z}{y}, Q^2\right) \tag{2.58}$$

Where σ_0 is the total cross-section for hadron production, D is the fragmentation structure function (analogous to a DIS structure function), and C_a are perturbatively calculable coefficient functions analogous to the splitting functions in DGLAP evolution. Through this relationship, it can be seen how the cross-section for hadron production in e^+e^- annihilation allows for direct measurement of the fragmentation functions. They exhibit a similar momentum sum rule constraint to the PDFs as well:

$$\int_0^1 dz \, z \, D_a(z, Q^2) = 1 \tag{2.59}$$

While the PDFs serve directly as inputs to Monte Carlo event generators used to simulate processes in QCD, the fragmentation functions typically do not. Instead, they are modelled as the combination of two phenomena: the parton shower, where a high-energy parton perturbatively radiates additional partons, and hadronization, where the parton configuration resulting from the shower hadronizes, producing a final set of outgoing hadrons. The final-state parton shower can be identified with the perturbative evolution of the fragmentation functions in 2.58, while hadronization describes the non-perturbative behavior.

2.3.1 Parton Showers: Fundamentals and Models

While factorization in QCD allows for a separation of scales between the hardest processes producing partons and the softer physics of hadronization, the two are connected by an intermediate step described by parton showers [94, 9]. Parton showering is the process by which a high-energy parton radiates additional, lower energy partons. Each radiated parton may then subsequently radiate as well, until the energies are suf ciently low that perturbative physics can no longer describe the splitting, and phenomenological models of the non-perturbative hadronization process are applied.

To understand the mechanics of parton showers, it is instructive to consider the simplest possible case: 4 4 ! @@ 6 proceeding via the production of @@ pair and subsequent radiation of a gluon by one quark. The cross-section as a function of of emission angle \ and momentum fraction I of the emitted gluon relative to the parent quark is [94]:

$$\frac{df_{@@6}}{d\cos\backslash\, dI} \,'\, f_{@@} \frac{4}{3}\frac{U_B}{2c}\frac{2}{\sin^2\backslash}\frac{1\,,\,{}^{1}1 \quad I^{o2}}{I} = f_{@@}\frac{U_B}{2c}\frac{2}{\sin^2\backslash}\%_{6@}{}^{1}I^{o} \tag{2.60}$$

This splitting probability is directly connected to the DGLAP splitting function, and these functions are generally important for calculating splitting probabilities in parton showers as well as in PDF evolution. It should be noted, though, that these splitting functions are only valid in the collinear limit [94], from which deviations will be relevant for describing parton showers. Expressions like 2.60 fully generalize, allowing for the general cross-section of some hard process with an additional parton, 9, to be written in terms of the hard process cross-section and splitting functions as [94]:

$$df_9 \,'\, f_0 \int_8^{\tilde{O}} \frac{U_B}{2c}\frac{d\backslash^2}{\backslash^2}dI\%_{98}{}^{1}I^{o}dq \tag{2.61}$$

For an emission at some azimuthal angle about the radiating parton, \. Expressions of this kind clearly diverge in three cases: ! 0, \ ! 0, and \ ! c. These divergences are referred to as the infrared (! 0) and collinear (! 0– c) divergences (IRC). Since the probability diverges

to produce a perfectly collinear or $! \quad 0$ parton, it is important to only measure "IRC-safe" observables, which are insensitive to splittings in these limits. IRC safety can be formally de ned for an observable which is a function of particle four-momenta, $O_=$, as [9]:

$$
\begin{aligned}
O_={}^1 ?_1 - \cdots - ?_B ? \cdots ?_={}^p &= O_={}_{,\,1}{}^1 ?_1 - \cdots - ?_B ?^1 1 \quad _{\_}{}^o ?_B - \cdots - ?_={}^p - _{\_}! \quad 0 \\
O_={}^1 ?_1 - \cdots - ?_0 ?_{,} \; ?_1 - \cdots ?_={}^p &= O_={}_{,\,1}{}^1 ?_1 - \cdots - ?_0 ? \; ?_1 - \cdots - ?_={}^p - ?_0 \; k \; ?_1
\end{aligned}
\tag{2.62}
$$

For IRC safe observables, calculations can be organized so the soft and collinear divergences cancel, yielding nite predictions. These divergences are resolved in parton shower calculations by introducing a resolution scale $\&_0$. Splitting probabilities are then only calculated for splittings where the two partons have relative transverse momentum greater than $\&_0$, meaning they are distinguishable from each other. One can then begin to write down splitting probabilities, neglecting the divergent region by restricting the bounds of integration using the cut-off scale. The probability dP_8 for branching for a parton of type 8 with virtuality between $@^2$ and $@^2 {}_{,} \; d@^2$, using this approach, is:

$$
dP_8 = \frac{U_B}{2c} \frac{d@^2}{@^2} \tilde{O}^{\,1}_{\;9} \; \frac{1 \; \&_0^2 \bullet @^2}{\&_0^2 \bullet @^2} \; dI\%_{98}
\tag{2.63}
$$

Then, the probability for the rst, hardest splitting to $@^2$ is the probability for no harder splittings to occur. This probability, known as the Sudakov form factor, can be written for a parton 8 with virtuality $\&^2$ as:

$$
{}_8{}^1 \&^2 - \&_0^{2o} = 4^{\;\frac{\&^2}{\&_0^2} \frac{d:^2}{:^2} \frac{U_B{}^1 :^{2o}}{2c} \tilde{O}^{\;1}_{\;\&_0^2 \bullet :^2}} \; dI\%_{98} I^o \; {}_{,} \; 8 = @ \; 4^{\;\frac{U_B}{2c} \ln^2 1 \&^2 \bullet \&_0^{2o}}
\tag{2.64}
$$

These Sudakov form factors provide the basis for an iterative, MC-based implementation of parton showers. It is important to recognize that the formulation of ${}_8$ relies on leading-order splitting functions, but by imposing a unitarity requirement on the splitting probability, it effectively sums splittings of all orders in U_B to determine the non-splitting probability. Thus, the application of parton showers is not strictly LO. This lack of a xed order in the calculation is important, since the

parton shower is responsible for transitions from kinematic regions where higher-order corrections are small to the $\& \ \&_{QCD}$ region, where such a xed-order calculation would have large errors. In general, errors do occur in the approximation used to determine the Sudakov form factor, which have a form that appears in the right-hand size of Equation 2.64. The Sudakov form-factor, along with many event-shape observables, is an expansion in U_B and logarithms of $\&^2 \bullet \&_0^2$ [95]:

$$\ln\left(\frac{\&^2 - \&_0^2}{\&^2 \bullet \&_0^2}\right) = \ln\left(\frac{U_B}{2c}\right) = \ln\left(\frac{\&^2 \bullet \&_0^2}{U_B}\right) + 6_1 U_B \ln\left(\frac{\&^2 \bullet \&_0^2}{U_B}\right) + 6_2 U_B \ln\left(\frac{\&^2 \bullet \&_0^2}{U_B}\right) + 6_3 U_B \ln\left(\frac{\&^2 \bullet \&_0^2}{U_B}\right) + \cdots \quad (2.65)$$

So, although the parton shower Monte Carlo approach will sum contributions at all orders in U_B, it only includes the leading-log terms of the expansion in Equation 2.65. These logarithms are potentially large, since partons may be emitted with virtuality much larger than the cutoff scale, $\&_0$, so the resummation of additional logarithms becomes important for precisely describing QCD nal-states [95]. It is also important to note that the coupling itself evolves, as is evident in the dependence of U_B in Equation 2.64. This evolution indicates that parton showering becomes much faster as $\&^2$ decreases throughout the showering process.

Because of running coupling effects, $\&_0$ cannot allow for emissions which are in a region with $U_B \quad 1$, meaning $\&_0$ must be substantially larger than $\&_{QCD}$. For this reason, $\&_0$ is not a cut-off in the typical sense of regulating QFT calculations, but rather a scale which applies non-perturbative corrections [94] proportional to $\&_0 \bullet \&^{?}$, where $? \quad 1$ is a constant depending on the parton shower algorithm and observable. These differences re ect how the parton shower connects the perturbative and non-perturbative descriptions relevant for calculating measurable QCD cross-sections.

Because of the competing soft and collinear divergences in the splitting probabilities, the ordering with which one simulates the branchings of a parton shower is important for maintaining consistency with the collinear approach that yields Equation 2.64. One can understand this ordering by imagining the case of two separate splittings from a single initiating quark: one soft, wide-angle gluon and one collinear, reasonably hard gluon. In the case that the hard splitting

occurs first, the soft gluon can be radiated from either parton, and the amplitudes for these two emissions interfere with one another. This interference can be shown [94] to be destructive, such that it provides the same result as emission by a single collinear quark. Since the interference between the two production mechanisms is not included in the algorithmic approach of repeating Equation 2.64 for subsequent splittings, one must order the splittings such that the soft, wide-angle gluon is emitted first. This picture is then consistent with the correctly accounting for interference effects in its emission.

The ordering of branchings is not necessarily always angular, although angular shower models [96] are often used to describe physical parton showers well. A popular alternative choice is the dipole shower [97], where the soft and collinear divergences are rewritten as a color dipole factorization. Many generators such as PYTHIA 8 [98] use other shower ordering variables, such as $?_T$ or virtuality, and in these cases, dedicated methods are applied to enforce consistency with QCD coherence [99].

2.3.1.1 Initial-State Parton Showers

In Section 2.1.5, it is discussed that the PDFs at higher values of Q^2 are calculated by evolving them at some starting scale via the splitting function. These splitting functions, however, produce an additional parton for each splitting, aside from the parton which evolves down to produce the DGLAP-evolved PDF. In an un-perturbed hadron, these additional partons all collapse back into the hadron, but when a hard-scattering process probes the hadron's PDF by pulling a parton out of it, all the additional partons associated with the evolution are emitted as accompanying radiation.

An important practical note about describing this initial-state radiation (ISR) is that the actual simulation of it works in reverse, starting with the hard parton probed in the PDF and tracing its DGLAP evolution backwards. This approach [100, 101] can be guided by the solution to the DGLAP equations and provides an ISR Sudakov form factor for emissions of virtuality Q^2 from a

radiating parton with virtuality $\&^2$ and momentum fraction G as:

$$\&^2 = \frac{d:^2}{:^2} \frac{UB^1:^2}{2c} \int_{\&_0^2 \cdot :^2}^{\&^2 \cdot :^2} d\% \, \frac{1 + \frac{1}{2}\left(1 - \frac{\&_0^2}{\&^2}\right)}{1 - \&_0^2 \cdot :^2} \qquad (2.66)$$

The same arguments about angular ordering and coherence of the parton shower still apply to ISR showers, and similar methods or organizing the sequential shower are useful. While the ISR emitted this way is conceptually separate from the final-state parton shower, it is not a separation which exists in QCD, only in the conceptual frameworks used to describe it [94].

