

NORTHWESTERN UNIVERSITY

Measurement of the Antineutrino Double-Differential Charged-Current
Quasi-Elastic Scattering Cross Section
at MINERvA

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Physics and Astronomy

By

Cheryl Patrick

EVANSTON, ILLINOIS

June 2016

© Copyright by Cheryl Patrick 2016

All Rights Reserved

ABSTRACT

Measurement of the Antineutrino Double-Differential Charged-Current

Quasi-Elastic Scattering Cross Section

at MINERvA

Cheryl Patrick

Next-generation neutrino oscillation experiments, such as DUNE and Hyper-Kamiokande, hope to measure charge-parity (CP) violation in the lepton sector. In order to do this, they must dramatically reduce their current levels of uncertainty, particularly those due to neutrino-nucleus interaction models. As CP violation is a measure of the difference between the oscillation properties of neutrinos and antineutrinos, data about how the less-studied antineutrinos interact is especially valuable. We present the MINERvA experiment's first double-differential scattering cross sections for antineutrinos on scintillator, in the few-GeV range relevant to experiments such as DUNE and NOvA. We also present total antineutrino-scintillator quasi-elastic cross sections as a function of energy, which we compare to measurements from previous experiments. As well as being useful to help reduce oscillation experiments' uncertainty, our data can also be used to study the prevalence of various correlation and final-state interaction effects within the nucleus. We compare to models produced by different model generators, and are able to draw first conclusions about the predictions of these models.

Acknowledgements

My journey from mild-mannered European software engineer to wild-haired, crabby thesis-writing machine has not always been a serene one, and I thank my friends, especially Chris and Claire, for their support, patience and cocktails. I thank my mum, Pat, and grandma, Joan, for being so kind about me running off to America to pursue this surprising new career path. Their love and encouragement has been wonderful.

My friends and colleagues at Fermilab have all been fantastic. Thanks to Debbie Harris and Kevin McFarland, our esteemed and beloved spokespeople, for sharing your advice and wisdom, and to the generator team, especially Tomek Golan, for making the models we use in this analysis. Thanks also to Anne Norrick, Emily Maher, Marianne Wospakrik, Dan Ruterbories, Minerba Betancourt, Phil Rodrigues, Mateus Carneiro, Rob Fine and Chris Marshall and the rest of my lovely MINERvA colleagues for your friendship and support, both emotional and practical.

Thanks to Leah Handel, Pam Villalovo and Leah Welty-Rieger for our ladies' lunches at Northwestern. They made Monday worth getting out of bed for (as did my 9am meeting).

To my committee, Heidi Schellman, Laura Fields, Giles Novak, Eric Dahl and André de Gouvêa - many thanks for your help, and I hope that this will be an interesting read.

Thanks to Teppei Katori and Joe Grange, whose detailed and comprehensive theses inspired me and taught me so much of what I needed to know for this analysis. Teppei's neutrino cross section newsletter has also kept me up to date on the latest developments. See you in London, Teppei! Thanks!

By a stroke of good fortune, I landed the world's best advisor, Heidi Schellman. She has been patient, funny, generous and very, very clever. Thanks for everything, Heidi. I couldn't have asked for better.

Laura Fields, wonderful, amazing Laura, I could not have done this without you. You have been a dear friend, a wise, insightful mentor, and the one who has talked me back to sanity when I thought all this was beyond me. A thousand thanks, Laura; I will miss you terribly.

I thank Bailey, who has listened to all of my practice talks and is the cat world's top neutrino cross section expert, and Rombles, who is too busy for physics, but has purrs and love when I need them the most.

Finally, my biggest thanks are for Brendan, who put up with all my grad-student drama and who has been nothing but sweet, supportive and perfect in every way. I love you so much, my sunshine. Thanks for everything.

Glossary

A9 This trigger message is sent from accelerator division before a beam spill to tell us to get ready to receive beam.

Arachne Our event viewer, in which you can watch events in the detector in real time, or scan through old events by run, subrun and gate number.

Argoneut A prototype liquid argon detector that used to be located between MINERvA and MINOS. This meant that the energy reconstruction of particles that had been matched in MINOS had to be corrected due to the extra mass between the two detectors.

Baffle Device protect equipment other than the target from the damaging effects of the proton beam

Bar OD equivalent of a strip. In each story , bar 1 will be upstream of bar 2.

Cal stage The calibration stage runs RawToDigitAlg. That does not make raw digits out of raw data (contrary to what the name suggests). Instead, it makes cal digits out of supdigits. Applies the gains, MEU, strip to strip and some attenuation calibrations.

CCQE Charged-current quasi-elastic scattering, when a neutrino scatters from a nucleon and exchanging a W boson. This turns the neutrino into a charged lepton (a muon, for our $\nu_m u$ beam) and a neutron into a proton or vice versa: $\nu_m u + n \rightarrow \mu^- + p$ or $\bar{\nu}_m u + p \rightarrow \mu^+ + n$

CCInclusive A charged-current analysis, where no particular interaction process is required. Thus it is looking for neutrino interactions that generate a muon.

Chain Up to 10 FEBs (each atop a PMT) are connected in a chain: a long network cable from the CROC-E runs to the first board in the chain, and then short jumper cables connect each board to the next. A chain of PMTs will be arranged vertically from bottom (FEB 1) to to (FEB 10) up either the east or west side of the top of the detector.

Channel Each PMT (and FEB) has 64 channels, which are connected to 64 different strips in the detector. Each corresponds to a Crate-CROC-Chain-Board-Pixel. Note: electrical engineers will sometimes use the word channel to refer to a chain (but they count from 1 instead of 0, so channel number will be chain number plus 1). This is very confusing. Do not do it.

Charged-current Any interaction wherein a neutrino exchanges a W boson, converting into its partner charged lepton.

Chvojka corollary Free coffee for people writing their theses!

Crate There are two crates, each of which contains CROCs. There are 8 CROC-Es in Crate 0 and 7 in Crate 1.

CRIM Helps to synch the timing of the CROC-Es. There are 4 CROC-Es per CRIM (two CRIMs per crate).

CROC-E Chain readout controller (ethernet). Up to four chains of FEBs are connected to a CROC, which polls the chains for readout.

Cross talk Current in a given channel can induce a small amount of current in the neighboring channel. The *weave* is used to protect us from false readings due to cross talk.

DAQ The system that receives raw data from the detector and stores it to disk.

dCache System for long-term data storage. Files copied to dCache will either be stored to disk or to tape, and can be accessed seamlessly using Samweb.

DIS Deep inelastic scattering - occurs at high Q^2 , where the neutrino scatters off a constituent quark in the nucleon, breaking it apart.

Downstream Further along the beamline, away from the target; MINOS is downstream of MINERvA.

ECAL Lead electromagnetic calorimeter downstream of the fiducial tracker volume and in the inner part of the outer detector. Designed to stop electromagnetic showers so that their energy can be measured.

ECL Electronic Control-room log, also known as Minerva Electronic Logbook. Used to log all shift tasks, hardware changes, or anything else that might affect the detector or data-taking.

Electromagnetic calorimeter See *ECAL*

Enstore The tape backup system. Files stored on Enstore can be accessed using Samweb.

Eroica The 2015 release of our reconstruction code.

FEB Electronics board attached on top of a PMT (one FEB per PMT) that outputs the signal from the PMT.

Fiducial volume The central scintillator tracker part of the detector.

Final-state interaction When an interaction with a nucleus knocks out a nucleon, this nucleon can re-interact with other particles in the nucleus. This is known as a final-state interaction or FSI.

Frame HCAL equivalent of a module. One frame per module.

Front-End Board See *FEB*

Frozen detector Before the full detector was completed, we ran with only the downstream part of the detector, from module 49 onwards (OD frame 51 onwards). This was known as the ‘frozen’ or ‘downstream’ detector.

FSI See *Final-state interaction*

GAUDI The C++ framework used to run our production and analysis jobs. Configured using options files, which live in Tools/SystemTests. Run using Gaudi.exe or SystemTestsApp.exe, to which you pass an options file - that file includes a list of algorithms you want to run, as well as various configuration parameters.

GEANT The program used to create our detector simulation

GENIE Our Monte Carlo generator

HCAL See *Hadron calorimeter*

Hadron calorimeter Iron calorimeter on the downstream and outside parts of the calorimeter. Designed to stop hadrons so that all their energy will be deposited and can be measured by the

Horn Parabolic magnets used to focus positive or negative pions (depending on current direction) produced when the proton beam collides with the beam target. These pions will decay to create our neutrino beam (they also create muons, which are filtered out by rock).

ID The inner detector, with respect to the beamline, including the scintillator tracker, nuclear targets, and downstream calorimeters.

IFDH Intensity frontier data handling. Use ifdh commands to do the equivalent of basic Linux commands on the Fermilab system, regardless of where the files are located. For example ifdh cp will allow you to copy files even if they are on the PNFS storage system.

Inclusive analyses count the number of interactions measured in a given situation, regardless of the reaction mechanism.

Inner detector See ID.

Keep-up processing Batch processing run on a daily cron job to process raw data into raw digits, pedestal-type raw digits to pedestal tables, and beam-type raw digits to supdigits (pedestal suppression). Also declares these files to SAM.

LI See *light injection*

Light Injection (*linjc*) is used to calibrate PMT gain. While there is no beam, the detector is flashed with a known amount of light from the LI boxes, and the response of each channel is measured. Runs set to *numil* mode alternate LI with beam spills.

linjc See *light injection*

logger The nearline computer, located underground. Raw data files are copied to the logger machine and then processed on *mnvnearline1-4*, to make the plots used in the shift room. Logger is in charge of reconstructing in semi-real time for the GMBrowser display.

Main Control Room The shift room for accelerator division personnel, who control our beam. Be sure to check in with them by calling x3721 at the beginning of your shift.

MCR See *Main Control Room*

MEU Overall energy level calibration. Equalizes the energy scale of a muon in data and Monte Carlo.

Michel electron The electron produced when a muon decays at rest.

Mirror plane Magnetic field shielding between MINERvA (unmagnetized) and MINOS (magnetized)

Module In the inner detector, a module consists of two planes of scintillator strips: one in the U or V direction, and one in the X. The U,V and X configurations are all at 60 ° to each other.

Muon monitor Four muon monitors are located upstream of the **MTM** Minerva timing module. The A9 triggers come to the MTM, which triggers the DAQ.

Nearline The computers that generate the shift plots that are used in the control room.

ν Energy of the incoming lepton minus energy of the outgoing lepton. Also the symbol for a neutrino.

numib A run that only consists of “numibeam” beam data, with no calibration data interspersed.

numil A run in which light injection (linjc) data is interspersed between the beam spills (numibeam data).

numip A run in which pedestal (pdstl) data is interspersed between the beam spills (numibeam data).

Nuclear target Passive materials interspersed between the active scintillator planes in the downstream part of the detector. MINERvA has graphite, lead, iron, water and liquid helium targets. Some planes are divided into sections of C, Pb and Fe.

OD the outer detector, around the sides of the fiducial tracker region.

Offline Offline processing does not take place in real time, but instead occurs in batches (keepup, calibration, reconstruction).

Online The online computers run the DAQ software.

Outer detector See *OD*

Paddle One active element of the veto wall, read out by two single-channel PMTs. The paddles are made of scintillator. 6 paddles make one wall, the veto wall has two walls.

pdstl See *pedestal*

Pedestal A measure of the amount of noise received by on our detector's channels when there is no light in the detector. We measure pedestals, typically between beam spills, in order to calibrate our data to remove this background noise.

Photoelectron When light from the detector's optical fibers arrives at the PMT, it hits a photocathode to produce photoelectrons via the photoelectric effect.

Photomultiplier tube (PMT)s receive light from the detector's optical fibers, which hit a photocathode to produce electrons. This signal is then amplified (typical gain is around 500,000) to produce the output signal. MINERvA has around 500 PMTs, each with 64 channels.

Plane Hexagonal sets of parallel scintillator strips that make up the detector. Arranged in X, U, or V configurations, which are at 60° to each other and all (almost) at right angles to the beam.

Playlist A list of MINERvA runs/subruns that correspond to a specific detector configuration. Analyses will typically process data from one or more playlists, depending on what the analysis is looking for (for example, an antineutrino analysis will look at a playlist of data taken in antineutrino mode).

Plex Because of *cross talk*, the detector is cabled so that adjacent strips are connected to distant PMTs and vice versa. This means that if cross talk is induced in the PMT channel next to one that received a real signal, the cross talk will show up in a distant

part of the detector and not be falsely assumed to be additional signal from the original particle. The plex defines the mapping of physical space to electronics space.

PMT See *photomultiplier tube*

PNFS It looks like a regular network filing system, but can actually include stuff that is backed up to tape. It is a front for dcache and Enstore.

POT normalization The number of protons on target (POT), meaning protons delivered from the main injector to the NuMI beam target that generates the neutrino beam, is a measure of how long the detector was exposed to beam, and how intense the beam was. The higher the POT, the more scattering events we would expect to measure. To compare event counts from two different data samples, we scale one sample by the ratio of the POT values for the target. To scale simulation to data, we scale by the ratio of the data POT to the simulated POT used to generate the Monte Carlo sample.

Q^2 The square of the four-momentum transferred to the nucleus in a neutrino interaction. This is a popular variable for differential cross section measurements, as different interaction mechanisms are favored at different values of Q^2 .

Raw data Data from the detector in the binary format in which it is stored on the DAQ. One file is stored per subrun.

Raw digits Decoded raw data, which has been through the first stage of keepup processing. Stored in a ROOT format. For each subrun, it the data is broken up into its pedestal, LI, and beam components. Pedestal type raw digits are used to make pedestal tables, which are used along with beam-type raw digits to make supdigits. LI raw digits are used to calculate gain constants, which are used in calibration stages.

Reco stage The reco stage of data processing takes cal digits and processes them into reco files (SAM data stream reconstructed-pool). These are the data files on which you can run an analysis tool.

Resonant An interaction that produces an excited state of a nucleon (typically the delta resonance Δ_{1232}). These typically decay to a pion and a nucleon.

Resurrection The 2013 release of our reconstruction code.

ROC W The shift room, located on the west side of the Fermilab atrium.

Rock muon A muon created by a neutrino from the beam interacting in the rock upstream of the detector. Creates a track from the front to the back of the detector. As muons behave as minimum-ionizing particles, these are used for calibration.

Run control Program used by shifters to start, stop and monitor the detector's running. Can be used to tell the detector what combination of beam, pedestal and light injection data it should take.

SAM Fermilab's metadata database, which stores information about experiment's data files (like file size, which run and subrun the file corresponds to, what version of code it was processed with etc). Accessed via the Samweb interface.

S2S See *strip to strip*

Scintillator The material used for our tracker, consisting of doped polystyrene. When a charged particle passes through the scintillator, it generates blue light, which is shifted to green by our wavelength-shifting fibers, and travels to our PMTs where it is converted to electrical current.

Slow control Program used by expert shifters to directly interface with the detector's hardware. Can be used to monitor hardware, upload configurations, and to directly read memory registers on the boards.

Spill Rather than being on constantly, our beam delivers a pulse of neutrinos approximately every 1.6s. This pulse is called a beam spill.

Story How far out you are in the HCAL.

Strip Long, triangular prism of scintillator, used to construct the active part of the inner detector.

Strip to Strip Calibration to equalize the energy response of the scintillator strips in the detector.

Subrun Our data runs are typically divided into around 40 subruns. A normal run will consist of something like 2 numip runs, 8 numil, and 30 numib. In the control room, you will see a run referred to as a 'run series' and a subrun referred to as a 'run'.

Supdigits Beam data in which pedestals have been suppressed. Generated from raw digits; input to the cal stage of processing.

SystemTests The directory where you store your options files for GAUDI jobs. It is nothing to do with testing systems.

Target Could refer to a *nuclear target* or the *beam target*.

Target 1 A nuclear target, consisting of iron and lead. Same mass as target 2.

Target 2 A nuclear target, consisting of iron and lead. Same mass as target 1.

Target 3 A nuclear target, consisting of graphite, iron and lead.

Target 4 A nuclear target, consisting of pure lead.

Target 5 A nuclear target, consisting of iron and lead. Thinner than targets 1 and 2.

Titan The original release of our reconstruction code.

Tower The 6 sides of the HCAL outer detector.

Veto wall The most upstream subdetector of MINERvA, used to tag rock muons for helium and target 1 analyses.

W Invariant mass of the final state hadrons in an interaction.

Weave Because of *cross talk*, the detector is cabled so that adjacent strips are connected to distant PMTs and vice versa. This means that if cross talk is induced in the PMT channel next to one that received a real signal, the cross talk will show up in a distant part of the detector and not be falsely assumed to be additional signal from the original particle.

Upstream Less far along the beam line, closer to the beam target. The veto wall is upstream of the MINERvA detector.

UROC Remote operations center - the setup that allows people to run shifts from a location other than the Fermilab shift room.

x Bjorken x. Fraction of the total nucleon momentum carried by the struck quark. Often used for DIS scattering.

x_F Feynman x.

YELL-DAQ Phone number to call to alert a MINOS expert

y Inelasticity

Dedication

Dedicated to my Dad, Richard Patrick, who I wish was here to read this. I miss you.

Table of Contents

ABSTRACT	3
Acknowledgements	5
Glossary	8
Dedication	20
List of Tables	25
List of Figures	32
Chapter 1. Introduction and motivation	54
1.1. Introduction	54
1.2. Neutrino oscillations	54
1.3. Long-baseline oscillation experiments	65
1.4. Importance of cross sections to oscillation experiments	70
1.5. Cross sections as a probe of nuclear effects	71
Chapter 2. Theory of quasi-elastic neutrino scattering on nuclei	73
2.1. Quasi-elastic scattering cross-sections	73
2.2. Quasi-elastic neutrino scattering	79
2.3. Scattering from heavy nuclei	91

	22
2.4. Long range correlations: the random-phase approximation	97
2.5. Multi-nucleon effects	98
2.6. Other neutrino-nucleon interaction processes, and final-state interactions	110
Chapter 3. The MINERvA Experiment	115
3.1. Introduction to MINERvA	115
3.2. The NuMI neutrino beam	115
3.3. The MINERvA detector	123
3.4. Light collection system	132
3.5. Data acquisition	134
3.6. Offline data processing and calibration	136
3.7. Simulation	153
3.8. Reconstruction	165
Chapter 4. Charged-current quasi-elastic event selection	177
4.1. Signal definition	177
4.2. Data and simulation samples	183
4.3. Event selection	187
Chapter 5. Double-differential cross section calculation	217
5.1. How the cross section is calculated	217
5.2. Background subtraction	219
5.3. Unfolding	225
5.4. Unsmearing the data	227
5.5. Efficiency and acceptance correction	235

	23
5.6. Cross section calculation from acceptance-corrected distribution	243
5.7. Cross section distributions	251
Chapter 6. Systematic uncertainties	267
6.1. Statistical and systematic uncertainties	267
6.2. Evaluating systematic uncertainties on our measurements	268
6.3. Flux uncertainties	282
6.4. Muon reconstruction uncertainties	285
6.5. Interaction model uncertainties	292
6.6. Final-state interaction uncertainties	306
6.7. Recoil reconstruction uncertainties	312
6.8. Propagation of uncertainty through the analysis stages	319
6.9. Detailed breakdown of uncertainties vs. p_{\parallel} and p_T	328
Chapter 7. Discussion of results	350
7.1. Alternative nuclear models	350
7.2. Comparison of quasi-elastic-like cross sections vs. p_T and p_{\parallel} with models	354
7.3. Comparison of true CCQE distributions with models	371
7.4. Conclusions from model comparisons	375
7.5. Comparison with previous and external results	378
Chapter 8. Conclusions	386
References	389
Appendix A. Preliminary tables of cross section measurements	413

	24
A.1. Quasi-elastic-like cross sections	413
A.2. True CCQE cross sections	423
Appendix B. Quasi-elastic kinematics	431
B.1. Q_{QE}^2 calculation in the quasi-elastic hypothesis	431
B.2. E_ν^{QE} in the quasi-elastic hypothesis	432
Appendix C. Ntuple branches	434
C.1. Quantities for cross section measurement	434
C.2. Reconstructed variables used in central value analysis	436
C.3. Variables used in systematic uncertainty variations (Monte Carlo only)	437
C.4. True GENIE information (Monte Carlo only)	441
C.5. Standard MINERvA information	443
Appendix D. Validating the unfolding method	444
D.1. Optimizing number of iterations	444
D.2. Ensemble test	445
D.3. Warped Monte Carlo unfolding test	449
Appendix E. Tables of systematic uncertainties	453
Vita	495

List of Tables

2.1	Values of M_A extracted from neutrino-nucleus scattering data	94
3.1	Composition by mass of a tracker plane (information taken from [37])	125
3.2	Composition by mass of nuclear targets, starting from the most downstream (information taken from [37]). The mass is for an area within a hexagon of apothem 85 cm, with a 2.5 cm cut on each side of the boundary between materials.	128
3.3	Approximate values of the resolution of various quantities, for a CCQE antineutrino sample	176
4.1	Information contained in the different ntuples	185
5.1	Summary of the scales by which the background fractions should be multiplied, in order to achieve the best data/simulation shape fit in the recoil distribution, along with uncertainty in the scale and χ^2 between data and simulation before and after. TFractionFitter introduces additional degrees of freedom to the fit to account for statistical fluctuation in the simulation template shapes, causing the χ^2 /DOF after the fit to appear lower than one would expect.	228

5.2	Quantities we divide by to convert our acceptance-corrected distributions to flux-integrated cross sections	243
6.1	Summary of techniques used to evaluate the individual systematic uncertainties	276
6.2	Summary of muon reconstruction uncertainties	285
6.3	Summary of background uncertainties	295
6.4	Summary of CCQE signal model uncertainties	302
6.5	Summary of hadron interaction uncertainties	307
6.6	Summary of recoil reconstruction uncertainties	313
6.7	Summary of which uncertainties are separately evaluated when generating different analysis histograms from Monte Carlo data	320
A.1	Measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).	414
A.2	Statistical uncertainty on the measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).	414
A.3	Systematic uncertainty on the measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).	415

- A.4 Differential quasi-elastic-like cross section $d\sigma/dp_T$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_T bins are in GeV. 415
- A.5 Differential quasi-elastic-like cross section $d\sigma/dp_{\parallel}$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_{\parallel} bins are in GeV. 416
- A.6 Measured energy-dependent quasi-elastic-like cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_{ν} (GeV). 416
- A.7 Statistical uncertainty on the measured quasi-elastic-like energy-dependent cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_{ν} (GeV). 417
- A.8 Systematic uncertainty on the measured energy-dependent quasi-elastic-like cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_{ν} (GeV). 417
- A.9 Flux-integrated quasi-elastic-like differential cross section $d\sigma/dQ_{QE}^2$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. The Q_{QE}^2 bins are in GeV^2 . 418

- A.10 Energy-dependent quasi-elastic-like cross section $\sigma(E_\nu)$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν bins are in GeV. 418
- A.11 Measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 419
- A.12 Statistical uncertainty on the measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 419
- A.13 Systematic uncertainty on the measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 420
- A.14 Reconstructed energy-dependent quasi-elastic-like cross section $\sigma(E_\nu^{QE})$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν^{QE} bins are in GeV. 420
- A.15 Ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 421
- A.16 Statistical uncertainty on the ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 421

- A.17 Systematic uncertainty on the ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV²), rows are bins of E_ν^{QE} (GeV). 422
- A.18 Measured double differential true CCQE cross section $d^2\sigma/dp_T dp_\parallel$. Units are 10⁻⁴¹cm²/GeV²/proton. Columns represent bins of p_T (GeV), rows are bins of p_\parallel (GeV). 424
- A.19 Statistical uncertainty on the measured double differential true CCQE cross section $d^2\sigma/dp_T dp_\parallel$. Units are 10⁻⁴¹cm²/GeV²/proton. Columns represent bins of p_T (GeV), rows are bins of p_\parallel (GeV). 424
- A.20 Systematic uncertainty on the measured double differential true CCQE cross section $d^2\sigma/dp_T dp_\parallel$. Units are 10⁻⁴¹cm²/GeV²/proton. Columns represent bins of p_T (GeV), rows are bins of p_\parallel (GeV). 425
- A.21 Differential true CCQE cross section $d\sigma/dp_T$, along with statistical and systematic uncertainties. Units are 10⁻⁴¹cm²/GeV/proton. The p_T bins are in GeV. 425
- A.22 Differential true CCQE cross section $d\sigma/dp_\parallel$, along with statistical and systematic uncertainties. Units are 10⁻⁴¹cm²/GeV/proton. The p_\parallel bins are in GeV. 426
- A.23 Measured energy-dependent true CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are 10⁻⁴¹cm²/GeV²/proton. Columns represent bins of Q_{QE}^2 (GeV²), rows are bins of E_ν (GeV). 426

- A.24 Statistical uncertainty on the measured energy-dependent true CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV). 427
- A.25 Systematic uncertainty on the measured energy-dependent true CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV). 427
- A.26 Flux-integrated true CCQE differential cross section $d\sigma/dQ_{QE}^2$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. The Q_{QE}^2 bins are in GeV^2 . 428
- A.27 Energy-dependent true CCQE cross section $\sigma(E_\nu)$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν bins are in GeV . 428
- A.28 Measured reconstructed energy-dependent cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 429
- A.29 Statistical uncertainty on the measured reconstructed energy-dependent true CCQE cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV). 429
- A.30 Systematic uncertainty on the measured reconstructed energy-dependent true CCQE cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are

	$10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).	430
A.31	Reconstructed energy-dependent true CCQE cross section $\sigma(E_\nu^{QE})$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν^{QE} bins are in GeV .	430
E.1	Absolute fractional flux uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	457
E.2	Absolute fractional muon reconstruction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	462
E.3	Absolute fractional primary interaction uncertainties (inelastic) vs. $p_{\parallel\mu}$ and $p_{T\mu}$	467
E.4	Absolute fractional primary interaction shape uncertainties (resonant and neutral current) vs. $p_{\parallel\mu}$ and $p_{T\mu}$	471
E.5	Absolute fractional signal model uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	476
E.6	Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	480
E.7	Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	485
E.8	Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	489
E.9	Absolute fractional recoil reconstruction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$	494

List of Figures

1.1	Particles of the standard model. Image credit: Fermilab.	55
1.2	Flavor mixing	61
1.3	Plot from NOvA showing the confidence level with which their observation of 3 electron neutrinos in the far detector disagrees with theory, for different values of δ_{CP} . For their primary technique (solid lines, corresponding to the LID likelihood-based selector technique), this disagreement confidence exceeds 90% for the inverted hierarchy between 0.1 and 0.5π . Taken from [23]. The secondary selector (Library Event Matching) shown in dashed lines does not agree with the current best fit values for the mixing angles.	66
1.4	Diagram of the T2K oscillation experiment, reprinted from http://www.t2k.org	67
1.5	Predicted ν_e appearance probabilities at DUNE, for three different values of CP violating phase δ_{CP} , reprinted from [13]	68
1.6	Exposure needed for DUNE to measure δ_{CP} for 75% of possible values of δ_{CP} , with different levels of systematic uncertainty. The blue hashed area shows the sensitivity with the current beam design, with the three lines representing how long DUNE must run with uncertainties from $5 \oplus 3$ to $5 \oplus 1\%$, where the two numbers refer respectively to the uncertainty on ν_μ	

	normalization and ν_e normalization relative to ν_μ and the antineutrinos. The dotted line shows the 3σ confidence level. The green colored area shows an equivalent for a new optimized design. Reprinted from [13]	72
2.1	Quasi-elastic antineutrino scattering	73
2.2	Electromagnetic electron-muon elastic (Mott) scattering	75
2.3	Mott scattering	76
2.4	Leading-order correction to $e^- \mu^-$ scattering	77
2.5	Elastic and quasi-elastic scattering of neutrinos from nuclei	80
2.6	Polynomial fit of form factor values measured by several experiments, from [136].	87
2.7	Pion corrections to the CCQE interaction	88
2.8	Flux-unfolded ν_μ and $\bar{\nu}_\mu$ CCQE cross sections per neutron in carbon, as a function of neutrino energy, from the MiniBooNE and NOMAD experiments, compared to the world average and MiniBooNE best-fit RFG predictions. Reproduced from [128]	94
2.9	Comparison of RFG and LFG momentum distributions, from [119]	96
2.10	Set of irreducible diagrams responsible for the polarization (RPA) effects in the 1p1h contribution to the W or Z self-energies, taken from [158]	97
2.11	The spreading function per neutron of electron- ^{12}C scattering, showing the range of energies \bar{E}_ν that would be reconstructed using quasi-elastic hypothesis for three corresponding E_ν^{true} values. The genuine quasi-elastic	

	(dashed lines) and the multi-nucleon (dotted lines) contributions are also shown separately. Reproduced from [151]	99
2.12	Scattering from a correlated pair of nucleons	100
2.13	Nucleon momentum distribution predictions for various nuclei, taken from [73]. Dotted lines show the RFG prediction, while the solid line shows the spectral-function prediction including correlated pairs. The points correspond to data from various experiments, detailed in the reference.	101
2.14	Ratio of the spectral function distributions from figure 2.13 to the corresponding distribution for the deuteron, also reproduced from [73]. The dashed line represents ${}^3\text{He}$, dotted ${}^4\text{He}$, dot-dashed ${}^{16}\text{O}$, long dashed ${}^{56}\text{Fe}$, and dot-long dashed nuclear matter.	102
2.15	Ratio of the cross sections for ${}^4\text{He}$, ${}^{12}\text{C}$, and ${}^{56}\text{Fe}$ vs Bjorken variable x_{Bj} to that for ${}^3\text{He}$, from CLAS[95].	103
2.16	The average fraction of nucleons in the various initial-state configurations of ${}^{12}\text{C}$, taken from [100]	104
2.17	The scaling functions f_T (top, red) and f_L (bottom, green), for various values of q , reproduced from [68]	107
2.18	Best-fit form of $\mathcal{R}_T(Q^2)$, fitted to data from [68] and [149]. Plot from [61].	108
2.19	The four types of diagram for MEC with one-pion exchange, based on a figure in [164]	109

2.20	Processes contributing to the total charged-current neutrino-nucleon scattering cross section, from [110]. ‘QE’ refers to quasi-elastic scattering, ‘RES’ to resonant pion production, and ‘DIS’ to deep inelastic scattering.	111
2.21	Non-quasi-elastic charged-current neutrino scattering processes	111
2.22	Comparison of MINERvA’s neutrino-scintillator scattering data with simulation with and without FSI effects (from [191])	113
3.1	The energy spectrum (flux) of antineutrinos used in this analysis. This flux is given in units of the number of antineutrinos per m ² per GeV per proton on target. This is the PPFX ‘Gen-2 thin’ flux, with neutrino-electron scattering constraint, including systematic uncertainties.	116
3.2	Schematic of the NuMI beam, reprinted from [98] (not to scale)	117
3.3	Schematic of the focussing horns, reprinted from [98] (not to scale)	119
3.4	Schematic of the MINERvA detector, taken from [37] (not to scale). The beam enters from the left side of figure 3.4b.	122
3.5	MINERvA scintillator strips and how they fit together, reprinted from [37]	124
3.6	Cartoon to demonstrate the X, U and V orientations of the scintillator strips in the $x - y$ plane. (Note that a real MINERvA plane has 127 strips)	125
3.7	Positions of the nuclear targets, taken from [37]. The beam passes from left to right in this diagram.	127

- 3.8 Schematic of the MINOS detector, taken from [14] (not to scale).
The grey area in figure 3.8b denotes the area covered by the partial scintillator planes used to instrument the upstream fine-sampling region on the left-hand side of figure 3.8a. The red hexagon (added by the author) indicates the position of MINERvA's inner detector in the beam line, as documented in [185]. The marked beam center is for MINOS; due to the beam's downwards angle, this does not correspond to the beam center in MINERvA. 130
- 3.9 Simplified flow chart of the stages of data processing and calibration 137
- 3.10 Variations in the mean gain of all of MINERvA's PMT pixels, over time. Reprinted from [37]. 144
- 3.11 Diagram showing how base position is calculated 146
- 3.12 Triangle base position shifts used to calculate plane alignment. From [37]. 146
- 3.13 Peak rock muon energy per centimeter was fitted for each plane. These peak values for each plane (with uncertainty from fitting) are plotted, to give a best-fit flat distribution, telling us the relative weighting for each plane. Plot reprinted from [37]. 148
- 3.14 Rock muon energy clusters in data and simulation (left) and the polynomial fit to data (right). Reprinted from [37]. 150
- 3.15 MEU factor $C(t)$ for MINERvA's low-energy run as it varies with time since MINERvA was switched on. (This plot shows only the

- full-detector data; day zero corresponded to the beginning of the earlier run with only a partial detector installed.) The red dashed lines denote the beginning and end of the ‘minerva5’ dataset corresponding to the data used in this analysis. Plot courtesy of J. Kleykamp, from [138]. 150
- 3.16 Time slewing vs hit PE, from [37] 151
- 3.17 An FEB’s high-, medium- and low-gain response to input charges in the test stand, reprinted from [37] 153
- 3.18 Contributions to flux uncertainty 154
- 3.19 Flux distributions and uncertainties for both neutrinos and antineutrinos, with and without the $\nu - e$ scattering constraint. In each figure the top plot shows the constrained (red) and unconstrained (black) distributions; the plot below shows the ratio of the constrained to unconstrained values. Reproduced from [104] 160
- 3.20 The energy hits from a readout gate are divided into time slices, corresponding to individual interactions or other events. Each time slice is indicated by a different color in the top image; black hits are below the energy threshold to be included in an event time slice. (This is data run 2160 subrun 1, gate 594). The event displays below show the energy deposited in the X view for the entire gate (left) and from the first slice only (the red time slice in the top image) (right). These images are from MINERvA’s event display utility, Arachne. [187] 167

- 3.21 Three example clusters through a plane of scintillator strips. a) Due to the triangular structure, even particles traveling perpendicular to the plane will almost always pass through at least two strips. b) Multiple particles from the same interaction may cause a larger cluster. c) A single particle may traverse several strips if it moves at a large angle from the beam direction. 168
- 3.22 An example of a track seed, made from three trackable clusters, in consecutive X planes, and forming a straight line 169
- 3.23 In example a) two track seeds combine to make a track candidate. In b), despite sharing a cluster, the seeds cannot be combined to a track candidate. 170
- 3.24 True - reconstructed x (left) and y (right) vertex positions , for antineutrino CCQE candidates 176
- 4.1 Background processes - non-quasi-elastic charged-current neutrino scattering 178
- 4.2 Final-state interactions 179
- 4.3 Endpoints of muon tracks generated in MINERvA interactions, plot courtesy of J. Ratchford 189
- 4.4 Number of non-muon tracks in data (points) and simulation (colors, POT-normalized to data). In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed

- by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red). All cuts except the track cut are applied. 191
- 4.5 Recoil energy region of the MINERvA inner detector (shown in blue). 192
- 4.6 Recoil energy for data (points) and simulation (colors, POT-normalized to data). In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red). All cuts except the recoil cut are applied. 193
- 4.7 Q_{QE}^2 -dependent recoil energy cut used in the previous MINERvA antineutrino CCQE cross section analysis [96]. Blue events are signal, red are background, for the previous analysis' signal definition of true GENIE CCQE. 195
- 4.8 Efficiency for non-CCQE events with a $CC0\pi$ signature, and distribution of these events in the recoil energy- Q_{QE}^2 phase space 196
- 4.9 Efficiency (top left), purity (top right) and the product of efficiency times purity (bottom) for $CC0\pi$ events with a maximum leading proton kinetic energy threshold as given by the x axis value, where

the recoil cut is as shown in figure 4.7, but additionally accepting all events with a reconstructed recoil energy below the y axis value. The black and white lines show the optimized values of a 120 MeV maximum proton kinetic energy in the signal, and a recoil energy shelf of 80 MeV.