2.3.2 Models of Hadronization

2.3.2.1 The Lund String Model

The Lund string model of hadronization [102, 103] is based on the observation described in Section 2.1.2 that the non-perturbative confining force between a color and anti-color charge, such as a $q\bar{q}$ pair, is described by a string-like color flux tube with energy proportional to their separation. In this picture, two color-connected partons propagating away from each other would thus be connected by a string storing linearly increasing energy as the partons separate further. Longer strings will store sufficient energy in their tension that they will break, producing new $q\bar{q}$ pairs from the vacuum, as depicted in the left panel of Figure 2.12. These hadrons are characterized by two adjacent vertices 8 and $8-1$, which are correlated through the hadron mass $<$, and the QCD string tension parameter Λ, in a simplified 1+1 dimensional picture as:

$$<_8^2 \cdot \Lambda^2 = (G_8 - G_{8-1})^2 - (G_8 - G_{8-1})^2 \qquad (2.67)$$

The string composed of these vertices will fragment into hadrons, the differential probability for which is given by [104]:

$$dP \propto \prod_{8=1} \# d^2?_8 \, X^2 \, ?_8^2 \, \exp\left(-X^2 \sum_8 ?_8 \, \%_{tot}\right) \qquad (2.68)$$

Figure 2.12: Left: The spacetime picture of string fragmentation in 1+1 dimensions in the Lund string model. Right: Formation of hadrons in the string fragmentation model by “popcorn production”. Both gures are reproduced from Ref. [98].

Where ?8 are the momenta of each outgoing hadron, 8%tot is the total momentum, is the area of the string in units of ^ prior to breaking (denoted on the left panel of Figure 2.12), and a tunable parameter. By imposing the requirement that string fragmentation must be the same result when fragmenting from left to right or right to left, one can derive the the light-cone momentum of each hadron (now in 3+1 dimensions) as:

$$5^{1}I^{o} / \frac{^{1}1 \quad I^{o0}}{I} 4 \frac{1^{1<2}, ?_{T}^{2o}}{I} \tag{2.69}$$

Where0 is an additional tunable parameter. Equation 2.69 is known as the Lund symmetric fragmentation function, and it provides the basis for the Lund string fragmentation model. While it is obvious how this process produces mesons, it is less clear how baryons are produced. This feature of hadronization is achieved through allowing diquark-diquark string breakings, in addition to simple @@pair production [105]. This feature is implemented in PYTHIA 8 through the “popcorn” model [106], where diquarks can be produced in pair production along the string without breaking it. This process is demonstrated as a series of steps a through d in the right panel of Figure 2.12. While these basic features describe most core principles of the Lund string fragmentation model, it has many complexities in its implementation to describe real data, which are detailed in Ref. [98].

2.3.2.2 Cluster Hadronization

While the Lund string model has been highly successful for describing hadronization processes across a variety of systems, other phenomenological models of hadronization based on alternative QCD-inspired pictures of this non-perturbative process are useful. The leading alternate hadronization model, cluster hadronization, is based on the precon nement picture of hadronization [107, 108, 109]. This model begins by non-perturbatively splitting all remaining gluons after the parton shower into@@pairs. Then, only color-connected quarks remain, and each color-singlet parton pair forms a cluster, where precon nement states that the mass distribution of these clusters is independent of the preceding perturbative process [107]. Clusters are then treated as excited hadron resonances and decayed to produce a collection of nal-state hadrons.

2.4 Properties of Relativistic Heavy Ion Collisions

Ultra-relativistic heavy ion collisions provide a unique and powerful laboratory for studying the properties of QCD. For the purposes of this thesis, such collisions will be divided into two varieties: “typical” collisions where the nuclei overlap and Ultra-Peripheral Collisions (UPCs) where they do not. Section 2.5 provides a comprehensive overview of the physics of UPCs, and this section will discuss the properties of more typical heavy ion collisions, where nuclear overlap leads to hadronic interactions between the two nuclei. In these processes, the overlap region between the two nuclei produces an enormous energy and baryon density in which a new phase of nuclear matter, Quark Gluon Plasma (QGP), [110] is produced [111].

These collisions are typically described through a geometrical model by Bjorken [112], where the two incoming Lorentz-contracted nuclei are effectively “pancakes” approaching each other at the speed of light. In this picture, the collision occurs in a at, cylindrical region with radius over a formation time, C_5, with a longitudinal length, I. The radius of the region is assumed to be the radius of the the nucleus, which is approximated with 1•2 1•3 fm. This spatial picture can then be related to the energy per unit rapidity by taking $H = dd\backslash\ '$ I•1 2C_5^o. By neglecting interactions

between the different colliding nucleons within the same nucleus, the total energy density,

$$n = \frac{}{+} = \frac{\frac{d\,)}{dH}\,dH}{Ic'^{\,2}}, \quad \frac{\frac{d\,)}{dH}\,\frac{I}{2G_5}}{Ic^{1}1\bullet 2^{\;1\bullet 3o2}} = \frac{1}{c2G_5^{1}1\bullet 2^{\;1\bullet 3o2}}\frac{d\,)}{dH} \tag{2.70}$$

The quantity $d\,) \bullet 3H$ can be directly measured in heavy ion collisions, and it is related to the temperature of the plasma. It was found to be 138 GeV in $\sqrt{B_{NN}} = 130$ GeV Au+Au collisions at RHIC by PHENIX [113] and 1737 GeV in $\sqrt{B_{NN}} = 2\bullet 76$ TeV Pb+Pb collisions at the LHC by ALICE [114]. If one assumes a formation time on the order of 1 fm at RHIC and 0•1 fm at the LHC, these measurements would correspond to energy densities of 5 GeV/fm and 130 GeV/fm, respectively. These energy densities are directly connected to the energy density of the plasma, where the temperature is expected to vary as $n^{1\bullet 4}$. Thus, even though the energy densities at RHIC and the LHC differ by over a factor of 20, the temperatures differ by not much more than a factor of 2.

2.4.1 Collision Centrality and Geometry

In relativistic heavy ion collisions, the impact parameter of the collision may vary such that the nuclei do not fully overlap. The size and shape of the overlap region in a collision of two nuclei, A and B, with impact parameter $4\ddot{Y}\,2'$, characterizes the total particle production and initial geometry of the interaction region. Particle production in the collision will exhibit a trivial geometric scaling with the total nuclear overlap, and probes sensitive to the formation of QGP will be quite sensitive to the shape and volume of QGP formed. Thus, it is helpful to de ne the nuclear overlap function, $)_{AB}{}^{1\,®}$, which describes the total ux of nucleons in the collision as a function of impact parameter, which depends on the nuclear geometry. The relationship between $)_{AB}$ and properties of heavy ion collisions is studied through the Glauber model [115, 116, 117].

The Glauber model requires a description of the radial nucleon density within a nucleus, for

which a Fermi function description is employed [116]:

$$d(r) = d_0 \frac{1 - F(r - r_0)^2}{1 - 4\frac{r - r_0}{a_0}} \tag{2.71}$$

Where d_0 is the nucleon density at the center of a nucleus, F characterizes deviations from a purely spherical nucleus, and a_0 is the nuclear "skin depth". The F parameter is typically 0 for the species of nuclei collided in heavy ion collisions (and certainly is for those relevant in this thesis), so setting F to 0 recovers the Woods-Saxon [118] form for the nuclear thickness. ^{208}Pb, the ion species used for the results in this thesis, the parameters of its Woods-Saxon distribution are [119] $r_0 = 6.62 \pm 0.06$ fm and $a_0 = 0.546 \pm 0.01$ fm. The Glauber approach then proceeds by defining the transverse nucleon density function for nucleus A, $T_A(\vec{b})$, as:

$$T_A(\vec{b}) = \frac{1}{} \int_{-\infty}^{\infty} d(\vec{b}, z)\, dz \Rightarrow \int T_A(\vec{b})\, d^2b = 1 \tag{2.72}$$

where the RHS indicates the normalization of $T_A(\vec{b})$ to unity. The transverse densities of the two nuclei are used to define T_{AB} by writing the product of the probability to simultaneously find a nucleon from each nucleus at a position $\vec{b}$ with respect to nucleus A:

$$T_{AB}(\vec{b}) = \int d^2b'\, T_A(\vec{b}')\, T_B(\vec{b} - \vec{b}') \tag{2.73}$$

A diagram of the geometry of a nucleus-nucleus collision used to calculate T_{AB} is shown in Figure 2.13. The quantity $T_{AB}(\vec{b})$ has units of inverse area, and since it represents the flux of nucleons from both nuclei per unit area, the probability of a nucleon-nucleon interaction can be written as $P_{NN} = T_{AB}\, \sigma_{NN}^{inel}$, where σ_{NN}^{inel} is the total inelastic nucleon-nucleon cross-section. Given this probability, the probability distribution to have ν nucleon-nucleon interactions as a function of impact parameter is described by a binomial distribution:

$$P_\nu(\vec{b}) = \binom{N_{coll}}{\nu} \left(T_{AB}(\vec{b})\, \sigma_{NN}^{inel}\right)^{\nu} \left(1 - T_{AB}(\vec{b})\, \sigma_{NN}^{inel}\right)^{N_{coll} - \nu} \tag{2.74}$$

Figure 2.13: A diagram of a hadronic heavy ion collision from a side view (left) and front view (right). The impact parameter of the nucleus-nucleus collision (is different from the impact parameter of the NN collision with respect to the nucleus A (or nucleus B ( ). Reproduced from Ref. [116].

This gives the probability of collisions, where can be as low as 0 and as high as , the latter being the case where all nucleons in one nucleus each scatter inelastically from every nucleon in the other nucleus. The probability of any inelastic interaction can thus be found by summing all = 1, which gives:

$$\frac{d^2 f^{inel}_{AB}}{d^2 1} = \tilde{O} \quad \% 1 = -1^{\circ} = 1 \quad 1 \quad)_{AB}{}^{1} 1^{\circ} f^{NN}_{inel} \tag{2.75}$$

This probability is very important in the context of UPCs, where events are required to have no hadronic interactions. For unpolarized nuclei, $)_{AB}{}^{1} 1^{\circ}$ can be written as a function of the scalar distance. This thesis will only consider the case of unpolarized collisions. Then, the total inelastic cross-section is:

$$f^{AB}_{inel} = \int_0^1 2c1 d1 \; 1 \quad 1 \quad)_{AB}{}^{1} 1^{o} f^{NN}_{inel} \tag{2.76}$$

We can then de ne $\#_{coll}{}^{1} 1^{o}$, the average number of nucleon-nucleon collisions as a function of

impact parameter, as:

$$\#_{coll}(b) = \sum_{i=1}^{\infty} \sigma_i(b) = \int T_{AB}(b)\,\sigma_{NN}^{inel} \tag{2.77}$$

Another useful quantity is the number of participants, $\#_{part}(b)$, which is defined as the total number nucleons from nuclei A and B which interact. It can be calculated by summing the participants from both nuclei:

$$\#_{part}(b) = \int \left[T_A(\vec{b}) \left(1 - \left(1 - T_B(\vec{b}-\vec{b})\,\sigma_{inel}^{NN}\right)^{\#_{coll}(b)}\right) + T_B(\vec{b}-\vec{b}) \left(1 - \left(1 - T_A(\vec{b})\,\sigma_{inel}^{NN}\right)^{\#_{coll}(b)}\right)\right] d^2b \tag{2.78}$$

All of the equations described relating to Glauber model so far have been in the optical limit, where one assumes continuous nucleon density distributions. Alternatively, the Glauber Monte Carlo approach [117] treats the Woods-Saxon distribution of nucleons within the nucleus as a probability distribution, shedding the requirement of continuous densities. Without this assumption, the MC approach is able to describe the features of actual collisions more accurately. Each Glauber MC event randomly samples b from the distribution of impact parameters, and the requirement for a collision to take place between any two nucleons separated by a distance d is:

$$d \le \sqrt{\frac{\sigma_{inel}^{NN}}{\pi}} \tag{2.79}$$

$\#_{coll}$ is simply the number of NN pairs which satisfy this condition and $\#_{part}$ is the number of nucleons for which this condition is satisfied at least once.

In order to connect this information about the Glauber model to measurements at colliders, the quantities $\#_{coll}$ and $\#_{part}$ must be connected to experimentally measurable observables. The most common two approaches are measuring forward neutron multiplicities or total transverse energy in some region of the detector. ATLAS uses the latter approach [120]. In this approach, the distribution of ΣE_T in the Forward Calorimeters is constructed for all hadronic collisions. Then, it is sub-divided into percentiles (centrality intervals), where a lower percentile is a more central,

higher overlap collision. A higher percentile is a more peripheral, lower overlap collision. The relationship between centrality and all of the different quantities described here is shown in Figure 2.14.

Figure 2.14: The relationship between different collision geometry and activity properties in Glauber MC from Ref. [116].

2.4.2 The Quark-Gluon Plasma

Following the development of the quark and parton models, it became clear that QCD would have a unique phase structure due to asymptotic freedom and con nement [121]. It was theorized that in the high-density, high-temperature matter produced in heavy ion collisions, a new state of matter could be formed where the quarks and gluons inside those hadrons would become de-con ned [110]. The study of this novel state of hot, dense nuclear matter, known as the Quark-Gluon Plasma (QGP), has been the cornerstone of heavy ion physics research for decades. The QGP is also believed to have been formed in the early universe, and developing an understanding of its phase structure and transitions would help substantially to understand models of the early universe [122].

Early experimental efforts to produce QGP began with the rst claim of evidence for its production in $\sqrt{s_{NN}} = 20$ GeV in Pb+Pb at the SPS at CERN [123]. The discovery of QGP was more

robustly con rmed by early results from all four experiments at the Brookhaven Relativistic Heavy Ion Collider (RHIC) [124, 125, 126, 127]. While the QGP was decisively discovered at RHIC, new probes only available at the LHC opened up new avenues for studying its properties, most notably the rst observation of jet quenching by ATLAS [128]. Through hard probes like jet quenching, the short-range, perturbative structure of QGP is accessible, which complements measurements of the long-range structure through phenomena such as collective ow [111]. After the discovery of QGP in heavy ion collisions, evidence was also observed for particle correlations which are characteristic of hydrodynamic behavior in ?? [129] and ?Pb collisions [130]. These results have led to a hypothesis that QGP is formed even in these small systems, although not all phenomena related to QGP formation have been observed in small systems. Most notably, no experimental evidence exists for jet quenching phenomena in ?? or ?Pb collisions [131].

Of particular relevance to this thesis are predictions that QGP exhibits the chiral magnetic effect [132], where chirality is induced in the QGP through the Adler-Bell-Jackiw axial anomaly [133]. It is predicted that in the presence of the strong magnetic eld produced in heavy ion collisions, the chiral magnetic effect will generate large currents in the QGP, producing long-lived electromagnetic elds inside the plasma. Results in this thesis presented in Section 7.2 set limits on these effects, and wide-ranging searches were conducted for the chiral magnetic effect by the STAR collaboration [134] at RHIC. No evidence has been found experimentally for the presence of these strong magnetic elds, though, and a recent theoretical explanation for the lack of measurable magnetic elds in the QGP is that Ohm's Law is broken in heavy ion collisions [135]. These calculations demonstrate that the QGP does not fully respond to electromagnetic elds, which suppresses the expected strength of magnetic eld effects by two orders of magnitude relative to the result when one assumes the validity of Ohm's law.

2.5 Photoproduction and Ultra-Peripheral Collisions

In ultra-peripheral heavy ion collisions [1, 2], the impact parameter of the collision is suf - ciently large (1 ¡ 2') that the nuclei do not overlap, and the probability of hadronic interactions

Accelerator	Collision Type	$\sqrt{s_{NN}}$	Maximum ω	Maximum $\gamma\omega$	Maximum $\sqrt{s_{\gamma\gamma}}$
RHIC	Au+Au	200 GeV	320 GeV	25 GeV	6 GeV
	p+Au	200 GeV	1•5 TeV	52 GeV	30 GeV
	pp	500 GeV	20 TeV	200 GeV	150 GeV
LHC	Pb+Pb	5•02 TeV	250 TeV	700 GeV	170 GeV
	p+Pb	8•16 TeV	1•1 PeV	1•5 TeV	840 GeV
	pp	14 TeV	16 PeV	5•4 TeV	4•2 TeV

Table 2.3: Parameters of the coherent photons emitted in UPCs for different existing collider facilities. Adapted from Refs. [1, 2]. For asymmetric collisions, the photon is assumed to be emitted from the higher Z nucleus. Energies are quoted in the target rest frame, by convention.

between the nuclei goes to zero. Instead, the dominant form of particle production arises from the intense electromagnetic elds which accompany highly charged ions moving at ultra-relativistic speeds [136]. These electromagnetic elds can be treated using the Weizsacker-Williams approach [137, 138], based on the ideas of Fermi [139]. This approach, also known as the method of virtual quanta, describes the electromagnetic elds as a ux of quasi-real photons. Since the intensity of the coherent photon ux goes as Z^2, the $Z = 82$ Pb ions at the LHC provide a massive increase in the rate for photon-induced processes.

The kinematics of these coherently emitted photons can be estimated through simple uncertainty principle considerations, where the maximum energy of an emitted photon in the nucleus rest frame is $\omega \sim 2\pi\hbar c/R \sim 25$ MeV for Pb. For ions moving at LHC energies, however, the longitudinal component of the photon energy is Lorentz boosted, while the transverse component is not, meaning $\omega \sim 2\gamma\hbar c/R \sim 80$ GeV, while the transverse momentum is $p_T \sim 2\pi\hbar c/R$. The photon Q^2 is also not Lorentz boosted, so the virtuality $Q \sim 2\pi\hbar c/R$, a factor of γ (2595 for Pb ions at the LHC) smaller than ω. The parameters of photon collision-system kinematics for different collisions and accelerator facilities are shown in Table 2.3. Additionally, the electric eld lines which spherically extend outwards in the nuclear rest frame are Lorentz-contracted to point radially outwards in the transverse direction. The magnetic eld, meanwhile, goes along the azimuthal direction, causing the photon ux to be linearly polarized in the transverse direction. The geometry of these collisions is shown in Figure 2.15

Figure 2.15: A side-view diagram of two colliding ions in an Ultra-Peripheral Collision from Ref. [1] (left) and a transverse view from Ref. [140] (right). The right diagram demonstrates the directions of electric and magnetic elds, along with the resulting polarization vectors (

2.5.1 The Equivalent Photon Approximation

While the typical photon kinematics can be estimated from basic uncertainty principle calculations, the Weizsacker-Williams formalism [137, 138, 139, 136] allows one to analytically derive the ux of photons coherently produced by an ion with radiusand charge14 , moving near the speed of light with relativistic boost1W= 11 V2o 12. This method is typically referred to as the Equivalent Photon Approximation (EPA) or method of virtual quanta. While in reality the photon-emitting nucleus has some nite radius, it is treated as a point charge at the origin for the sake of this derivation, with electric and magnetic elds [136]:

$$H^{1}\mathcal{C} = /4 \frac{W1}{^{1}1^{2} \text{¸} WE\text{-}C^{2o3\bullet 2}} \tag{2.80}$$

$$G^{1}\mathcal{C} = V \; H^{1}\mathcal{C} = /4 \frac{VW1}{^{1}1^{2} \text{¸} WE\text{-}C^{2o3\bullet 2}} \; ' \; /4 \frac{W1}{^{1}1^{2} \text{¸} WE\text{-}C^{2o3\bullet 2}} \tag{2.81}$$

$$I^{1}\mathcal{C} = \; /4 \frac{WEC}{^{1}1^{2} \text{¸} WE\text{-}C^{2o3\bullet 2}} \tag{2.82}$$

The time-varying electromagnetic elds are equivalent to two pulses (and) of plane polarized radiation, the rst travelling in the direction, co-moving with the charged particle, and the second in the direction, moving away from the charged particle. These equivalent pulses have corresponding frequency spectra given by:

$$\frac{2}{2c}\left|\tilde{E}_1(\omega)\right|^2 = \frac{2}{2c}\left|\tilde{E}_2(\omega)\right|^2 \tag{2.83}$$

Here, is the energy per unit area per unit frequency and denotes the Fourier transform of the electric eld. Direction corresponds to the direction, and directions and correspond to the and directions, respectively. The application of these Fourier transforms yields:

$$\frac{2}{2c}\left|\tilde{E}_1(\omega)\right|^2 = \frac{1}{c^2}\frac{\left(\frac{1}{4\pi\varepsilon_0}\right)^2}{V^2\gamma^2}\frac{1}{\omega^2}\left(\frac{\omega\gamma}{VE}\right)^2 K_1^2\left(\frac{\omega\gamma}{VE}\right) \tag{2.84}$$

$$\frac{2}{2c}\left|\tilde{E}_2(\omega)\right|^2 = \frac{1}{c^2}\frac{\left(\frac{1}{4\pi\varepsilon_0}\right)^2}{V^2\gamma^2}\frac{1}{\omega^2}\left(\frac{1}{V^2}\right)\left(\frac{\omega\gamma}{VE}\right)^2 K_0^2\left(\frac{\omega\gamma}{VE}\right) \tag{2.85}$$