		197
4.10	Purity of our reconstructed sample	200
4.11	Relationship between E_ν^{QE} and Q_{QE}^2 in the quasi-elastic hypothesis, and muon kinematic variables p_T and p_{\parallel} . Blue dashed lines show constant values of E_ν^{QE} , increasing in 1 GeV increments from 1 to 10 GeV. Green dashed lines show constant values of Q_{QE}^2 , increasing in 0.2 GeV ² increments from 0.2 to 2 GeV ² .	203
4.12	Reconstructed event counts in data and simulation vs. muon p_T (normalized to data POT).	207
4.13	Reconstructed event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).	208
4.14	Reconstructed event counts in data, and pull between data and simulation, vs. muon p_T and p_{\parallel}	209
4.15	Reconstructed event counts vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)	210
4.15	Reconstructed event counts vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)	211

4.16	Reconstructed event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).	212
4.17	Reconstructed event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).	213
4.18	Reconstructed event counts in data, and pull between data and simulation, vs. Q_{QE}^2 and E_ν^{QE}	214
4.19	Reconstructed event counts vs. Q_{QE}^2 , in bins of E_ν^{QE} (continued in next figure)	215
4.19	Reconstructed event counts vs. Q_{QE}^2 , in bins of E_ν^{QE} (continued)	216
5.1	Stages of cross section analysis	218
5.2	Distribution of signal and background events vs. muon transverse and longitudinal momentum. Simulation is POT-normalized to data. In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red).	220
5.3	Background tuning: $p_T < 0.15$ GeV	223
5.4	Background tuning: $0.15 < p_T < 0.25$ GeV	224
5.5	Background tuning: $0.25 < p_T < 0.4$ GeV	225

5.6	Background tuning: $p_T > 0.4$ GeV; $p_{\parallel} < 4$ GeV	226
5.7	Background tuning: $p_T > 0.4$ GeV; $p_{\parallel} > 4$ GeV	227
5.8	Scales by which the background fractions should be multiplied, in order to achieve the best data/simulation shape fit in the recoil distribution	228
5.9	Background-subtracted event counts in data and simulation vs. muon p_T (normalized to data POT).	229
5.10	Background-subtracted event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).	230
5.11	Cartoon to demonstrate how a small fractional migration from each bin leads to a smearing effect on the distribution's shape	231
5.12	Monte Carlo events distribution versus reconstructed and true muon kinematics	231
5.13	Migration matrix for the p_{\parallel} / p_T distribution. The x axis corresponds to reconstructed bins, the y to true.	232
5.14	Unfolded event counts in data and simulation vs. muon p_T (normalized to data POT).	236
5.15	Unfolded event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).	237
5.16	Unfolded event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).	238

5.17	Unfolded event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).	239
5.18	Efficiency \times acceptance of our reconstructed sample	242
5.19	Acceptance- and efficiency-corrected event counts in data and simulation vs. muon p_T (normalized to data POT).	244
5.20	Acceptance- and efficiency-corrected event counts in data and simulation vs. muon $p_{ }$ (normalized to data POT).	245
5.21	Acceptance- and efficiency-corrected event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).	246
5.22	Acceptance- and efficiency-corrected event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).	247
5.23	The energy spectrum (flux) of antineutrinos used in this analysis. This flux is given in units of the number of antineutrinos per m^2 per GeV per proton on target.	249
5.24	Mapping of the true E_ν (in the binning of the flux histogram) to E_ν^{QE} reconstructed in the quasi-elastic hypothesis, from our truth-level simulation ntuple	250
5.25	Differential quasi-elastic-like cross section in data and simulation vs. muon p_T (normalized to data POT).	252
5.26	Differential quasi-elastic-like cross section in data and simulation vs. muon $p_{ }$ (normalized to data POT).	253

5.27	Double-differential quasi-elastic-like cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)	254
5.27	Double-differential quasi-elastic-like cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)	255
5.28	Flux integrated quasi-elastic-like differential cross section $d\sigma/dQ_{QE}^2$ in data and simulation (normalized to data POT).	256
5.29	Energy-dependent total quasi-elastic-like cross section $\sigma(E_\nu)$ in data and simulation(normalized to data POT).	258
5.30	Energy-dependent total quasi-elastic-like cross section $\sigma(E_\nu^{QE})$ in data and simulation(normalized to data POT).	259
5.31	Double differential quasi-elastic-like cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued in next figure)	260
5.31	Double differential quasi-elastic-like cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued)	261
5.32	True CCQE cross section measurements in data and simulation. As before, the differential cross sections are flux-integrated.	262
5.33	Double-differential flux-integrated true CCQE cross section $d_\sigma^2/dp_T dp_\parallel$ vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)	263

5.33	Double-differential flux-integrated true CCQE cross section $d^2\sigma/dp_T dp_{\parallel}$ vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)	264
5.34	True CCQE cross section $d\sigma(E_{\nu})/dQ_{QE}^2$, in bins of E_{ν} (continued in next figure)	265
5.34	True CCQE cross section $d\sigma(E_{\nu})/dQ_{QE}^2$, in bins of E_{ν} (continued)	266
6.1	Summary of fractional uncertainties on the cross section	280
6.2	Summary of fractional shape uncertainties in the cross section	281
6.3	Summary of fractional uncertainties on the final cross section due to flux	283
6.4	Summary of fractional uncertainties on the cross section due to muon reconstruction	286
6.5	Summary of fractional uncertainties on the final cross section due to background interaction model uncertainties	301
6.6	Summary of fractional uncertainties on the final cross section due to CCQE model uncertainties	305
6.7	Summary of fractional uncertainties on the final cross section due to final-state interaction uncertainties	312
6.8	Summary of fractional uncertainties on the final cross section due to recoil energy reconstruction	318

6.9	Summary of fractional uncertainties on raw Monte Carlo distributions. Many of these will be irrelevant to the final cross section result; some will affect background subtraction, unfolding, and efficiency correction.	321
6.10	Summary of fractional uncertainties on raw Monte Carlo background distributions, as a fraction of the total background event count in simulation. The data-driven background scaling means that the effect of these on the cross section will be reduced.	322
6.11	Summary of fractional uncertainties reconstructed data event count distributions, after tuned backgrounds have been subtracted	323
6.12	Summary of fractional uncertainties on data distributions, after backgrounds have been subtracted and the distribution has been unfolded	324
6.13	Summary of muon reconstruction uncertainties on data distributions, after backgrounds have been subtracted and the distribution has been unfolded	325
6.14	Summary of fractional uncertainties on acceptance-corrected data distributions	326
6.15	Summary of fractional uncertainties on the final data cross section distributions	327
6.16	Absolute fractional uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)	330
6.16	Absolute fractional uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)	331

6.17	Flux uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	333
6.17	Flux uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	334
6.18	Muon reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	336
6.18	Muon reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	337
6.19	Background interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	339
6.19	Background interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	340
6.20	CCQE signal model uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	342

6.20	CCQE signal model uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	343
6.21	Final-state interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	345
6.21	Final-state interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	346
6.22	Recoil reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).	348
6.22	Recoil reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).	349
7.1	Cartoon of Fermi potential, from [119]	351
7.2	Fractional difference between NuWro's cross sections for local Fermi gas and spectral function models, compared to the global Fermi gas. As seen in these comparisons, the LFG model enhances the cross section; spectral functions reduce it. Note that these effects are most pronounced at neutrino energies lower than the minimum (around 1.5 GeV) analyzed in this study. Plot reproduced from [119]	352

- 7.3 Additional fractional contribution of the Nieves and TEM multi-nucleon effects, compared to the quasi-elastic cross section, as predicted by NuWro. Reprinted from [119]. 353
- 7.4 Comparison of measured quasi-elastic-like cross section with the **Global Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects (continued in next figure) 356
- 7.4 Comparison of measured quasi-elastic-like cross section with the **Global Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects (continued) 357
- 7.5 Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects, to GENIE (continued in next figure) 359
- 7.5 Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects, to GENIE (continued) 360
- 7.6 Comparison of measured quasi-elastic-like cross section and **Global Relativistic Fermi Gas** models, **with RPA**, with and without 2p2h effects, to GENIE (continued in next figure) 361
- 7.6 Comparison of measured quasi-elastic-like cross section and **Global Relativistic Fermi Gas** models, **with RPA**, with and without 2p2h effects, to GENIE (continued) 362

7.7	Comparison of measured quasi-elastic-like cross section with the Local Relativistic Fermi Gas model, with RPA , with and without 2p2h effects, to GENIE (continued in next figure)	364
7.7	Comparison of measured quasi-elastic-like cross section with the Local Relativistic Fermi Gas model, with RPA , with and without 2p2h effects, to GENIE (continued)	365
7.8	Comparison of measured quasi-elastic-like cross section with the Spectral function model, with and without 2p2h effects, to GENIE (continued in next figure)	366
7.8	Comparison of measured quasi-elastic-like cross section with the spectral function model, with and without 2p2h effects, to GENIE (continued)	367
7.9	Comparison of measured quasi-elastic-like cross section and GENIE with 26% 2p2h effect, to GENIE with no 2p2h (continued in next figure)	368
7.9	Comparison of measured quasi-elastic-like cross section and GENIE with 26% 2p2h effect, to GENIE with no 2p2h (continued)	369
7.10	Comparison of our measured quasi-elastic-like cross section $d\sigma/dp_T$ with GENIE without 2p2h, and with 10, 26 (the standard amount) and 40% 2p2h per CCQE interaction, simulated by the Nieves MEC model.	370
7.11	Comparison of quasi-elastic-like cross section $d\sigma/dp_T$ to GENIE for data and models. a) Global Fermi Gas without RPA; b) Local Fermi	

	Gas without RPA; c) Global Fermi Gas with RPA; d) Local Fermi Gas with RPA; e) Spectral Functions; f) GENIE with and without MEC	372
7.12	Comparison of quasi-elastic-like cross section $d\sigma/dp_{\parallel}$ to GENIE for data and models. a) Global Fermi Gas without RPA; b) Local Fermi Gas without RPA; c) Global Fermi Gas with RPA; d) Local Fermi Gas with RPA; e) Spectral Functions f) GENIE with and without MEC	373
7.13	Comparison of $d\sigma/dp_T$ and $d\sigma/dp_{\parallel}$ measurements with GENIE's (red) and NuWro's (blue) implementations of the Global Fermi Gas without RPA, for both the QE-like and true CCQE signal definitions	374
7.14	Comparison of measured true-CCQE cross section and GENIE with 26% 2p2h effect (blue), to GENIE with no 2p2h (red) (continued in next figure)	376
7.14	Comparison of measured true-CCQE cross section and GENIE with 26% 2p2h effect (blue), to GENIE with no 2p2h (red) (continued)	377
7.15	Data / model comparison of the 2013 $d\sigma/dQ_{QE}^2$ measurement, taken as a ratio to GENIE, from [96], but scaled to use the updated PPFX1 NuMI flux [36], by the method of [105]	379
7.16	Comparison between 2013 CCQE $d\sigma/dQ_{QE}^2$ for a true CCQE signal, and the current measurement of $d\sigma/dQ_{QE}^2$ for a true CCQE signal. The two measurements use a different integrated flux.	380

7.17	Double differential $\bar{\nu}_\mu$ scattering cross section from MiniBooNE[30]	382
7.18	NOMAD's antineutrino-carbon CCQE cross section measurements (red) compared with those from previous bubble chamber experiments (black). The data were consistent with a global relativistic Fermi gas model, with $M_A = 1.06 \pm 0.14$ GeV. Reprinted from [146]	383
7.19	Preliminary plot showing the cross section $\sigma(E_\nu)$ from this analysis (black) compared with measurements from MiniBooNE [30] (red squares) and NOMAD [146] (blue triangles) and GENIE's CCQE-like prediction (black dashed line). Our measurements show statistical (inner tick) and total (full error bar) uncertainties.	384
B.1	Quasi-elastic scattering kinematics	431
D.1	Average pull of the 66 bins in 10 unfolded Monte Carlo sub-samples, vs number of Bayesian unfolding iterations	445
D.2	Pull averaged over the 10 unfolded Monte Carlo sub-samples, after four iterations. Note that the bins with the largest pulls correspond to those with very low acceptance (for example, at high p_T and low p_{\parallel} , where the high-angle trajectory makes MINOS matching close to impossible)	446
D.3	Pull between the average statistical uncertainty of 10 Monte Carlo sub-samples, and the width of a Gaussian fit to a histogram of the number of events in each sample	447

D.4	Pull between the average statistical uncertainty of 10 unfolded Monte Carlo sub-samples, as reported by the unfolding procedure, and the width of a Gaussian fit to a histogram of the number of events in each sample	447
D.5	Pull between mean unfolded and true values in each bin.	448
D.6	Ratio of the RMS of the 10 unfolded values to the average uncertainty reported by the unfolding procedure	449
D.7	Ratio of the warped to the central value for two of the warped subsamples	450
D.8	Average pull between unfolded and true values of the warped samples, vs number of iterations	450
D.9	Pull between mean unfolded and true values of the warped samples in each bin.	451
D.10	Ratio of the RMS of the 10 unfolded values in each bin to the average uncertainty reported by the unfolding procedure	452

CHAPTER 1

Introduction and motivation

1.1. Introduction

This thesis describes a measurement of a charged-current double-differential antineutrino scattering cross section on carbon scintillator, at the MINERvA detector at Fermilab. Cross section measurements such as this are a crucial input for neutrino oscillation experiments, and also provide a probe for studying the structure and behavior of atomic nuclei. In this chapter we will introduce the phenomenon of neutrino oscillations and the unanswered physics questions that they raise. We will explain why an accurate and precise cross section model is so important to answer these questions, and describe the current state of the art. In the following chapters, we will give an introduction to the theory of neutrino-nucleus cross sections, and describe the detector hardware and software tools we use to measure them. We will then give a step-by-step explanation of how the antineutrino cross section was measured, and discuss the sources of uncertainty in this measurement and their effects on our final distribution. Finally, we will discuss how our measurement compares with the predictions of various different nuclear models, and how it compares with cross sections measured by other experiments.

1.2. Neutrino oscillations

1.2.1. Neutrinos in the Standard Model

In the Standard Model (Figure 1.1), there are three flavors of neutrinos: the electron neutrino ν_e , the muon neutrino ν_μ , and the tau neutrino ν_τ . All three flavors are electrically neutral, interacting only via the weak interaction. The Standard model predicts them to be massless. Each neutrino flavor is characterized by the fact that it is produced in conjunction with its charged lepton partner;

the electron, muon or tau lepton. In these reactions the number from each generation of lepton is conserved, if a particle of a given generation has a value of 1, and an antiparticle of the same generation has a value -1.

For example, the beta decay of a neutron produces an electron (electron-generation number +1) and an electron antineutrino (electron-generation number -1):

$$n \rightarrow p + e^- + \bar{\nu}_e$$

but it will never produce, for example, an electron (electron-generation number +1) and a *muon* antineutrino (muon-generation number -1, but electron-generation number 0):

$$n \not\rightarrow p + e^- + \bar{\nu}_\mu$$

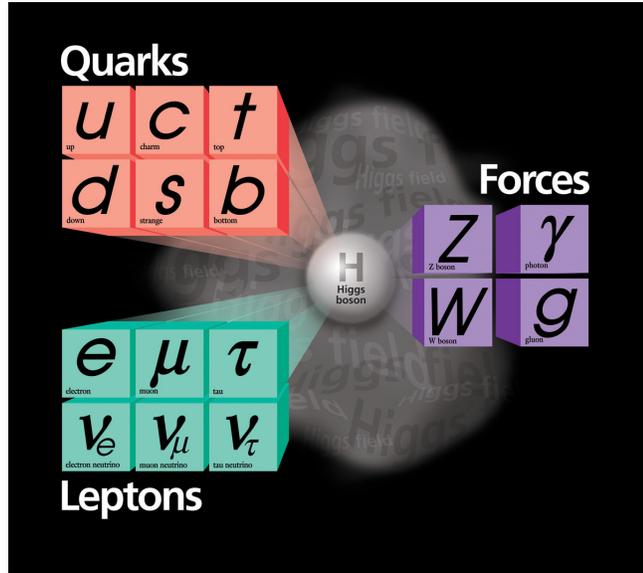


Figure 1.1. Particles of the standard model. Image credit: Fermilab.

Meanwhile, charged pion decay produces anti-muons or muons, and thus, muon neutrinos or antineutrinos:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$

Because they are neutral particles, none of our detectors can directly observe neutrinos themselves. However, in the rare cases that neutrinos do interact, these same rules of lepton flavor conservation are observed. Because of this, we can infer what type of neutrino has interacted in a detector by looking at what charged lepton is created in that interaction. Just as before, a muon neutrino will produce a muon in the final state, an electron neutrino will generate an electron, and so on.

So, looking at the beta-decay equation above, we see that if we have a neutrino source produced by beta decay, we will expect it to be a source of electron antineutrinos, and thus, we will expect to see a positron in our final state (electron-generation lepton number = -1, just like the $\bar{\nu}_e$). If we have a neutrino beam produced by π^+ decay, we expect it to be a beam of muon neutrinos, so we will be looking for a muon in our final state (muon-generation lepton number remaining at 1, as it is for ν_μ).

However, there have been some cases where a beam of muon neutrinos has produced an electron in the final state, rather than a muon. The only neutrino that can create an electron is an electron neutrino ν_e — it is forbidden by the Standard Model for a ν_μ to interact and produce an electron. The alternative possibility is that the ν_μ has somehow transformed into a ν_e on its way to the detector, but this is also forbidden by the Standard Model. This oscillatory behavior was proposed by Pontecorvo [174][172], and was confirmed by the Super Kamiokande experiment [112], which determined that the relative flux of muon and electron neutrinos generated by cosmic ray interactions

in the upper atmosphere had an angular dependence, indicating that the rate at which neutrinos changed from one flavor to another was dependent on the distance they had traveled since creation. This shows that the effect must be something that occurred as the neutrino propagated, rather than at the point of interaction.

These *oscillations* between neutrino flavors have been observed by several different experiments. In this section, we will show how quantum mechanics supplies an explanation for this strange phenomenon through the mixing of flavor states to form mass eigenstates, and discuss how neutrino mass plays an important role in oscillatory phenomena. In addition, we will introduce some of the current experiments that are studying neutrino oscillations, explain the new standard model parameters they are hoping to pinpoint, and how neutrino scattering cross sections play a part in this search.

1.2.2. Flavor state mixing

Boris Kayser et al.’s book “The physics of massive neutrinos” [135] provides an excellent description of the theory of neutrino oscillations, which is outlined here.

We know (or, at least, we have so far observed) that all interactions involving a neutrino involve a particular flavor or “weak interaction” eigenstate - that is, each flavor of neutrino ν_l is coupled to its equivalent lepton l : electron neutrino to electron and so on. However, if neutrinos are massive, it is possible that the freely propagating neutrino states, each of which must have a definite mass, are not in fact these flavor eigenstates, but a different set of states ν_m . In this case, flavor eigenstates would be a linear combination

of these individual mass eigenstates:

$$\nu_l = \sum_m U_{lm} \nu_m$$

and conversely, you could also express a mass eigenstate as a combination of flavor states:

$$\nu_m = \sum_{l'} U_{l'm}^* \nu_{l'}$$

Note that this mixing of flavor states to form mass states was first observed in the quark sector, where small amounts of cross-generational couplings were seen, leading Glashow, Iliopoulos and Maiani [118] to propose that instead of a d quark (mass state), the weak interaction coupled to a combination of d and s quarks, defined by the Cabibbo angle θ_C : $d' = d \cos \theta_C + s \sin \theta_C$, leading to a small component of strangeness in decays. As further quarks were discovered, this was extended to produce the CKM (Cabibbo-Kobayashi-Maskawa) [67][139] matrix combining the mass states into weak interaction flavor states:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}$$

where the current best measurements of the magnitudes of the CKM matrix V are [163]:

$$|V_{ij}| = \begin{pmatrix} 0.97425 \pm 0.00022 & 0.2253 \pm 0.0008 & 0.00413 \pm 0.00049 \\ 0.225 \pm 0.008 & 0.986 \pm 0.016 & 0.0411 \pm 0.0013 \\ 0.0084 \pm 0.0006 & 0.0400 \pm 0.0027 & 1.021 \pm 0.032 \end{pmatrix}$$

The equivalent unitary matrix for neutrinos U_{lm} is known as the neutrino mixing matrix, or the *Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix* [173][148]. This is the neutrino-sector analogy of the quark-sector CKM matrix. We can substitute this into the plane-wave wave function for a neutrino propagating through space and time to give

$$\psi(x, t) = \sum_m U_{lm} \nu_m e^{ip_\nu x - iE_m t}$$

where the energy E_m is related to the neutrino's momentum p_ν and the mass M_m of the eigenstate by the special relativity relation

$$E_m^2 = p_\nu^2 + M_m^2$$

(where we are using natural units, meaning $\hbar = c = 1$). If we assume that (as is seen to be the case for neutrinos), the particle is moving at high speed, close to the speed of light, such that $p_\nu \gg M_m$, we can Taylor expand to get

$$E_m \simeq p_\nu + \frac{M_m^2}{2p_\nu}$$

Additionally, we can approximate that, as it has a speed $\simeq 1$ in our units, $x \simeq t$, giving us

$$\psi(x) \simeq \sum_m U_{lm} \nu_m e^{-i(M_m^2/2p_\nu)x}$$

or, substituting in our expression for ν_m as a combination of flavor states,

$$\psi(x) = \sum_{l'} \left[\sum_m U_{lm} e^{-i(M_m^2/2p_\nu)x} U_{l'm}^* \right] \nu_{l'}$$

So our wave function is actually a superposition of all the flavor states $\nu_{l'}$. After a neutrino created with flavor l has traveled a distance x , it has an amplitude to have evolved into any flavor l' , which is given by the coefficient of $\nu_{l'}$.

So, for example, the probability that a muon neutrino is created, but that, after traveling a distance x , it is detected as an electron neutrino can be calculated by considering the amplitude that a muon neutrino is of mass state m_1 , and multiplying it by the amplitude that an m_1 that has traveled a distance x will be an electron neutrino. You must then add the equivalent amplitudes for other mass states m_2 and m_3 . However, as we're adding amplitudes rather than probabilities, just as in the famous Young's two-slit experiment, we get a kind of "interference pattern" or oscillation of probabilities, which are equivalent to the absolute value squared of the amplitude. Mathematically, we can see this:

$$\begin{aligned}
P(\nu_l \rightarrow \nu_{l'}, x) &= \left[\sum_{m'} U_{lm'}^* e^{-i(M_{m'}^2/2p_\nu)x} U_{l'm'} \right] \times \left[\sum_m U_{lm} e^{-i(M_m^2/2p_\nu)x} U_{l'm}^* \right] \\
&= \sum_m |U_{lm}|^2 |U_{l'm}|^2 \\
&\quad + \sum_{m' \neq m} \text{Re}(U_{lm} U_{lm'}^* U_{l'm'} U_{l'm}^*) \cos\left(\frac{M_m^2 - M_{m'}^2}{2p_\nu} x\right) \\
&\quad + \sum_{m' \neq m} \text{Im}(U_{lm} U_{lm'}^* U_{l'm'} U_{l'm}^*) \sin\left(\frac{M_m^2 - M_{m'}^2}{2p_\nu} x\right)
\end{aligned}$$

Note the sinusoidal behavior in the quantity $\frac{M_m^2 - M_{m'}^2}{2p_\nu} x$. This leads to a characteristic oscillation length $L_{mm'}$, corresponding to the ratio of the momentum and the difference between the squares of the masses:

$$L_{mm'} = 2\pi \frac{2p_\nu}{M_m^2 - M_{m'}^2} = 4\pi \frac{p_\nu}{\Delta M_{mm'}^2}$$

Note that if the matrix U is not real, we will have CP violation effects in neutrino oscillations; in other words, $P(\nu_l \rightarrow \nu_{l'}) \neq P(\bar{\nu}_l \rightarrow \bar{\nu}_{l'})$.

In summary:

- For neutrinos to oscillate, they must have mass
- The masses of the different states must be different ($m \neq m'$) — otherwise, no oscillation takes place
- The mass eigenstates are a mixture of flavor eigenstates and vice versa
- Mixing is dependent on the *difference* in the squares of masses between the eigenstates, not on the absolute masses of the states
- There is a characteristic oscillation length proportional to the neutrino's momentum
- There may be CP violation

1.2.3. Mixing fractions and the PMNS matrix

As shown previously, oscillations will only occur if mass eigenstates consist of a mixture of flavor states. Oscillation experiments probe the extent to which this mixing occurs. The approximate mixtures are illustrated in the figure 1.2, where yellow indicates electron neutrinos, pink muon, and brown tau. These mixture ratios are functions of the PMNS matrix. The mass splittings between the mass states are also indicated.

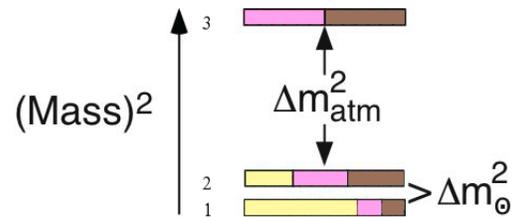


Figure 1.2. Flavor mixing

These mixture ratios are functions of the PMNS matrix. The mass splittings between the mass states are also indicated.

Mass splittings

It should be noted that, as the neutrino oscillations depend only on the squared mass difference between the mass states, rather than on their absolute values, it is unknown whether the states are arranged as shown in the the diagram, with state 3 being far more massive than states 1 and 2 (*normal hierarchy*), or whether it is in fact states 1 and 2 that are heavier, with state 3 being the lightest (*inverted hierarchy*). However, normal hierarchy would intuitively seem the more likely scenario, as this would correspond to electron neutrinos tending to be lighter, while muon and tau neutrinos would be heavier, as is the case for their corresponding charged leptons.

The mass splittings between the states are determined experimentally. The current best fit value for the larger (“atmospheric”) mass splitting $\Delta M_{23}^2 = 2.42 \pm 0.06 \times 10^{-3} eV^2$ [163]. This is from a combined fit to measurements made by Daya Bay [43] which looks at $\bar{\nu}_e$ disappearance over a short baseline in the flux of electron antineutrinos produced by a nuclear reactor; T2K [3] and MINOS [21], which measure ΔM_{23}^2 by looking at muon neutrino disappearance from an accelerator-generated ν_μ beam over a long baseline.

The current world best fit value for the smaller (“solar”) mass splitting ΔM_{12}^2 or $\Delta M_{\odot}^2 = 7.53 \pm 0.18 \times 10^{-5} eV^2$ [163]. This is extracted from a fit to data from the KamLAND long-baseline reactor neutrino oscillation experiment [116] constrained by measurements of solar neutrino oscillations [74][1][54][131][31], and from short baseline reactor experiments Double Chooz [8], Daya Bay [42] and RENO [33] as well as the accelerator experiments T2K [2] and MINOS [17]. Notice that the solar mass splitting is comparable to the uncertainty on the atmospheric splitting.

Absolute mass scale

The mass splittings measured in neutrino oscillations give no indication of the absolute masses of the neutrinos. It is hoped that the mass of the lightest neutrino can be determined by examining the spectrum of low-energy beta-decays. There are several experiments hoping to establish neutrino masses, most notably KATRIN, currently under construction in Karlsruhe, Germany. KATRIN is designed to look at the beta decay of tritium, with a sensitivity of 0.2eV. [155]

Mixing angles

The approximate values of the PMNS matrix are (from [13]):

$$\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} = \begin{pmatrix} 0.8 & 0.5 & 0.1 \\ 0.5 & 0.6 & 0.7 \\ 0.3 & 0.6 & 0.7 \end{pmatrix}$$

However, this is known not to be precise, and experiments are currently underway to determine the matrix elements more accurately. Although it appears that there are 9 parameters to be determined, this is not in fact the case.

We assume that the PMNS matrix is unitary (it is possible that it would not be, if there were one or more sterile neutrinos, but these have not been observed). Furthermore, we know that the total probability of a neutrino being in *some* flavor eigenstate must be equal to 1. From this, it is possible to parameterize the mixing matrix with just four parameters: three *mixing angles*, θ_{12} , θ_{23} and θ_{13} , and a single phase δ_{CP} (this phase creates an imaginary part in some of the matrix elements, that would indicate the presence

of CP violation). Using these parameters, the matrix looks like [198]

$$\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - s_{23}c_{12}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{23}s_{12}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{23}c_{12}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}$$

Here, c_{ij} and s_{ij} are shorthand for $\cos\theta_{ij}$ and $\sin\theta_{ij}$, respectively. Neutrino oscillation experiments attempt to determine these mixing angles and phase shift, as well as the mass differences between the different states. ¹

The current best fits for the mixing angles are (from [163]):

- θ_{13} is around 9° ; ($\sin^2\theta_{13} = 2.19 \pm 0.12$ from $\bar{\nu}_e$ disappearance at reactor experiments Double Chooz [9], RENO[32] and Daya Bay [43]).
- θ_{23} is close to 45° , the maximal mixing. It is unknown which quadrant it falls into ($\sin^2 2\theta_{23} = 0.999_{-0.018}^{+0.001}$, from T2K [3])
- θ_{12} has been measured to be approximately 33° . (Best fit $\sin^2\theta_{12}$ is 0.304 from KamLAND [116]).

1.2.4. CP-violating phase

The parameter δ_{CP} is the CP- (charge-parity symmetry) violating phase. If δ_{CP} is non-zero, there will be differences in oscillatory behavior between neutrinos and antineutrinos.

¹Note: there exists a theory (as yet unconfirmed) that neutrinos could be *Majorana* particles - that is, that they are their own antiparticles, with the neutrino and antineutrino of a given flavor differing only in their helicity. If this is found to be the case, the mixing matrix would be modified by multiplying by additional phase-shift parameters. The Majorana nature of neutrinos is not detectable in oscillation experiments; instead neutrinoless double beta decay experiments such as SuperNEMO [45] are needed to establish whether neutrinos are indeed Majorana.

The value of δ_{CP} is as yet unknown; measuring this phase is one of the top priorities of current and future neutrino oscillation experiments.

The MINOS collaboration have attempted to fit their oscillation data to extract a value of δ_{CP} . The fit gives a 90% confidence level that $0.05\pi < \delta_{CP} < 1.2\pi$ if $\theta_{23} > \pi/4$ and mass hierarchy is normal; they place no restrictions for inverted hierarchy or $\theta_{23} < \pi/4$ [21]. Global fits from Forero et al. [109], using accelerator data from MINOS and T2K and reactor data from experiments such as KamLAND and Daya Bay, they extract a best fit of $\delta_{CP} = 1.34^{+0.64}_{-0.38}\pi$ for normal hierarchy, and $\delta_{CP} = 1.48^{+0.34}_{-0.32}\pi$ for inverted. NOvA's first electron appearance results disfavor $0.1\pi < \delta_{CP} < 0.5\pi$ for inverted hierarchy, with a 90% confidence level.

1.3. Long-baseline oscillation experiments

The major unanswered questions in neutrino physics can be summarized as follows:

- Is the mass hierarchy normal (the small mass splitting is between the two lightest neutrinos) or inverted (the small mass splitting is between the two heaviest neutrinos)?
- What is the value of δ_{CP} ?
- In what octant does θ_{23} fall?
- What is the mass of the lightest neutrino?
- Are neutrinos Majorana or Dirac fermions?
- Does the data fit the three mass and three flavor state mixing paradigm we have described, or is this just part of a more complex theory?

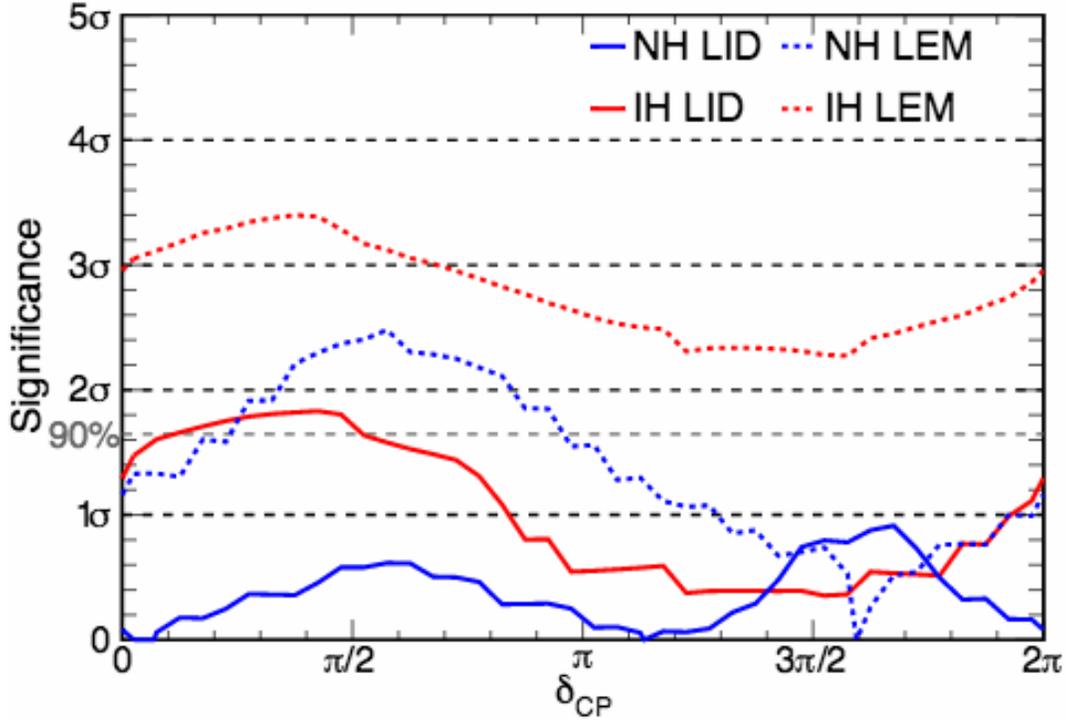


Figure 1.3. Plot from NOvA showing the confidence level with which their observation of 3 electron neutrinos in the far detector disagrees with theory, for different values of δ_{CP} . For their primary technique (solid lines, corresponding to the LID likelihood-based selector technique), this disagreement confidence exceeds 90% for the inverted hierarchy between 0.1 and 0.5π . Taken from [23]. The secondary selector (Library Event Matching) shown in dashed lines does not agree with the current best fit values for the mixing angles.

Long-baseline accelerator-based oscillation experiments are particularly suited to address the first three of these questions, looking at the mass hierarchy, CP-violating phase, and θ_{23} octant. As all of these affect the oscillation probabilities, experiments will fit for combinations of the parameters.

Figure 1.4 shows the setup of T2K, a typical long-baseline oscillation, in Japan. An intense neutrino beam is generated by a particle accelerator (accelerators generate beams of ν_μ or $\bar{\nu}_\mu$, though there will typically be some $\nu_e / \bar{\nu}_e$ contamination. See [98] or section

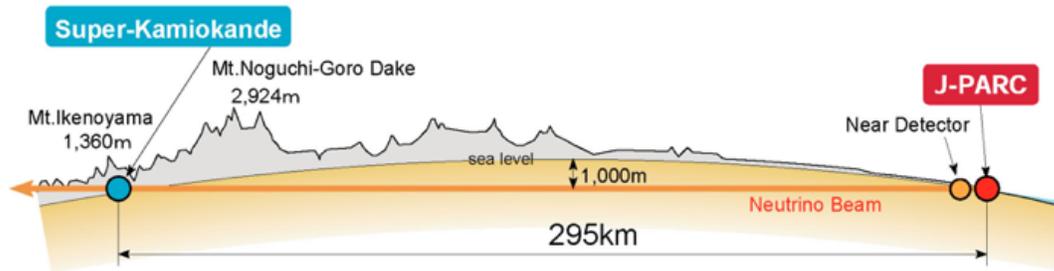


Figure 1.4. Diagram of the T2K oscillation experiment, reprinted from <http://www.t2k.org>

3.2 for a description of the NuMI beam at Fermilab, used by the NOvA and MINOS oscillation experiments, as well as by MINERvA).