The functions are the th Bessel functions of the second kind. Since is suppressed by a factor of , it is clear that for ultra-relativistic particles, the contribution from dominates, which demonstrates that the radiation is linearly polarized along the direction of motion. Since the impact parameter is not known in a given UPC, the form integrated over impact parameter is much more useful. To derive it, we restrict to impact parameters $\rho \geq \rho_{min}$ and integrate the sum of the intensities:

$$\begin{aligned}\left|\tilde{E}(\omega)\right|^2 &= 2c\int_{\rho_{min}}^{\infty}\rho\,d\rho\,\left[\left|\tilde{E}_1(\omega)\right|^2 + \left|\tilde{E}_2(\omega)\right|^2\right] \\ &= \frac{2}{c}\frac{\left(\frac{1}{4\pi\varepsilon_0}\right)^2}{V^2\gamma^2}\left[K_0(\xi)K_1(\xi) + \frac{V^2\gamma^2}{2}\left(K_1^2(\xi) + K_0^2(\xi)\right)\right]\end{aligned} \tag{2.86}$$

We have de ned the variable $\xi = \lambda_{1\,min} \cdot \gamma / W$. After deriving the intensity, we can derive the number density of virtual quanta (photons) with:

$$\lambda_1 \cdot d\lambda_1 = \hbar\omega_W \lambda_1 \cdot d\lambda_1 \cdot \tag{2.87}$$

This slightly modi es Equation 2.86 (taking $\omega_W = \hbar\lambda_1$) to get:

$$\frac{d\#}{d\omega_W} = \frac{2}{c}\frac{\alpha/4\pi^2}{V^2 2}\frac{1}{\omega_W} \xi K_0(\xi) K_1(\xi) - \frac{V^2\xi^2}{2}\left[K_1^2(\xi) - K_0^2(\xi)\right] \tag{2.88}$$

Then, if one simpli es by taking the ultra-relativistic limit ($V \to 1$) and using natural units ($\hbar = c = 1$), the expression becomes:

$$\frac{d\#}{d\omega_W} = \frac{U_{EM}/\pi^2}{\omega_W} \xi K_0(\xi) K_1(\xi) - \frac{\xi^2}{2}\left[K_1^2(\xi) - K_0^2(\xi)\right] \tag{2.89}$$

This result is still parameterized in terms of $\lambda_{1\,min}$, which is included in the de nition of ξ. In order to have a fully practical result for the photon ux, one must choose a suitable $\lambda_{1\,min}$ which, for the case of coherent emission from a nucleus, is typically R_2. It is also useful to have a form for the photon ux retaining explicit λ_1-dependence, which follows directly from Equations 2.86 and 2.87:

$$\frac{d\#}{d\omega_W d\lambda_1} = \frac{\alpha/\pi^2 U_{EM}}{c^2 V^2}\frac{\xi^2}{\omega_W \lambda_1^2}\left[K_1^2(\xi) - \frac{1}{W^2}K_0^2(\xi)\right] \tag{2.90}$$

Note that λ_1 in this context is the impact parameter with respect to the emitted photon, not the nucleus-nucleus impact parameter. In addition to this standard formulation for the equivalent photon ux, further modi cations must be made to achieve the level of precision required to predict observables in recent UPC measurements, such as those presented in this thesis. Different process-dependent modi cations are made which improve on certain assumptions of the EPA procedure, and the methods through which those modi cations occur are discussed Sections 2.5.3.2 and 2.5.3.4.

2.5.2 The Structure of Quasi-Real Photons

In UPCs, the photons emitted differ from those relevant for DIS, since the photons are nearly real and can be treated theoretically as real photons. In DIS, the virtuality of the photon prevents it from exhibiting the full, complex structure of real photons [141], which is much more relevant in UPCs. This complicated structure arises from the QED coupling of the photon to quarks, through the W@vertex, which allows it to uctuate to@@pairs. In the case of high-virtuality photons, this @@system has an intrinsic hard scale, which allows the uctuation to be described solely through perturbative methods. For low-virtuality photons, however, this system is governed by non-perturbative QCD dynamics and thus has its own non-perturbative structure, similar to that of a hadron described in Section 2.1.5.

Descriptions of the hadronic structure of the photon which describe low-energy processes have existed for a long time, actually pre-dating QCD, through the Vector Meson Dominance (VMD) [142, 143] picture. In this approach, the photon can uctuate to any light vector meson with the same quantum numbers:% = 1 , & EM = = (= 0. This picture is often called Generalized VMD (GVMD) when the photon is allowed to uctuate to heavier vector mesons as well. In principle the GVMD approach could describe the wave-function of the photon at high energies as well as low energies, but matching the many high-mass states onto the partonic structure observed at high& 2 is theoretically challenging. Instead, a more common approach is the one proposed in Ref. [141], where the photon wave-function can be written as:

$$jW\dot{V} = \frac{p}{/_3} jW i \; , \; \tilde{O}_{+=d^0-l-q} \frac{4}{5_+} j+i \; , \; \frac{4}{5_{@@}} j@@ \tag{2.91}$$

The different components of the wave-function are the bare photon, jW i, the vector meson components j+i , and the perturbative @@ uctuations. The corresponding normalization constants are / 3, 5+ , and 5@@ In this description, the VMD part of the wave-function is explicitly restricted to only states with masses light enough that < + o & 1. Ref. [141] explicitly describes these three contributions to the photon structure as the direct, VMD, and “anomalous”. For studying& 2 high-

processes, it often makes sense to refer to the hadronic components as one uni ed contribution, known as the resolved photon.

The coef cients 5+ are directly related to the transition probability to different light meson states, so they can be determined by ts to low-energy scattering experiments. In order to avoid double-counting between the vector meson and perturbative parts near the transition, a cut-off must be introduced as function of mass. This cut-off scale, ?0, is taken in [141] to be ?0 = < q•2 0•5 GeV. ?0 sets the mass of hadronic uctuations at which the photon is assumed to be completely non-perturbative, 4• 5@@ = 0. For higher virtuality, this perturbative component of the wave-function evolves as:

$$\frac{4}{5_{@@}}^{2} = \frac{U_{EM}}{2c}\frac{\#}{3}\ 2^{\tilde{O}}_{@}\ 4^{2}_{@}\ \ln\ \frac{`^{2}}{?^{2}_{0}} \qquad (2.92)$$

In general, the majority of the photon wave-function is the bare photon (1 / 3 ŸŸ 1), but because the cross-sections for hadronic interactions are much larger than the QED cross-sections, all three components play an important role in photon-hadron scattering. Since results in this thesis focus on high momentum-transfer jet production, we will now focus speci cally on this regime of the photon structure. In this region, the k2 of the scattering process regulates perturbative calculations of the cross-section, so the partonic content of the photon can be described in the language of structure functions and PDFs. Following from Equation 2.91, the PDFs of the photon can be expressed as the sum of three parts:

$$5_{0\bullet W}{}^{1}G-\&^{2o} = 5^{dir}_{0\bullet W}{}^{1}G-\&^{2o}\ ,\ 5^{VMD}_{0\bullet W}{}^{1}G-\&^{2o}\ ,\ 5^{anom}_{0\bullet W}{}^{1}G-\&^{2o} \qquad (2.93)$$

Where the direct photon PDF, as one would expect, has only contributions from W

$$5^{dir}_{0\bullet W}{}^{1}G-\&^{2o} = /\ _{3}X_{0W}X^{1}1\quad G^{o} \qquad (2.94)$$

Typical photon PDF descriptions combine the two hadronic terms into one resolved photon PDF:

$$5^{res}_{0\bullet W}{}^{1}G-\&^{2o} = 5^{VMD}_{0\bullet W}{}^{1}G-\&^{2o} \text{ , } 5^{anom}_{0\bullet W}{}^{1}G-\&^{2o} \Rightarrow 5_{0\bullet W}{}^{1}G-\&^{2o} = 5^{dir}_{0\bullet W}{}^{1}G-\&^{2o} \text{ , } 5^{res}_{0\bullet W}{}^{1}G-\&^{2o} \quad (2.95)$$

The quark components of the resolved photon PDFs have a modi ed evolution equation [141], compared to standard DGLAP evolution (see Section 2.1.5):

$$\&^{2}\frac{m\,5^{res}_{@W}{}^{1}G-\&^{2o}}{m\&^{2}} = \frac{U_{B}{}^{1}\&^{2o}}{2c}\int_{G}^{1}\frac{dH}{H}\,5^{res}_{@W}{}^{1}H-\&^{2o}\%_{@@}\frac{G}{H} \text{ , } 5^{res}_{6\bullet W}{}^{1}H-\&^{2o}\%_{@6}\frac{G}{H}$$
$$\text{ , } 34^{2}_{@}\frac{U_{EM}}{2c}\,G^{2} \text{ , } {}^{1}1 \quad G^{p2} \quad (2.96)$$

The rst term is recognizable as the standard DGLAP evolution, where quarks or gluons with momentum fraction H ¡ G may contribute to the PDFs at G via perturbative splittings. The second term is unique to the evolution of the photon PDF, and it comes from $@@$ splittings. The gluon PDFs of the photon follow an un-modi ed DGLAP evolution, as described in Equation 2.41. It is important to note that, from the perspective of DGLAP evolution, no distinction is made between partons arising from the VMD component of the photon wave-function or the “anomalous” part, so they must be combined into one resolved photon PDF.

In determining the photon PDFs, then, the only remaining question becomes nding the values of $5^{res}_{0\bullet W}{}^{1}G-\&^{2o}_{0}$ for some reference scale, from which it can be evolved. The typical approach taken by leading photon PDF models [144, 145] is the same as for the PDFs of hadrons, where experimental measurements are used to constrain the PDFs at a variety of scales and DGLAP evolution is used to interpolate between them. These ts rely on measurements from $e^{+}e^{-}$ colliders which tag one electron to have $\&^{2} = 0$, emitting a quasi-real photon. These WW collisions effectively allow for DIS from a photon, directly probing the photon structure function. Evaluations of the photon PDF are much less sophisticated than those of the proton PDFs or nPDFs, largely due to the lack of available data and the challenging nature of measurements which probe the photon PDFs. Regardless, several parameterizations exist for the photon PDFs, the leading ones being GRV NLO

Figure 2.16: The photon PDFs from leading NLO ts for the up quark (left) and the gluon (right) from Ref. [148]. While the gluon PDF falls rapidly as 0, the up quark PDF has a maximum near 0•6 0•7, demonstrating the valence structure of the photon. While the models have some substantial disagreements in scale, the shapes are reasonably consistent.

[146], GRS NLO [147], CJKL LO [145], and CJK NLO [148]. These different parameterizations are compared in Figure 2.16.

2.5.3 Particle Production in Ultra-Peripheral Collisions

2.5.3.1 Rapidity Gaps

In the typical picture of hadronic collisions, some number of partons from one hadron collide with partons from the other hadron, and these partonic collisions provide the dominant source of particle production. In addition to the particle production resulting from the hard-scattering process, though, there is also additional particle production resulting from the “underlying event”. In the string model, discussed in Section 2.3.2.1, this underlying event has two sources: the breaking of color ux tubes connecting the struck hadron at beam rapidity with outgoing partons and additional (multi-parton) interactions, separate from the hard-scattering process. These processes deposit low-? hadrons throughout the region between the beam rapidity and the outgoing partons,

Figure 2.17: Feynman diagrams showing the photoproduction of a dijet pair in direct (left) or resolved (right) W A UPCs. In the case of the direct photon, no particles are produced between the photon-going nucleus and the hard-scattering, while the resolved photon remnant partially lls the gap in that case.

and which is characteristic of particle production in hadronic collisions.