A near detector, located at the facility that generates the neutrino beam, records the spectrum of neutrinos before oscillation. The beam then travels through the earth to a far detector. The distance of the far detector is matched to the beam energy, with the aim of it corresponding with the first or second oscillation maximum for one of the mixing angles. The spectrum of neutrinos is examined at the far detector, to determine what fraction of muon neutrinos have disappeared, and what fraction of electron neutrinos have appeared. (Tau neutrinos may also have appeared; these are, however, more difficult to detect). Figure 1.5 shows the predicted ν_e appearance spectrum for the proposed DUNE oscillation experiment.

There are several current running long-baseline neutrino oscillation experiments. MINOS (Main Injector Neutrino Oscillation Search) [94] shares the NuMI beam with MINERvA. This beam has a broad energy spectrum which peaked around 2.5 GeV for its low-energy run until 2012. It is now running in a medium-energy configuration (MINOS+) with a peak energy around 7 GeV. The far detector is situated on the beam

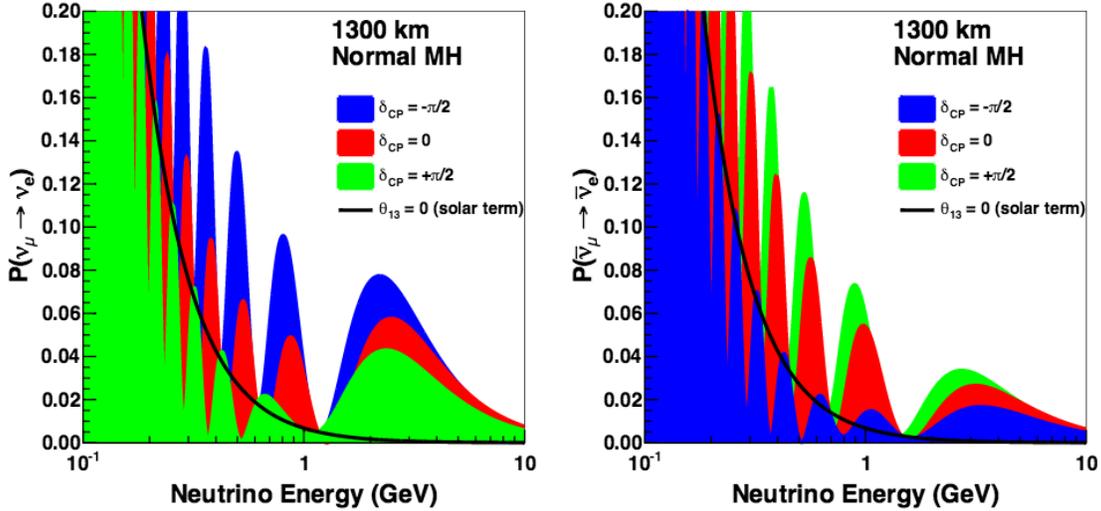


Figure 1.5. Predicted ν_e appearance probabilities at DUNE, for three different values of CP violating phase δ_{CP} , reprinted from [13]

axis, 450 miles from Fermilab at the Soudan mine in Minnesota. MINOS has taken both neutrino and antineutrino data, and has published many analyses, measuring both muon neutrino disappearance and electron neutrino appearance. [21]. This has enabled them to measure θ_{23} and to constrain θ_{13} and δ_{CP} . MINOS has also performed sterile neutrino searches [15] and studied atmospheric neutrinos [22]. The MINOS detectors are both made of magnetized iron as a neutrino target, with polystyrene scintillator for detection.

NOvA, completed in 2014, also shares the NuMI beam, but its far detector is located off the beam axis, 500 miles from Fermilab, at Ash River, Minnesota. By being off axis, NOvA's event rate is lower than that of MINOS, but its energy spectrum much narrower. Both of its detectors are made of liquid scintillator, stored in extruded PVC shells. NOvA has recently submitted its first measurements of electron neutrino appearance [23] and muon neutrino disappearance [24] to the arXiv.

T2K's (Tokai to Kamioka) neutrinos travel 295km from J-PARC to the Super-Kamiokande detector at Kamioka in Japan. Super-Kamiokande is a water Cherenkov detector, situated 2.5° off the beam axis. T2K has two near detectors: INGRID, located on the beam axis, and ND280, off axis at the same angle as the Super-K detector. T2K has measured muon neutrino disappearance and electron neutrino appearance, and was instrumental in confirming that θ_{13} was non-zero. It also has the world's leading measurement of θ_{23} [3]. T2K is now recording antineutrino data. The two near detectors have also been used to make cross section measurements, which will be discussed later.

The Oscillation Project with Emulsion-tRacking Apparatus (OPERA) uses a neutrino beam from CERN in Switzerland, with its far detector at the Laboratori Nazionali del Gran Sasso (LNGS) in Gran Sasso, Italy. OPERA's detector, which consists of bricks of lead plates interleaved with photographic emulsion film, allows it to see tau neutrinos; so far 5 ν_τ candidates have been observed. [26] This confirms $\nu_\mu \rightarrow \nu_\tau$ oscillations (responsible for the ν_μ disappearance at MINOS and other oscillation experiments) with a significance of 5.1σ .

Two next-generation oscillation experiments are currently in the planning phase. The DUNE (Deep Underground Neutrino Experiment) [123] wide-band neutrino beam will travel 800 miles from Fermilab to the Sanford Underground Research Facility in South Dakota. The proposed DUNE detector will be a 40,000 tonne liquid argon time projection chamber. DUNE's physics goals are to measure CP violation, mass hierarchy and the θ_{23} octant.

Hyper-Kamiokande, a proposed new megaton water Cherenkov detector to replace Super-Kamiokande as a T2K upgrade, aims to be able to determine δ_{CP} to a precision of

19° or less for any values of δ_{CP} , and to establish CP violation at 5σ significance for 58% of the parameter space [175].

1.4. Importance of cross sections to oscillation experiments

While the oscillation formula can give a prediction of the probability of neutrino oscillation for a given baseline and neutrino energy, to convert this to an expected number of events, one must also know the initial flux of unoscillated neutrinos, as well as the probability that a neutrino (oscillated or otherwise) will interact within the detector. Thus an accurate cross section model for neutrino scattering from heavy nuclei is vital for experiments to compare their event counts to models' predictions and thus extract physics information such as the value of δ_{CP} and the mass hierarchy.

In order to meet their physics goals, it is important that next generation experiments keep systematic uncertainties to a minimum. DUNE has a goal of 2% uncertainty on its measurements [123]. Figure 1.6 shows the time (and therefore operating cost) savings of reducing the uncertainties: it details the exposure needed to measure CP violation for 75% of possible values of δ_{CP} . We see that, to discover CP violation with 3σ significance is predicted to take approximately 1000kt-MW-years with their standard reference design if they can achieve 5% uncertainty on the total normalization and 1% on the relative normalization of ν_e to the other neutrinos (top line of blue hashed area); with $5 \oplus 2\%$ (middle line), it will take 1250kt-MW-years. Thus reducing uncertainty is key to saving time and expense.

The sources of uncertainty come from neutrino flux spectrum, which can be reduced by comparing near and far detector measurements; fiducial volume identification, which

will be small ($< 1\%$) in such a large detector, energy scale, and interaction models for neutrinos and antineutrinos on nuclei. The goal of the measurement in this analysis is to reduce uncertainty on the interaction models, by constraining them with data from our cross section measurements. According to [13], T2K currently has 5.3% interaction model uncertainty; to meet its physics goals, DUNE must reduce this to 2%. As DUNE intends to measure the CP-violating phase by comparing $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ scattering, antineutrino cross sections, which have been studied less extensively than their neutrino counterparts, are particularly necessary.

1.5. Cross sections as a probe of nuclear effects

As well as reducing model uncertainties for oscillation experiments, cross section measurements can be used to investigate nuclear physics. Neutrino scattering on free nucleons is well understood (the charged-current quasi-elastic scattering discussed in this paper was modeled in 1972 by Llewellyn Smith [145], and this model works well for scattering from hydrogen and deuterium). However, in heavy nuclei, interactions between the nucleons in the nucleus affect the scattering behavior. By examining the cross section distributions and comparing them to various models of these nuclear effects, we are able to increase our knowledge of the nature and strength of these behaviors. As very different processes affect cross sections in similar ways, a double-differential cross section, where we see how the distribution behaves in a two-dimensional phase space, provide vital additional information to help distinguish between these models. The theory of neutrino-nucleus scattering, as well as introductions to several different nuclear models, can be found in chapter 2.

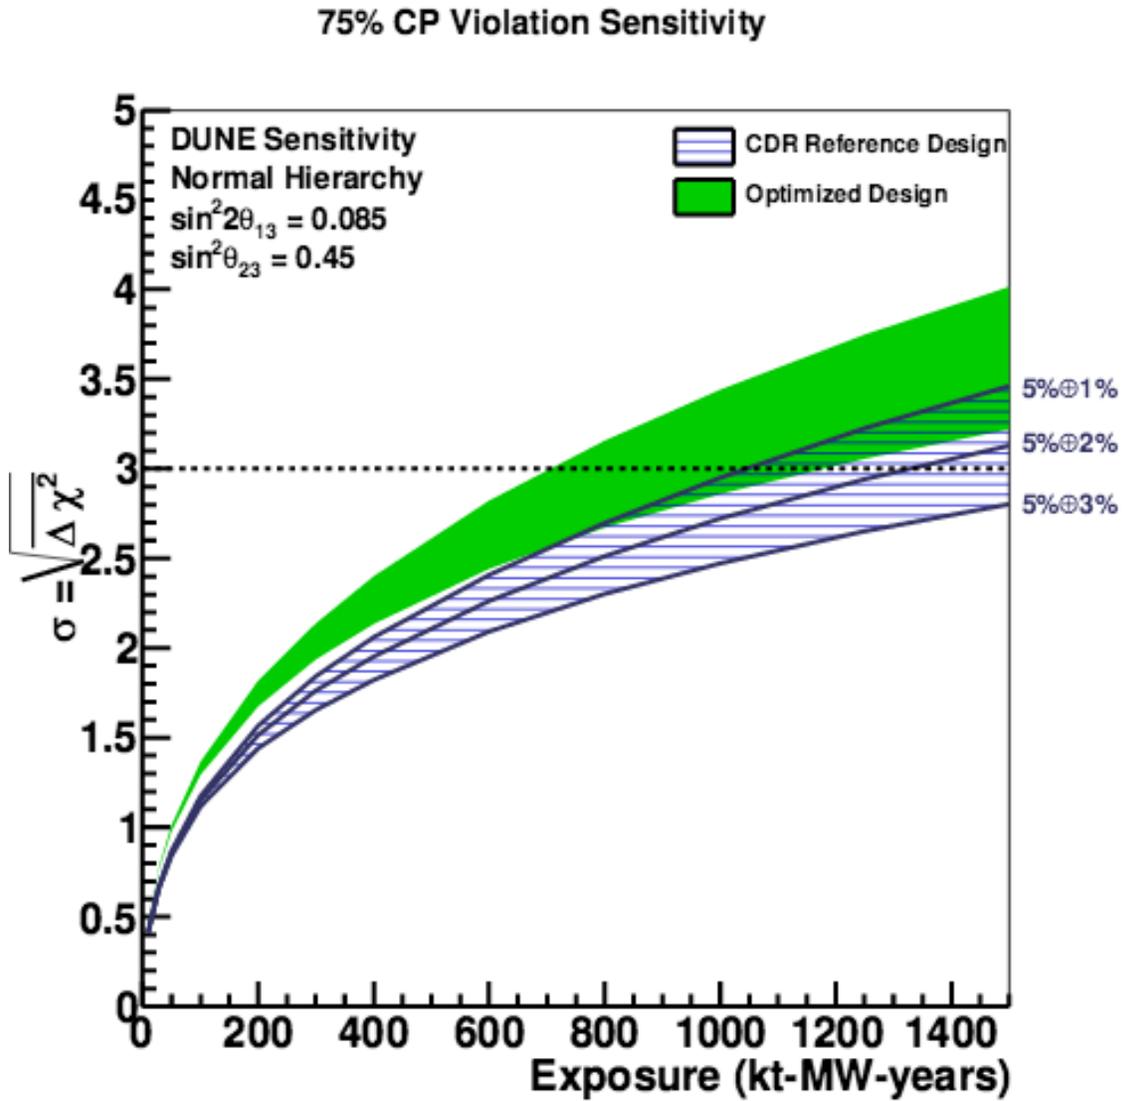


Figure 1.6. Exposure needed for DUNE to measure δ_{CP} for 75% of possible values of δ_{CP} , with different levels of systematic uncertainty. The blue hashed area shows the sensitivity with the current beam design, with the three lines representing how long DUNE must run with uncertainties from $5 \oplus 3$ to $5 \oplus 1\%$, where the two numbers refer respectively to the uncertainty on ν_μ normalization and ν_e normalization relative to ν_μ and the antineutrinos. The dotted line shows the 3σ confidence level. The green colored area shows an equivalent for a new optimized design. Reprinted from [13]

CHAPTER 2

Theory of quasi-elastic neutrino scattering on nuclei

2.1. Quasi-elastic scattering cross-sections

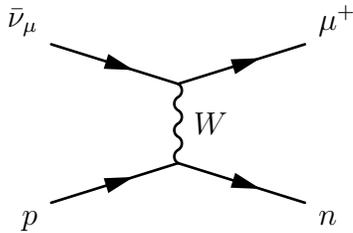


Figure 2.1. Quasi-elastic antineutrino scattering

This analysis looks at charged-current quasi-elastic (CCQE) scattering of muon antineutrinos on the material of MINERvA’s tracker region, which is made up of strips of doped polystyrene scintillator, with a titanium dioxide coating. The percentage composition by mass of this material is detailed in table 3.1; the main con-

stituents are carbon and hydrogen atoms, of which there are almost equal numbers.

A simplistic description of CCQE scattering is that the incoming antineutrino interacts with a target proton within the nucleus, exchanging a W boson to knock out a neutron, also leaving a positively charged muon in the final state:

$$\bar{\nu}_\mu + p \rightarrow \mu^+ + n$$

(see figure 2.1). It should be noted that, in the quasi-elastic case, the neutrino can be considered to be scattering off of the nucleon, rather than off of one of its constituent quarks. (Scattering from the quarks can also occur; this is known as “deep inelastic scattering”, and dominates at higher neutrino energies than does quasi-elastic scattering).

In the case of pure quasi-elastic scattering, it is possible to reconstruct certain characteristics of the interaction using only the kinematics of the outgoing charged lepton - particularly useful as muons tend to be relatively easy to reconstruct in current neutrino detectors. In particular, the incoming neutrino energy and the four-momentum transfer Q^2 can be estimated.

When calculating scattering amplitudes, we must remember that nucleons are not point-like particles, but that they have finite size and complex internal structure.

Additionally, we should note that the target material used in this experiment consists largely of carbon; thus, the protons from which our antineutrinos scatter are frequently bound within a nucleus consisting of twelve nucleons. The nucleons within a nucleus interact with each other in complicated ways which are not fully understood. This can affect the initial state of the target proton in a scattering experiment, as well as modifying the final state as the ejected neutron may interact with other nucleons while escaping the nucleus. It is also suspected that incident neutrinos may interact with bound multi-nucleon states within the nucleus. These effects are complicated and not yet fully understood, yet can cause significant modifications to the free-nucleon scattering cross section. In this section, we will explain the basic theory of quasi-elastic scattering from a free nucleon, followed by descriptions of contemporary theories that are used to model nuclear effects.

2.1.1. Elastic scattering

In elastic scattering, energy is conserved, and both the incoming particle and the target retain their structure. A well-known, and relatively easy-to-study, example of this would be electromagnetic scattering of electrons from protons, mediated by the exchange of a

virtual photon. “Quasi-elastic” scattering shares many similarities with elastic scattering, with the exception that in the quasi-elastic case, it is a weak interaction involving a charge exchange, mediated by a W boson. In attempting to understand the theory behind quasi-elastic scattering, it may be easiest to first consider the elastic case before examining how this is modified for the quasi-elastic.

2.1.1.1. Mott scattering. The simplest example of electromagnetic scattering involves two point-like particles, such as an electron and a muon, known as “Mott scattering” [159]

$$\mu^- + e^- \rightarrow \mu^- + e^-$$

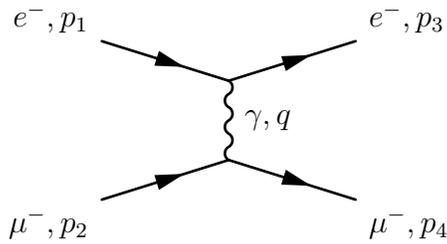


Figure 2.2. Electromagnetic electron-muon elastic (Mott) scattering

The Feynman diagram for this process, at tree level, is mediated by a photon, as shown in Figure 2.2. The spin-averaged differential cross-section for this interaction can be calculated analytically. The following calculation follows the method of [129].

The differential cross-section is given by :

$$\frac{d\sigma}{d\Omega} = \left(\frac{\hbar}{8\pi m_\mu c} \right)^2 \langle |\mathcal{M}|^2 \rangle$$

where the matrix element \mathcal{M} is given by the Feynman rules as:

$$\mathcal{M} = -\frac{g_e^2}{(p_1 - p_3)^2} [\bar{u}^{s_3}(p_3)\gamma^\mu u^{s_1}(p_1)] [\bar{u}^{s_4}(p_4)\gamma_\mu u^{s_2}(p_2)]$$



Figure 2.3. Mott scattering

where the $u^{s_n}(p_n)$ refer to the Dirac spinors corresponding to the spin s_n and momentum p_n of each fermion n . Using the method of traces to average over initial spins and sum over final spins, we get:

$$(2.1) \quad \langle |\mathcal{M}|^2 \rangle = \frac{g_e^4}{q^4} L_{\text{electron}}^{\mu\nu} L_{\mu\nu}^{\text{muon}}$$

where

$$(2.2) \quad L_{\text{electron}}^{\mu\nu} \equiv 2\{p_1^\mu p_3^\nu + p_1^\nu p_3^\mu + g^{\mu\nu}[(m_e c)^2 - (p_1 \cdot p_3)]\}$$

and equivalently for the muon, replacing the electron momenta and mass with muon equivalents.

Substituting equation 2.2 into equation 2.1, we get:

$$(2.3) \quad \langle |\mathcal{M}|^2 \rangle = \frac{8g_e^4}{(p_1 - p_3)^4} \left[(p_1 \cdot p_2)(p_3 \cdot p_4) + (p_1 \cdot p_4)(p_2 \cdot p_3) \right. \\ \left. - (p_1 \cdot p_3)(m_\mu c)^2 - (p_2 \cdot p_4)(m_e c)^2 + 2(m_e m_\mu c^2)^2 \right]$$

As the muon mass $m_\mu \gg$ electron mass m_e , we can simplify by treating the electron as massless, and the muon as remaining stationary throughout the collision (see figure 2.3). Following the derivation in [129], we get the cross section distribution for electron-muon electromagnetic scattering:

$$(2.4) \quad \frac{d\sigma}{d\Omega_{\text{Mott}}} = \left(\frac{\alpha \hbar}{2\mathbf{p}^2 \sin^2 \frac{\theta}{2}} \right)^2 \left[(m_e c)^2 + \mathbf{p}^2 \cos^2 \frac{\theta}{2} \right]$$

where $g_e = \sqrt{4\pi\alpha}$ and the fine structure constant $\alpha \approx 1/137$.

Higher order diagrams, such as that shown in Figure 2.4 also contribute to the cross-section, but as each vertex contributes a factor of $\alpha_{QED} \approx \frac{1}{137}$, these make little difference to the total cross-section. A similar process may also take place via the weak interaction, in which case the mediator is a Z boson rather than a photon. This process, however, is heavily suppressed below the 90GeV Z resonance peak.

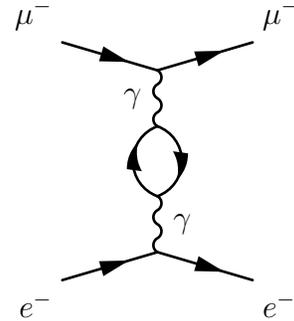


Figure 2.4. Leading-order correction to $e^- \mu^-$ scattering

2.1.1.2. Electron-proton scattering. Like the muon, the proton is a fermion; if it were point-like, electron-proton scattering would have exactly the same form as that just shown for Mott scattering:

$$\langle |\mathcal{M}|^2 \rangle = L_{\text{electron}}^{\mu\nu} L_{\mu\nu}^{\text{proton}}$$

with $L_{\text{electron}}^{\mu\nu}$ defined as in equation 2.2 and $L_{\mu\nu}^{\text{proton}}$ defined similarly, by replacing electron momenta and mass with those of the proton.

However, it is known that the proton is not a point charge; it is a composite particle with a finite size, consisting of strongly-interacting partons: three valence quarks (uud) plus the “sea” of gluons and quark-antiquark pairs. At high energies, typically of the several-GeV range, electrons can scatter off of these individual partons (“deep inelastic scattering”, $e^- + p \rightarrow e^- + X$ in which the final products are not simply the electron and proton with which we started, but a shower of hadrons). At lower energies, of the order of 1 GeV, we tend to see elastic scattering. The proton is even heavier than the muon, so as with the previous example of Mott scattering, we can assume an electron scatters from a stationary proton, and that the electron’s mass can be neglected. However, the interaction between the virtual photon and the proton is now no longer a simple QED vertex, as in the Mott scattering example, meaning that our matrix element must now be modified. Again, we follow the method of [129]. The matrix element becomes:

$$\langle |\mathcal{M}|^2 \rangle = L_{\text{electron}}^{\mu\nu} K_{\mu\nu}^{\text{proton}}$$

where $K_{\mu\nu}^{\text{proton}}$ is some second-rank tensor that models the unknown interaction between the proton and the mediating photon. Furthermore, the tensor can only depend on the four-momentum transferred to the proton, q , and the incoming proton momentum p (the outgoing momentum is simply $p + q$), as these are the only momenta that contribute at the proton vertex.

Bearing in mind that in order to contribute to the cross-section, terms in $K_{\mu\nu}^{\text{proton}}$ must, as with $L_{\text{electron}}^{\mu\nu}$, be symmetric, the most general form is given below (m_p is the proton

mass):

$$(2.5) \quad K_{\text{proton}}^{\mu\nu} = -K_1 g^{\mu\nu} + \frac{K_2}{(m_p c)^2} p^\mu p^\nu + \frac{K_4}{(m_p c)^2} q^\mu q^\nu + \frac{K_5}{(m_p c)^2} (p^\mu q^\nu + p^\nu q^\mu)$$

(There is no K_3 - this is reserved for a parity-violating term that is not relevant to electromagnetic scattering, which conserves parity. The weak interaction, however, violates parity, and thus for neutrino scattering, an additional form factor is included, of the form $-i(\epsilon^{\mu\nu\alpha\beta} p_\alpha q_\beta)/2m_p^2$, as detailed in [145]). For elastic scattering, these form factors are functions only of the squared four-momentum transfer q^2 ; for inelastic scattering, the form factors would also be functions of energy transfer ν . The quantity q^2 is used frequently in calculations and measurements involving this type of interactions; more frequently we refer to $Q^2 = -q^2$, as this quantity is always positive. It can be shown that the $K_i(q^2)$ are not all independent; K_4 and K_5 can be eliminated, giving:

$$K_{\text{proton}}^{\mu\nu} = -K_1 \left(g^{\mu\nu} + \frac{q^\mu q^\nu}{q^2} \right) + \frac{K_2}{(m_p c)^2} \left[\frac{(p_1 \cdot p)(p_3 \cdot p)}{(m_p c)^2} + \frac{Q^2}{4} \right]$$

2.2. Quasi-elastic neutrino scattering

Neutrinos, having no electric charge, do not undergo electromagnetic interactions, meaning that elastic scattering via virtual photon exchange, as shown in the previous section, does not take place. However, neutrinos do undergo weak interactions, and it is thus possible for neutral-current elastic scattering to take place via exchange of a Z boson, as shown in figure 2.5a.

Figures 2.5b and 2.5c show the charged-current quasi-elastic process for an antineutrino and a neutrino respectively. These processes produce a charged lepton in the final

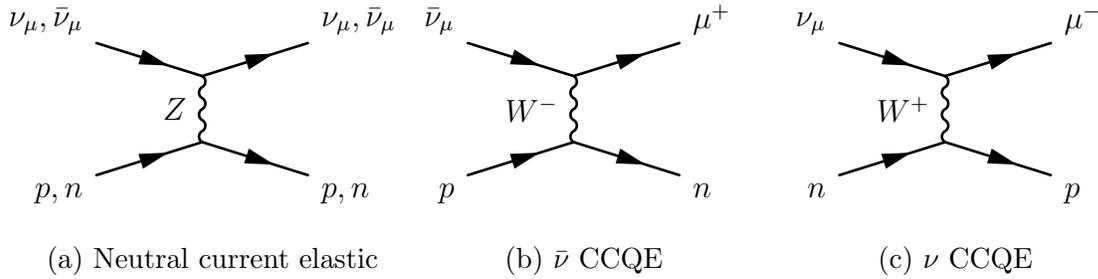


Figure 2.5. Elastic and quasi-elastic scattering of neutrinos from nuclei

state. (In this example, we show scattering of muon neutrinos, but the process can equally take place with tau or electron neutrinos, in which case a tau lepton or electron respectively would be produced in the final state). These leptons can be detected, and their charge and momentum analyzed, not only in the MINER ν A detector, but also in those of numerous neutrino experiments, including those that look for oscillations. In this case, the mediating particle is the charged W boson, which causes a neutrino to change to its charged leptonic partner, while simultaneously changing the flavor of the target nucleon. Neutrinos interact with neutrons, with a W^+ being exchanged from the lepton to the hadron :

$$\nu_l + n \rightarrow l^- + p$$

while antineutrinos interact with protons, with a W^- being exchanged from the lepton to the hadron (or, equivalently, a W^+ from the hadron to the lepton):

$$\bar{\nu}_l + p \rightarrow l^+ + n$$

This analysis will principally concentrate on the antineutrino interaction, which is illustrated in figure 2.5b.

CCQE interactions are of particular interest to oscillation experiments for several reasons. They dominate at energies in the GeV range, a common energy range for neutrino beams (for example, the Fermilab booster beam used by MiniBooNE has a mean energy of 0.5 GeV [77]; T2K's beam is centered at 0.6 GeV [78]; and MINOS[94] and NOvA[47] are situated, along with MINERvA, in the NuMI beam [98] which, in its low energy configuration, had a mean energy around 3 GeV, and now delivers 1-3 GeV neutrinos to NOvA's off-axis detector and a broad-spectrum beam peaking around 6 GeV to MINOS and MINERvA). Additionally, in the case of a pure quasi-elastic interaction, we are able to use conservation of energy and momentum parallel and perpendicular to the neutrino's direction to reconstruct both the energy of the incoming neutrino, E_ν ; and the negative square of the 4-momentum transferred from the leptonic to the hadronic system, Q^2 . Q^2 is an important quantity because nucleon form factors (as explained above, and shown again below) are functions of Q^2 .

$$(2.6) \quad E_\nu^{QE} = \frac{m_n^2 - (m_p - E_b)^2 - m_\mu^2 + 2(m_p - E_b)E_\mu}{2(m_p - E_b - E_\mu + p_\mu \cos \theta_\mu)}$$

$$(2.7) \quad Q_{QE}^2 = 2E_\nu^{QE}(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2$$

These formulae are valid for a quasi-elastic interaction antineutrino incident upon a proton at rest within a nucleus, with a binding energy E_b . (Note that this assumption of a stationary proton with a constant binding energy is an approximation; this will be discussed in subsequent sections.) The interaction produces a positively-charged muon and a recoil neutron. Under the quasi-elastic assumption, no energy is lost to the rest of

the nucleus - its only effect is to provide the binding energy that lowers the initial state energy of the stationary proton. In these formulae, E_ν and E_μ are the neutrino and muon energy respectively. Muon momentum is represented by p_μ , and θ_μ represents the angle between the outgoing muon and the incoming neutrino. Four-momentum transfer q is the difference in 4-momentum between the final state muon and the initial-state neutrino. As the neutrino mass is negligible (less than 1eV), we take $m_\nu = 0$, meaning $E_\nu = |\vec{p}_\nu|$. The muon mass is represented by m_μ . We recall that $E^2 = m^2 + p^2$ in natural units (where the speed of light is set to 1).

Muons typically behave as minimum-ionizing particles in detectors, meaning that their kinematics are relatively easy to reconstruct. This makes this interaction especially appealing for oscillation experiments that wish to compare measured to theoretical cross-sections.

2.2.1. The Llewellyn-Smith model for quasi-elastic cross-section

As with the elastic electron-proton scattering discussed earlier, we are unable to make a precise analytical calculation of the neutrino-nucleon quasi-elastic cross-section; due to the internal structure of the nucleon, our cross-section depends on nucleon form factors. Recall the form-factor parametrization given earlier for the proton, with the additional parity-violating term:

$$(2.8) \quad K_{\text{proton}}^{\mu\nu} = -K_1 g^{\mu\nu} + \frac{K_2}{(m_p c)^2} p^\mu p^\nu - i \frac{K_3}{2m_p^2} \epsilon^{\mu\nu\alpha\beta} p_\alpha q_\beta + \frac{K_4}{(m_p c)^2} q^\mu q^\nu + \frac{K_5}{(m_p c)^2} (p^\mu q^\nu + p^\nu q^\mu)$$

In 1972, C. Llewellyn-Smith [145] used these form factors to calculate the differential quasi-elastic cross-section. He regroups the form factors in the following way:

$$(2.9) \quad \frac{d\sigma}{dQ^2}_{QE} \left(\begin{array}{l} \nu_l n \rightarrow l^- p \\ \bar{\nu}_l p \rightarrow l^+ n \end{array} \right) = \frac{M^2 G_F^2 \cos^2 \theta_C}{8\pi E_\nu^2} \left\{ A(Q^2) \mp B(Q^2) \frac{s-u}{M^2} + C(Q^2) \frac{(s-u)^2}{M^4} \right\}$$

where:

G_F : is the Fermi coupling constant, $1.166 \times 10^{-5} GeV^{-2}$

M : is the nucleon mass; $M_{proton} = 938.27 MeV/c^2$; $M_{neutron} = 939.57 MeV/c^2$

θ_C : is the Cabibbo angle, 13.04°

s, u : are the Mandelstam variables; $s - u = 4ME_\nu - Q^2 - m_l^2$

E_ν : is the incoming neutrino energy which, in the quasi-elastic hypothesis, can be calculated from the angle and energy of the final state lepton.

Q^2 : is the square of the four-momentum transferred from the lepton to the hadron which, in the quasi-elastic hypothesis, can be calculated from the angle and energy of the final state lepton.

(Constants from [163].) The coefficients A, B and C are functions of the nuclear form-factors:

$$(2.10) \quad A(Q^2) = \frac{m_l^2 + Q^2}{M^2} \left\{ \left(1 + \frac{Q^2}{4M^2} \right) |F_A|^2 - \left(1 - \frac{Q^2}{4M^2} \right) F_1^2 \right. \\ \left. + \frac{Q^2}{4M^2} \left(1 - \frac{Q^2}{4M^2} \right) (\xi F_2)^2 + \frac{Q^2}{M^2} Re(F_1^* \xi F_2) - \frac{Q^2}{M^2} \left(1 + \frac{Q^2}{4M^2} \right) (F_A^3)^2 \right. \\ \left. - \frac{m_\mu^2}{4M^2} \left[|F_1 + \xi F_2|^2 + |F_A + 2F_P|^2 - 4 \left(1 + \frac{Q^2}{4M^2} \right) ((F_V^3)^2 + F_P^2) \right] \right\}$$

$$(2.11) \quad B(Q^2) = \frac{Q^2}{M^2} Re [F_A^* (F_1 + \xi F_2)] - \frac{m_l^2}{M^2} Re \left[(F_1 - \tau \xi F_2) F_V^{3*} - \left(F_A^* - \frac{Q^2}{2M^2} F_P \right) F_A^3 \right]$$

$$(2.12) \quad C(Q^2) = \frac{1}{4} \left\{ F_A^2 + F_1^2 + \tau(\xi F_2)^2 + \frac{Q^2}{M^2} (F_A^3)^2 \right\}$$

The form factors are associated with different physics processes, and all but F_A are known to a good level of approximation from other processes, as explained below.

Of these, F_1 and F_2 are vector form factors, F_P pseudoscalar, and F_A axial vector. The form factors F_V^3 and F_A^3 correspond to second-class currents which, if they exist at all, have a small enough effect that they can be neglected [194].

2.2.2. Form factors in the Llewellyn-Smith model

2.2.2.1. Conserved Vector Current and the vector form factors. The weak interaction has both a vector and an axial vector component, leading to the factor $\gamma_\mu(1 - \gamma^5)$ in the weak vertex between elementary particles. However, the internal strong interactions between the nucleon's components (valence quarks, gluons, and sea quark-antiquark pairs which are constantly being created and recombining) modify this vertex factor; the term $(1 - \gamma^5)$ can now be written as $(c_V - c_A \gamma^5)$, where c_V is a weak vector charge and c_A a weak axial charge. The Conserved Vector Current (CVC) hypothesis [117][102] claims that the vector component will be the same for charged-current interactions (such as neutrino-nucleon scattering) as it is for the equivalent electromagnetic interaction (charged lepton-nucleon scattering). In other words, $c_V = 1$ exactly. This is because, in the Standard Model, the vector parts of the weak charged current and its hermitian conjugate, along with the pure-isovector electromagnetic current, form an isospin triplet. This hypothesis has been confirmed by various experiments, such as [186].

Thus that the vector form factors can be measured via electron-nucleon scattering, a simpler experiment due to the relative ease of producing a mono-energetic electron beam.

Using [145] [102] to relate the CCQE vector form factors to the Sachs form factors [91][181] G_E^V and G_M^V (explained below), one gets (for scattering from a free nucleon):

$$F_1(Q^2) = \frac{G_E^V(Q^2) + \frac{Q^2}{4M^2} G_M^V(Q^2)}{1 + \frac{Q^2}{4M^2}}$$

$$\xi F_2(Q^2) = \frac{G_M^V(Q^2) - G_E^V(Q^2)}{1 + \frac{Q^2}{4M^2}}$$

where G_E^V and G_M^V are related to the electron scattering form factors by

$$G_E^V(Q^2) = G_E^p(Q^2) - G_E^n(Q^2)$$

$$(2.13) \quad G_M^V(Q^2) = G_M^p(Q^2) - G_M^n(Q^2)$$

The electric (G_E^p, G_E^n) and magnetic (G_M^p, G_M^n) form factors of the proton and neutron can be measured via electron scattering. At $Q^2 = 0$, the electric form factors simply correspond to the electric charge of the proton (1, in Heaviside particle physics units) and neutron (0), while the magnetic form factors correspond to the magnetic moments μ_p and μ_n respectively.

At $Q^2 > 0$, the behavior of the form factors must be determined experimentally (from electron scattering). Budd, Bodek and Arrington[65] fitted measured cross-sections to various theoretical models. The simplest model commonly used experiments is a dipole approximation, based around the assumption that the nucleon's charge has an exponential

distribution $\rho(r) = \rho_0 e^{-Mr}$, whose Fourier transform gives the dipole form:

$$(2.14) \quad \begin{aligned} G_D(Q^2) &= \frac{1}{\left(1 + \frac{Q^2}{4M_V^2}\right)^2} \\ G_E^p(Q^2) &= G_D(Q^2) & G_E^n(Q^2) &= 0 \\ G_M^p(Q^2) &= \mu_p G_D(Q^2) & G_M^n(Q^2) &= \mu_n G_D(Q^2) \end{aligned}$$

In [65], M_V^2 is taken to be 0.71 GeV².

The dipole approximation is thought to hold well at low Q^2 , where the key issue is the spatial structure of the nucleus; however, they are less accurate at higher Q^2 values, where quark structure plays a more important role. Budd, Bodek and collaborators have since proposed various modifications to this model[65][63][58], the most recent being the so-called BBBA07 parameterization[58]. This is based around a parameterization by Kelly[136] of the nucleon form factors. For G_M^p , G_E^p and G_M^n , these use a ratio of two polynomials in $\tau = \frac{Q^2}{4M^2}$ (M is nucleon mass), in which the degree of the denominator is two powers of Q^2 greater than that of the numerator, in order to replicate the Q^{-4} behavior of the form factors observed at large Q^2 . A pre-factor of magnetic moment μ_n or μ_p is included for G_M^n and G_M^p respectively (see figure 2.6).