In the context of UPCs, it is important to recognize that no color connection exists between the outgoing partons and the photon-emitting nucleus. For this reason, no hadronic particle production would be expected in the rapidity interval between the photon-going beam rapidity and the outgoing partons, a feature which aids in distinguishing UPCs from hadronic interactions [149]. This feature of large rapidity gaps in particle production is also characteristic of diffractive processes, since the pomeron has zero color charge [9]. However, for the case of resolved photons, there are color connections between the photon remnant and the outgoing partons, causing particle production to partially ll the rapidity gap. Figure 2.17 shows the Feynman diagrams for resolved and direct photoproduction of jets, demonstrating these features.

The ATLAS collaboration has demonstrated [149, 150] that rapidity gaps are a useful method of distinguishing photo-nuclear from hadronic events. Additionally, many UPC measurements use exclusivity requirements [1, 2] to select photonuclear events, requiring gaps to populate the entire detector, aside from the selected nal-state particles. In order to optimally perform these selections, several gap quantities are useful, as shown in Figure 2.18. In particular, two gap quantities are de ned: the sum of gaps and the edge gap. The edge gap is (the difference in rapidity be-

Figure 2.18: A diagram showing an example photonuclear jet production event the gap quantities relevant to inclusive photonuclear measurements. The blue circles are jets, and the black circles are particles produced in the event. The edge gaps are the darker shaded intervals, and both green intervals compose $_W$ [. Figure by A. Angerami.

tween the most forward particle on a given side of the detector and the detector edge on that side. Meanwhile, the sum of gaps ([ ) is de ned as the sum of all gaps between particles greater than 0•5 units in rapidity which are on a given side of the hard-scattering process, de ned by the $_{min}$, [$_{max}$ of the jets in the event. This gap quantity is particularly useful because it is both resilient to detector noise and allows for a resolved photon remnant to partially ll the gap without rejecting events.

2.5.3.2 Electromagnetic Dissociation and Nuclear Break-up

The UPC process with the highest cross-section, by a large margin, is electromagnetic dissociation [1, 2], since the photon spectrum falls so steeply with energy. EM dissociation occurs when a low-energy photon emitted by one nucleus is absorbed by the other, causing a Coulumb excitation. These excitations are dominated by the Giant Dipole Resonance (GDR) [151], where the neutrons and protons collectively oscillate against each other. The GDR excitation is the most

common because it has the same quantum numbers as the photon, and its primary decay channel is the emission of a single neutron, where higher-order excitations cause more neutron emission. Since these neutrons are emitted with very low energy in the rest frame of the emitting nucleus, they are produced at beam rapidity with no transverse momentum, allowing them to be measured by detectors near the beamline.

The total EM dissociation cross-section, according to the RELDIS model [152], is 22 kb, over 30 times as large as the hadronic Pb+Pb cross-section. These predictions are consistent with measurements of EM dissociation at the LHC from ALICE [153]. It is important to note that secondary photon emissions factorize from other photon emissions, even when both photons are emitted by the same nucleus [154], so even though the rate for these processes is very large, it does not interfere with other photoproduction processes.

In addition to forward neutron emission resulting from the Coulomb excitation of nuclei and their subsequent decay, nuclei may also break apart in UPCs due to higher energy photon-nucleus interactions. The nal-states produced near mid-rapidity in these processes are discussed in much more detail in Section 2.5.3.3, but in addition to the particle production at mid-rapidity in inclusive photonuclear scattering, these processes have a very high probability to break up the struck nuclei. Due to the relatively small nuclear binding energy (10 MeV) compared to the photon energy (¡ 1 GeV), these processes are expected to break up the nucleus into fragments, resulting in the emission of some (potentially large) number of forward neutrons. The probability of this break-up is very high (¿ 95%), and it is predicted to be 100% by RELDIS [152].

Forward neutron emission topology is a key signature of photonuclear events, and three different topologies are helpful for classifying heavy ion collisions: 0n0n (no forward neutrons on either side), 0n1n (forward neutrons on exactly one side), and 1n1n (forward neutrons on both sides). For many photonuclear measurements, tagging events with a forward neutron topology provides a way to help identify the events of interest and un-ambiguously determine which nucleus emits the photon. These topologies, however, may be in uenced by additional neutron production from independent Coulomb excitations of the photon-emitting nucleus, where the additional

forward neutron emission causes an event which should be accepted by the 0=-= requirement to instead fall into the -=-= category.

The 0=-= requirement, which is crucial to measurements of photonuclear UPCs, thus imposes an additional requirement of not having additional Coulomb excitation of the photon-emitting nucleus. Since both processes depend on impact parameter, this requirement modi1es the distribution measured in photonuclear events non-trivially, changing the effective photon ux in 0=-= events. This modi cation can be modelled by imposing the 0=-= condition while integrating the photon ux over impact parameter, rather than simply applying the form integrated over 1 ¡ 2' . This procedure begins with writing the cross-section for some photonuclear process ¸ ! W ! - as:

$$f^{AB}_{EPA} = \int d\,{}_W d^2B\, d^2 1 \frac{d\#}{d\,{}_W d^2B} f_{W!-}{}^{1}\,{}_W - B^o X^2 {}^1 B^{\circledR} \quad {}^{\circledR}B \quad 1^{\circledR} \tag{2.97}$$

Here, the photonuclear collision kinematics are given by B, the impact parameter of the collision with respect to the photon emitting nucleus, B_s, the impact parameter of the collision within the struck nucleus, and $B^{\circledR}$ the impact parameter of the two-nucleus system. These kinematics are shown in Figure 2.19. The 0=-= requirement is then applied by restricting the region of integration through the inclusion of two probability distributions as a function of $1^{\circledR}$. These distributions correspond to the probability of breakup due to a hadronic interaction, $\%_{no\,had}{}^1 1^o$, and the probability of breakup due to Coulomb excitation, $\%_{no\,EM}{}^1 1^o$. Following the approach used in STARLIGHT [155], the former probability can be written as:

$$\%_{no\,had}{}^1 1^o = \left(1 - \,)_{AA}{}^1 1^o f^{inel}_{NN}\right)^2 \approx 4\,)_{AA}{}^1 1^o f^{inel}_{NN} \tag{2.98}$$

Where $)_{AA}{}^1 1^o$ is the nuclear overlap function de ned in Section 2.4.1 and f^{inel}_{NN} is the total inelastic nucleon-nucleon cross-section. The approximation uses the large-limit to simplify the expression considerably. For nuclear breakup due to additional electromagnetic interactions, it is important to only account for events where the additional neutron production would cause an event

Figure 2.19: A diagram of a photonuclear UPC from a side view (left) and front view (right). The impact parameter of the nucleus-nucleus collision ($\vec{b}$) is different from the impact parameter of the W A collision with respect to the photon-emitting nucleus ($\vec{b}_s$) (or the struck nucleus $\vec{b}_t$). Inspired by Figure 3 from Ref. [116].

to fail the 0=-= requirement. STARLIGHT models the impact-parameter dependence of additional breakup probabilities as:

$$\hat{\sigma}_{\mathrm{inel}} = \int d^2b \, \frac{d\sigma_{\mathrm{inel}}}{d^2b} \, f_{\mathrm{inel}}(W_{\gamma A}) \tag{2.99}$$

Here, $f_{\mathrm{inel}}(W_{\gamma A})$ is a parameterization of the energy dependence of the total inelastic photon-nucleus cross-section, particularly for GDR excitations. The quantity $\hat{\sigma}_{\mathrm{inel}}$ is unitarized in order to form the probabilities for each breakup topology, but it is not itself a probability. Thus, the “no

Figure 2.20: Left: The different probabilities relevant for calculating the nuclear breakup probability as a function impact parameter. Right: The photon uxes which arise from modeling W A ! jets with or without the full calculation of the breakup probability. The sum of the 0nXn and 0n0n contributions is is de ned to sum % $_{no\ had}$ as does the sum of % $_{EM}$ and % $_{no\ EM}$. Figures by A. Angerami.

EM" breakup requirement can be de ned as:

$$\%_{no\ EM} = \%_{no\ had}\ 4\ \%\ 1\ \ 4\ \%\ ,\ 4\ 2\%\ \tag{2.100}$$

The various probabilities which de ne these different conditions are shown in the left panel of Figure 2.20. The last piece required for fully modeling photonuclear interactions is considering the spatial extent of the struck nucleus. In order to do so, we assume a factorization of the cross-section:

$$f_{W!\ -\ ^1\ W} - B^o = f_{WN!\ -\ ^1\ W^o)\ \ ^1B\ o} \tag{2.101}$$

The nuclear thickness function, , is de ned fully in Section 2.4.1. Then, combining all of the pieces, the total cross-section can be written as:

$$f^{AB}_{eff} = \ ^1\ d\ _W d^2B\ d^2 1 \frac{d\#}{d\ _W d^2B} f_{WN!\ -\ ^1\ W^o)\ \ ^1B\ o\%_{no\ had}\ ^1 1^o\%_{no\ EM}\ ^1 1^o X^{2\ 1}B\ \ \ \ B\ \ \ \ 1^{\ o} \tag{2.102}$$

Where it is particularly helpful to de ne the new effective ux $\Phi^{eff}(\omega)$ as:

$$\Phi^{eff}(\omega) = \int d^2b\, d^2B\, P_{no\ had}(b)\, P_{no\ EM}(b) \left(\frac{d\sigma}{d\omega\, d^2B}\right)_{\omega} \Phi(\omega, b, B) \qquad (2.103)$$

These quantities are often formulated in terms of the photon energy fraction, $\omega = \frac{E_\gamma}{E_{beam}}$, where E_{beam} is the energy of the beam. In terms of the modi cation of the photon ux due to these different corrections, compared to the simple EPA ux derived in Section 2.5.1, is shown in the right panel of Figure 2.20. These modi cations to the ux improve agreement with experimental measurements that impose a $\Theta = \pi$ requirement, but they imperfectly model these breakup effects, as is shown in Section 5.6.

2.5.3.3 Photonuclear Scattering

Photonuclear scattering processes, one of the main topics of this thesis, can be classi ed by the momentum-transfer of the hardest scattering process. Minimum bias events are characterized by softer physics, dominated by the VMD part of the photon wave-function, while events requiring high-energy or high-mass nal-states probe the harder parts of the photon wave-function, governed by photon PDFs. While the results of this thesis primarily focus on the case of hard-scattering, measurements of minimum bias photonuclear interactions at the LHC [149, 156] provide useful insights into studying inclusive photonuclear processes. Measurements by ATLAS [149] demonstrated the usefulness of applying selections on gap quantities and forward neutrons in order to select a high-purity sample of photonuclear events. This measurement also found that the particles produced in these events exhibit correlated behavior which is characteristic of the production of Quark-Gluon Plasma. Theoretical predictions [157] indicate that these ow-like event characteristics should only be possible in resolved photon events, where the hadronic uctuation of the photon provides the initial-state geometry necessary for elliptic ow. Similar measurements by CMS [156] in $W\gamma$ collisions do not nd evidence of such effects, but they are limited by the lack of a full subtraction of non- ow backgrounds.

In the case of photonuclear hard-scattering processes, Figure 2.17 demonstrates leading-order Feynman diagrams for direct and resolved photoproduction. The two outgoing hard partons from these collisions produce jets which, if measured, probe the kinematics of the partons which enter into the scattering. This mechanism was proposed [4] in order to study both the nPDF modi- cations in a novel kinematic range and the diffractive parton distribution functions. An older, preliminary measurement by ATLAS [150] demonstrated the feasibility of these studies, introduc- ing several important variables for mapping the jet kinematics onto the kinematics of incoming partons. These kinematics are de ned in terms of the jet system mass and rapidity:

$$<_{jets} \quad \tilde{O}_{8} \quad {}^{2} \quad , \quad \tilde{O}_{8} \quad {}^{2} \quad {}^{1\bullet 2} \quad - H_{jets} \quad \frac{1}{2} \ln \left(\frac{{}^{1}_{8} \quad {}^{8}, \quad ?_{I-8}}{{}^{1}_{8} \quad {}^{8} \quad ?_{I-8}} \right) \tag{2.104}$$

Where the index 8 sums over each jet in the event. A proxy for the k^2 of the scattering process comes from the linear sum of the individual jet transverse momenta:

$$T \quad \tilde{O}_{8} \quad ?_{T-8} \tag{2.105}$$

Two additional variables G_A and I_W are also used to characterize the jet kinematics. In the limit of 2 ! 2 scattering kinematics, they correspond to the parton momentum fractions. These observ- ables are de ned as:

$$G_A \quad \frac{<_{jets}}{B_{NN}} 4^{H_{jets}} \tag{2.106}$$

$$I_W \quad G_W H_W = \frac{<_{jets}}{B_{NN}} 4^{H_{jets}} \tag{2.107}$$

Where G_A is the struck parton momentum fraction with respect to the emitting nucleus, and I_W is the momentum fraction of the parton entering into the hard-scattering from the photon-going direction. In the case of direct photons, this is the photon energy fraction (H_W), but for resolved photons ($G_W < 1$), it is the momentum fraction of the partonic constituent of the resolved photon, with respect to the photon-emitting nucleus. The system rapidity is oriented such that the photon is going in the, I direction, which is referred to as “collision system coordinates”, as opposed to

Figure 2.21: Projected reductions in nPDF uncertainties for the gluon PDF resulting from the inclusion of data on photonuclear jet production, from Ref. [158]. The left panel shows the projected uncertainty reduction if all point-to-point errors are 5% of the cross-section, while the right panel shows the reduction for 10% errors. In both cases, no knowledge of the correlation in the uncertainties is assumed.

detector coordinates. Note that while the correspondence between these variables and the parton kinematics is only fully applicable for 2 ! 2 processes, the de nitions include all jets in the event, accounting for the possibility of multi-jet events.

The lack of a nalized measurement or full correction to theoretically calculable results in Ref. [150] has prevented the use of the preliminary data for constraining nPDFs until now. Still, results comparing NLO pQCD calculations to these preliminary results [159] have demonstrated good initial agreement with the data. Further studies have been conducted [158, 160] demonstrating the potential of these measurements to constrain nuclear parton distributions. Figure 2.21 demonstrates the potential reduction of uncertainties with different assumptions about the experimental error. These projections do not, however, include the impact of any correlated uncertainty, which is important for extracting the full information about nPDFs from the results presented in this thesis. It has also been proposed [161] that because of the high breakup rate in photonuclear jet events due to Coulomb excitations, higher $W_{\gamma N}$ events probe a different range of impact parameters within the struck nucleus than low $W_{\gamma N}$ events, when a θ=-= requirement is applied. This property could, in theory, be used to constrain the dependence of nuclear PDF modi cations on impact parameter in future studies.

2.5.3.4 Two-Photon Processes

Since the photon uxes from both nuclei in heavy ion collisions have intensities which scale as $/^{4}$, the total cross-section for two-photon fusion processes in Pb+Pb collisions is enhanced by a factor of $82^{4} = 4 \cdot 5 \quad 10^{7}$, relative to processes of the same kinematics in pp collisions. This massive increase in the cross-section for $W_{\gamma}W_{\gamma}$ processes in heavy ion collisions allows them to serve as a laboratory for precision QED measurements, studies of physics beyond the standard the Standard Model, and, as discussed further in Chapter 6, a novel electromagnetic probe of the QGP. Following from the expression for the single photon ux, the effective ux for two-photon processes can be written as [2]:

$$\frac{d\#_{WW}}{d:_{1}d:_{2}} = \int_{1_{1}i'}^{1} d^{2}1_{1} \int_{1_{2}i'}^{1} d^{2}1_{2}\#^{1}:_{1} - 1_{1}^{o}\#^{1}:_{2} - 1_{2}^{o}\%_{no\ had}{}^{1}j\,1_{1}^{®} \quad 1_{2}^{®}j^{o}\%_{n}{}^{1}1_{1}^{®} \quad 1_{2}^{o} \qquad (2.108)$$

Where $\%_{n}$ is the probability of a given forward neutron con guration. Because the initial-state transverse momentum is so small in $W_{\gamma}W_{\gamma}$ production of dileptons, the relevant variables for studying the properties of these systems are related to the relative de ection of the two leptons. The two most basic forms of these observables are the acoplanarity, U, and asymmetry, c, de ned as:

$$U \quad 1 \quad j\,q_{1} \quad q_{2}j \cdot c - \qquad (2.109)$$

$$j?_{T1} \quad ?_{T2}j \cdot {}^{1}?_{T1}\,,\ ?_{T2}{}^{o} - \qquad (2.110)$$

where q_{1-2} and $?_{T1-2}$ represent the azimuthal angles and the transverse momenta of each of the two leptons, respectively. The variable c is also a useful as a modi cation to the acoplanarity, where the two are the same except for a factor of $?_{T}$:

$$:_{?} \quad \frac{1}{2}\ ?_{T_{1}}\,,\ ?_{T_{2}}\ {}^{1}c \quad j\,q_{1} \quad q_{2}j^{o} = cU?_{T} - \qquad (2.111)$$

Figure 2.22: Feynman diagrams from Refs. [2, 162] showing the different processes where leptons are produced via two-photon fusion in UPCs. The right-most and center diagrams show leading-order production mechanisms for electron and muon pairs, while the latter shows the case of nal-state radiation of a photon by one lepton.

The : ? of a dimuon pair effectively represents the transverse momentum "kick" of one muon relative to the other. It is particularly useful since, for lower ?e pairs, the same added transverse momentum will result in a larger shift in b than for higher ?T pairs. The same is not true for ?, since it scales with the pair ?T, allowing it to incorporate these effects more consistently across ranges in ?T.

An additional method of quantifying the relative de ection of dileptons is the total transverse momentum of the dilepton system. Rather than adding the scalar transverse momenta, taking the magnitude of the vector sum allows it to quantify the momentum de ection in the transverse plane:

$$q \frac{}{\%\,\&\,?} \quad q \frac{}{j\,?\,\&\,?_{T_1} \quad \&\,?_{T_2} j^2} \tag{2.112}$$

The physics of UPCs and the equivalent photon ux have been extensively measured through two-photon fusion processes producing dileptions, often called the Breit-Wheeler process [163]. Examples of the Feynman diagrams which correspond to these processes are shown in Figures 2.22. Both ATLAS and CMS have studied the production of dimuon pairs in UPCs, including detailed studies of the dependence of the pair kinematics on forward neutron topology [5, 164]. The CMS collaboration used forward neutron topology as a proxy for the impact parameter, demonstrating with a phenomenological tting procedure that the "core" of the distribution has a signif-

Figure 2.23: Left: The core+tail fit of dimuon acoplanarity distributions by the CMS collaboration from Ref. [5]. Right: The coherent+dissociative photoproduction fit of di-electron acoplanarity distributions by the ATLAS collaboration from Ref. [162]

icant trend with neutron topology. While the CMS result suggested that this effect is due to the impact parameter dependence of the initial photon, ATLAS observed a similar modification in the acoplanarity distributions but found that the behavior could be described by an increase in the contribution of dissociative photoproduction. These dissociative processes occur when the photon is emitted incoherently by a nuclear constituent, causing it to break up but not necessarily exhibit any dependence. ATLAS presented further evidence for this description in studies of two-photon production of pairs [162], where a dissociative contribution described the tails of the distributions quite well. A comparison of these two methods of describing the acoplanarity distributions is shown in Figure 2.23.

The STAR collaboration has also investigated the properties of di-lepton pair kinematics using collisions [165]. Their studies demonstrated the impact of the photon's longitudinal polarization on the electron kinematics. In the left panel of Figure 2.24, it shows that the angle of one electron in the pair with respect to the angle of both electrons exhibits a modulation which is well-described by modeling of the photon polarization. STAR [165] and ATLAS [7] have also measured the dependence of di-lepton pair broadening on collision centrality for non-UPC events. Both results observe a substantial modification to the deflection of back-to-back pairs in more

Figure 2.24: Left: The dependence of di-electron production on the angle of a single electron with respect to the pair measured by the STAR collaboration. Right: The STAR measurement of the different distributions for UPC and peripheral Au+Au collisions Both gures are reproduced from Ref. [165].

central collisions. The STAR result is shown in the right panel of Figure 2.24, and Section 7.2 provides a more detailed measurement of these properties using a larger, updated dataset. These observations have been explained by two possible mechanisms: de ection of the leptons by magnetic elds in the QGP [166, 167] or impact parameter-dependent modi cations of the photon spectrum [168, 169, 170, 171].

The hypothesis that the de ection of muon pairs could arise due to interactions with magnetic elds makes several other predictions about the dependence of these effects on muon kinematics [166, 170]. First, it is predicted that magnetic broadening of the dimuon distributions should depend on the rapidity separation of the two muons. If they have similar rapidity, the magnetic eld would de ect them the same, largely cancelling the effect. Additionally, magnetic elds in the QGP would have a direction strongly correlated with the second-order event plane angle, $_2$, [170], where the second-order event plane is de ned as the plane formed by the impact parameter vector and beam axis [172]. Thus, any modi cation to these distributions due to magnetic eld effects should depend strongly on the orientation of the dimuon pair with respect to the event plane. Each of these hypotheses is studied in detail in the results of this thesis in Section 7.2.