As insufficient data was available to model G_E^n in this way, the Galster form-factor[115] $G_E^n = \frac{A\tau}{1+B\tau} G_D$ was used instead. The polynomial coefficients were obtained through fits to data as described in [136]. The Kelly form factors are multiplied by scaling functions that affect their behavior at high Q^2 , having been fitted to match data from neutrino-deuterium quasi-elastic scattering. (Deuterium was chosen in order to minimize the effect of unknown nuclear corrections).

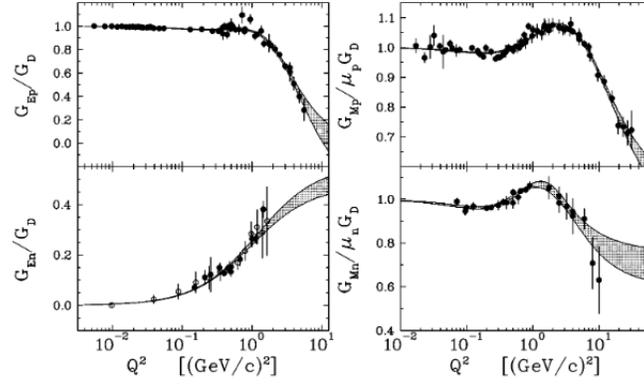


Figure 2.6. Polynomial fit of form factor values measured by several experiments, from [136].

2.2.2.2. The pseudoscalar form factor. We now turn our attention to the axial component of the weak interaction.

Following the formalism of [80], we note that the axial part of the proton-neutron weak matrix element has the form

$$(2.15) \quad \langle p | A^{\dagger\mu} | n \rangle = -\cos\theta_C \bar{u}_p \left[F_A \gamma^\mu \gamma_5 + F_P \frac{q^\mu}{M} \gamma_5 \right] u_n$$

The first “axial” term corresponds to the basic leading order diagram for the weak interaction, plus vertex corrections such as that shown in figure 2.7a. The second term corresponds to pseudoscalar interactions, the most significant being that shown in figure 2.7b, where a charged pion (a pseudoscalar particle) with the same charge as the exchanged W boson, is exchanged.

In this second case, our diagram now has two vertices: one corresponding to $p + \pi^- \rightarrow n$ and another which is a rotation of the weak decay of the pion: $\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$. Connecting

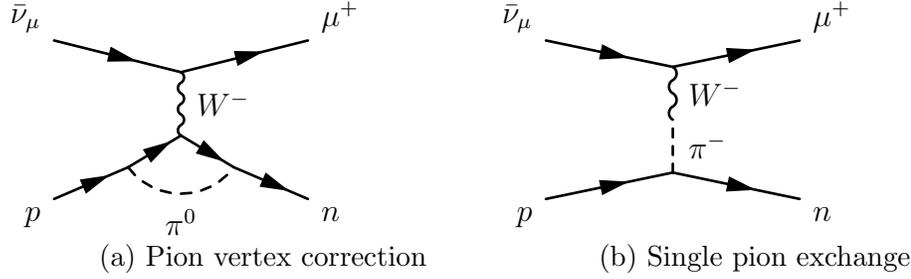


Figure 2.7. Pion corrections to the CCQE interaction

these we have a charged pion propagator. Thus we get an amplitude that looks like:

$$\begin{aligned}
 \mathcal{A} &= [p + \pi^- \rightarrow n \text{ vertex}] \times [\pi^- \text{ propagator}] \times [\pi^- \rightarrow \mu^- \bar{\nu}_\mu \text{ vertex}] \\
 (2.16) \quad &= [ig_0 \sqrt{2} \bar{u}_p \gamma^5 u_n] \times \left[\frac{1}{q^2 - m_\pi^2} \right] \times \left[\cos \theta_C \frac{G_F}{\sqrt{2}} i f_\pi q_\mu (\bar{u}_{\text{muon}} \gamma^\mu (1 - \gamma_5) u_{\nu_\mu}) \right]
 \end{aligned}$$

where g_0 is the coupling constant for the $p + \pi^- \rightarrow n$ interaction and f_π is the charged pion decay constant. This would give

$$(2.17) \quad \langle p | A^{\dagger\mu} | n \rangle = -\cos \theta_C \bar{u}_p \left[F_A \gamma^\mu - \frac{q^\mu}{M} \frac{g_0 f_\pi}{q^2 - m^2} \right] \gamma_5 u_n$$

The pseudoscalar part (the part multiplied by just γ^5) gives us our pseudo scalar form factor F_P :

$$(2.18) \quad F_P(q^2) = -\frac{g_0 f_\pi}{q^2 - m_\pi^2}$$

While the axial current is not conserved; it can be considered “partially conserved”; in that it would be conserved in the limit that $m_\pi \rightarrow 0$. Thus, taking a divergence of our axial current in that limit, we get:

$$\begin{aligned}
(2.19) \quad 0 &= \lim_{m_\pi \rightarrow 0} \frac{\partial \mathcal{A}^\mu}{\partial x^\mu} \\
&= \lim_{m_\pi \rightarrow 0} \bar{u}_p \left[F_A(M_p + M_n) - \frac{f_\pi g_0}{q^2 - m_\pi^2} \frac{q^2}{M} \right] \gamma_5 u_n \\
&= \bar{u}_p [2MF_A(0) - f_\pi g_0/M] \gamma_5 u_n
\end{aligned}$$

(We approximate $M_p \approx M_n \equiv M$). This gives us the ‘‘Goldberger-Treiman relation’’ [122]

$$(2.20) \quad f_\pi g_0 = 2M^2 F_A(0)$$

Finally, substituting back to 2.18, we can relate our pseudo scalar form-factor to the axial form factor:

$$(2.21) \quad F_P = -\frac{f_\pi g_0}{q^2 - m_\pi^2} = \frac{2M^2}{Q^2 - m_\pi^2} F_A$$

Note that the expression for F_P only appears in the Llewellyn-Smith cross-section formula with a multiplier of $(M_l/M_N)^2$. Thus its effect is small except in the case of tau neutrino interactions.

2.2.2.3. The axial form factor. The weak interaction (by which neutrinos interact) has a (V-A) form [102], including both a vector and an axial component; the electromagnetic interaction, meanwhile, is pure vector. For this reason, it is very difficult to measure the strength due to the axial component by looking at electromagnetic electron-nucleon scattering (there is a small contribution from parity-violating weak neutral-current scattering, but this is very hard to detect). The axial component, represented by the axial form factor F_A , is therefore measured through either neutrino-nucleon scattering or pion

electro-production. As with the vector form factor, it is usual to use a dipole approximation,

$$(2.22) \quad F_A(Q^2) = \frac{g_A}{\left(1 + \frac{Q^2}{M_A^2}\right)^2}$$

The constant g_A , the value of the axial form-factor at $Q^2 = 0$, has been measured through beta-decay experiments [193] to be 1.2756(30)[150], (MINERvA's nominal Monte Carlo simulation, GENIE [93], uses 1.2670) leaving one free parameter, the axial mass M_A . Bubble chambers (in which neutrinos scatter from deuterium) [50] [35] [57] and pion electro-production [71] experiments have measured the value of M_A on free or quasi-free nucleons.

An average value of $M_A = 1.014 \pm 0.014 \text{ GeV}/c^2$ was extracted by Bodek et al. [57] in 2008, from various neutrino-deuterium scattering experiments. Kuzmin et al. [140] performed global fits to neutrino and antineutrino scattering data on hydrogen, deuterium and several materials including heavy nuclei, using alternative models for the vector form factors. The extracted values were all consistent with Bodek et al.'s world average. Our Monte Carlo simulation, GENIE, [93] uses $M_A = 0.99_{-15\%}^{+25\%} \text{ GeV}$. While this model works well for neutrino-deuterium scattering, it has been shown to break down when scattering from heavier nuclei, as will be discussed in the following section. A summary of the history of measurements of M_A can be found in [114].

2.3. Scattering from heavy nuclei

2.3.1. The Relativistic Fermi Gas model

Protons and neutrons are fermions. This means that they must obey the Pauli exclusion principle, which forbids two identical fermions from occupying identical states. Because of this, multiple protons or neutrons in a nucleus each have a number distribution dictated by Fermi-Dirac statistics:

$$(2.23) \quad n_i = \frac{1}{e^{\beta(E_i - \mu)} + 1}$$

where n_i denotes the number of protons or neutrons in a given energy state, E_i is the energy of the state, μ is the chemical potential and $\beta = 1/kT$ where k is Boltzmann's constant and T is the temperature. As temperature reduces to absolute zero, this results in a distribution where all energy states are filled up to the Fermi energy $E_F = \mu_{(T=0)}$ while all states above E_F are empty. As temperature rises, the distribution smears, with some states above E_F being filled, and some below becoming empty. We can thus think of the nucleus as a gas consisting of nucleons moving in "Fermi motion", each with some energy and momentum satisfying the Fermi-Dirac distribution.

In the Relativistic Fermi Gas (RFG) model, proposed by Smith and Moniz[184], quasi-elastic scattering from a nucleon in a nucleus is treated as if the incoming lepton scatters from an independent nucleon (the "impulse approximation") as in the Llewellyn-Smith formalism; however, in the case of the RFG, the target nucleon is not stationary, but has a momentum consistent with the Fermi distribution. Thus the cross section for scattering off the nucleus is replaced by a coherent sum of cross sections for scattering off of each

individual nucleon, with the remaining nucleus (depleted by 1 nucleon) as a spectator. In this case, we can conclude that, with a four-momentum transfer q , energy transfer ν , nucleon mass M , and nucleon initial and final momenta p_i and p_f respectively:

$$\begin{aligned}
 &\text{Initial nucleon kinetic energy, } \text{KE}_i = \vec{p}_i^2 / 2M \\
 &\text{Final nucleon kinetic energy, } \text{KE}_f = \vec{p}_f^2 / 2M \\
 &\hspace{15em} = (\vec{q} + \vec{p}_i)^2 / 2M \\
 &\text{Energy transfer, } \nu = \text{KE}_f - \text{KE}_i \\
 (2.24) \hspace{15em} &= Q^2 / 2M + \vec{q} \cdot \vec{p} / M
 \end{aligned}$$

(Remember $Q^2 = -q^2$ is always positive.) Thus we would expect the distribution of ν at fixed Q^2 to be centered around $\bar{\nu} = Q^2 / 2M$, with a width corresponding to the average momentum in the direction of energy transfer, which is a function of the Fermi momentum. Fitting these distributions yields a measurement of the Fermi momentum, which for carbon-12 has been measured to be 221 ± 5 MeV [192]. This is the value used by our Monte Carlo event generator, GENIE.[93]

A further consideration of the Relativistic Fermi Gas model is the concept of Pauli blocking. This is a consequence of nucleons being fermions and thus being required to obey the Pauli exclusion principle. Thus, a struck nucleon cannot be raised to an momentum state that is already occupied; that is, it must have a final-state momentum above k_F . This has the effect of, for a given energy transfer, setting a lower limit on the possible energy range of target nucleons for which an interaction is allowed. Thus, for a pure Fermi distribution where all states up to the Fermi limit, and none above it, are occupied, the

range of energies allowable to a target nucleon is:

$$(2.25) \quad \begin{aligned} E_{max} &= \sqrt{k_F^2 + m_N^2} \\ E_{min} &= \sqrt{k_F^2 + m_{N'}^2} - E_B - \nu \end{aligned}$$

Where for a quasi-elastic interaction on a proton, m_N is the proton mass, $m_{N'}$ is the neutron mass and E_B the proton binding energy (30 MeV in carbon); for interaction on a neutron, m_N is the neutron mass, $m_{N'}$ is the proton mass and E_B the neutron binding energy (34 MeV in carbon). As before, k_F is the Fermi momentum and ν the energy transfer. In a real nucleus, in which there is not a strict Fermi momentum cutoff, the Pauli blocking mechanism is more complex; this has been studied in papers such as [144]. In our Monte Carlo generator, GENIE, Pauli blocking is implemented via a modification to the Fermi momentum [93].

2.3.1.1. Limitations of the RFG model. Figure 2.8 shows measurements of CCQE ν_μ and $\bar{\nu}_\mu$ scattering cross sections on carbon. (MiniBooNE subtracted the $\bar{\nu}_\mu$ -hydrogen component of their cross section). The plot includes results from the NOMAD experiment at CERN [146], which operated in the 3-100 GeV range, as well as lower energy results from MiniBooNE [30] at around 1 GeV. In each case, the results were fitted to the Relativistic Fermi Gas model, extracting best fit parameters of the axial mass M_A and Pauli blocking strength κ . In the case of NOMAD, the data were a good fit for the world average M_A measured on bubble chambers, giving a value of $1.05 \pm 0.02(stat) \pm 0.06(sys)$ GeV/c². MiniBooNE, however, extracted a value of $M_A = 1.35 \pm 0.17$ GeV/c² for scattering from

Experiment	Material	Energy (GeV)	Q^2 cut (GeV ²)	M_A (GeV)
K2K [127]	Oxygen	0.3-5	$Q^2 > 0.2$	1.20 ± 0.12
K2K [92]	Carbon	0.3-5	$Q^2 > 0.2$	1.14 ± 0.11
MINOS [88]	Iron	≈ 3	None	1.19 ± 0.17
MINOS [88]	Iron	≈ 3	$Q^2 > 0.2$	1.26 ± 0.17
MiniBooNE [29]	Carbon	≈ 1	None	1.35 ± 0.17
MiniBooNE [29]	Carbon	≈ 1	$Q^2 > 0.25$	1.27 ± 0.14
NOMAD [146]	Carbon	$\approx 3 - 100$	None	$1.05 \pm_{0.06(sys)}^{0.02(stat)}$
T2K [4]	Carbon	≈ 1	None	$1.26_{-0.18}^{+0.21}$
T2K [4]	Carbon	≈ 1	None (shape)	$1.43_{-0.22}^{+0.28}$

Table 2.1. Values of M_A extracted from neutrino-nucleus scattering data

2.3.1.2. Local Fermi Gas model. The local Fermi gas model is an extension to the relativistic Fermi gas, in which a local density approximation [160] [153] is used, so that instead of using a constant average field for the whole nucleus, the momentum distribution is instead dependent on a nucleon's position within the nucleus. For a relativistic Fermi gas:

$$(2.26) \quad p_F^{\text{RFG}} = \frac{\hbar}{r_0} \left(\frac{9\pi N}{4A} \right)^{1/3}$$

where A is the atomic mass, r_0 is a constant 1.25 ± 0.20 fm, where the nuclear radius $R = r_0 A^{1/3}$ in the Relativistic Fermi Gas model. N is the number of neutrons. (This gives the Fermi momentum for neutrons; for protons, substitute Z , the number of protons, for N). For the local Fermi gas, this is modified [119] to:

$$(2.27) \quad p_F^{\text{LFG}}(r) = \hbar \left(3\pi^2 \rho(r) \frac{N}{A} \right)^{1/3}$$

By selecting a form for $\rho(r)$ that corresponds to a uniform sphere of charge, the LFG simplifies to the RFG. However, more accurate distributions are known from electron scattering. [83]

Figure 2.9 shows the difference in Fermi momentum between the LFG and RFG models. This model has been used as the basis for many more complicated treatments of nuclear effects [161][41].

2.3.1.3. Spectral functions. The Hamiltonian for a large nucleus is so complicated that it is impractical to try to solve the many-body Schrödinger equation for the entire nucleus. However, if a mean field is used to replace the sum of the individual

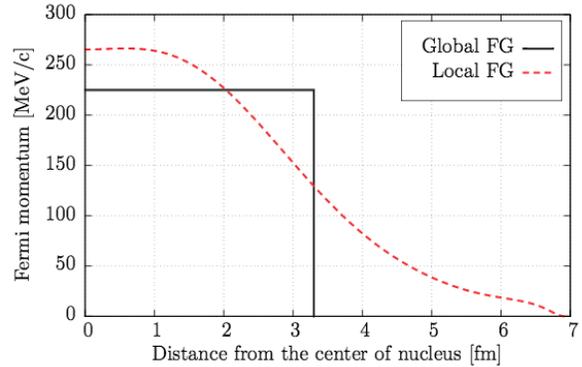


Figure 2.9. Comparison of RFG and LFG momentum distributions, from [119]

interactions, a shell model can be created. In this case, a spectral function $S(\vec{k}; w)$ can be defined such that the density of momentum states $\rho(k)$ is given by $\rho(k) = \int_{-\infty}^{E_F} d\omega S(\vec{k}; w)$, as explained in [147]. The spectral function $S(k; w)$ represents the probability of finding a nucleon with momentum \vec{k} and removal energy w within the nucleus. Spectral functions can be used to improve the Relativistic Fermi Gas model [70]. A realistic spectral function, however, must also include correlation effects. For small nuclei (atomic masses up to 4), these can be calculated using many-body theory; beyond this, the calculation becomes extremely complex. In order to generate a spectral function for heavier nuclei that includes correlations, Benhar et al. [55] use the mean-field approach to model the

single-particle contribution, while the correlated part is calculated for infinite nuclear matter. A local-density approximation is used to translate this to an appropriate value for a finite nucleus.

The effect of correlations on the spectral function prediction will be explained in the following section.

2.4. Long range correlations: the random-phase approximation

The random-phase approximation (RPA) is a method of incorporating long-range correlations between nucleons.[158] It is based around the phenomenon, observed in β -decay experiments, that the electroweak coupling can be modified by the presence of strongly-interacting nucleons in the nucleus, when compared to its free-nucleon coupling strength - similarly to the screening of an electric charge in a dielectric. The effect of nuclear polarization on the W and Z can be modeled by a chain of diagrams as shown in figure 2.10, where a solid line represents a nucleon and a double line an isobar resonance. It should be noted that long-range correlations affect the 1p1h (1-particle-1-hole) response - in other words, this is not a multi-nucleon knockout effect.

The RPA approach affects cross section predictions at low energy transfers (and low Q^2), where a quenching of the axial current reduces the cross section compared to the RFG prediction.

2.5. Multi-nucleon effects

The Relativistic Fermi Gas does not address multi-nucleon effects - effects due to interactions between pairs or larger groups of nucleons within the nucleus. In this section we will explain how, if nucleons are correlated, they no longer obey the assumptions of

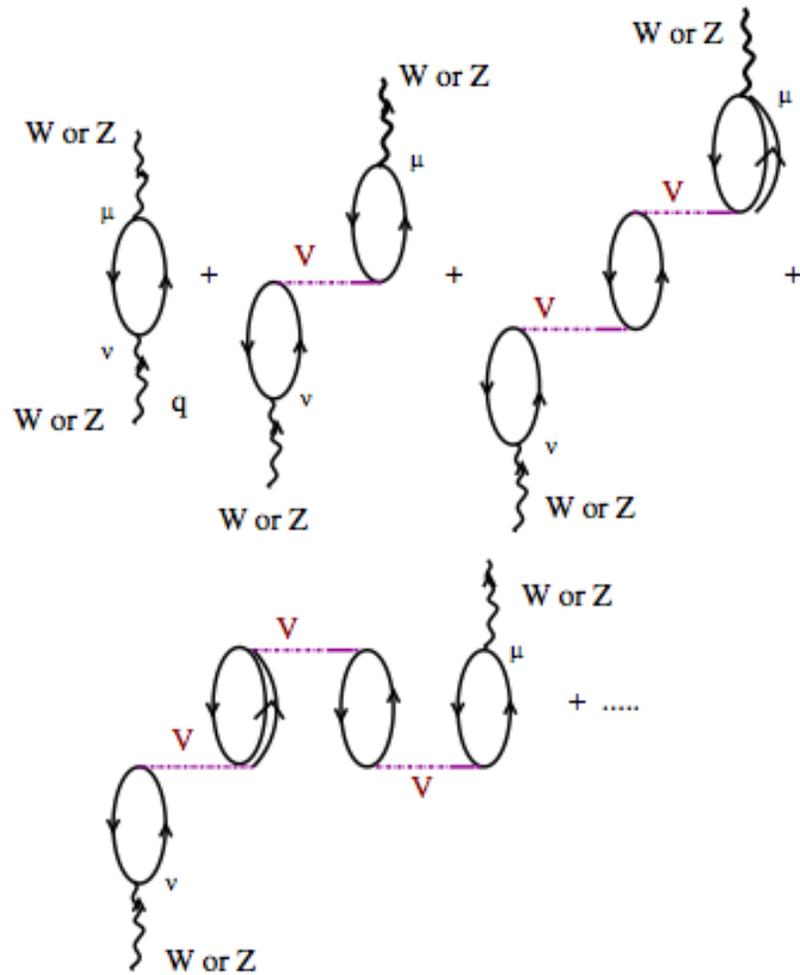


Figure 2.10. Set of irreducible diagrams responsible for the polarization (RPA) effects in the $1p1h$ contribution to the W or Z self-energies, taken from [158]

the quasi-elastic hypothesis and Fermi Gas distribution, meaning that we will reconstruct their energies incorrectly. For three example energies, figure 2.11, reprinted from [151], demonstrates the predicted extent to which a true energy will be smeared when reconstructed using the CCQE hypothesis on a system including multinucleon effects, described using Martini et al.'s model.

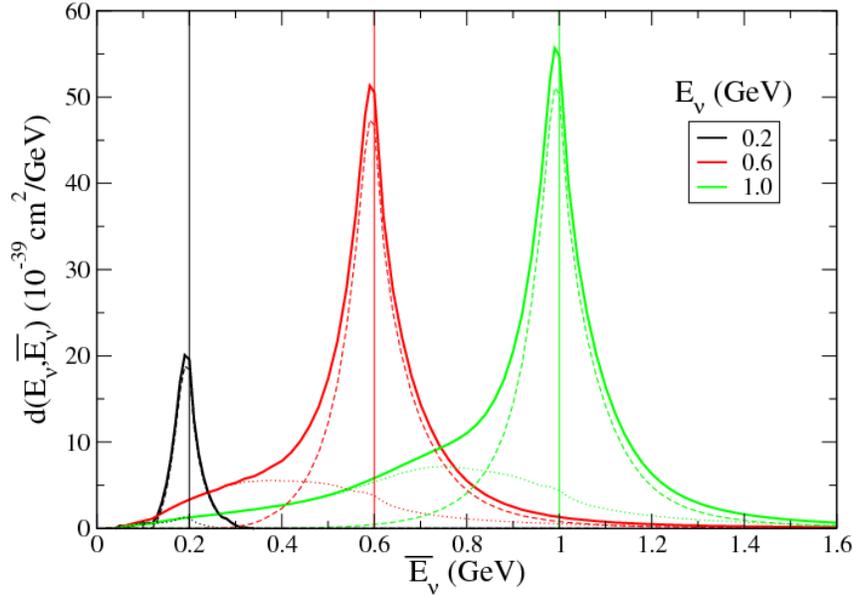


Figure 2.11. The spreading function per neutron of electron- ^{12}C scattering, showing the range of energies \bar{E}_ν that would be reconstructed using quasi-elastic hypothesis for three corresponding E_ν^{true} values. The genuine quasi-elastic (dashed lines) and the multi-nucleon (dotted lines) contributions are also shown separately. Reproduced from [151]

2.5.1. Short-range correlations

The Relativistic Fermi Gas model, which treats the individual nucleons as independent, does not adequately take into account the nature of the nuclear force, which has a short range with a repulsive core [46]. This short-range interaction between two (or more) very close-together nucleons can give the individual nucleons very high momenta, far above the Fermi momentum k_F , leading to a high-momentum tail when compared to the mean-field prediction of the Relativistic Fermi Gas model.

Figure 2.13, reprinted from [73], illustrates such a high-momentum tail, modeled using a spectral function, in contrast to the mean-field prediction of the RFG model; the model including the correlations has better agreement with data. Note that, while individual

particles within a correlated pair may have very high momentum, the pair as a whole has a low total momentum relative to k_F ; this corresponds to a pair of nucleons moving back-to-back. This effect is poorly understood, but can be examined with high-energy probes, which can be used to isolate scattering from these high-momentum nucleons.

The nucleon-nucleon force has a tensor component, which is responsible for the isosinglet np pair with an isospin of 0, as well as a central isotriplet ($I = 1$) component, which can contribute pp , nn and np pairs. When the relative momentum of the nucleons in the pair is between 300 and 600 MeV/c, and the total momentum is

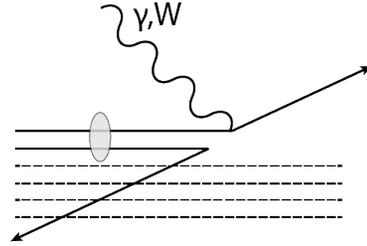


Figure 2.12. Scattering from a correlated pair of nucleons

small, the tensor component appears to dominate, meaning that we would expect pairs to mainly consist of a proton and a neutron [182]. As this is the form of the deuteron, we expect the high-momentum tail for different nuclear materials to have the same form as that of deuterium, in this momentum range (at higher momenta, where the short-range repulsive core dominates, we would expect this to break down). Electron-scattering experiments have demonstrated this constant ratio of nucleus-deuteron ratio, as shown in figure 2.14, reprinted from [73]. This shows the ratio of several nuclei's momentum distributions to that of the deuteron; all of these become almost flat at momenta $k > 2fm^{-1}$.

An alternative way to think of this high-momentum tail is to consider the Bjorken variable $x_{Bj} = Q^2/2m_N\nu$, which corresponds to the fraction of a nucleon's momentum carried by a struck quark. (Q^2 is the squared four-momentum transfer, m_p the mass of the nucleon, and ν the energy transfer.) In the case of elastic or quasi-elastic scattering off

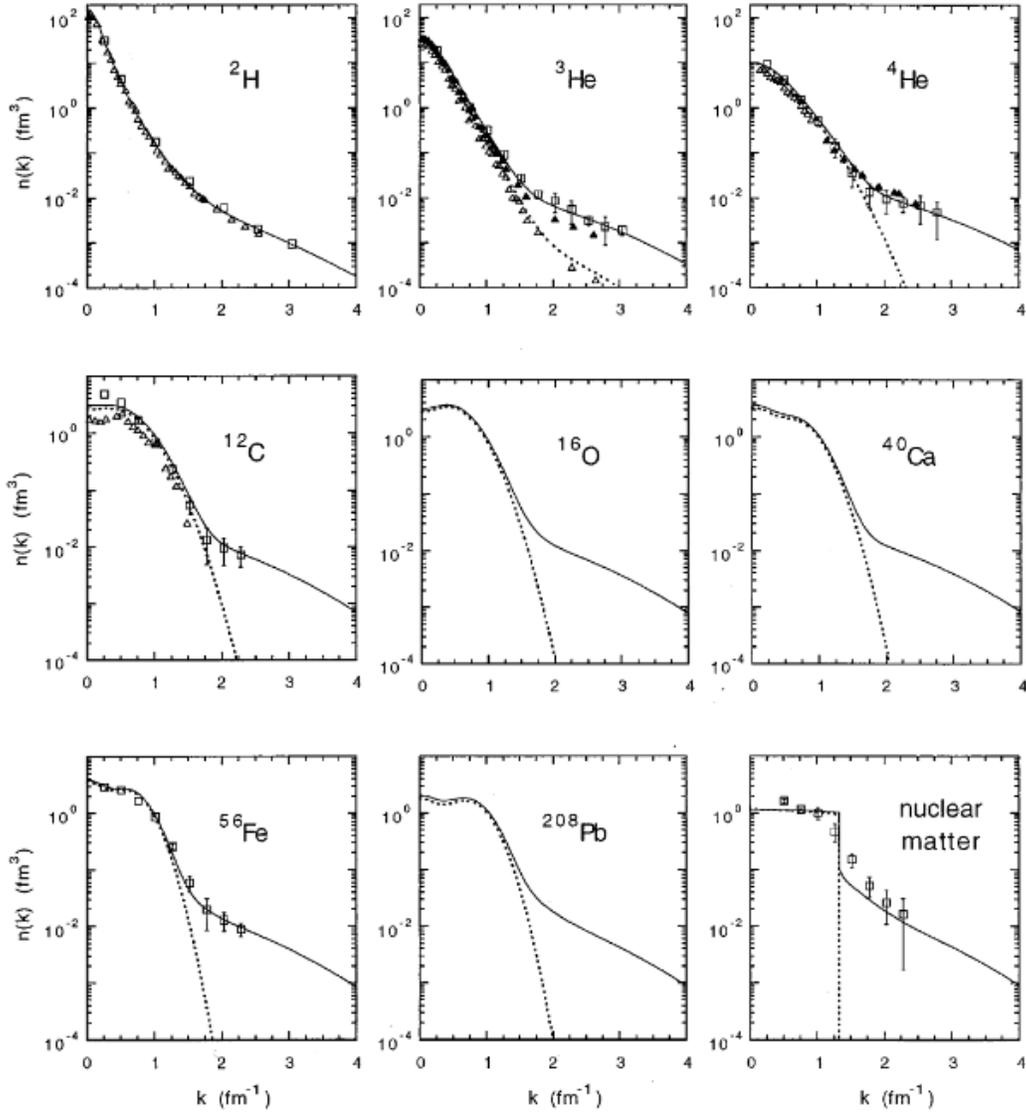


Figure 2.13. Nucleon momentum distribution predictions for various nuclei, taken from [73]. Dotted lines show the RFG prediction, while the solid line shows the spectral-function prediction including correlated pairs. The points correspond to data from various experiments, detailed in the reference.

of a nucleon at rest, where we scatter of the entire nucleon, $x_{Bj} = 1$ - the nucleon carries all of its own momentum. But if we scatter off of a pair, the struck nucleon can actually have a higher momentum than the pair's total, effectively corresponding to $x_{Bj} > 1$. A

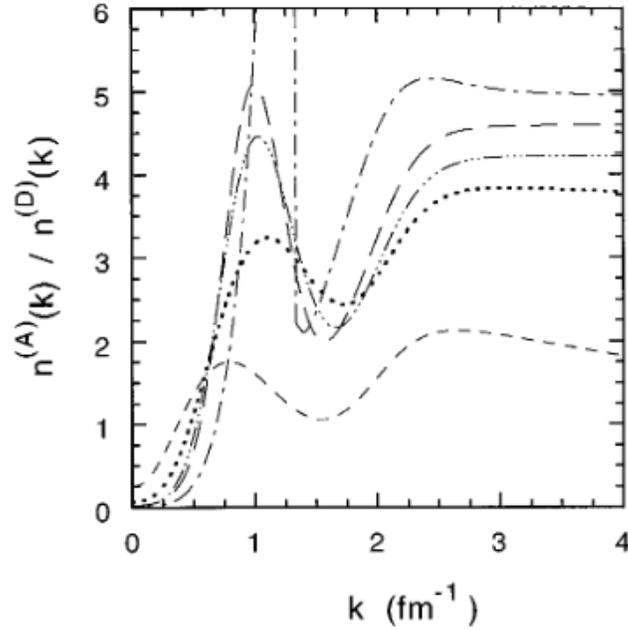


Figure 2.14. Ratio of the spectral function distributions from figure 2.13 to the corresponding distribution for the deuteron, also reproduced from [73]. The dashed line represents ${}^3\text{He}$, dotted ${}^4\text{He}$, dot-dashed ${}^{16}\text{O}$, long dashed ${}^{56}\text{Fe}$, and dot-long dashed nuclear matter.

comparison of inclusive electron-scattering cross sections vs x_{Bj} for different materials[97] shows that the ratio of this cross section to that for deuterium reaches a flat plateau for all materials when $x_{Bj} > 1.5$ (note that $x_{Bj} = 2$ is the maximum for scattering off a deuteron, which has an atomic weight of 2). At lower x_{Bj} , there is still significant contribution from the mean-field part of the distribution, whereas above $x_{Bj} = 1.5$, this is overwhelmed by the correlation contribution.

Figure 2.15 shows this distribution. Plateaus indicate 2-nucleon and 3-nucleon correlation zones, above $x_{Bj} = 1.5$ and 2.25 respectively, in electron-nucleus scattering. By integrating the momentum distribution in these area, the CLAS collaboration [95] were

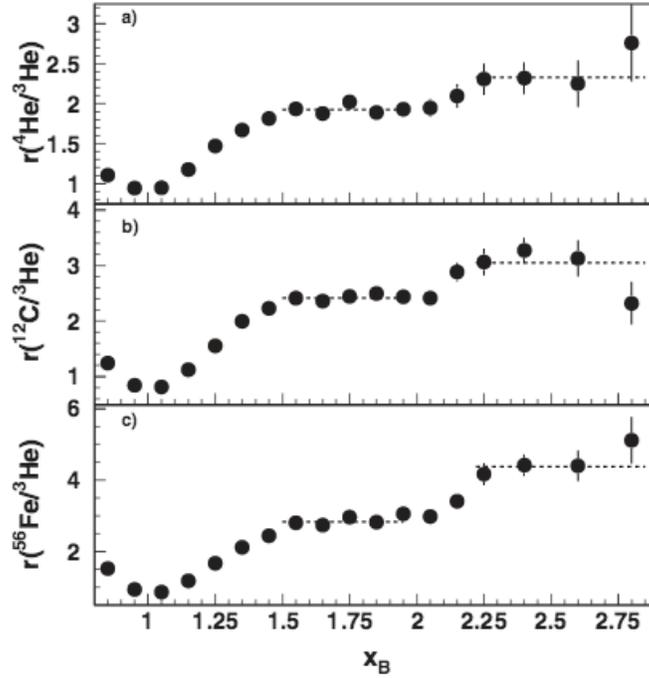


Figure 2.15. Ratio of the cross sections for ^4He , ^{12}C , and ^{56}Fe vs Bjorken variable x_{Bj} to that for ^3He , from CLAS[95].

able to determine the fraction of nucleons that were in short-range correlated pairs. For carbon, this probability was calculated to be approximately 20%.

High-momentum knock-out experiments, in which high-energy electrons are scattered on a nucleus, show that, when an interaction has missing energy - that is when the proton ejected from the nucleus has less energy than would be expected from conserving energy and momentum - that energy is typically found in a second ejected nucleon[99]. This is consistent with knocking out the partner nucleon, when scattering from a nucleon that is in a correlated pair. Subedi et al. at JLab's Hall A [100] used custom neutron and proton detectors to analyze the makeup of these pairs for electromagnetic scattering of electrons from protons in carbon. They concluded that the ratio of np to pp pairs ejected

was 18 ± 5 . This supports the theory that the dominant contribution to these short-range correlations is from the tensor component, which includes only the isosinglet np state. Including the findings of [95], and assuming isospin symmetry between nn and pp pairs, they conclude that the carbon nucleus consists of 80% uncorrelated nucleons, 18% np pairs and 1% each of nn and pp pairs.

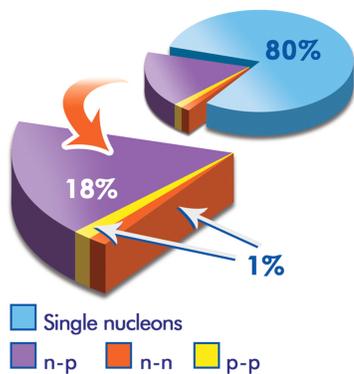


Figure 2.16. The average fraction of nucleons in the various initial-state configurations of ^{12}C , taken from [100]

Even in heavier nuclei, which tend to have more neutrons than protons, np pairs dominate, meaning that a proton has a greater chance of being in a correlated pair than does a neutron. This was demonstrated with the CLAS detector at JLab, which looked at two-nucleon knock-outs in carbon, aluminum, iron and lead. [182]

Note that the consequence of this for neutrino-nucleus charged-current quasi-elastic scattering would be that we would expect a neutrino scattering from a correlated np pair to produce two protons (one created by the charged current interaction with a neutron, plus its proton partner). Correspondingly, charged-current quasi-elastic antineutrino scattering with a correlated pair would tend to produce two neutrons.

Bodek and Ritchie's modification to the Relativistic Fermi Gas model [59] adds a high-momentum tail to the RFG's momentum distribution, based on the nucleon-nucleon correlation function as explained in [124]. This modified version of the RFG is the nuclear

model adopted by the GENIE Monte Carlo generator v.2.8.4, which is used in this analysis' cross-section calculation for background subtraction, unsmearing, and acceptance correction of our data measurements. While this method attempts to provide a realistic initial momentum distribution, it does not include any model for the ejection of paired nucleons, which may affect our final state.

It should be noted that there is also a small contribution from correlated groups of three nucleons, which take the form of a ${}^3\text{He}$ nucleus. These 3N correlations can have $x_{Bj} > 2$, and in carbon, contribute to only approximately 0.5% of interactions[95].

2.5.2. Transverse enhancement model

The cross section for electron scattering on a nucleus can be separated into transverse and longitudinal parts (that is, into parts where the exchanged photo is polarized transversely and longitudinally), giving [52]

$$(2.28) \quad \frac{d\sigma}{d\Omega d\nu} = \sigma_{\text{Mott}} \frac{E'}{E} \left[\frac{Q^4}{\vec{q}^4} R_L(Q^2, \nu) + \left(\frac{Q^2}{2\vec{q}^2} + \tan^2 \frac{\theta}{2} \right) R_T(Q^2, \nu) \right]$$

The longitudinal response function R_L is proportional to the nucleon's electric form factor, G_E , while the transverse response function R_T is proportional to the magnetic form factor G_M . We can define a scaling function f such that

$$(2.29) \quad \begin{aligned} f_L(Q^2, \nu) &\propto \frac{R_L(Q^2, \nu)}{G_E^2} \\ f_T(Q^2, \nu) &\propto \frac{R_T(Q^2, \nu)}{G_M^2} \end{aligned}$$

For deep inelastic scattering, where the incoming neutrino or electron scatters off a quark within the nucleon, the structure functions have been found to be functions only of the Bjorken scaling variable $x_{Bj} = Q^2/2m_N\nu$ and Q^2 , a process known as x-scaling. This allows one to relate the cross section distribution to the momentum distribution of the quarks in the nucleon. In analogy to this, one can look for a variable y which performs a similar function for scattering from nucleons within a nucleus. As the form factors are functions of q and ν , this y must also be a function of these variables. Day et al. [82] derive this as:

$$(2.30) \quad y(q, \nu) = \left[(m_A + \nu) \sqrt{\Lambda^2 - m_{A-1}^2 W^2 - q\Lambda} \right] / W^2$$

Where $W \equiv \sqrt{(m_A + \nu)^2 - q^2}$ is the center-of-mass energy, and $\Lambda \equiv (m_{A-1}^2 - m^2 + W^2)/2$. In doing so, they assume the plane-wave input approximation, that is, that a nucleon of momentum \vec{k} in a nucleus of mass m_A (not to be confused with axial mass M_A) absorbs a photon with momentum \vec{q} , causing the nucleon to be ejected with a momentum $(\vec{k} + \vec{q})$ and the residual nucleus, now with mass m_{A-1} , to recoil with momentum $-\vec{k}$.

The validity of the y scaling method was tested on electron scattering data from Saclay by Finn et al.[106], and by Carlson et al.[68], who plotted $f_L(y)$ and $f_T(y)$ for various values of q . This data (see figure 2.17, where the scaling variable is referred to as ψ') showed that, while f_T and f_L each independently obey this scaling relationship (f_T diverges with q at high y , but this effect is due to the onset of resonant interactions, and is not relevant to the quasi-elastic behavior), the f_T shows an enhancement with respect to f_L . Because meson-exchange currents are known to affect the transverse response, it is assumed that the longitudinal response is due to a single-nucleon scattering effect,

and that the transverse enhancement is due to multi-nucleon effects, with the difference between the two scaled curves providing some measure of the strength of these effects.

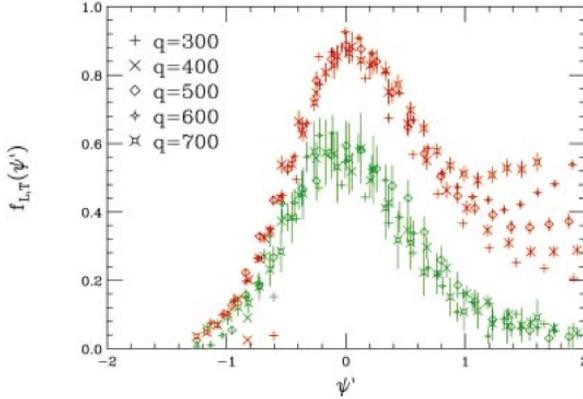


Figure 2.17. The scaling functions f_T (top, red) and f_L (bottom, green), for various values of q , reproduced from [68]

Bodek, Budd and Christie [61] have isolated the transverse enhancement contribution. At low Q^2 they used the results of Carlson et al. [68], who calculated ratios of transverse to longitudinal response; at high Q^2 , they used a similar measurement from the JUPITER experiment [149].

Inclusive cross section data was fitted in four components: transverse and longitudinal quasi-elastic contributions, inelastic contributions, and the transverse enhancement to the QE. From this fit, they extracted a transverse enhancement ratio

$$(2.31) \quad \mathcal{R}_T = \frac{\text{QE}_{\text{transverse}} + \text{TE}}{\text{QE}_{\text{transverse}}}$$

This ratio $\mathcal{R}_T(Q^2)$ was parameterized in the form $\mathcal{R}_T = 1 + A Q^2 e^{-Q^2/B}$ where $A=6.0$ and $B=0.34 \text{ (GeV/c)}^2$. This best-fit parametrization is shown, along with the data, in figure 2.18.

A prediction for the effect on neutrino-nucleus scattering was made by applying a multiplier corresponding to \mathcal{R}_T to the magnetic form factors of the nucleons. By applying this correction to the Relativistic Fermi Gas model, they were able to produce cross

section predictions that were consistent with the low-energy neutrino-scattering results from MiniBooNE[77], and with the high-energy results from NOMAD[146]. Comparisons of the $d\sigma/dQ^2$ for quasi-elastic neutrino[107] and antineutrino[96] scattering in MINERvA also showed better agreement with this model than with the Relativistic Fermi Gas.

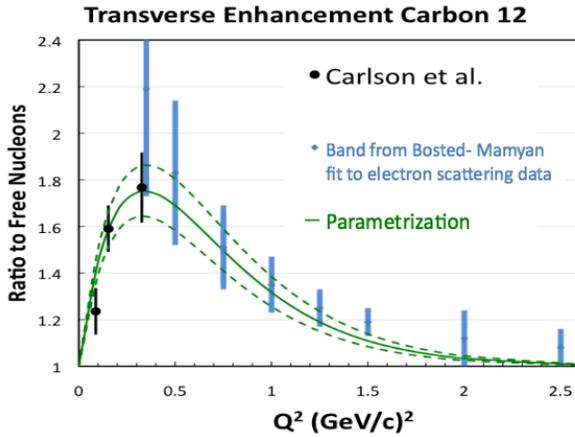


Figure 2.18. Best-fit form of $\mathcal{R}_T(Q^2)$, fitted to data from [68] and [149]. Plot from [61].

neutrino and electron scattering. Having done so, the strength of this in electron-scattering is extracted, compared to theory, and then that same correction is applied to the neutrino channel. This approach was adapted to take into account meson-exchange currents, as these occur in the transverse channel.

Note that the transverse enhancement and SuSA models are empirical models that ascribe an effect observed in electron scattering to multi-nucleon effects, and then attempt to predict what effect it might have on neutrinos. Conversely, meson-exchange current models attempt to predict multi-nucleon effects in neutrino scattering from first principles.

A somewhat similar approach was taken in a phenomenological model based on super-scaling (SuSA) of longitudinal electron-nucleus scattering data [40]. This is based around the idea of finding a scaling variable that is independent of both momentum transfer and nucleus target, for both

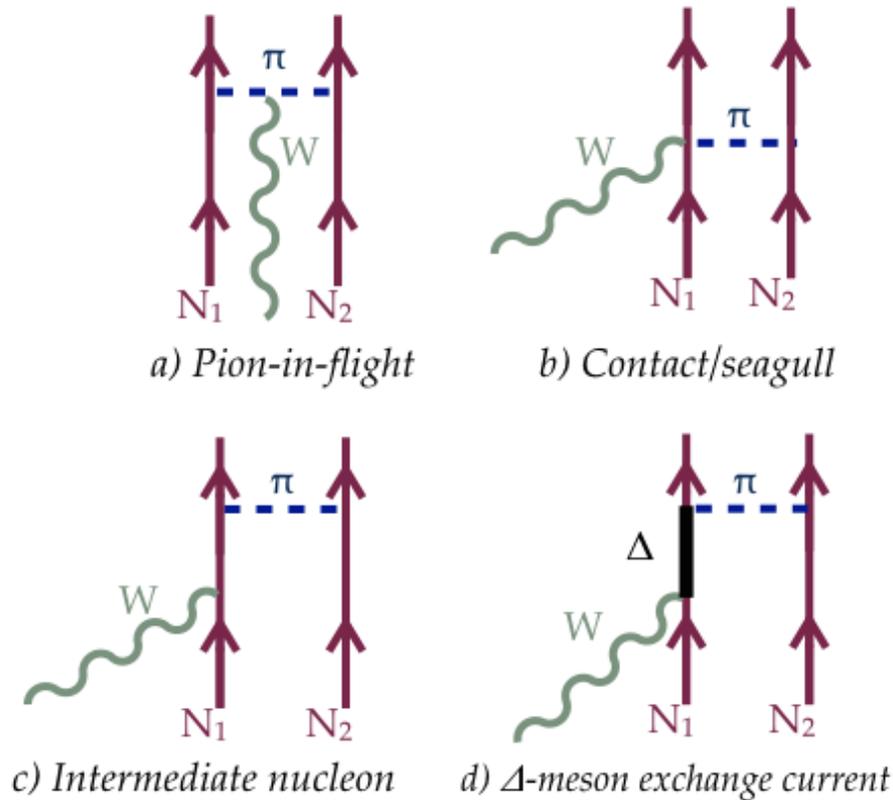


Figure 2.19. The four types of diagram for MEC with one-pion exchange, based on a figure in [164]

2.5.3. Meson-exchange currents

An alternative way of thinking about the binding between pairs of nucleons is to consider them as being bound by the exchange of virtual mesons [164], a process known as meson-exchange currents (MEC). In this case, it would be possible for the exchanged vector boson to couple to the exchanged meson, or for it to couple to a nucleon only if a virtual meson is exchanged. The longest-range diagrams will correspond to those in which a single pion (the lightest meson) is exchanged. They can be divided into four categories, as shown

in figure 2.19: a) pion-in-flight diagrams, where the boson couples to the exchanged pion b) the contact or ‘seagull’ diagram, including the pseudovector nucleon-pion coupling, c) diagrams including an intermediate virtual nucleon which exchanges the pion, and d) a similar diagram where the intermediate is an isobar resonance, $\Delta(1232)$.

By evaluating some set of these diagrams, several cross section models have been produced. The IFIC model of Nieves et al. [162] and the Lyon model (Marteau/Martini) [152] are based on a local Fermi gas and both include RPA polarization effects. However, the Lyon model does not model pion-in-flight, pion pole or contact diagrams. Additionally, the IFIC model is only valid for momentum transfers below 1.2 GeV/c, while the Lyon model is not fully relativistic. An alternative calculation using the Giessen Boltzmann-Uehling-Uhlenbeck framework (GiBUU) [143] includes the local Fermi gas and a spectral-function nuclear model, but does not include RPA.

2.6. Other neutrino-nucleon interaction processes, and final-state interactions

Quasi-elastic scattering is not the only possible form of charged-current interaction between a neutrino and a nucleon. Figure 2.20, reproduced from [110], shows how other processes come into play as neutrino energy E_ν increases. The lines on the plots (which represent charged-current neutrino and antineutrino scattering cross sections respectively) represent the predictions of the NUANCE neutrino interaction generator [69] for the quasi-elastic (QE), resonant (RES) and deep inelastic scattering (DIS) processes, as well as the total charged-current inclusive cross section.

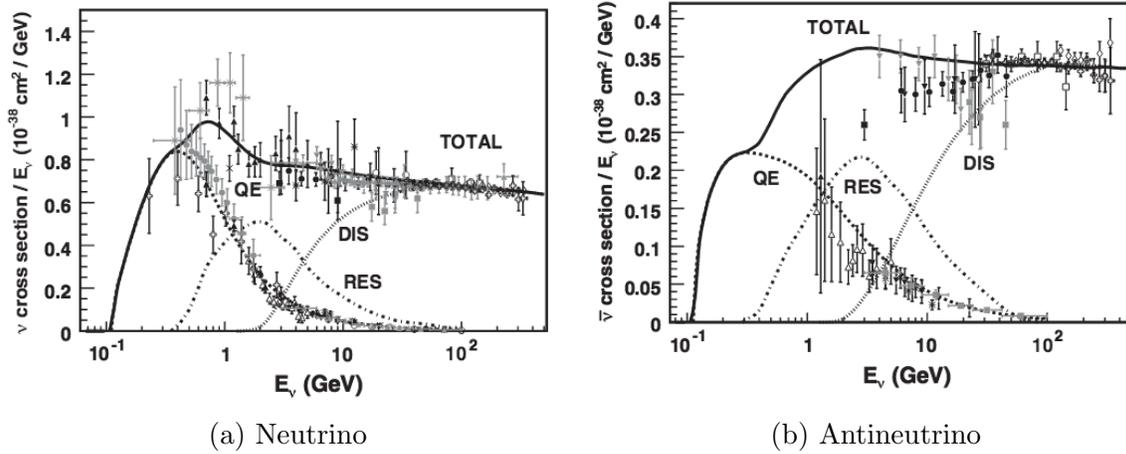


Figure 2.20. Processes contributing to the total charged-current neutrino-nucleon scattering cross section, from [110]. ‘QE’ refers to quasi-elastic scattering, ‘RES’ to resonant pion production, and ‘DIS’ to deep inelastic scattering.

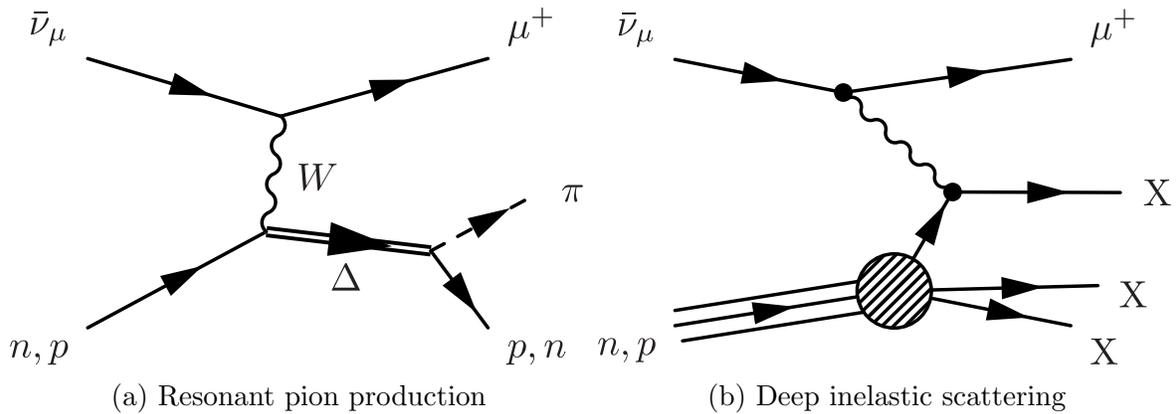


Figure 2.21. Non-quasi-elastic charged-current neutrino scattering processes

Figure 2.21 shows the Feynman diagrams for the resonant and DIS interactions. In resonant single-pion production (figure 2.21a), a neutrino scatters on a nucleon with the exchange of a W boson, as in quasi-elastic scattering. However, if the exchanged W carries sufficient energy, the nucleon can be raised to an excited state, typically one of the Δ_{1232} resonances, spin $\frac{3}{2}$ states comprising u and d quarks. This unstable resonance will decay to a pion and nucleon. For antineutrinos, there are three possible charged-current

resonant pion production processes:

$$\bar{\nu}_\mu p \rightarrow \mu^+ p \pi^-$$

$$\bar{\nu}_\mu p \rightarrow \mu^+ n \pi^0$$

$$\bar{\nu}_\mu n \rightarrow \mu^+ n \pi^+$$

(Compare with the quasi-elastic process $\bar{\nu}_\mu p \rightarrow \mu^+ n$). In charged-current deep inelastic scattering, the neutrino exchanges a W boson with one of the constituent quarks within the nucleon, producing a complex hadronic final state in addition to the μ^+ . (Naked quarks cannot exist outside of a hadron, therefore they rapidly recombine to create this hadronic shower).

These processes are of concern to us because, when they occur on nucleons within a nucleus, their final states can mimic those of quasi-elastic interactions. This is because of the phenomenon known as “final-state interactions” or FSI. Hadrons produced by interactions within the nucleus must traverse the rest of the nucleus in order to reach the final state. In some cases, the hadronic products of the initial interaction will rescatter or be absorbed, altering the kinematics and multiplicity of the hadronic final state. Of particular concern when measuring a quasi-elastic cross section is the case in which a pion is produced, but is then absorbed, leaving a quasi-elastic-like final state of a single muon and neutron. (Note: it is actually not possible for a pion to simply be absorbed by the nucleus, leaving no other products than the original nucleus; conservation of momentum and energy require that at least two low-energy nucleons must be produced. However, these nucleons are very difficult to detect, and may not exit the nucleus themselves.)

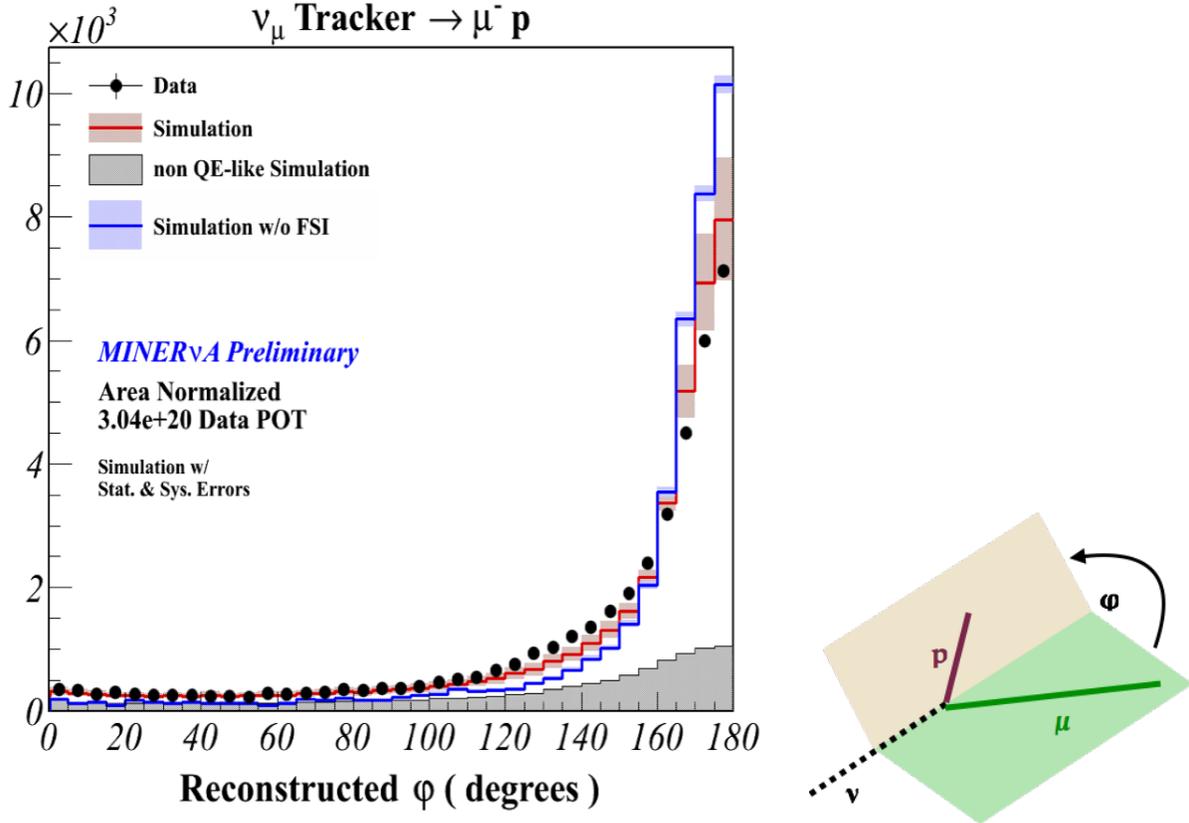


Figure 2.22. Comparison of MINERvA's neutrino-scintillator scattering data with simulation with and without FSI effects (from [191])

A recent measurement of MINERvA's quasi-elastic-like neutrino-scintillator scattering cross section demonstrates the importance of modeling FSI effects when measuring CCQE cross sections [191]. Figure 2.22 shows the distribution angle between the neutrino-muon and neutrino-proton plane for fully-reconstructed quasi-elastic-like events ($\nu_\mu n \rightarrow \mu^- p$) on MINERvA's scintillator tracker. This distribution is then compared to the predictions of the GENIE event generator [93] with and without final-state interactions enabled. Each of the simulated distributions is normalized to the area beneath the data distribution; this serves to remove contamination from uncertainties in the measurement of the neutrino

flux, allowing the shapes of the distributions to be compared. It can be clearly seen that the data far more closely the simulation including FSI effects; however the agreement is not exact, indicating that FSI modeling can be improved.

CHAPTER 3

The MINERvA Experiment

3.1. Introduction to MINERvA

MINERvA (Main INjector ExpeRiment ν -A) is a dedicated neutrino-nucleus scattering cross-section experiment. It is situated in the NuMI neutrino beam at Fermi National Accelerator Laboratory (FNAL, or Fermilab), in Batavia, Illinois. MINERvA detects neutrino interactions using a tracker consisting of strips of doped polystyrene scintillator, arranged in three orientations to enable three-dimensional track reconstruction. Upstream of the central tracker region, planes of scintillator strips are interspersed with passive nuclear targets consisting of several different materials, allowing MINERvA to study how scattering cross section distributions depend on the composition of the target nucleus. The near detector for the MINOS experiment [94] (Main Injector Neutrino Oscillation Search), located two meters downstream of MINERvA, serves as a muon spectrometer. Its magnetized detector provides data on the charge and momentum of muons exiting the back of MINERvA.

3.2. The NuMI neutrino beam

Fermilab's NuMI (Neutrinos at the Main Injector) [98] beam serves the oscillation experiments MINOS [94] and NOvA [169], as well as MINERvA. NuMI delivers a high-intensity, broad-spectrum neutrino beam, consisting mainly of muon neutrinos or muon antineutrinos, depending on configuration. For the study described in this thesis, the

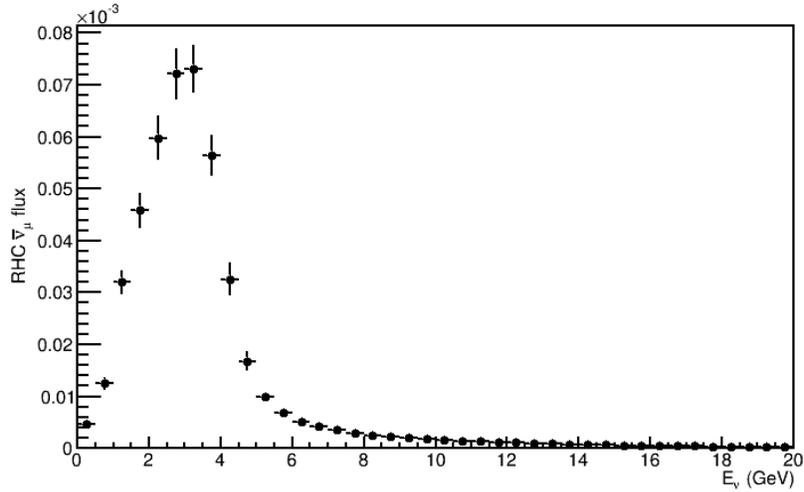


Figure 3.1. The energy spectrum (flux) of antineutrinos used in this analysis. This flux is given in units of the number of antineutrinos per m^2 per GeV per proton on target. This is the PPFX ‘Gen-2 thin’ flux, with neutrino-electron scattering constraint, including systematic uncertainties.

peak beam energy was around 3 GeV, and used the antineutrino-rich setup. This section explains the NuMI beam as it relates to this analysis, and is summarized from [98].

In order to generate the neutrino beam, 120 GeV/c protons from Fermilab’s Main Injector accelerator are fired into a narrow graphite target 1 meter long. Interactions between the protons and the graphite produce pions and kaons, which are focused by a pair of parabolic horns. The direction of the electrical current in the horns allows a beam of neutrinos or antineutrinos to be selected. The mesons produced are allowed to decay in a decay pipe, producing muons and neutrinos. An absorber removes any remaining hadrons from the beam; then over 200 meters of rock filter out muons, leaving a beam of neutrinos. A schematic of the beam components is shown in figure 3.2. The energy spectrum of the antineutrino beam used in this study is shown in figure 3.1.

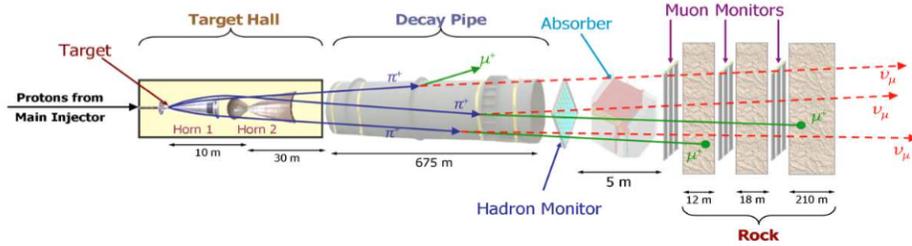


Figure 3.2. Schematic of the NuMI beam, reprinted from [98] (not to scale)

3.2.1. Proton beam

Fermilab's proton source produces H^- ions, which are accelerated to 400 MeV by the linear accelerator (Linac). [75] From there, the Booster synchrotron [132] converts them to protons, accelerates them to 8 GeV, and groups them into $1.6\mu s$ bunches with 53MHz spacing. Six Booster bunches are then injected into the Main Injector synchrotron, which accelerates them to 120 GeV. A *spill* of protons, corresponding to either five or six of these bunches (depending on what other experiments are running) is extracted to the NuMI beam approximately every 2.2 seconds, with a spill time of 8-10 μs . By the end of MINERvA's low energy run in 2012, 3.6×10^{13} protons on target (POT) were delivered in each spill. In the full low-energy run from March 2010 to April 2012, MINERvA recorded data corresponding to a total of 4.0×10^{20} POT in neutrino mode, and 1.7×10^{20} in antineutrino mode. [37]

The proton beam is aimed at the NuMI target, located in the Target Hall, 41m underground at Fermilab. The beam is oriented downwards at an angle of 3.343° (58.87 milliradians) to the horizontal; this is to align the beam to the MINOS[94] far detector, 734 km away and half a mile underground, in the Soudan mine in Minnesota. As the

MINERvA detector is horizontal, that means that there is a 3.343° discrepancy between the z axis of the detector and the beam direction.

3.2.2. Meson production

In the Target Hall, the beam is collimated using a device called a baffle. This a narrow tube, designed to degrade the beam if it is mis-steered, as its 400kW of power could destroy expensive components. Upon exiting the baffle, the beam hits the NuMI target; at this point the beam spot has a diameter of 1.1-1.2 mm. The protons interact, producing pions and kaons.

The NuMI target is 95.38 cm long, and consists of 47 20mm-long graphite fins, with a cross-sectional area of 15×6.4 mm, 0.3mm apart. This small area was chosen to prevent re-interaction of secondary mesons. These “tertiary interactions” are a source of systematic uncertainty on our final cross-section measurement, as they affect how well we know our final antineutrino flux. The position of the target along the beam line is configurable; when the data analyzed here was collected, it was inserted 50.4cm into the first focussing horn. This position was chosen to maximize the flux of neutrinos in the 1-3 GeV energy range, the energy that will maximize oscillation probability between the MINOS near and far detectors.

3.2.3. Focussing horns

The pions and kaons produced in the target are focused by a pair of aluminum horns (figure 3.3), each approximately 3m long. The horns are pulsed with a high half-sine wave with high current (185kA for the data taken here; other currents were used in

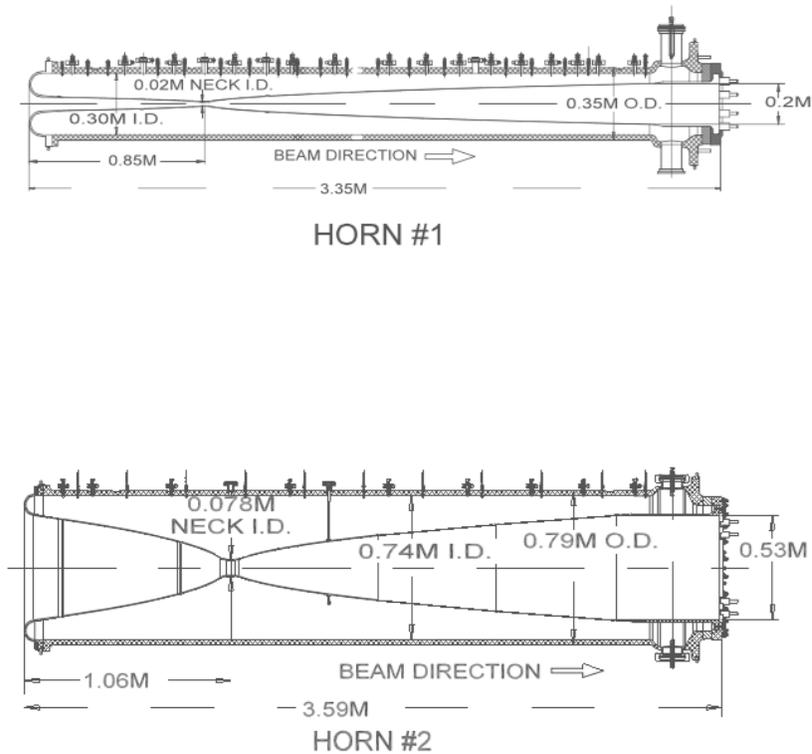


Figure 3.3. Schematic of the focussing horns, reprinted from [98] (not to scale)

“special runs” to generate different beam energy profiles), to create a toroidal magnetic field of up to 3T between the inner and outer conductors. The direction of the current selects whether positive or negative pions will be focused; in forward horn current mode (FHC, +185 kA), we focus π^+ , leading to a neutrino beam, whereas with reverse horn current (RHC, -185 kA), we select π^- and thus antineutrinos.

The horns act as lenses for the beam; the parabolic shape of the inner conductor has the effect of generating a magnetic field between the inner and outer conductors that focuses charged particles of one sign so that they are directed towards the beam direction. Particles inside the inner conductor experience very little magnetic field (just a small residual field due to irregularities in the horn shape); those between the conductors

experience a field proportional to $1/R$; however, due to the shape of the horn, the amount of time they spend in this field is proportional to R^2 , meaning that the correcting transverse momentum is proportional to R . Thus particles of the preferred sign are focussed towards the beam direction, with the most correction being received by the particles with the greatest deviation; meanwhile, particles of the opposite sign are defocussed and removed from the beam. Note that any wrong-sign particles that pass inside the inner conductor will experience no magnetic field and will thus remain in the beam, leading to contamination. As these particles, with their large forward momenta, tend to be the most energetic, wrong-sign contamination is greatest at high energies. It is a particular issue in the antineutrino beam, as the collision of protons on the NuMI target produces more π^+ than π^- .

The efficiency of beam focussing at different energies is a source of systematic uncertainty on our cross section measurement, as it affects the flux distribution of incoming neutrinos.

3.2.4. Meson decay

Particles exiting the horns pass next through the NuMI decay pipe. The composition of the beam at this point is mainly pions, along with some kaons and some residual protons that passed through the target without interacting. The decay pipe is 675 m long (approximately the length needed to accommodate the decay of a 10 GeV pion), and 2 m in diameter. Ideally, this interior of this pipe would be a vacuum to prevent re-interaction; due to the impracticality of maintaining this, it is instead filled with helium gas, which protects against corrosion.

Charged pions decay via the interaction $\pi^+ \rightarrow \mu^+ \nu_\mu$ (and $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$) with a branching ratio of 99.99%. [163] However, charged kaons have several significant decay modes that produce neutrinos (its other main decay modes are to pions):

$$K^+ \rightarrow \mu^+ \nu_\mu (63.5\%)$$

$$K^+ \rightarrow \pi^0 e^+ \nu_e (5.1\%)$$

$$K^+ \rightarrow \pi^0 \mu^+ \nu_\mu (3.4\%)$$

3.2.5. Absorption

3.2.5.1. Hadron monitor and absorber. Downstream of the decay pipe is the hadron absorber. This large structure of aluminum steel and concrete removes pions and kaons that have not decayed, as well as the remaining non-interacting protons, from the beam; neutrinos and muons pass through virtually unaffected. As well as cleaning the beam, the absorber has a safety function, reducing radiation levels in groundwater and in areas accessed by personnel.

A hadron monitor is located 80cm upstream of the absorber to monitor the hadronic content of the beam, around 80% of which is 120 GeV protons that did not interact. The monitor consists of a 7x7 section ionization chamber, covering an area of approximately 1 m×1 m. It is used to monitor the status of the NuMI target (these had to be replaced several times throughout the low energy run, due to radiation damage), and to check the horn position, primary beam angle, and the position of the beam spot.

3.2.5.2. Muon shield and monitors. Downstream of the absorber, muons are absorbed by 240 m of dolomite rock, to leave a beam of neutrinos. However, neutrino

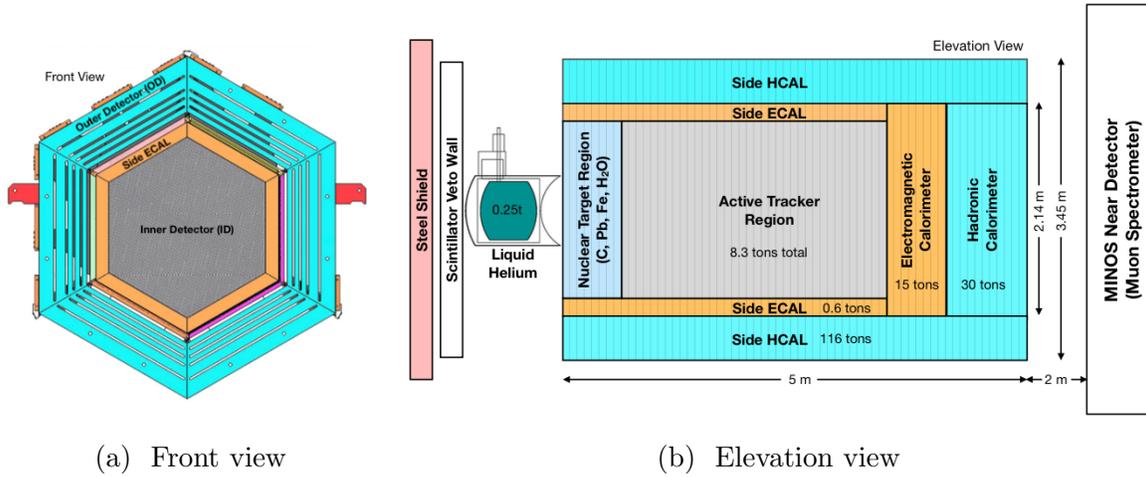


Figure 3.4. Schematic of the MINERvA detector, taken from [37] (not to scale). The beam enters from the left side of figure 3.4b.

interactions in the rock generate more muons, which sometimes have enough energy to reach the MINERvA detector. These “rock muons” can be identified with the help of MINERvA’s veto wall. As muons behave as minimum-ionizing particles in MINERvA, they are used for energy scale calibration.

Four chambers have been carved in the rock of the muon shield. For the period included in this analysis, three of these contained muon monitors, positioned at the distances where 4 GeV, 10 GeV and 20 GeV muons would range out. Each of these monitors, which are 2.3 m square, consists of a 9×9 array of ionization chambers. They are used to monitor the shape of the muon beam, which is expected to approximately mimic the shape of the neutrino beam. In 2014, a fourth muon monitor was approved, at a distance corresponding to 40 GeV muons [48].