Two methods exist in the literature for explaining the centrality-dependent modi cation of

W, W! » pair de ections via initial-state effects. The rst, described by Refs. [173, 168], uses a quantum-mechanical description of the matrix to study quantum mechanical interference effects which generate a substantial impact parameter dependence of the initial photon kinematics. The method of Refs. [170, 171] treats correlations between impact parameter and photon transverse momentum more generally, by introducing the photon Wigner distribution and using it to generalize the EPA approach to include independent modi cations to the photon spectrum. Each of these methods predicts a suppression for very low and : ?, not a simple broadening, which is consistent with results shown in Section 7.2.

The approach of Refs. [173, 168] begins by recognizing that, in order to generalize EPA to include photons generated inside a nuclear target, the total relevant charge must be modi ed. This is achieved by including the nuclear form factor, 1; 2o, in the de nition of the electromagnetic four-potentials. This correction using the form factor is also applied in STARTRAP, but it does not include any other corrections described here. The classical expression for the cross-section then exhibits a dependence on the form factor:

$$
\begin{aligned}
f = 16 \frac{1/4\,o^4}{1 4c^{o2}} \int d^2 1 \int \frac{dF_1}{F_1} \frac{dF_2}{F_2} \frac{d^2 :_{1?}}{1 2c^{o2}} \frac{d^2 :_{2?}}{1 2c^{o2}} \frac{d^2 @_?}{1 2c^{o2}} \\
\frac{1 \quad :_1^{2o}}{:_1^2} \frac{1 \quad :_2^{2o}}{:_2^2} \frac{1 \quad :_1^{02o}}{:_1^{02}} \frac{1 \quad :_2^{02o}}{:_2^{02}} 4 \; 8^{®} @_? \\
\times 1 :_{1?}^{®} \quad :_{2?}^{®} o 1 :_{1?}^{®0} \quad :_{2?}^{®0} o f_B 1 F_1 - F_2 o \; , \; 1 :_{1?}^{®} \quad :_{2?}^{®} o 1 :_{1?}^{®0} \quad :_{2?}^{®0} o f_{?B} 1 F_1 - F_2 o 1/4
\end{aligned}
\tag{2.113}
$$

Where the f_B and $f_{?B}$ are the scalar and pseudo-scalar contributions to the cross-section, de ned in Ref. [173]. The four-momenta of the ingoing $(:_1, :_2)$ and outgoing $(:_1^0, :_2^0)$ photons are [168]:

$$
\begin{aligned}
:_1 = 1 F_1 - :_{1?} \frac{F_1}{E} o \; - \; :_2 = 1 F_2 - \%_? o \; - \; :_{1?} \frac{F_2}{E} o \\
:_1^0 = 1 F_1 - :_{1?} \; @_? \frac{F_1}{E} o \; - \; :_2^0 = 1 F_2 - :_{2?} \; , \; @_? \frac{F_2}{E} o
\end{aligned}
\tag{2.114}
$$

The photon kinematics are formulated in terms of the transverse momentum difference in the incoming pair, $\%_?$ and the outgoing pair $@_?$. The crucial element of this approach is the complex

exponential term which arises in equation 2.113, coupling the impact parameter to the outgoing pair total transverse momentum. This expression, however, does not model the observed de ections well. The authors then extend it by combining it with a QED calculation of the cross-section with this equivalent photon ux. In this case, the impact parameter dependence from the complex exponential term does model the behavior of the distribution well and can be attributed to a quantum mechanical interference effect in the combination of the b-dependent ux and the QED cross-section.

An alternative approach to predict the cross-section [170] relies on introducing the Photon Wigner distribution [171] in order to describe these modi cations. The full, 5-dimensional form of the Wigner distribution is [170, 174]:

$$G\,5_W\,{}^1G-):-1_?{}^o = \int \frac{d^2\,?}{{}^1 2c^{o2}} 4^{8\,?\,1?} \int \frac{db\,d^2A_?}{{}^1 2c^{o3}} 4^{8G\%b\;8:)\;A_?} \; - \frac{?}{2}\;,?\;0 - \frac{A_?}{2}\;,?\;b - \frac{A_?}{2}\;-\frac{?}{2} \tag{2.115}$$

The integrand of the rst integral is a case of the generalized transverse momentum distribution (GTMD), which can be written as a decomposition of polarized and unpolarized GTMDs:

$$^{89}{}^1G-):-\;?{}^o = \frac{X^{89}}{2} G\,5_W{}^1G-):-\;?{}^o\,,\;\frac{:\,8:\,9}{®\;\;®} \;\frac{X^{89}}{2} G\;{}_W{}^1G-):-\;?{}^o \tag{2.116}$$

Where 5_W and $_W$ are the unpolarized and linearly polarized photon GTMDs, $?$ is the Fourier conjugate momentum of the transverse impact parameter. In the classical limit, the GTMD reduces to the classical photon ux, accounting for the nuclear form factor by taking the Fourier transform of the classical distribution. This Fourier transform introduces an oscillation function of the intrinsic photon transverse momentum, where the shape of the uctuation depends on the value of ®. This effect is demonstrated in Figure 2.25, which shows the shape of this uctuation for several different impact parameters. The differential cross-section can then be written down using the photon Wigner distributions, incorporating this b-dependence of the ux.

In addition to allowing for precision studies of QED and potential new probes of the QGP, two-

Figure 2.25: The oscillations in the photon Wigner distribution as a function of transverse momentum for different values of impact parameter from Ref. [170].

photon UPCs also provide a collision system where new searches can be performed for physics beyond the Standard Model (BSM). Two processes which are expected to be sensitive to BSM physics are [2] the production of g g pairs via two-photon fusion and W, W! W, W light-by-light scattering. The left two panels of Figure 2.26 show example diagrams of these two processes. Tau production allows for direct measurement of the anomalous magnetic moment of the g, is predicted to have sensitivity to the same BSM physics as [175], with effects which scale by the ratio of lepton masses squared, a factor of 280 more sensitive. This process has been measured in UPCs by the ATLAS and CMS collaborations [175, 176], and the ATLAS result has shown initial sensitivity competitive with the best limits from LEP. A comparison of these limits from ATLAS is shown in the left panel of Figure 2.27. The light-by-light scattering process, meanwhile, is sensitive to the existence of massive axion-like particles (ALPs) through the diagram shown in the right panel of Figure 2.26. Measurements by CMS and ATLAS [177, 178] have set new limits on these processes by measuring the photon kinematics produced via light-by-light scattering, and these limits are shown in the right panel of Figure 2.27.

2.5.3.5 Photo-diffractive Scattering

Photo-diffractive scattering processes occur when the photon emitted by one nucleus in a UPC interacts diffractively with the other nucleus, either through a pomeron exchange or some higher-

Figure 2.26: Feynman diagrams of different two-photon fusion processes sensitive to physics beyond the Standard Model from Refs. [175, 178].

Figure 2.27: Left: Constraints on Ω_g set by ATLAS measurements of two-photon production of g, g pairs compared to existing limits from world data, from Ref. [175]. Right: World limits on the existence of axion-like particles where ATLAS and CMS measurements of light-by-light scattering contribute to a substantial region of mass vs. coupling, from Ref. [178].

order reggeon exchange [2]. For high photon energies, though, diffractive scattering processes are dominated by the pomeron contributions. This interaction can be either coherent, where the nucleus remains intact, or incoherent, where the nucleus is excited or dissociated. In the case of coherent photoproduction, the pomeron is emitted coherently by the entire nucleus, while the incoherent case involves emission by some nuclear constituent. The pomeron exchange can be treated as a two-gluon exchange in a perturbative picture, which provides an intuitive reason for the feature that this process is predicted to be sensitive to the square of the gluon density [179].

The most common measurements in UPCs are, by a large margin, measurements of the photo-diffractive production of exclusive vector mesons. The cross-sections for these processes are often quite large [1], on the order of several hundred microbarns in some cases. Since these states are exclusive, the kinematics of the vector meson entirely determine the kinematics of the hard-scattering process. This relationship yields a similar de nition of the struck pomeron's energy relative to the emitting nucleus [2]:

$$G_{\mathbb{P}} = \frac{\text{"}_{+}}{\sqrt{B_{NN}}} 4^{H} \tag{2.117}$$

Here, "_{+} is the mass of the vector meson, and the denotes an important feature of these collisions: because of the single-particle nal state and lack of nuclear breakup, the case where one nucleus emits the photon is indistinguishable from the case where the other nucleus emits the photon. Two vector meson states have been studied most extensively in UPCs: ρ^0 production and J/ψ production. Measurements of coherent ρ^0 production [180, 181, 182] allow one to spatially image the target via the Fourier transform of the ρ^0 distribution due to its large cross-section. J/ψ production is not as common as ρ^0 production, but its larger mass allows for the application of pQCD methods to probe the gluon densities at very low x. Many measurements of J/ψ photoproduction, both coherent and incoherent, have been conducted for this purpose [183, 184, 185, 186, 187, 188].

More recent measurements [189, 190] have attempted to resolve the directional ambiguity in Equation 2.117 via a proposal [191, 192] to measure J/ψ photoproduction with different neutron topologies. These different topologies modify the photon ux, which allows one to dis-entangle the cross-section contribution from each con guration away from mid-rapidity. The picture of

Figure 2.28: A Feynman diagram for direct dijet photoproduction (left) and an NLO QCD prediction for the cross-section of photo-diffractive jet production as a functi on of the jet system rapidity (right). Both gures are reproduced from Ref. [196].

•k production probing the gluon density squared is complicated, however, by other recent results showing large NLO corrections [193], which indicate that the quark contributions are much more important than expected. Proposed methods to remedy this surprising reduction in sensitivity to the gluon densities include measurements of photoproduction [194] or •k photoproduction in O+O collisions [195].

While the exclusive case for photo-diffractive scattering has been intensely studied through vector meson production, inclusive photo-diffractive scattering processes have not been explored deeply in UPC measurements. In particular, the case of photo-diffractive jet production, where a high-energy photon strikes a parton within the pomeron, probing the diffractive PDF [196], is of particular interest. An example Feynman diagram for this process is shown in Figure 2.28, along with an NLO QCD prediction [196] for the differential cross-section of this process. STARLIGHT 8 has also recently implemented coherent diffractive jet production [197], allowing the structure of these events to be studied at leading order in much greater detail.