3.3. The MINERvA detector

The MINERvA detector is described in detail in [37]; this section summarizes its main features, with particular focus on the components relevant to this analysis. As shown in figure 3.4, the main body of the detector consists of hexagonal “modules”. Each of these is composed of an inner detector (ID) and an outer detector (OD). Looking along the beam line, the ID consists first of active scintillator planes interspersed with passive nuclear targets, followed by a tracking region of pure scintillator, a downstream electromagnetic calorimeter (ECAL), then a hadronic calorimeter (HCAL). The outer detector is mainly composed of a heavy steel frame, interspersed with scintillator bars, which serves both for calorimetry and as a support structure for the detector. Outside of this are the electronics and light collection systems. Upstream of the main detector are a steel shield, veto wall, and liquid helium target. The MINOS near detector, which serves as a muon spectrometer, is located 2m downstream of the back of the MINERvA detector.

3.3.1. The MINERvA coordinate system

MINERvA uses a Cartesian coordinate system to refer to positions and directions within the detector. This is defined such that the z axis is horizontal, pointing downstream along the detector axis (from left to right in figure 3.4b). The y axis is vertical (from bottom to top in figure 3.4a) and the x axis horizontal pointing to beam left (right to left in figure 3.4a), with the origin of the $x - y$ plane in the center of the inner detector. The z axis is chosen so that $z = 1200\text{cm}$ at the front face of the MINOS near detector.

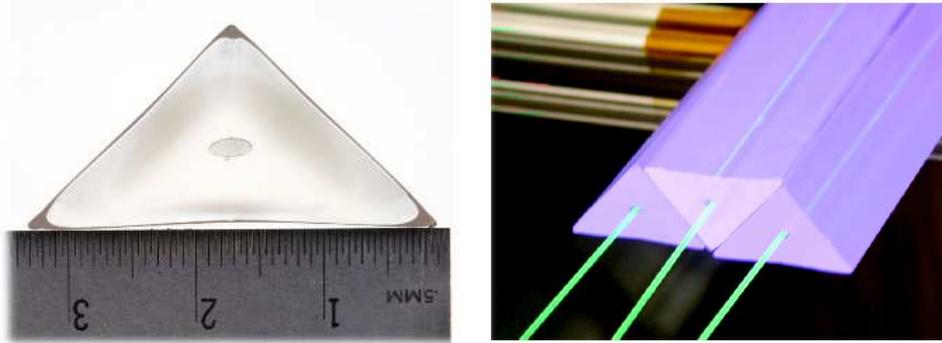


Figure 3.5. MINERvA scintillator strips and how they fit together, reprinted from [37]

As explained in section 3.2.1, the beam points slightly downwards towards the MINOS far detector in Soudan, MN. Thus the beam access is in the $y-z$ plane, with a downwards angle of 3.343° .

3.3.2. Tracking modules in the inner detector

The modules in the active tracking region (the region of the detector in which the interactions studied in this analysis take place) are composed entirely of scintillator planes. Planes of the same design are also interspersed with the passive nuclear targets in the upstream region, and with the calorimeter materials in the ECAL and HCAL.

A scintillator plane is made up of 127 “strips” of doped polystyrene scintillator, with a titanium dioxide coating. The strips have a triangular cross section 17 ± 0.5 mm and 33 ± 0.5 mm wide. The strips are arranged in an alternating orientation as shown in figure 3.5, in order to ensure that any charged particle passing through the plane will produce scintillation in at least two strips. The lengths of the individual strips vary from 122 to 245 cm, depending on their positions in the hexagonal detector. The strips are

Material	Percentage (%)
Hydrogen	7.42
Carbon	86.6
Oxygen	3.18
Aluminum	0.26
Silicon	0.27
Chlorine	0.55
Titanium	0.69

Table 3.1. Composition by mass of a tracker plane (information taken from [37])

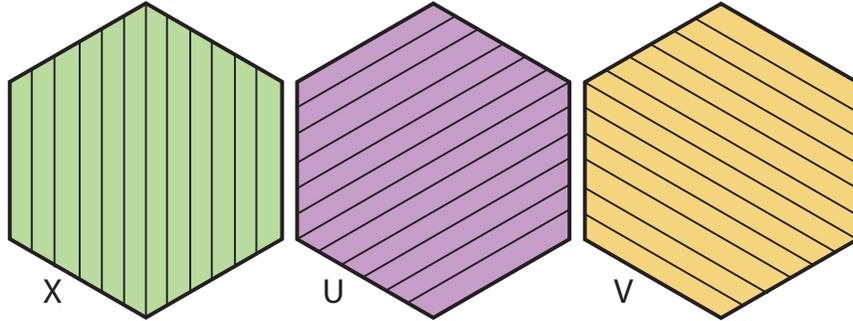


Figure 3.6. Cartoon to demonstrate the X, U and V orientations of the scintillator strips in the $x - y$ plane. (Note that a real MINERvA plane has 127 strips)

glued together with epoxy, and the planes are then covered in Lexan to prevent light leakage between one plane and the next. While the polystyrene is a hydrocarbon with a CH structure, the tracker also includes the TiO_2 coating, dopant and epoxy, leading to a composition as shown in table 3.1. More details can be found in [37].

A 2.6 ± 0.2 mm hole drilled down the center of each strip contains a wavelength-shifting fiber, sealed in optical epoxy. The light collection system, including the function of the scintillator and wavelength-shifting fibers, will be explained in section 3.4.

Each plane is installed in one of three orientations, X, U or V. In the X orientation, the strips are vertical (parallel to the y axis) meaning that scintillation in a given strip gives information about the x position of a charged particle passing through the plane. Planes with a U or V orientation have the strips oriented at 60° clockwise or counterclockwise, respectively, to the vertical (see figure 3.6). By including planes in each orientation within the detector, we are able to reconstruct three-dimensional tracks. (Two orientations would actually be sufficient to generate a 3-d track; the third provides a check, especially useful in the case of multiple crossing tracks).

Each module in the active tracker region consists of two planes of scintillator strips. These alternate between UX and VX configurations, with the X orientation always being downstream of the U or V. The central tracking region, in which this study is based, contains a total of 62 modules.

A 2mm-thick lead collar, colored orange in figure 3.4a, covers the outer 15cm of each scintillator plane, on the downstream side; this is designed to contain electromagnetic showers in the ID, acting as a side electromagnetic calorimeter.

3.3.3. Calorimeter modules in the inner detector

The furthest-downstream portion of the detector contains electromagnetic and hadronic calorimeters. The ten electromagnetic calorimeter (ECAL) modules are similar to tracker modules, except that the 2mm-thick lead collar is replaced by a 2mm-thick sheet of lead covering the entire module. In order to ensure that every ECAL module has lead absorber directly upstream, a transition module between the tracker and the ECAL has a 2mm sheet of lead on its downstream side. The dense lead absorber increases the likelihood of

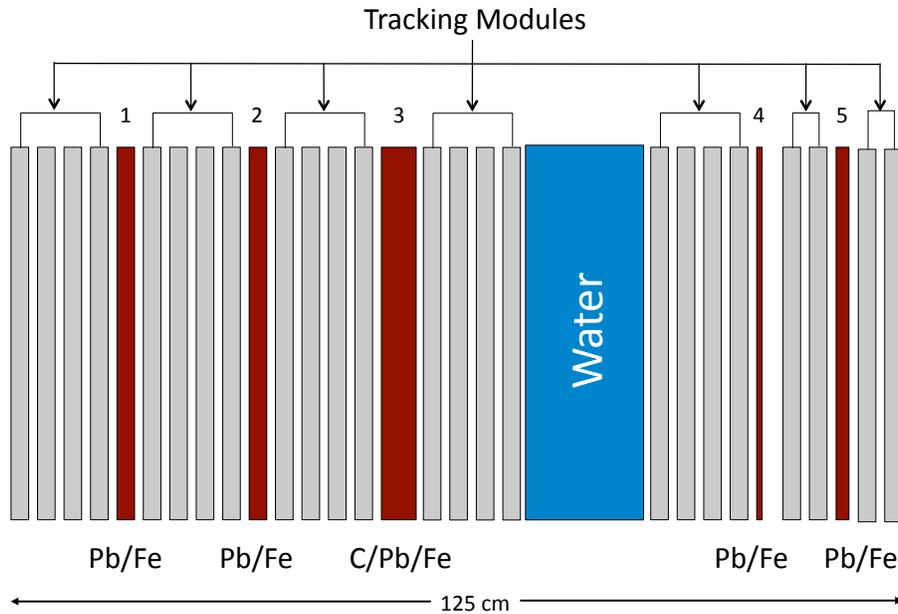


Figure 3.7. Positions of the nuclear targets, taken from [37]. The beam passes from left to right in this diagram.

photon and electron interactions, which can then be tracked by the fine-grained scintillator modules of the ECAL.

The 20 hadronic calorimeter (HCAL) modules, downstream of the ECAL, contain only one plane of scintillator, followed by a 2.54cm-thick plane of steel, to encourage hadronic interactions. The plane orientations alternate in an X-V-X-U pattern.

3.3.4. Nuclear target region

The region upstream of the tracker, as seen in figure 3.4b, contains passive nuclear targets, which are used to test the A -dependence (that is, the dependence on the nuclear mass) of neutrino scattering cross sections.

Target	z position (cm)	Thickness (cm)	Material	Mass (kg)
1	452.5	2.6	Steel	322
			Lead	263
2	470.2	2.6	Steel	321
			Lead	263
3	492.3	2.6	Steel	158
			Lead	107
			Graphite	160
H ₂ O	528.4	17-24	Water	452
4	564.5	0.8	Lead	225
5	577.8	1.3	Steel	162
			Lead	134

Table 3.2. Composition by mass of nuclear targets, starting from the most downstream (information taken from [37]). The mass is for an area within a hexagon of apothem 85 cm, with a 2.5 cm cut on each side of the boundary between materials.

The target region consists of five solid nuclear targets, plus a water target. Between each of these are tracker modules, exactly analogous to those in the active tracker region. The nuclear targets are not used for the analysis described in this thesis, which only measures cross sections in the tracker.

The detector includes five hexagonal planes of solid targets (see figure 3.7). A water target was later added; this was not installed when the data for this analysis were taken. The solid targets are composed of various combinations of lead, steel and graphite. The thicknesses and orientations of the targets vary. The thinner targets allow study of low-energy interactions, while the thicker ones, placed further upstream, provide a higher event rate. The nuclear targets are summarized in table 3.2.

A 1m³ cryostat is located directly upstream of the active detector. In the latter part of the low energy data run, this was filled with liquid helium; however it was empty at the time the data for this analysis was collected.

3.3.5. Veto wall

The veto wall is a structure placed upstream of the helium target. It consists of alternating planes of steel and scintillator (5cm steel, 1.9 cm scintillator, 2.5 cm steel, 1.9 cm scintillator). The purpose of this is to identify “rock muons”: muons generated through neutrino interactions with the dolomite rock upstream of the target, which enter the front of the detector and pass all the way through. These rock muons may otherwise be confused with muons produced in charged-current neutrino interactions in the first plane of the detector. The veto wall also removes any hadrons remaining in the beam. This structure was a late addition to the detector setup, and was not in place when the data analyzed in this thesis was collected.

3.3.6. The outer detector

The outer detector (OD) is located on the six sides of the hexagonal modules. Its steel frame construction serves as both a supporting structure for the detector modules, and as a hadronic calorimeter. Each MINERvA module consists of an ID and an OD component. The OD is colored blue in figure 3.4a.

The steel of the outer detector is interleaved with “bars” of scintillator. The steel enables us to contain hadronic showers generated in the ID; the scintillator enables us to measure the energy produced by these hadrons.

3.3.7. The MINOS near detector

The Main Injector Neutrino Oscillation Search (MINOS)[94], the original experiment in the NuMI beamline, has been running since 2005. Its extensive program of analysis has

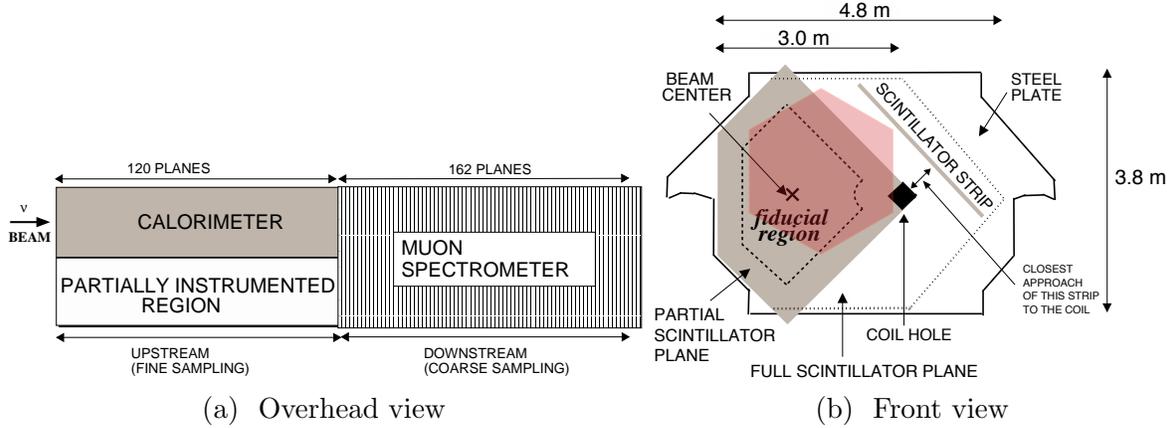


Figure 3.8. Schematic of the MINOS detector, taken from [14] (not to scale). The grey area in figure 3.8b denotes the area covered by the partial scintillator planes used to instrument the upstream fine-sampling region on the left-hand side of figure 3.8a. The red hexagon (added by the author) indicates the position of MINERvA’s inner detector in the beam line, as documented in [185]. The marked beam center is for MINOS; due to the beam’s downwards angle, this does not correspond to the beam center in MINERvA.

included measurements of θ_{23} [21] through ν_μ [18] and $\bar{\nu}_\mu$ [16][20] disappearance, and of θ_{13} [19] through ν_e appearance, as well as searching for sterile neutrinos [15]. The MINOS near detector (henceforth referred to as ‘MINOS’) is located 2.1 meters downstream of MINERvA, and is used as a muon spectrometer. MINOS is of key importance to this analysis, as in order to identify antineutrino charged-current events, we require that the muon produced is matched as a μ^+ in MINOS.

The 1kTon MINOS near detector [156], shown in figure 3.8, is composed of 2.54 cm-thick steel planes, interspersed with 1 cm-thick layers of scintillator. The scintillator planes are formed from 4.1 cm-wide parallel strips, with orientation of the strips alternating between $+45^\circ$ and -45° to the vertical in successive planes. The first 120 planes are instrumented for fine sampling; in this region, every fifth steel plane is followed by a

fully-instrumented scintillator plane, while all other steel planes are followed by a partially-instrumented scintillator plane. The area of the partially-instrumented plane is shown in grey in figure 3.8b, while the additional area of the fully-instrumented plane is shown in white. The coarse-sampling region, further downstream, has only the fully-instrumented scintillator every five planes; there are no partial scintillator planes in this region.

The MINOS detector is magnetized by a coil that runs in a loop passing through the detector (see the coil hole in 3.8b). This generates a toroidal field with an average strength of 1.3 T. This field causes charged particle tracks to curve; the direction of curvature indicates the particle's charge, while its radius of curvature can be used to estimate the particle's momentum. If a particle ranges out within the calorimeter region, the range of the particle can also give a momentum estimate. Both of these methods are used to obtain the muon momenta used in this analysis; thus uncertainties on the MINOS reconstruction and simulation contribute to our systematic uncertainty on muon energy scale.

The requirement of a muon charge-matched in MINOS significantly aids our purity, by removing almost all wrong-sign neutrino events. The price for this is a limitation on our angular acceptance, as muons must be sufficiently forward-going to intercept the front of the MINOS detector. They must also have sufficient energy to traverse any material between the MINERvA and MINOS detectors. While this decreased acceptance is also dependent on the position within the MINERvA detector where an interaction took place (muons produced at the downstream end are more likely to reach MINOS), the approximate result of the MINOS-matching restriction is that we can only reconstruct

events with a muon energy above 1.5 GeV and an angle of less than 20° with respect to the beam direction.

3.4. Light collection system

This section describes how MINERvA measures the energy of neutrino interaction products. As a charged particle traverses a scintillator strip, it generates light, which is transmitted via a wavelength-shifting fiber to one of the 64 channels of a photomultiplier tube (PMT). The readout from the PMTs is processed by front-end electronics boards, mounted on top of the PMTs, which are in turn mounted on the outside of the MINERvA detector. The steps of this process are described below.

3.4.1. Scintillation and light transmission

MINERvA's 32,000 scintillator strips are extruded from polystyrene doped with 1% (by weight) 2,5-diphenyloxazole (PPO) and 0.03% (by weight) 1,4-bis(5-phenyloxazol-2-yl)benzene (POPOP). PPO acts as a primary scintillator, absorbing energy from the particles and emitting it as ultraviolet light with a wavelength of 357nm, with POPOP serving to absorb the light produced and re-emit it at a violet wavelength (410nm) [87]. This combination of materials is also used in the MINOS detectors [156], for which it was extensively evaluated.

The light produced in the strip is read out by a 1.2 mm diameter, 175 ppm Y-11 doped, S-35, multicladd wavelength-shifting (WLS) optical fiber inserted in a hole passing along the length of the strip (see figure 3.5). These fibers shift the wavelength of the light to green. Light is totally internally reflected within the fibers, and thus passes along them

to the end of the strip, where it is connected by a Fujikura-DDK optical connector to a 1.2 mm clear optical fiber. In order to minimize light loss, the other end of the WLS fiber is mirrored.

Each connector services eight WLS-clear fiber pairs. The eight fibers are bundled together to form a single optical cable, surrounded by a light-tight sheath, as detailed in [37]. The clear fibers transmit light to the photomultiplier tubes. They vary in length from 1.08 m to 6 m.

3.4.2. Photomultiplier tubes

Light delivered by the optical fibers is read out by Hamamatsu H8804MOD-2 photomultiplier tubes (PMTs), as described in [37]. The full detector has 507 PMTs, each of which consists of 64 pixels. Each PMT services eight optical cables, which enter the PMT at a fiber-end face plate at the end of the cylindrical optical box housing the PMT. The fibers from the cables are then connected into an 8×8 array of pixels on a “fiber cookie”. Fibers from two adjacent cables are arranged in a checkerboard-style “weave” pattern over two rows of the pixel array. This is to ensure that fibers from adjacent detector strips do not correspond to adjacent PMT pixels. Thus, if cross talk occurs - that is, if a signal in one pixel induces a current in the adjacent pixel, this will not then show up in adjacent detector strips and be misconstrued as a physical effect. The low-level signals in strips not adjacent to true activity, resulting from cross talk with the fiber weave in place, can be identified and disregarded when reconstructing particle activity in the detector.

Scintillation light delivered through the optical fibers impinges on a photocathode, releasing photoelectrons via the photoelectric effect. The electrons are accelerated along

a chain of 12 dynodes, releasing further electrons at each dynode. The resulting electrons are collected at an anode, with a typical gain of around half a million electrons per photoelectron.

3.5. Data acquisition

The MINERvA data acquisition system is described in detail in [171]. The following gives a summary of the key elements.

3.5.1. Readout electronics

A front-end board or FEB is attached to the outside of each optical box; this reads out the pulse height for the 64 channels of the corresponding PMT, as well as providing high voltage to the PMT via a Cockroft-Walton high-voltage generator. The output of the PMTs is read out by six “Trip-T” application-specific integrated circuit chips. Three chips read the high-, medium- and low-gain output respectively, on 32 of the 64 PMT channels.

Up to ten FEBs are daisy-chained together into a readout *chain*, connected to a *CROC-E* controller (Chain ReadOut Controller - Ethernet). Each CROC-E can support four chains, and temporarily stores data “frames” received from the FEBs on those chains. The CROC-E also received timing information from a CRIM (CROC-E interface module), which services up to four CROC-Es. As we receive our neutrino beam in timed pulses, our readout window is triggered from the NuMI beam timing signal. The CROC-Es and CRIMs are divided between two VME *crates*.

3.5.2. The readout gate

The NuMI beam delivers neutrinos for a $10\mu\text{s}$ spill every 2.2 seconds. A readout *gate* is opened, allowing us to take data, 500 ns before the spill is expected, and closed 5.5 ms after it ends. We refer to these as NuMI beam gates. Additionally, we have the ability to take pedestal gates (in which we measure the background electronics noise with no beam present) or light injection (LI) gates (in which we flash the PMTs with LED light) for calibration purposes. When the beam is on, a pedestal or LI gate can be opened in the gap between two beam gates. A typical MINERvA run, lasting around 12 hours, will record two subruns consisting of a total 1500 gates of beam data interspersed with pedestals, eight of 1500 gates of beam/LI, and 30 consisting of 700 beam gates.

A *frame* of data containing the high-voltage state, hit timing and hit pulse heights is read out from each FEB on the chain when the gate ends, and stored temporarily on the CROC-E. The data acquisition system (*DAQ*) then reads the data from the CROC-Es. The DAQ checks each FEB's data to see whether it contains hits above a discriminator threshold (on the high-gain TriP-T channel, as this is the most sensitive); if so, it reads the data from that FEB, which can contain up to 7 timed hits plus an eighth untimed hit corresponding to the total integrated charge in the remaining part of the gate after the last timed hit. The firing of a discriminator triggers chips corresponding to the low, medium and high gains to push data 20 clock ticks later (a tick is 9.4 ns). After the data is pushed, the TriP-T chips must reset. This process takes 20 clock ticks, and during this period of *dead time*, the 32 channels corresponding to those chips are unable to read data from any further activity in the detector.

This information corresponds to the low-energy data taking, as used for this analysis; the detector is now taking medium energy data, using a different configuration.

3.5.3. The MINERvA DAQ

MINERvA has a set of three computers running its data acquisition (DAQ) software: one per crate plus one master computer. These machines run Scientific Linux Fermi; the DAQ software is written in C++. The DAQ computers are situated underground next to the MINERvA detector and readout crates. Data from the FEBs is collated in memory and written to a binary file on disk at the end of each subrun (roughly every 20 minutes, depending on run series setup). No external network connection is required to ensure that the data is written to disk. Twice a day, the newly-produced files are copied to Fermilab's bluearc servers above ground, where they undergo offline processing and cataloguing.

The *nearline* monitoring system reads data a frame at a time from the DAQ head node and processes each event in near-real time from binary format to the raw digits format, compatible with the ROOT[64] program. This data is then further processed on the nearline machine to produce monitoring plots and event displays, which are displayed on the control room system, allowing the MINERvA shifter to monitor the detector.

3.6. Offline data processing and calibration

The raw data files generated by the DAQ contain the timing and pulse height information from the detector in a binary format. This data then undergoes several stages of processing in order to convert it to a more easily readable format for processing, and to calibrate the data in such a way that the ADC counts read out from the FEBs are

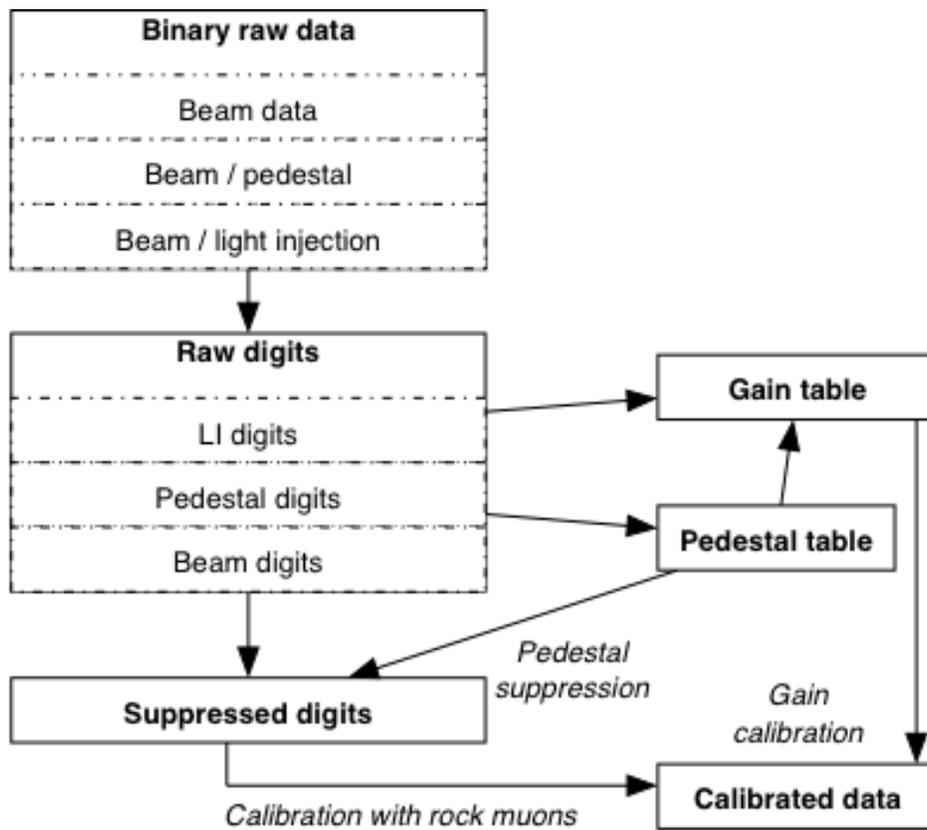


Figure 3.9. Simplified flow chart of the stages of data processing and calibration

converted to measurements of energy deposited in the detector. A summary of the main components of this process is shown in figure 3.9. This section outlines the key components of the processing and calibration process; a full explanation of the calibration can be found in [37].

3.6.1. Keep-up processing

The *keep-up* processing refers to a suite of tasks scheduled to run automatically once or twice a day, without human intervention. It encompasses the processing of raw data to raw digits, pedestal table generation, and pedestal suppression. The procedure for each

of these is explained below: the principles behind pedestal calibration will be described in section 3.6.1.1.

In each case, a BASH script calls a corresponding python program that manages the processing and metadata management. The BASH script also reports any warnings generated via email to experts. The `dailyProcessingKeepup` script begins by identifying raw data files in need of processing, by checking the directory on the bluearc server in which the files from the DAQ have been deposited. To prevent problems due to excessive files in a single directory, the files are arranged in a directory structure whereby they are separated by the type of run (beam, for example, or combined beam and pedestal) and then into a separate directory containing all the files for a single run, contained in a parent directory structure dividing them into groups of 100 runs. The keep-up script scans this structure for files that have not yet been declared to the metadata database, SAM [133]. These will typically be files that have been delivered from the DAQ since the last keep-up run; however this check allows for recovery in the case of other issues. Any undeclared file will be declared to SAM. The SAM database stores metadata about individual files, as well as tracking their locations on disk and/or tape, allowing access to the files without knowledge of their physical locations. Examples of metadata stored include the run and subrun information, data tier (raw data, reconstructed Monte Carlo simulation, calibrated data etc) and validation information such as file size and checksum. We are also able to add custom information, recording such things as the run configuration and detector setup at the time of the run.

Once the files have been declared, unprocessed files (both the newly declared, and any other files which were previously declared, but for which the processing could not be

completed, for example due to a grid outage) are processed into raw digits. A locking mechanism is implemented to ensure that duplicate grid jobs are not generated concurrently for the same file. The processing involves submitting a grid job for each file, which runs the BuildRawEventAlg algorithm for each data stream (pedestal, light injection or beam). This algorithm is written using the GAUDI framework [51]. The data from the FEBs is parsed into the ROOT [64] format, creating a file for each data stream that can be used for further processing. The resulting ROOT file, known as a raw digits file, is written to Fermilab's dCache [111]. This is a distributed disk storage system fronted by a single filesystem tree, allowing a file to be stored anywhere in a pool of data server nodes, while being accessed from a Linux-style directory structure independent of its physical storage location. This file system location is stored in the SAM database. All files written to dCache are backed up to the Enstore tape filing system via the FTS file transfer system.

3.6.1.1. Pedestal suppression. Similarly scheduled procedures manage pedestal table creation (from raw digits from the pedestal data stream) and pedestal suppression (where pedestal tables from the calibration database are used to calibrate raw digits from the beam data stream). In these cases, the initial raw digits in need of processing are found by querying SAM, in which a flag is set to indicate which files have already been processed, and which are in need of reprocessing.

For two subruns at the beginning of each run (in the current standard run configuration), pedestal data is taken, interspersed with the beam data, giving us a total of 1500 gates of pedestal data. This records the background level of electronic noise present in

the detector when the beam is not firing. It can also include readings from charged particles passing through the detector, but not produced by the beam; typically, these would be cosmic ray muons. Outliers such as these are removed using the Pierce’s Criterion technique [180]. Once the outliers have been removed, each channel is found to have a steady (less than 2% variation) rate of pedestals (there is around a 7% variation between channels). The mean pedestal rate for each channel is subtracted from the ADC counts in the raw digits beam data before performing any further calculations. This is carried out by an automated pedestal-suppression keep-up job, which turns the beam-type raw digits into suppressed digits. Pedestals are also subtracted from the light-injection data used to calculate the gains. In each case, the most recent pedestal data taken before a given LI or beam subrun is used to provide up-to-date information. We refer to the period between consecutive sets of pedestal data as an interval of validity (IoV). As pedestals are typically taken in the first two subruns of each data run, the pedestals from those two subruns are used for all of the run’s data; the interval of validity corresponds to one run, or around 12 hours. A more detailed explanation can be found in [37]

Pedestals are calculated by the pedestal table production keep-up processing job, which runs automatically each day. This outputs pedestal tables in a text file format. Once a month, these pedestal tables are subjected to manual data quality checks, and once any problem tables have been removed, the tables are uploaded to the MINERvA conditions database, which stores the pedestal level for each channel/IoV combination.

3.6.2. Data calibration

The keep-up processing translates the ADC counts from the FEBs into an accessible ROOT format, and suppresses the pedestals, leaving only the ADC counts generated by actual beam events. These counts must then be calibrated in order to convert them to measurements of energy deposited in a given strip i . This is accomplished using the following formula (taken from [37]).

$$(3.1) \quad E_i = [C(t) \cdot S_i(t) \cdot \eta_i^{att} \cdot e^{l_i/\lambda_{clear}} \cdot G_i(t) \cdot Q_i(\text{ADC})] \times \text{ADC}_i$$

where:

- E_i is the estimated energy deposited in strip i
- $C(t)$ is the overall energy scale of the detector
- $S_i(t)$ is a relative energy scale correction for strip i
- η_i^{att} corrects for attenuation in a strip as a function of location
- $e^{l_i/\lambda_{clear}}$ corrects for the attenuation in an optical fiber of length l_i
- $G_i(t)$ is the gain of the PMT pixel corresponding to strip i
- $Q_i(\text{ADC})$ is the ADC-to-charge conversion factor of the FEB channel corresponding to strip i

Remember that each strip corresponds to a single pixel of a PMT, which corresponds to a channel on an FEB (see section 3.4). Section 3.6.2.1 will explain how the gain factor $G_i(t)$ is calculated as a function of time, using light injection data. Section 3.6.2.2 will

discuss the rock muon calibrations that give us $C(t)$, $S_i(t)$ and η_i^{att} , and section 3.6.2.3 will explain how the Q_i is calibrated for each FEB before it is installed in the detector.

3.6.2.1. Gain calibration. The gain of a PMT is the number of electrons produced at the anode per photoelectron (PE) generated at the cathode. The gain response of a given PMT channel changes gradually as the PMT ages, meaning that in order to correctly calculate the energy response from a given ADC count, we need to monitor the gains regularly and ensure that we use a recent measurement. The method used to calculate gains is explained in detail in [37], and summarized below.

In order to measure the gain, light injection (LI) data is used. In a MINERvA data-taking run consisting of 40 subruns (typically lasting around 12 hours), runs 3-10 are combined beam/LI runs, consisting of 750 beam gates interspersed with 750 gates of LI. Thus, for a given run, 6000 gates of LI are available.

During an LI gate, the light collection system is flashed with 472 nm blue light from AlGaInP LEDs in MINERvA's light box (described in [37]). This light is delivered via clear optical cables to the optical boxes, where it is spread to each pixel using a polypropylene diffuser. By delivering a small but known amount of light (from 1 to a few PEs, configurable via a serial connection to the light box), and measuring the resulting ADC, it is possible to calculate the gain. Over time, we can define a channel's gain as

$$(3.2) \quad g = \frac{\bar{Q}}{\lambda e}$$

where \bar{Q} is the mean charge measured at the anode, and λ is the number of photoelectrons reaching the first dynode of the chain. Using the method of [176], we can calculate

a probability distribution $P_n(q)$ of measuring a charge q at the anode (for a PMT with n dynodes, each amplifying according to a Poisson distribution, with an amplification proportional to the number of incident PEs). This distribution $P_n(q)$ has a width σ given by summing in quadrature the widths of the pedestal, the incoming photoelectron distribution and the uncertainty introduced by the broadening of the dynode chain:

$$(3.3) \quad \sigma^2 = \sigma_P^2 + \lambda g^2 e^2 + \lambda g^2 e^2 w^2$$

In this equation σ_P is the width of the pedestal distribution, and w^2 is defined by

$$(3.4) \quad w^2 \equiv \sum_{j=1}^n \left(\prod_{i=1}^j \frac{1}{g_i} \right)$$

Where the individual dynode gains g_i can be multiplied to get the total gain g . Substituting $\lambda = \bar{Q}/eg$ from equation 3.2 and rearranging, we get an equation for the gain:

$$(3.5) \quad g = \frac{\sigma^2 - \sigma_P^2}{\bar{Q}(1 + w^2)e}$$

Thus we are able to calculate the gain from the anode charge distribution and the pedestal distribution, as long as we know the parameter w . We find that by using the properties of photomultiplier tubes as given in the PMT handbook [137], which relates the gain g_i to the voltage V_i at each dynode.

$$(3.6) \quad g_i \propto V_i^\alpha$$

where $\alpha \approx 0.75$. Thus, using our knowledge of the high voltages to the PMT, we are able to calculate the average gain.

A gain measurement is made for each MINERvA run, using the combined data from all LI subruns in that run (6000 gates in the current configuration). The gain of each channel is stored in the MINERvA conditions database, and is used to calibrate all runs in its interval of validity - that is, all runs until the next good gain table. Before storing the data, quality checks are run - these typically look at around a month's worth of data and involve checking how gains are distributed over the course of the month - outliers in the distribution can be indicative of a bad gain table (perhaps produced during some kind of detector maintenance) that should be discarded. In this case the interval of validity of the previous gain table would be extended. These quality checks can also be used to identify PMTs with dead channels, or with gains that apparently vary wildly from run to run - this could be indicative of problems with either the PMT or the readout electronics, and provides a timely warning when failing hardware will soon need to be replaced.

Gains in the MINERvA detector are typically of the order of 6×10^5 . As shown in 3.10, the average gain increases with time, due to PMT aging. The sharp peaks occur when the high voltage to the PMTs is reset; the PMTs take around a day to stabilize. After the period of outage around day 160, the high voltages were recalibrated, causing a drop in overall gain. Suggested high-voltage values are produced by the gain calibration procedure, and are chosen to equalize gains across the detector. This is done by choosing a value for each PMT that will set its lowest-gain 8 pixels to have a fixed average gain of 4.38×10^5 .

This calibration contributes the value $G_i(t)$ to the calibration equation 3.1.

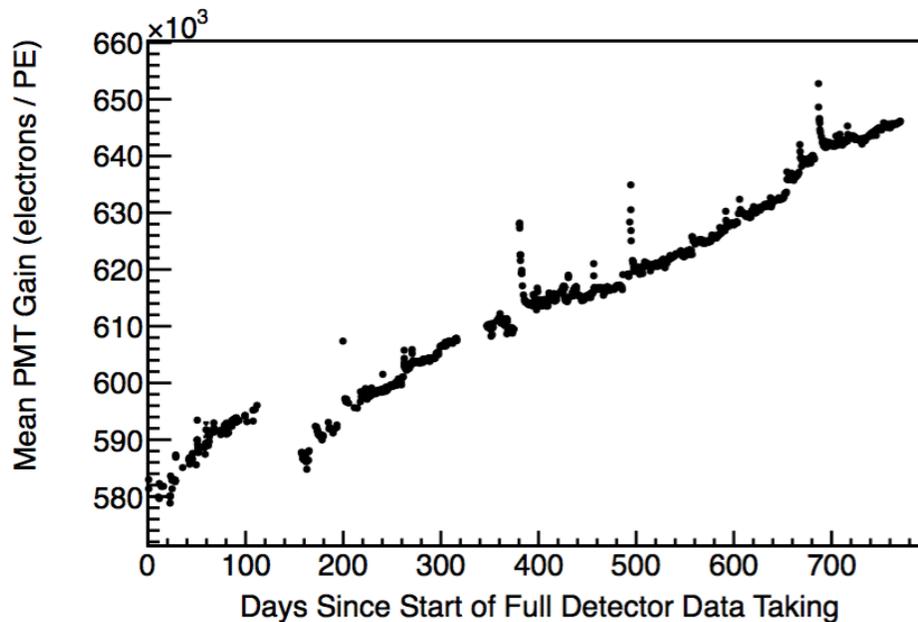


Figure 3.10. Variations in the mean gain of all of MINERvA’s PMT pixels, over time. Reprinted from [37].

3.6.2.2. Calibration with rock muons. The NuMI beam [98] travels through 240 m of dolomite rock before reaching the MINERvA detector. While the main purpose of the rock is to filter out muons generated in the interactions that produce the neutrino beam, a small number of neutrinos from the beam will interact with the nuclei in the rock, producing further muons. On average, one of these so-called “rock muons” crosses the MINERvA detector from front to back every two beam pulses.

These “rock muons” , which pass through the detector as minimum-ionizing particles depositing a known amount of energy per centimeter, can be used to correct for scintillator plane alignment, variations in light yield between different scintillator strips, and for overall energy scale. They can also be used to calibrate the timing. As many of these calibrations require us to measure the response of every single detector strip to enough

rock muons to generate a statistically significant sample, these calibrations cannot be performed on a daily basis like the gain and pedestal corrections; this is not problematic as none of these are expected to change rapidly over time. Instead, calibrations use data taken from longer intervals of the order of months. The periods are typically chosen to correspond to hardware changes in the detector that could affect the calibration constants.

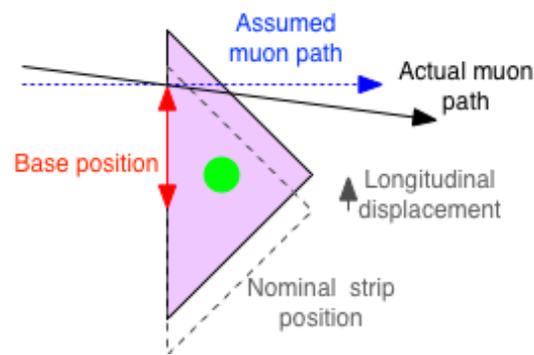


Figure 3.11. Diagram showing how base position is calculated

Plane alignment

The MINERvA inner detector consists of planes of 127 interlocking scintillator strips with triangular cross section. It is possible for these planes to be misaligned in two ways - first, they can be translated longitudinally, meaning that the point of each triangle is shifted by the same distance to one side. Secondly, they can be rotated about the z axis. In order to check for longitudinal displacement, the amount of energy deposited in each strip is plotted vs the position relative to the strip's nominal triangle base (the point at which we believe the triangle's point to be). As rock muons tend to travel in the direction of the beam, we assume a normal incidence; in that case, the energy deposited by the muons, which is proportional to the amount of material they travel through, should peak at the

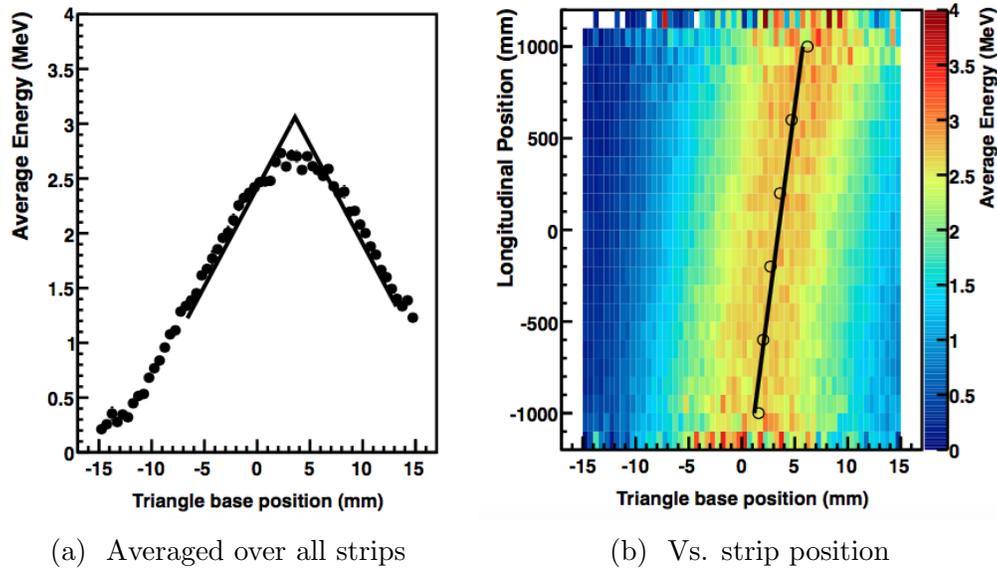


Figure 3.12. Triangle base position shifts used to calculate plane alignment. From [37].

triangle's point and reduce linearly to zero at its edges. By plotting energy deposited vs the distance from the nominal peak, averaged over all strips in the plane, we can see if there is an offset. Figure 3.12a shows this for an example plane; when the triangular strip shape is fitted, we see that this particular plane is aligned approximately 3 mm to the right of its nominal position.

If, instead of averaging these energy deposition patterns for all strips, we make a 2-dimensional histogram, plotting the energy vs both the offset from the nominal base and the longitudinal position of the strip in the plane, we are able to determine whether a rotation around the z axis causes this shift to vary as we move across the plane. In figure 3.12b, we see that the illustrated plane does indeed have a small rotation, as evidenced by the fitted line connecting the peak position in different strips; if there were no rotation, this line would be vertical.

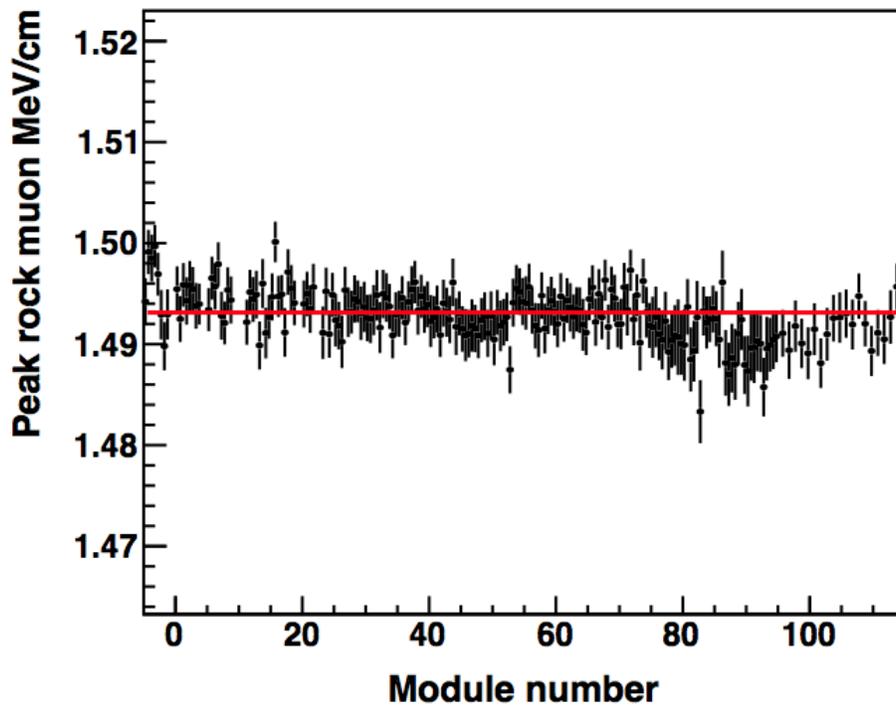


Figure 3.13. Peak rock muon energy per centimeter was fitted for each plane. These peak values for each plane (with uncertainty from fitting) are plotted, to give a best-fit flat distribution, telling us the relative weighting for each plane. Plot reprinted from [37].

Plane misalignment is corrected for during the reconstruction phase.

Strip-to-strip calibration

Rock muons can also be used to calibrate the relative light yields of the detector's individual strips. Variations in light yield can be caused by tiny differences in the strips' extrusion and assembly process. They are accounted for by measuring the average energy deposited in each strip per rock muon, a process which is repeated when dead channels have been identified and removed from the calculation. The peak energy is then equalized

for each plane. The resulting constants are then multiplied to give a strip-to-strip weighting for each strip. These are all renormalized to ensure that the average weighting is 1, guaranteeing that the overall energy scale is unaffected. More details of this calibration can be found in [37].

This calibration contributes the value $S_i(t)$ to the calibration equation 3.1.

Absolute energy scale

The absolute energy scale (in “muon energy units” or MEU) is calibrated using a sample of rock muons whose tracks in MINERvA have been matched to corresponding MINOS tracks. The range and curvature of the MINOS tracks is used to provide an energy estimate for the muon in MINOS; this is corrected for projected energy loss in the MINERvA detector to give an estimate for its energy at the front of MINERvA. As the behavior of muons in scintillator is well understood (see [94]), this set of muons is used to produce a GEANT4 [27] simulation modeling how much energy such a set of muons would be expected to deposit in the detector. The energy of each reconstructed energy cluster (using a trial MEU factor) from the muon tracks (see section 3.8 for an explanation of how energy clusters are reconstructed) is then plotted in both data and simulation. The peak region is fitted to a degree-5 polynomial in each case. (See figure 3.14). By comparing these fitted distributions, an improved MEU factor can be extracted. This is then multiplied by the slope of a plot of true vs. reconstructed energy cluster energy from the simulation to correct for reconstruction efficiency issues.

As this does not require multiple readings in every strip, and because the overall energy scale has been found to change with time, the interval of validity for the MEU is

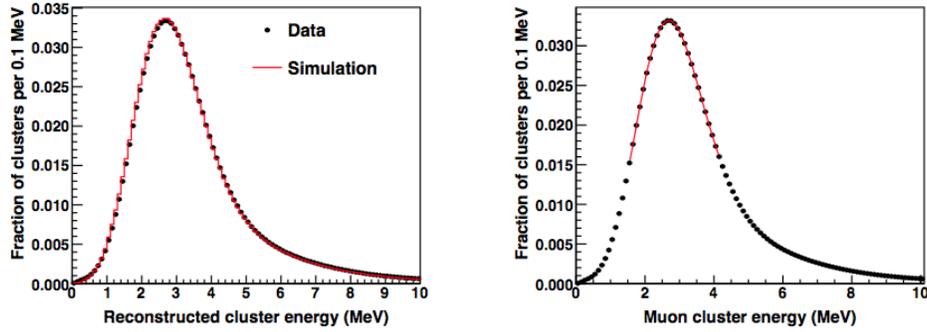


Figure 3.14. Rock muon energy clusters in data and simulation (left) and the polynomial fit to data (right). Reprinted from [37].

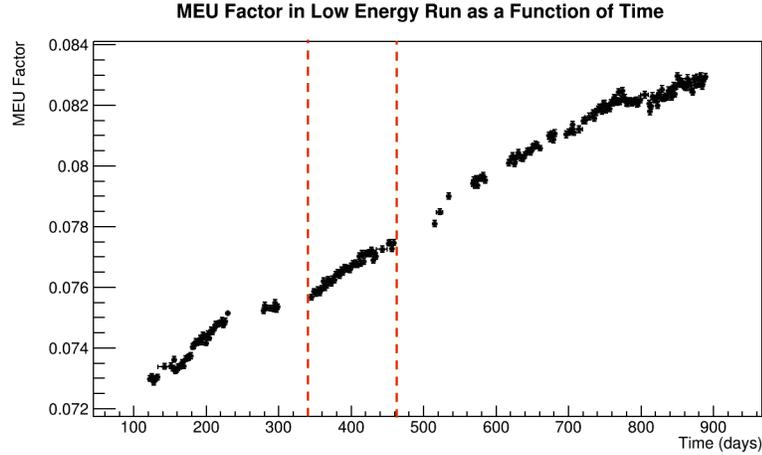


Figure 3.15. MEU factor $C(t)$ for MINERvA’s low-energy run as it varies with time since MINERvA was switched on. (This plot shows only the full-detector data; day zero corresponded to the beginning of the earlier run with only a partial detector installed.) The red dashed lines denote the beginning and end of the ‘minerva5’ dataset corresponding to the data used in this analysis. Plot courtesy of J. Kleykamp, from [138].

two days - much shorter than for the other rock-muon calibrations. The main reason for the change in energy scale in the low energy MINERvA run period was due to scintillator aging, which decreased light yield. A cooling system has since been installed, improving this issue.

This calibration contributes the value $C(t)$ to the calibration equation 3.1. As with the other calibrations, more details can be found in [37].

Timing calibration

There is a time delay between a charged particle entering a scintillator strip, and the charge being read out from the channel corresponding to that strip. This is due to a combination of scintillator decay time, the length of the optical fiber that the light must traverse, and data transfer differences due to an individual FEB's position in a chain (and the chain's position in the CROC-E).

While the length of optical fibers is known and can be corrected for, the other effects must be

calibrated. Again, we use rock muons. After correcting for muon time of flight and optical fiber length, we have an estimate of when the muon should be detected in each strip along its track. The offsets between expected and measured time are measured as a function of the number of photoelectrons produced (time slewing, due to scintillator decay, is known to be a function of energy). The resulting time slewing is parametrized as a function of the number of photoelectrons (figure 3.16. Because timing is dependent on the hardware, it must be recalibrated whenever an FEB or PMT is changed).

By summing all the timing information for a given FEB, the readout timing delay can be isolated.

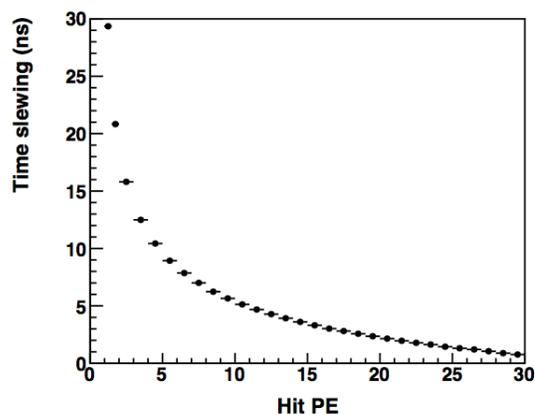


Figure 3.16. Time slewing vs hit PE, from [37]

3.6.2.3. *Ex situ* calibrations. In addition to the ongoing time-dependent calibrations, some one-off calibration processes are performed outside of the detector. All of these are also described in [37], but are summarized briefly below.

Optical cable attenuation

The optical fiber attenuation constant λ_{clear} in the calibration equation 3.1 was measured on a dedicated test stand, where clear optical cables were connected between an LED source box and a readout box to measure their response to a known amount of light. This yielded a measurement of $\lambda_{clear} = 7.83$ m.

Scintillator strip attenuation

The module mapper was an apparatus used to test each module of the MINERvA detector prior to installation. As well as checking for dead or problematic channels, this device measured the attenuation constant for each individual scintillator strip within the module. This test was done using two ^{137}Cs radioactive sources, which were moved in a predefined scan pattern over the module. The response was measured by a custom data-acquisition system connected to a series of PMTs. By reconstructing each strips response as a function of position along the strip, the attenuation factor η_i^{att} in the calibration equation 3.1 was determined for each strip.

FEB constants

Before installation in the detector, every FEB undergoes a series of tests, including measuring its response to charge. A known amount of charge is injected to four non-adjacent

FEB channels (to protect against cross talk). The FEB being tested is connected to a data acquisition system similar to the MINERvA DAQ, and the responses of the low, medium and high gain channels are recorded and plotted vs the input charge (figure 3.17 shows the response for one FEB). The response on each channel is fitted to three linear segments, the parameters of which are entered into the conditions database when the FEB is installed into the detector. When converting an ADC count to charge, the high-gain channel will be used, unless the ADC count is above the point where its response saturates, in that case, it will use the medium-gain channel, or if that has saturated, the low-gain.

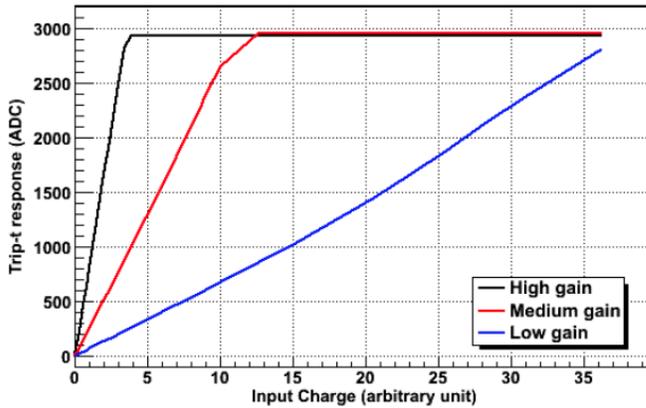


Figure 3.17. An FEB’s high-, medium- and low-gain response to input charges in the test stand, reprinted from [37]

This calibration contributes the value Q_i to the calibration equation 3.1. As with the other calibrations, more details can be found in [37].

3.7. Simulation

We have explained how MINERvA is able to detect and record energy deposited by charged particles. In order to turn this data into measurements of physics processes that can be interpreted by the community, we must go through several further steps. The data recorded must be used to reconstruct particle tracks, and interpreted to determine

what types of interaction occurred. To do this, we need to simulate the behavior of different types of particles in our detector, so that we can identify the particles produced and estimate their energies. As different interaction types can give identical signatures in MINERvA, we must use our knowledge of the theory of neutrino interactions to predict what fraction of reconstructed interactions were likely to have been produced by the process we are studying, and what fraction were produced by background processes; we do this using a Monte Carlo simulation of the possible interaction types produced by neutrinos with a NuMI-like energy spectrum interacting with a MINERvA-like detector. As we cannot measure the energy spectrum of the incoming neutrinos, we must also simulate the NuMI beam to determine our incoming neutrino flux. These simulation and reconstruction processes are described below.

3.7.1. NuMI flux simulation

The NuMI neutrino beam is described in section 3.2. This beam has been simulated to provide an estimate of the flux of neutrinos incident upon the MINERvA detector. The components that go into producing this flux simulation are summarized below, and explained in detail in [36].

3.7.1.1. Hadron production. Hadron production cross sections for the NuMI proton beam on the graphite target are simulated using the GEANT 4 simulation [27] software with the G4numi package, which uses the FTFP_BERT (FRITIOF Precompound - Bertini cascade) inelastic scattering model. This is constrained with proton-carbon scattering data from CERN's NA49 experiment [39] and cross-checked against results from the lower-energy experiment NA61 [10]. NA49 used a 158GeV proton beam (as opposed

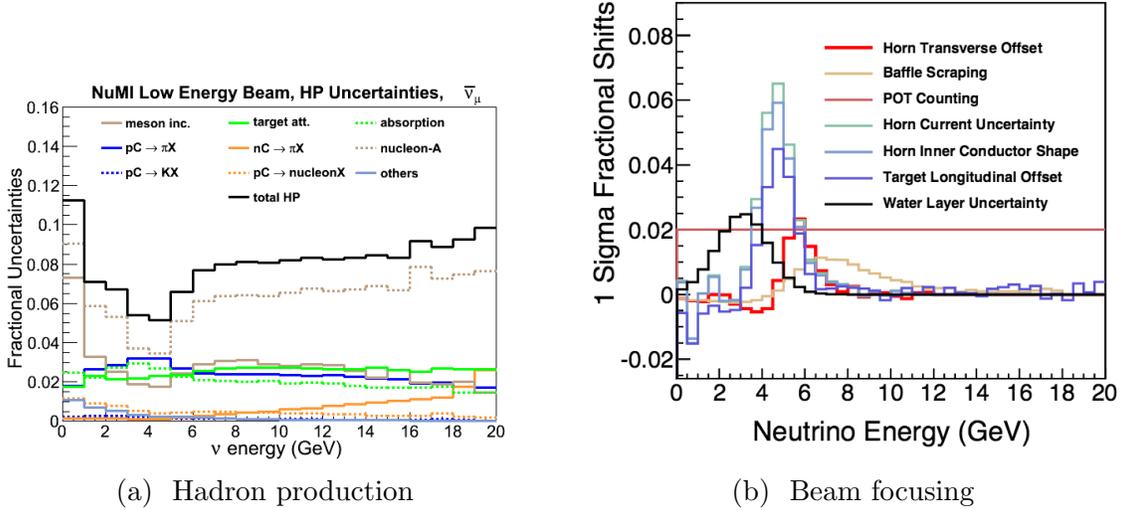


Figure 3.18. Contributions to flux uncertainty

to the 120GeV NuMI beam) incident on a short graphite target (as opposed to NuMI’s long rod-shaped target). NA49’s data is scaled to NuMI energies using the FLUKA Monte Carlo simulation, which assumes Feynman scaling. NA49 data is used where the Feynman scaling variable $x_F < 0.5$, where

$$(3.7) \quad x_F = \frac{2p_L}{\sqrt{s}}$$

where s is the Mandelstam variable corresponding to the squared center of mass energy and p_L is the forward momentum. For $x_F > 0.5$, measurements from the Fermilab Single Arm Spectrometer are used [53]; NA49’s measurement takes precedence where data overlaps. It is also used to re-weight kaon production cross sections for $x_F < 0.2$, and nucleon production for $x_F < 0.95$. Because the NA49 data is taken at a different proton beam energy, it must be rescaled to the NuMI beam energy, which is done using the Feynman scaling technique detailed in [196], supplemented by the FLUKA Monte Carlo simulation

[101] [62]. This re-weighting is detailed in [38]. NA49 cross sections agree with the FTFFP simulation to about $\pm 10\%$ for antineutrino production.

For $0.2 < x_F < 0.5$, the NA49 pion yields are scaled using the K/π ratios measured on a thin carbon target at the MIPP experiment [142].

The PPFX (Package to Predict the FluX) package, released in 2015, is used to implement the reweighting scheme described above. For this analysis, we use PPFX version 1. It includes uncertainties on the hadron production cross sections, as well as on attenuation of the pions, kaons and protons due to re-interaction in the target, or with the materials of the horn and decay pipe (not carbon). Additionally, there is uncertainty due to K^0 production and for the estimated contribution of isoscalar conjugate of the $pC \rightarrow \pi X$ interaction, $nC \rightarrow \pi X$, which has not been directly measured.

PPFX accounts for uncertainty in several components of hadron production, evaluated using the many-universe method, where uncertainties are evaluated by looking at how simulated distributions vary when input parameters are varied within their uncertainties (explained in more detail chapter 6). Their relative contributions to the total flux estimate are shown in figure 3.18a. These uncertainties include:

- **Meson incident correction:** interactions with mesons as the projectile, for which little data is available
- $pC \rightarrow \pi X$: uncertainty on pion production cross section, based on NA49 [39] measurements scaled with FLUKA [101][62]
- $pC \rightarrow KX$: uncertainty on kaon production cross section, based on NA49 [39] pion production and scaled with MIPP ratios [142]

- **Target attenuation:** accounts for the uncertainty on the total interaction cross section of nucleons and mesons on carbon
- $nC \rightarrow \pi X$: neutron interaction cross section estimated through isoscalar symmetry
- $pC \rightarrow \text{nucleon}X$: uncertainty on nucleon production, based on NA49 measurements [39]
- **Absorption:** similar to target attenuation, but for the materials of the horn and decay pipe
- **Nucleon-A:** interactions of protons and neutrons for materials and energy ranges for which no experimental data is available, including proton-carbon interactions for which there is no external data. This is predicted by a nuclear mass-dependent scaling method.

3.7.1.2. Beam focusing. Two magnetic horns (described in [98]) are used to focus pions and kaons produced in the proton-carbon target interaction. These horns take a maximum current of 200kA, and can be run in a forward or reverse current configuration to favor neutrino or antineutrino production respectively. Whether a given particle is focused sufficiently such that it will produce a neutrino that hits the MINERvA detector depends on its initial momentum and angle, as well as on its charge. For this analysis the horn current was set up so as to prioritize the focusing of pions that give an antineutrino beam energy peak around 3-3.5 GeV. The horn system is modeled in GEANT4 [27] using the g4numi package.

Parameters affecting the beam focusing are listed below (taken from [38]), with their relative contributions shown in figure 3.18b.

- **Horn transverse offset:** There is a 0.3 mm uncertainty on horn 1 position and 0.5 mm on horn 2, as detailed in [98]
- **Baffle scraping:** at the tails of the beam transverse position distribution, the beam may hit (‘scrape’) the walls of the baffle, as explained in [98]. There is a 0.25% uncertainty [170] on how much of the beam scrapes the baffle.
- **POT counting:** number of protons on target delivered by the NuMI beam is known to 2% [170].
- **Horn current uncertainty:** uncertainty on the current delivered to the horns (nominally -185 kA for this analysis) has 1% uncertainty [170].
- **Horn inner conductor shape:** two different implemetations of the inner conductor shape model change give flux differences similar to changing the horn current by 0.8%. We use 100% of this as an uncertainty on the inner conductor shape. [38]
- **Target longitudinal offset:** Target position changed at different times during the low-energy run, affecting the falling edge of the focusing peak. This accounts for residual uncertainty on the offset.
- **Water layer uncertainty:** A 1 ± 0.5 mm layer of water on the inside surface of the inner conductor cools the horns [38]. We simulate a 1 mm layer and use the difference between that and 0.5 mm as an uncertainty.

3.7.1.3. Neutrino-electron scattering constraint. This procedure to constrain the flux is explained in detail in [167]. While nucleons in nuclei are composite particles experiencing complex nuclear effects, the electrons surrounding the nuclei are point particles. The cross section for neutral current neutrino-electron scattering $\nu_\mu + e \rightarrow \nu_\mu + e$ is well

understood in electroweak scattering theory.

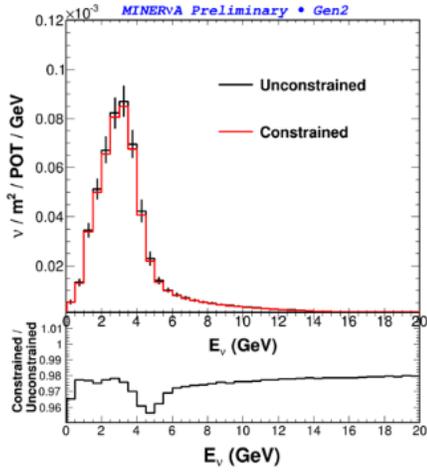
$$(3.8) \quad \frac{d\sigma}{dy}_{\nu_\mu e \rightarrow \nu_\mu e} = \frac{G_F^2 s}{\pi} \left[\left(\frac{1}{2} - \sin^2 \theta_W \right)^2 + \sin^4 \theta_W (1 - y)^2 \right]$$

Where y is the inelasticity E_e/E_ν , the Mandelstam variable $s \approx 2m_e E_\nu$, G_F is the Fermi constant and θ_W is the Weinberg angle.

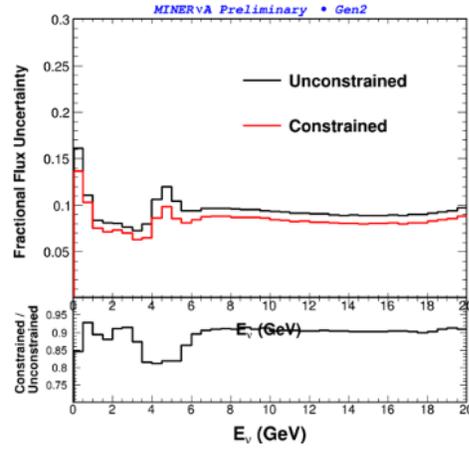
While it is a much rarer process than scattering from the nucleons in the nuclear material, MINERvA's low-energy run yielded 135 ± 17 ν - e scattering events from 3.5×10^{20} protons on target, with a predicted background of 30 events [166]. (The beam consisted of 93% ν_μ ; the component of $\bar{\nu}_\mu$ and traces of ν_e and $\bar{\nu}_e$ were also accounted for.) As the cross section is well known, discrepancies between data and Monte Carlo predictions will instead be due to mis-modeling of the flux distribution. Thus, the final-state distribution of electron energies can be used to constrain both the overall normalization and the shape of the neutrino flux. This procedure is explained in [166] and is implemented using Bayes' theorem, which relates the probability of a given flux model M given an observed neutrino-electron scattering measurement $N_{\nu e \rightarrow \nu e}$ to the probability of the flux model and the probability of seeing such a measurement given that the model was correct:

$$(3.9) \quad P(M|N_{\nu e \rightarrow \nu e}) \propto P(M)P(N_{\nu e \rightarrow \nu e}|M)$$

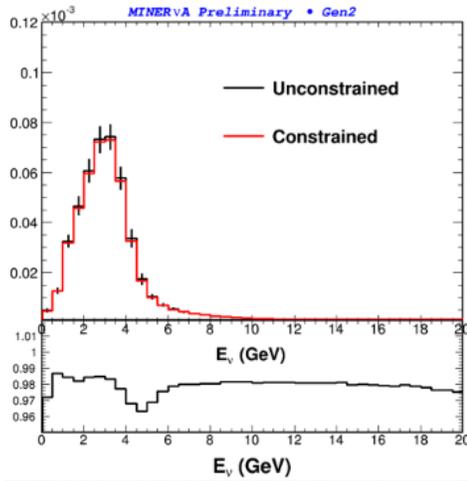
where $P(N_{\nu e \rightarrow \nu e}|M) \propto e^{-\chi_M^2}/2$, χ_M^2 being the chi-squared statistic between the observed and predicted electron spectra. By weighting each universes's distributions by $e^{-\chi_M^2}$, the flux is constrained. While an equivalent measurement is not available for $\bar{\nu}_\mu e$ scattering, the known correlations between the neutrino and antineutrino fluxes are used to translate this method to the antineutrino distribution. As shown in 3.19, applying the



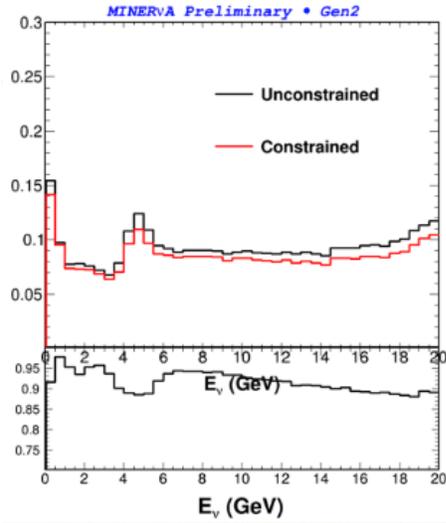
(a) Neutrino flux



(b) Neutrino flux uncertainty



(c) Antineutrino flux



(d) Antineutrino flux uncertainty

Figure 3.19. Flux distributions and uncertainties for both neutrinos and antineutrinos, with and without the $\nu - e$ scattering constraint. In each figure the top plot shows the constrained (red) and unconstrained (black) distributions; the plot below shows the ratio of the constrained to unconstrained values. Reproduced from [104]

constraint results in approximately a 2% decrease in the antineutrino flux prediction, and a 10% reduction in flux uncertainty, changing the flux uncertainty from around 9% to 8%.

3.7.2. Simulation: the GENIE Monte Carlo event generator

MINERvA uses the neutrino interaction event generator GENIE (**G**enerates **E**vents for **N**eutrino **I**nteraction **E**xperiments) [93] version 2.8.4 to model physics processes within the detector. GENIE is a Monte Carlo event generator, capable of simulating interactions between any neutrino or antineutrino flavor and any nuclear target, over a range of energies from the MeV to PeV scale. It concentrates specifically on the few-GeV range pertinent to accelerator-based oscillation experiments. GENIE’s simulated event distributions are used in this analysis to estimate background levels and efficiencies. From an input neutrino flux distribution, GENIE uses a system of random number generation to simulate neutrinos from MINERvA’s energy spectrum, and then uses its physics models corresponding to the relative probabilities of different interactions, along with a simulation of the MINERvA detector to generate a simulated interaction chain for each neutrino. A description of the MINERvA geometry, modeled in GEANT 4, tells GENIE what materials make up the detector, and where they are positioned.

In order to estimate the magnitude of uncertainty due to the physics models used in GENIE, weights are then also generated to correspond to the change in probability of the given event occurring in the case that a certain input parameter of the physics models (for example, the cross section for an event producing a pion) were increased or decreased. These weights are used to evaluate the effect of model uncertainties on our distributions, as explained in chapter 6.

This section will summarize the physics models used by GENIE, with particular emphasis on those that are most significant to this analysis.

Nuclear model. Chapter 2 explains in detail how important nuclear effects are when scattering from nuclear material, and details several commonly-used nuclear models. GENIE models the nucleus using the Relativistic Fermi Gas model [184] incorporating the Bodek-Ritchie high-momentum tail [59] that simulates short-range correlations. For carbon, the Fermi momentum is taken as $k_F = 0.221$ GeV/ c . A mass density is determined for any nucleus by interpolating between those for which there is data; in this case, a fit is made to a parametrization of data from review articles. Pauli blocking is also applied. A factor is also included to model Bjorken x_{Bj} -dependent effects such as the EMC effect, shadowing and anti-shadowing; however these effects are not relevant to quasi-elastic scattering, where $x_{Bj} = 1$.

Quasi-elastic scattering model. GENIE models quasi-elastic cross sections following Llewellyn Smith's prescription [145], as explained in section 2.2.1. This parametrizes the cross section as a function of the squared four-momentum transfer Q^2 , where the pre-factors depend on the nucleon form factors. Vector form factors are modeled by default using the BBBA05 model [63]. The simulated distributions used in this analysis use BBBA05; however, the Sachs [91] dipole form factor model is used to evaluate the effect of uncertainty in the vector form factor model (see chapter 6 for more details). For the axial vector form factor f_A , a dipole form is used, with $g_A = f_A(0) = 1.2670$ and axial mass $M_A = 0.99$ GeV/ c^2 .

Background models. In section 2.6, we introduced the resonant and deep inelastic scattering (DIS) processes. These processes are typically backgrounds to the quasi-elastic scattering analysis; however, they can also contribute to the quasi-elastic-like cross section through the effect of final-state interactions - in particular, if a pion produced by a resonant

interaction is absorbed, leaving a final-state that mimics a quasi-elastic. Uncertainty in the modeling of these processes also affects the uncertainty on our cross section distributions, as we rely on GENIE’s predictions of resonant and DIS interaction rates when subtracting the fraction of events that we believe to correspond to background processes.

GENIE uses the Rein-Sehgal model [178] to simulate baryon resonance production, which provides cross sections for 16 different unambiguous resonance states. All resonance parameters are updated to the latest best values. As with quasi-elastic scattering, resonant scattering includes an axial mass parameter, M_A^{RES} , which GENIE takes to be 1.12 GeV/c². Resonant events are the dominant background for this analysis.

DIS cross sections are calculated with an effective leading order model with a low- Q^2 modification from Bodek and Yang [60]. Hadronic showering is modeled with the AGKY model [197]. More detail is given in chapter 6 where we discuss how this model affects our systematic uncertainties. The Bodek-Yang model also describes other low-energy non-resonant pion production processes. At the energy range of this analysis, and after our reconstruction cuts, the contamination from DIS is minimal.