Predictions for diffractive jet photoproduction are particularly challenging because of a phenomenon known as factorization-breaking. Although QCD factorization has been proven for diffractive hard-scattering [38], it was demonstrated [39] at HERA [198, 199] that factorization

Figure 2.29: Estimates of the fraction of photonuclear jet events which come from photo-diffractive scattering as a function of G_A from Ref. [200]. The predictions vary substantially with the choice of factorization-breaking correction, but the contribution is limited to the very small G_A region.

is, in fact, broken for diffractive jet photoproduction. This factorization-breaking arises from resolved photons re-interacting with the particles from the hard-scattering, while direct photons preserve factorization. These effects are quite large, causing NLO pQCD predictions which do not account for it to overestimate the cross-section by almost a factor of 2 [39]. For this reason, the relative rate of photo-diffractive jet production, a potential background to the measurements in this thesis, is not easily calculated. Two estimates [200] for the relative rate of coherent diffractive dijet photoproduction are shown in Figure 2.29, where the effects of factorization-breaking cause large changes in the rate for these events. These calculations indicate that it is only substantial at very small values of G_A, and these events are characteristically different from inclusive jet production due to large gaps on both sides of the event, allowing them to be removed from measurements via gap requirements.

2.6 Monte Carlo Generators and Models

2.6.1 Leading Order Generators

2.6.1.1 Pythia

PYTHIA 8 [98] is a multi-purpose Monte Carlo event generator which has been extensively used for decades, in a variety of collision systems: e^+e^- at LEP, $p\bar{p}$ at the Tevatron, ep at HERA,

Figure 2.30: The schematic event structure of a $\sigma(?? \to ??)$ event in PYTHIA 8 from Ref. [98]. The generator simulates the process with a hardness-ordering, moving from the hard-scattering matrix element in the center out to fully hadronized outgoing particles.

and ?? at a number of hadron colliders, including RHIC and the LHC. PYTHIA 8 implements a massive array of physics processes, including all leading-order perturbative Standard Model processes, along with a huge number of theoretical BSM processes. Figure 2.30 demonstrates the event generation process for PYTHIA 8, which implements a "hardness-ordering", simulating each aspect of the event generation with sub-processes in order of the relevant metric of hardness for that sub-process ($\hat{s}$, Q^2, , etc.).

In physics simulation relevant for this thesis, PYTHIA 8 begins with the hard-scattering process and additionally computes multi-parton interactions (MPI), which are a phenomenological description of additional semi-perturbative hadronic scatterings between additional partons in the two colliding objects [201]. These MPIs generate "minijets", which contribute to the p_T of the event and can be color-connected to partons from the hard-scattering during hadronization. After

simulation of MPI, PYTHIA 8 implements a parton shower (see Section 2.3.1) native to the generator [99], which uses p_T-ordering in order to interleave MPIs with the parton shower of the hard-scattering process. Finally, the hadronization step in PYTHIA 8 uses the Lund string model [102, 103] (see Section 2.3.2.1). Additionally, it handles color flow in hadronization by taking the $N_c \to \infty$ limit, allowing the generator to neglect color interference effects which are proportional to $1/N_c^2$. It implements color reconnection [202] to define the starting strings in the fragmentation process.

An update to PYTHIA 8 since version 6 is that it allows for a modular implementation of photon fluxes for studying the physics of photoproduction processes [98]. For samples used in the measurement of photonuclear jet production described in this thesis, this functionality was very important, since a user-generated photon flux could be supplied according to the formulation in Section 2.5.1. PYTHIA 8 uses an implementation of the CJKL [145] LO photon PDFs, and it includes a full initial-state parton shower for resolved photons. Additionally, MPIs are also implemented in PYTHIA 8 with resolved photons, along with beam remnants, which allow for excellent modeling of rapidity gap distributions. PYTHIA 8 also implemented coherent photo-diffractive dijet production [197], which allows it to be used to study an important background process described in Chapter 5.

2.6.1.2 HERWIG

HERWIG++ [203, 204] is a multi-purpose leading order event generator, which is the C++ implementation of HERWIG 7 (Hadron Emission Reactions With Interfering Gluons). HERWIG++ simulates much of the same physics as PYTHIA 8, with leading-order implementations of hard-scattering matrix elements. The default parton shower implemented in HERWIG++ is angular-ordered [96], in order to directly take QCD interference into account. It also implements an alternative parton shower using the dipole model with local recoils to account for coherence effects [97]. HERWIG++ uses a similar MPI method to PYTHIA 8, except that it does not interleave MPI with the parton shower, keeping the partons produced in each step separate until the hadroniza-

tion step. HERWIG++ implements a cluster hadronization model (see Section 2.3.2.2), which is based on the precon nement property of QCD [205] shared with its default angular-ordered parton shower [107].

2.6.2 Next to Leading Order Generators

2.6.2.1 POWHEG

POWHEG [206, 207, 208] (POsitive Weight Hardest Emission Generator) is a specialized Monte Carlo event generator used to simulate the hard-scattering matrix elements for NLO processes. Since POWHEG is not a multi-purpose generator, it must rely on other generators such as PYTHIA 8 or HERWIG++ to implement other aspects of the physics in full event generation, such as the parton shower or hadronization. As discussed in Section 2.3.1, the parton shower is not a strictly leading-order process, since it sums contributions at all orders in α_s, to some $\ldots$ xed order in the $\ln \ldots$ coef cient logarithms. NLO calculations, meanwhile, compute the cross-sections to a $\ldots$ xed order in α_s, but they sum all orders in the logarithmic coef cients. Physically, these two approximations are complementary: NLO corrections account for hard wide-angle splittings, while parton showers primarily generate emissions in the soft and collinear limit. Since these two methods will double-count certain terms, though, POWHEG must introduce a method to systematically remove overlap in the two descriptions, allowing for results which take advantage of the physics described by both.

In general, an NLO cross-section calculation can be written schematically as [209]:

$$d\sigma = \ldots \quad (2.118)$$

Where $\ldots$ is the Born-level cross-section, $\ldots$ is the Born phase-space, $\ldots$ is the virtual part of the NLO correction, $\ldots$ is the real part of the NLO correction, and $\ldots$ is the phase-space for additional real emissions. Note that these calculations have divergences in both $\ldots$ which are typically handled by cancelling them explicitly in NLO calculations, meaning that the separation of virtual and real corrections for the purpose of combining them with a parton shower is non-trivial. The

core idea of the POWHEG method for combining NLO calculations with parton showers is to always generate the hardest emission first and then use the parton shower MC to simulate all subsequent emissions. Treatment of the first emission with full NLO accuracy requires defining the Matrix Element Corrected (MEC) parton shower cross-section:

$$\frac{d\sigma^{MEC}}{dH\,d\Phi_T} = d\Phi_{rad}\,\Phi\left(\Phi_1, \Phi_0, \Phi_T'\right) \tag{2.119}$$

Here, Φ_{rad} is the radiation phase-space, where $d\Phi = d\Phi_1 \cdot d\Phi_{rad}$. The MEC parton shower also involves applying a correction to the Sudakov Form Factor:

$$\Delta_T(\Phi_T) = \exp\left(-\int d\Phi_{rad}\,\Phi_T(\Phi_1, \Phi_{rad}, \Phi_T)\right) \tag{2.120}$$

Which introduces a correction to the no-splitting probability accounting for the ratio of the real NLO correction and Born cross-section. Then, the POWHEG method involves systematically decomposing the singularity structure of the real NLO correction, using the coefficient of the divergent term to correct the Sudakov form factor and including the regular portion directly in the parton shower cross-section. Then, the hardest splitting generated by this NLO-corrected splitting function is designed to cancel the divergences in its NLO virtual corrections with the divergences of the real part.

Through this method, POWHEG allows for the generic implementation of NLO QCD cross-sections with parton showers via a systematic approach to prevent double-counting. In addition to matching the hardest emission of the shower, vetoes on certain emissions and subtraction schemes involving negative weights are often necessary to remove all double-counting contributions. The POWHEG method has the advantage of minimizing the need for negative weights (hence its name), but it cannot remove them completely.

2.6.2.2 Sherpa

SHERPA [210] is a multi-purpose Monte Carlo event generator which can simulate either leading-order or NLO processes. Unlike POWHEG, it is a multi-purpose generator which implements other physics processes necessary for simulating full events, rather than only matrix elements. It implements a different matching scheme from POWHEG, following the MC@NLO method [211, 212], with a built-in dipole parton shower [97]. It implements both the cluster and string hadronization models, with the option to select a particular approach. Its string hadronization model directly uses the implementation built into PYTHIA.

2.6.3 STARlight

STARLIGHT [155] is a specialized generator which implements the Equivalent Photon Approximation (see Section 2.5.1) for photon-nucleus or two-photon scattering processes. It implements a calculation for the probability of nuclear break-up due to Coulomb excitation as a function of impact parameter based on photonuclear scattering data. The photon ux for impact parameters smaller than twice the nuclear radius is modi ed by the nuclear form factor, and xed impact parameters for collisions are supported, but correlations between the impact parameter and photon $?_T$ are not considered. The probability of no hadronic interaction is implemented via the Glauber model [116, 117]. Additionally, it only calculates photonuclear reactions for the diffractive photoproduction of vector mesons, so it does not implement an integration over the thickness of the nuclear target which is relevant for inclusive photonuclear reactions.

Chapter 3: Experimental Apparatus and Techniques

3.1 Particle Collisions at the Large Hadron Collider

3.1.1 The LHC Injection Chain

The Large Hadron Collider [213] was built around the existing accelerator complex at CERN in the tunnel originally used for LEP, the Large Electron-Positron collider [214]. It takes advantage of existing, lower-energy colliders such as the Proton Synchrotron (PS) and Super Proton Synchrotron (SPS) in order to serve as injectors to the LHC, and the full complex is shown in Figure 3.1. The LHC injection chain [215] is a very complex process, the details of which are beyond the scope of this thesis. However, the schematics of the process described here are useful for understanding the basics of how ion beams enter the ATLAS detector with 2.6 TeV of energy per nucleon.

The injection process for Lead-208 ions begins with isotopically pure $^{208}_{82}Pb$ entering into an electron cyclotron resonance ion source [216], which effectively heats the Lead into a plasma and extracts the ions from that plasma for acceleration. The ions are accelerated in Linac 3 up to 4.2 MeV/n. The ion beam is then injected into the Low Energy Ion Ring (LEIR), stripping it of additional electrons, after which LEIR accelerates the ions up to 72 MeV/n. At this point, the beam forms ”bunches“ of ions, where the each bunch contains the vast majority of the beam charge and their spacing is set by the accelerator RF frequency. The LEIR bunches each long pulse from Linac 3 into 2 bunches of 2.25 $\times 10^8$ ions each. Then, the bunches are injected 2 at a time into the PS, where they are split into 4 bunches of 1.2 $\times 10^8$ ions per bunch with 100 ns bunch spacing and accelerated to 5.9 GeV/n. These ions are then stripped of their remaining electrons as they are injected into the SPS, where they are accelerated to 177 GeV/n with 9 $\times 10^7$ ions/bunch. Finally, these bunches are injected to the LHC, where they are accelerated up to the target energy with a target intensity of 7 $\times 10^7$ ions/bunch.