Final-state interactions. Section 2.6 discusses the effect of final-state interactions: re-scattering that sometimes occurs as a hadronic interaction product traverses the nucleus. GENIE simulates the re-scattering of nucleons and pions in the nucleus using the INTRANUKE-hA intranucleon hadron cascade package [89]. This works by tracking particles’ progress through the nucleus in steps of 0.05 fm. After each step, a mean free path λ for the hadron in question is calculated based on the local nucleon density (a function of position in the nucleus), $\rho_N(r)$, the hadron’s energy E_h , and the isospin-averaged

interaction cross section for the propagating hadron σ_{hN} :

$$(3.10) \quad \lambda(r, E_h) = \frac{1}{\sigma_{hN} E_h \rho_N(r)}$$

This is used to determine a probability that the simulated particle will ‘interact’; a roll is then made to determine whether or not an interaction will take place at the step in question. If so, an interaction type is determined (absorption, charge-exchange scattering etc) based on their relative cross sections. Once the interaction type is determined, final-state particles will be generated. Each individual interaction product is only allowed to re-interact once; in reality, there could be multiple re-interactions. Where possible, cross sections are determined from data; where no data exists for nucleon or pion scattering at a given material and energy, the cross section must be extrapolated.

3.7.3. Detector simulation with GEANT4

Once GENIE has generated a neutrino scattering interaction, and has produced a set of final-state particles, we use the GEANT4 toolkit [27] v9.4.02 to simulate how these particles propagate through the material of the detector. A detailed model of the detector’s geometry, including coated scintillator strips, absorbers and all the nuclear targets is generated. GEANT4 has an extensive range of physics interaction models for many particles on all types of nuclei, and can thus model particles’ passage through this simulated detector, predicting the energy that will be deposited in each scintillator strip. In MINERvA, the Bertini Cascade model [130] is used to simulate hadronic interactions. The optical and electronics systems are also simulated, which allows this energy deposition to be converted to a simulated readout that can be analyzed as if it were MINERvA data.

3.8. Reconstruction

The calibrated data measurements, and our Monte Carlo simulation, each provide us with the energy depositions generated by (real or simulated) neutrino scattering interactions in the MINERvA detector. The next stage of processing involves running reconstruction algorithms on this data in order to analyze the patterns in the energy deposits, and to identify particle tracks. A general reconstruction procedure is run on all of MINERvA's data and simulation; this does such things as identifying muon tracks and matching them to tracks in MINOS, and identifying energy formed into distinct clusters, as well as counting the dispersed energy. This reconstructed data is saved in the form of ROOT ntuple files, which are then made available for further analysis-specific processing.

As this study analyzes interactions based on the kinematics of a MINOS-matched muon produced in a quasi-elastic interaction, muon reconstruction is of key importance. In order to reconstruct a muon, we must:

- (1) Divide a gate's data into *time slices* corresponding to individual interactions or events (eg a rock muon passing through)
- (2) Identify energy *clusters* within a time slice
- (3) Group clusters to generate *track* candidates
- (4) Identify which track represents the muon, and identify the *interaction vertex*
- (5) Identify muon tracks in MINOS
- (6) Match the MINERvA track to a MINOS track to reconstruct its charge and energy

These steps will be explained below; more detail can be found in [37].

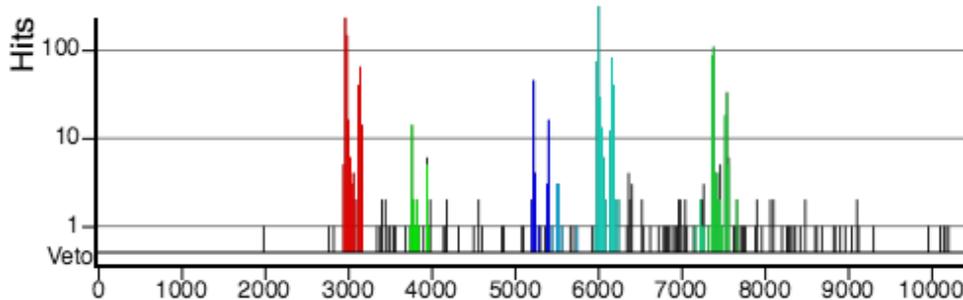
3.8.1. Time slicing

Figure 3.20 shows how the hit profile of a $10\mu\text{s}$ gate is divided up into time slices. The x axis shows the time (in ns) within the gate; the y axis indicates the number of hits. The time slicing procedure attempts to identify groups of hits that occur at approximately the same time, and are thus likely to correspond to a single interaction or other event (such as a rock muon or cosmic ray passing through the detector). In the figure, 12 candidate events are isolated.

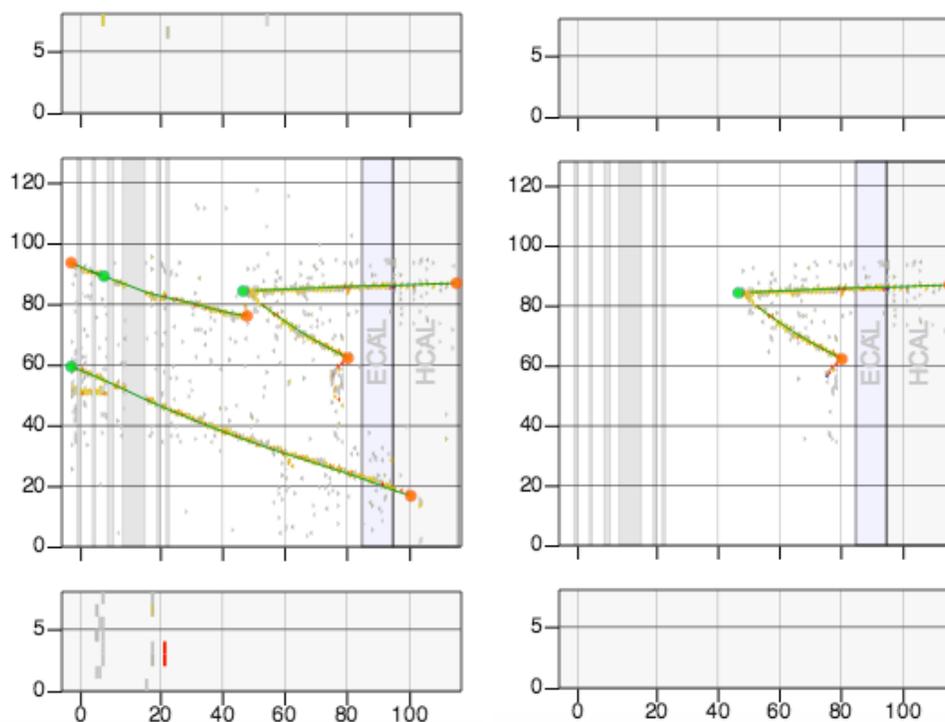
In order to create time slices, an offline time slicing algorithm scans through the spill time in 80ns blocks. When it identifies a block in which sufficient hits (summed over the entire detector) have fired the discriminator to indicate 10 photoelectrons, a new time slice begins. The following 80ns blocks are added to the time slice until the number of hits drops below the 10-PE threshold (hits that do not fire the discriminator are included if they share a Trip-T with hits that do). When the threshold is no longer met, the algorithm continues to scan until it again finds a time block with 10 PE or more, when it opens the next time slice. Note that the time slicing sums all hits across the detector that fire the discriminator - it does not take the spatial distribution of the hits into account.

3.8.2. Clustering

Once a time slice has been identified, the next step is to generate ‘clusters’ of hits that are adjacent in space. A cluster consists of adjacent strips showing activity. Figure 3.21 shows a plane with three example clusters - the presence of non-activated strips between the activity separates the three. The energy-weighted central position of each cluster is



(a) Hits from a single gate divided into time slices



(b) All hits in the gate (X view)

(c) Hits in time slice 1 (X view)

Figure 3.20. The energy hits from a readout gate are divided into time slices, corresponding to individual interactions or other events. Each time slice is indicated by a different color in the top image; black hits are below the energy threshold to be included in an event time slice. (This is data run 2160 subrun 1, gate 594). The event displays below show the energy deposited in the X view for the entire gate (left) and from the first slice only (the red time slice in the top image) (right). These images are from MINERvA's event display utility, Arachne. [187]

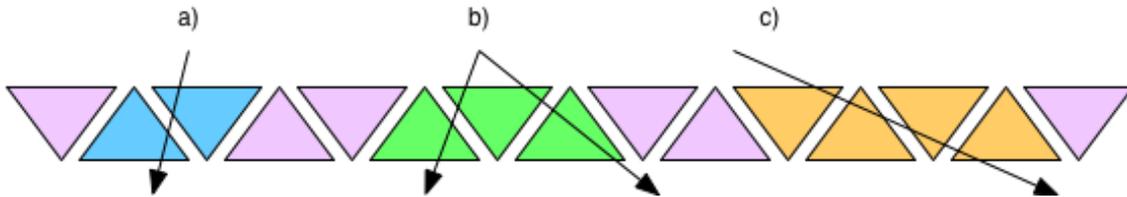


Figure 3.21. Three example clusters through a plane of scintillator strips. a) Due to the triangular structure, even particles traveling perpendicular to the plane will almost always pass through at least two strips. b) Multiple particles from the same interaction may cause a larger cluster. c) A single particle may traverse several strips if it moves at a large angle from the beam direction.

found; its time is set to the time of the highest-energy hit. Clusters are then divided into 5 categories:

- (1) *Low energy*: less than 1 MeV
- (2) *Trackable*: 1-12 MeV total, maximum four hits, one (or more, if adjacent) hit with more than 0.5MeV
- (3) *Heavily ionizing*: not trackable, but with over 1MeV total, and with 1-3 (adjacent) hits with more than 0.5 MeV
- (4) *Supercluster*: more than 1 MeV (but not trackable or heavily ionizing), or with 5 or more hits
- (5) *Cross talk*: composed of hits on PMT pixels next to pixels associated with a particle interaction, suggesting it has been generated by optical or electronic cross talk, rather than by an actual particle

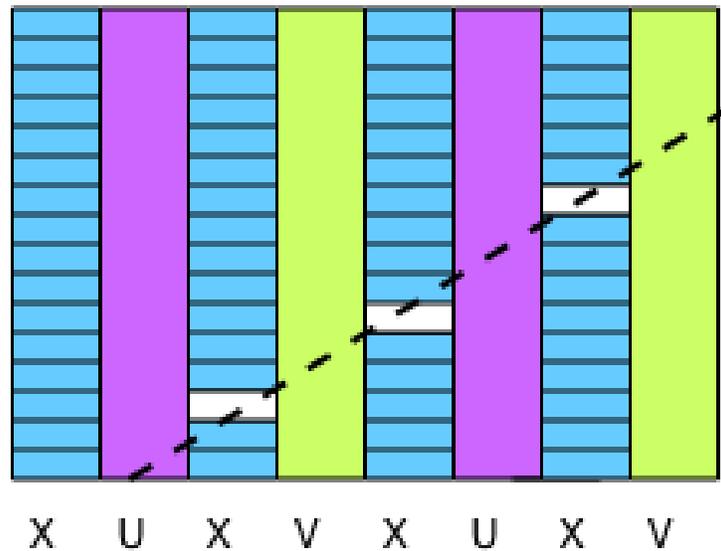


Figure 3.22. An example of a track seed, made from three trackable clusters, in consecutive X planes, and forming a straight line

3.8.3. Track reconstruction

Clusters are formed into tracks, three dimensional objects representing particles' trajectories through the detector. The first step of the track reconstruction algorithm is to combine a time slice's clusters into 'seeds':

- Three trackable or heavy ionizing clusters
- ...in consecutive planes of the same (X,U or V) orientation
- ...that fit to a straight line

Note that this limits the ability to reconstruct very sideways-going seeds, which will not have enough of a forward component to form clusters in three different planes of the same orientation. This leads to an inability to reconstruct tracks with angles between 70° and 110° to the beam direction. This restriction does not affect this analysis, as the only track we require is a MINOS-matched muon; to reach MINOS, muons must be far

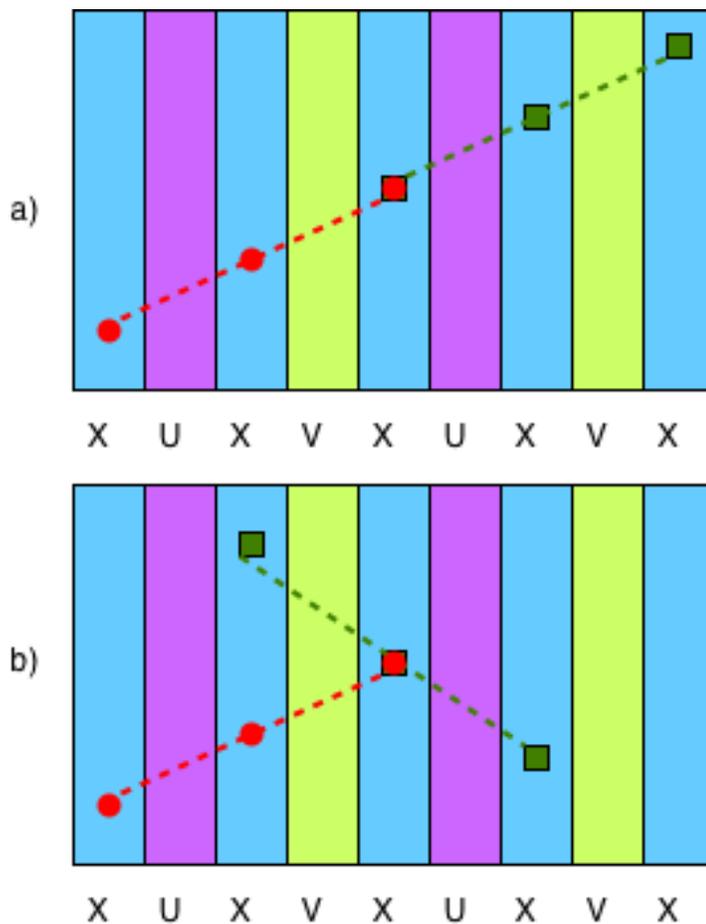


Figure 3.23. In example a) two track seeds combine to make a track candidate. In b), despite sharing a cluster, the seeds cannot be combined to a track candidate.

more forward-going, and we are limited to a maximum angle of around 20° to the beam direction.

Multiple seeds in the same orientation are merged into a track candidate if they contain a common cluster, have a consistent direction, and don't include more than one cluster in the same plane (see figure 3.23). A given seed can only be in one track candidate. Once candidates are built, they can be merged according to the same criteria as seeds. This

allows a track candidate to continue to follow its trajectory even if there are dead pixels in one of the planes along the track.

Once the track candidates are formed in each plane, they are formed into 3-dimensional tracks using two algorithms. The first looks for combinations of three tracks, one from each orientation (X, U and V) that overlap in the z direction and fall along the same three-dimensional line. (Note that, because of this requirement, the shortest possible track that can be generated by this method is eleven planes long. The muon tracks characteristic of the quasi-elastic interactions sought in this analysis must be longer than this, in order to satisfy the requirements that they both originate in the fiducial volume, and exit the back of the detector in order to continue to MINOS.)

After this algorithm is run, a second algorithm attempts to form the remaining pairs of the candidates, with different plane orientations (X/U, X/V, or U/V) into three-dimensional lines, and to fill in these lines with untracked clusters. If sufficient clusters fit to the line, it is promoted to a track. This method can find shorter tracks (nine planes), as well as tracks that are obscured in one orientation.

Once tracks are identified, they are fitted with a Kalman filter routine, and additional untracked clusters (including superclusters) in planes adjacent to the track. This allows tracks to intersect or to be extrapolated through areas of high activity (such as a hadron shower) by sharing the energy of superclusters they encounter.

Once tracks have been created, the longest track of over 25 clusters is designated the ‘anchor track’. The downstream end of this track is designated the primary vertex. For the interactions studied in this analysis, the anchor track corresponds to the muon, and its origin is the neutrino interaction vertex. The vertex identification is to be improved by

identifying other tracks emerging from the vertex point, and using those to re-estimate its exact position. Note that for this analysis, we reject all events in which any tracks other than the muon are identified.

Once it has been found, the anchor track, is ‘cleaned’ to ensure that it emulates the minimum-ionizing particle behavior expected of a muon, by removing superclusters and extra energy near the vertex that are inconsistent with known muon characteristics.

MINERvA also includes a short tracker algorithm; however this is not used to identify muons, the only tracks that are of interest to us.

3.8.4. Charge identification and energy measurement with MINOS

For this analysis, we require that an interaction produces a muon that can be matched with a track in MINOS. As explained in section 3.3.7, MINOS is situated downstream of MINERvA, and is magnetized, enabling us to identify the muon’s charge. This is vital to filter out negative muons produced by neutrino contamination of the antineutrino beam. MINOS’s magnetic field causes a muon’s track to be deflected. When the data for this analysis were taken, the NuMI horns were set up with reversed current, to produce an antineutrino-enhanced beam. Accordingly, during this period, MINOS’s coil was connected in such a direction that positive muons would be deflected towards the coil, focusing them and providing them with a maximum trajectory through the detector. Negative muons, conversely, would be deflected away from the coil.

In MINOS, as in MINERvA, time slices are selected by looking at hits clustered in space and time. The hits in a given time slice are then formed into clusters, which are grouped into tracks if their positions are correlated. The track’s path is the estimated

using a Kalman filter; unlike in MINERvA, MINOS tracks curve due to the detector's magnetic field. The track fitting algorithm outputs an estimates of the track's momentum calculated from its curvature, and an estimate calculated from its range [86].

MINOS tracks are matched to MINERvA muons when activity is measured in the last five planes of MINERvA, and a track starts in the first four planes of MINOS within 200ns of the MINERvA track time. The MINERvA track is extrapolated forwards to where it would intercept the first MINOS plane, and the MINOS track is extrapolated back to the last plane of MINERvA. If, in each case, the extrapolated track intercepts within 40 cm of the track in the other detector, the tracks are considered a match. Failing this, tracks may be matched if the point of closest approach between the two tracks is within 40cm. (This was of particular value in early measurements, not used in this analysis, for which the ArgoNEUT [44] detector was positioned between MINOS and MINERvA; muons could scatter from the material of ArgoNEUT, altering their direction).

As mentioned in 3.3.7, the upstream section of MINOS is instrumented as a calorimeter (see figure 3.8), while the downstream part serves only as a tracking spectrometer. For muons that stop in the calorimeter region, their energy can be measured by summing the total energy deposited, using the method developed by MINOS and explained in [12]. This 'range' measurement leads to an uncertainty of 2% on the momentum.

For muons that do not stop in the calorimeter region, the momentum must be measured by the track's curvature K in the magnetic field, which is the inverse of its radius of curvature R , in cm. This is related to the momentum by

$$(3.11) \quad K \equiv \frac{1}{R(cm)} = \frac{0.3B}{P}$$

where B is the magnetic field in kGauss and P is the component of momentum perpendicular to the magnetic field, in MeV. The curvature method is less precise than the range method, adding an additional 2.5% uncertainty on the momentum for muons with momenta below 1 GeV and 0.6% for those with higher momenta, to be added in quadrature to the range uncertainty, as stated in [37].

3.8.5. Recoil energy reconstruction

We refer to final-state energy not associated with the lepton (for this analysis, the μ^+) track as “recoil energy”. The recoil system can include non-muon tracks as well as non-tracked energy deposits. For this analysis, we reject all events with tracks other than the muon (as our signal is for events resembling the $\bar{\nu}_\mu p \rightarrow \mu^+ n$ final state, where we allow only the muon, plus neutrons and sub-120MeV kinetic energy protons in the final state; neutrons are uncharged, and the protons allowed in our signal have too little energy to generate tracks). We do, however, look at the summed recoil energy, and make a Q^2 dependent cut on this in order to identify quasi-elastic-like events. The procedure for measuring the recoil energy is summarized below, and explained in detail in [37].

The recoil energy - all the energy not associated with the muon track - is summed calorimetrically across the detector, using the formula

$$(3.12) \quad E_{\text{recoil}}^{\text{cal}} \equiv \alpha \sum_i C_i^{sd} E_i$$

This is a weighted sum of the energy E_i deposited in each sub-detector i . The energy associated with a given interaction is the total sum of all clusters except those due to cross talk and those associated with the muon track. The clusters must have a hit time

between 20 ns before and 35 ns after the first hit on the muon track (used to identify the interaction time).

Each sub-detector is weighted by a calorimetric constant C_i^{sd} , in order to take account of the different materials therein. (For this analysis, we look at the recoil in both the scintillator tracker and in the electromagnetic calorimeter (ECAL), which consists of scintillator and lead absorber). The calorimetric constants have the form

$$(3.13) \quad C^{sd} = \frac{E_{\text{abs}} + E_{\text{scint}}}{f \times E_{\text{scint}}}$$

Where E_{abs} and E_{scint} are the energy lost in a single plane of absorber and scintillator respectively when a minimum-ionizing particle travels through it perpendicular to the plane, and f is the active fraction of the scintillator plane. The constants measured for scintillator, ECAL and HCAL are $C^{scint} = 1.22$, $C^{ECAL} = 2.013$ and $C^{HCAL} = 10.314$.

The overall scale α in equation 3.12 is not relevant to this analysis, as we use the simulation to assess the recoil energy, and apply the same reconstruction cuts to simulation and data.

Uncertainties on the recoil energy are complex, as they depend on the detector composition, the scintillator's response to different particles, and the response of the optical and electronics systems. We estimate the uncertainties by using a test beam to measure the response of a MINERvA plane to different particles [76]. More information about recoil uncertainties and how they affect this analysis is included in section 6.7.

3.8.5.1. Detector resolution. We can use the simulation to evaluate our detector's resolution. To do this, we subtract the reconstructed value of a quantity in the simulation (such as muon angle, or vertex position in a given direction) from the true value generated

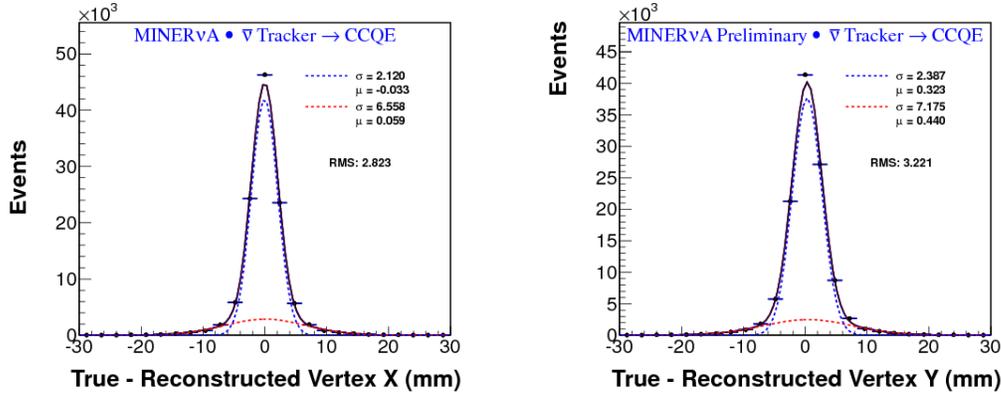


Figure 3.24. True - reconstructed x (left) and y (right) vertex positions , for antineutrino CCQE candidates

by the Monte Carlo generator, for each interaction. Making a histogram of these plots and fitting it (the shape fits well to a double Gaussian) allows us to get an estimate of the resolution from the width of the central Gaussian. Figure 3.24 shows an example of how the vertex position resolution in the $x - y$ plane is estimated. The plot in x has a width of 2.120 mm; that in y has a width of 2.387 mm. Thus we can conclude that the vertex position resolution is less than 3 mm in each direction. (Note that the z resolution is far less simple, as it is limited by the plane structure of the detector.)

Approximate values of the resolution in various quantities are shown in table 3.3. For detailed plots and tables of resolutions, see [103].

Quantity	Resolution
Vertex x	3 mm
Vertex y	3 mm
Vertex z	10 mm
Muon angle	0.5°
Muon energy	0.1 to 0.6 GeV (best resolution at low energies and angles)

Table 3.3. Approximate values of the resolution of various quantities, for a CCQE antineutrino sample

CHAPTER 4

Charged-current quasi-elastic event selection**4.1. Signal definition****4.1.1. The CCQE interaction**

This analysis measures the charged-current quasi-elastic cross-section for anti-neutrinos incident on the plastic scintillator (CH) tracker region of the MINERvA detector. In the free-nucleon case, this interaction would be represented by:

$$\bar{\nu}_\mu + p \rightarrow \mu^+ + n$$

The interaction is explained in detail in section 2.2. A true charged-current quasi-elastic (CCQE) antineutrino scattering interaction should produce a characteristic signature: an outgoing positive muon plus a recoil neutron. Thus, in general, our reconstruction should attempt to identify final states with this configuration, while rejecting events from other processes.

Specifically, we need to remove resonant events, wherein a pion (or kaon) is produced from the decay of a resonant state such as one of the Δ_{1232} resonances, spin $\frac{3}{2}$ states comprising u and d quarks, as shown in figure 4.1a. We also attempt to reject deep inelastic scattering, where scattering off one of the component quarks in the nucleon produces a hadronic shower. (See section 2.6 for an explanation of these processes).

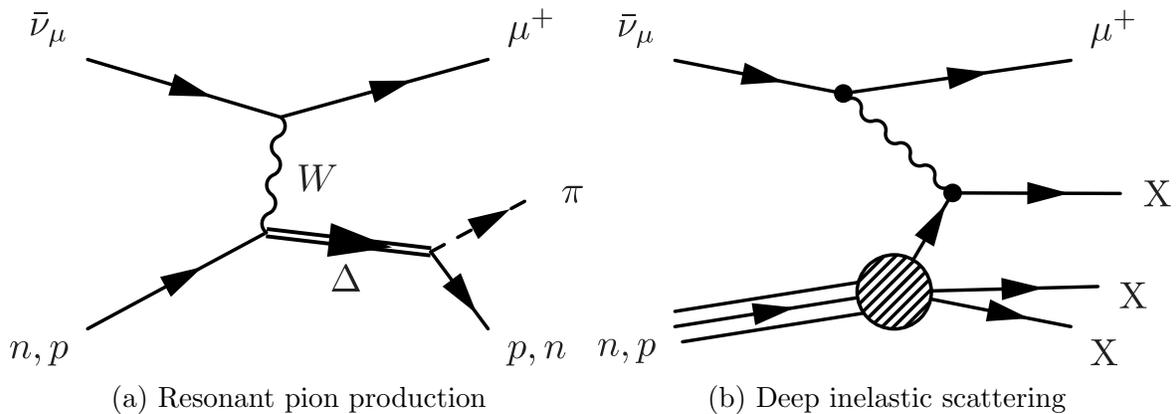
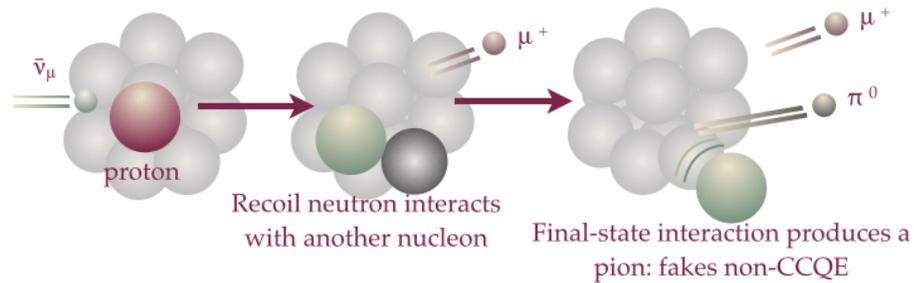


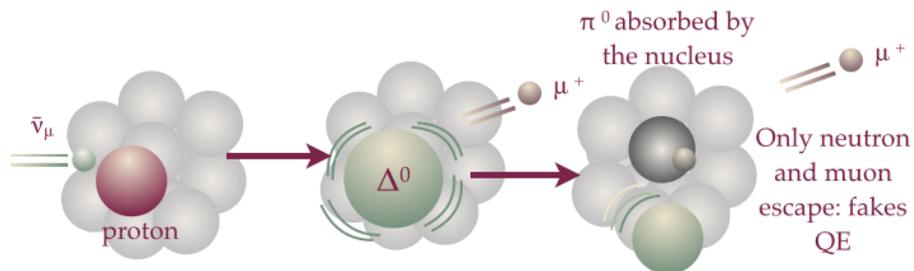
Figure 4.1. Background processes - non-quasi-elastic charged-current neutrino scattering

However, we know from electron scattering experiments (and see hints from our own previous neutrino-scattering results) that correlated pairs of nucleons can exist within the nucleus (see section 2.5). A quasi-elastic scatter from a correlated pair can cause the ejection of additional nucleons. The situation is further complicated by the possibility of final-state interactions, hereafter referred to as FSI, (figure 4.2), in which hadrons produced in an initial interaction may re-interact as they propagate through the nucleus. In this way, for example, a neutron produced in a quasi-elastic interaction could produce a pion by interacting with another nucleon as it exits the nucleus (figure 4.2a). Conversely, a pion produced in an initial resonant interaction may re-interact and become absorbed within the nucleus, leading a quasi-elastic-like final state of just the muon and neutron (figure 4.2b). There is thus no direct one-to-one correlation between final states (which we can attempt to detect) and the initial interaction type that we are attempting to identify. (The theory behind final-state interactions is discussed in section 2.6.)

As the definition of what constitutes a quasi-elastic interaction on a heavy nuclear target is somewhat ambiguous, our goal is to define a signal which:



(a) Resonant process with a QE-like final state



(b) QE process with pion in the final state

Figure 4.2. Final-state interactions

- Corresponds as well as possible to the CCQE process
- Can be identified in the MINERvA detector with well-understood efficiency (the fraction of signal events that can be successfully reconstructed) and purity (the fraction of reconstructed events that correspond to signal, rather than background)
- Is clearly defined, so that theorists can use models to generate cross section predictions for the specified signal definition

Bearing this in mind, we consider the pros and cons of alternative definitions, as listed below.

4.1.2. True CCQE signal definition

For MINERvA’s initial CCQE cross-section publications [96][107], we measured the cross section for interactions that were initially quasi-elastic (that is, no resonant or deep inelastic scatters), regardless of the final-state particles produced. Thus, a quasi-elastic scatter in which the resulting neutron undergoes FSI, producing a pion, would be considered a signal event under this definition. A resonant event in which the pion is absorbed in the nucleus would be a background event, as would resonant events that produce low-energy pions not removed by the recoil energy cut.

This signal definition has the advantage of including contributions only from one physical interaction process (CCQE, with no contribution from resonant or DIS events). Its disadvantage is that it has no defined final-state signature: signal events may or may not, for example, include pions in the final state. When reconstructing events, the only information available to us is the composition of this final state. We can determine which events in our simulation should be considered signal by looking at those in which GENIE tags as CCQE, and in turn use this to correct our data distribution. However, this is very dependent on GENIE’s CCQE model - a model which our work in [96] and [107] has already shown to be incomplete.

4.1.3. Quasi-elastic-like signal definition

“Quasi-elastic-like” events are those whose final-state signature matches that of a quasi-elastic scatter. This definition, however, is open to interpretation. For a free-nucleon antineutrino quasi-elastic scattering event, the final state would include a positively-charged

muon and a single neutron. However, when scattering from a heavy nucleus, the possibility of nucleon pair correlations in the initial state leads us to suspect that a neutrino scattering quasi-elastically from a correlated pair could eject both partners [100]. If the correlations observed in electron-nucleus scattering are correct, this would lead us to expect that we will scatter from a single proton (leading to a neutron in the final state) around 80% of the time, from an np pair (leading to a possible second final-state neutron) 18% of the time, with a small fraction of 1-2% of scattering events from pp pairs, which could produce a low-energy proton. Thus there is an argument for including events with an additional nucleon in the final state under the definition of quasi-elastic-like.

Furthermore, there is a historical basis for a quasi-elastic-like definition of $CC0\pi$ - that is, a final state of an appropriately-charged muon, plus nucleons. No other types of hadrons are allowed; neither are photons. In this case, backgrounds include resonant events producing a pion that is not caught by the recoil cut; however, resonant events where the pion is absorbed will now become signal, even if they contain multiple high-energy nucleons (pion ‘absorption’ produces a minimum of two nucleons in addition to that already generated by the resonant event; thus events of this type will have three or more nucleons in the final state). This signal definition made sense for the equivalent measurement at MiniBooNE[30], whose mineral oil Cherenkov detector, while able to detect muons and pions, was not sensitive to nucleons. (It was also used for T2K’s CCQE-like neutrino-scattering measurement [5]. T2K has only recently begun taking antineutrino data, and has not yet amassed sufficient statistics to analyze antineutrino scattering cross sections). MINERvA’s scintillator tracker, however, is able to resolve

proton tracks above a kinetic energy of 110 MeV [191], and to detect the energy deposited by charged particles.

Our principal signal (that from an antineutrino scattering quasi-elastically from a non-correlated proton) has a signature of a single neutron (and positive muon). The electron-scattering data suggests that the majority of events will be of this type, with most antineutrino-correlated pair scatters also giving a neutrons-only final state. Therefore, our cuts preferentially select for events with this signature. For this study, we define our signal for quasi-elastic-like events to be $CC0\pi$ with no protons above a measurable kinetic energy threshold. We take this threshold to be 120MeV. Protons with energies below this threshold will be contained within our exclusion region around the vertex, for which we do not measure the recoil energy, as explained in section 4.3. The energy threshold's value was selected by optimizing the product of purity and efficiency, and will be explained in section 4.3.

As we are measuring a cross section on scintillator, we require the interaction to have taken place within the detector's fiducial volume: the tracking region of the detector that is formed of scintillator strips (see figure 3.4).

Thus our quasi-elastic-like signal definition is antineutrino scattering events:

- whose interaction vertex is in the fiducial volume
- with a final-state μ^+
- with any number of final-state neutrons
- with any number of final-state protons under 120MeV of kinetic energy
- with no other mesons, photons (de-excitation photons, which have energies less than 10MeV, are allowed), or heavy baryons

4.2. Data and simulation samples

4.2.1. Run periods

For the anti-neutrino double-differential analysis, we use data taken from the “minerva5” playlist. This was taken in MINERvA’s low-energy run, with the reverse horn current setting that favors antineutrinos. It corresponds to MINERvA runs 2650 to 2856, recorded between November 5 2010 and February 24 2011. The data was taken in the low-energy beam mode with a peak neutrino energy of around 3 GeV, as shown in figure 3.1. Simulated data (hereafter referred to as Monte Carlo) was taken from the ”minerva5” Monte Carlo playlist, with run numbers from 50200-50249. Data and Monte Carlo were processed using the v10r8p8 (Eroica) version of the MINERvA analysis framework and reconstructed using the CCQEAntiNuTool analysis tool. The data sample corresponds to 1.041×10^{20} protons on target (POT), the total number of protons delivered by the Main Injector accelerator to the NuMI beam target over the entire run period. The Monte Carlo simulation corresponds to 9.515×10^{20} POT; in this case the protons are, of course, simulated, rather than delivered by the actual accelerator.

Many of the plots in the following sections will compare data and simulation. As the Monte Carlo simulates more protons than were actually delivered in our data run period, it is natural that it will also produce correspondingly larger event counts. We thus show *POT-normalized* distributions when comparing data and simulation; this means that the simulation is scaled by the ratio of the data POT to the simulated Monte Carlo POT.

4.2.2. Data and simulation ntuple formats

Both data recorded by the detector and simulated data generated by the GENIE event generator [93] is stored in the ROOT [64] file format. These files contain tree-like structures known as *ntuples*, which store information about each interaction event. We generate four types of ntuple files:

- **Data** ntuples, with reconstructed information about real interactions in the detector
- **Reconstructed Monte Carlo** ntuples, with reconstructed and generator-level information about simulated interactions
- **Truth** ntuples, with generator-level information about all simulated interactions, including those that we could not reconstruct (for example where we were not able to reconstruct a MINOS-matched muon, so could not measure the properties that depend on muon kinematics). The reconstructed Monte Carlo ntuples contain only the subset of the events in the Truth ntuples that we were able to reconstruct.
- **Meta** ntuples, which contain information about the contents of the other ntuples, such as the number of POT

For simulation, it is important to recognize the difference between *true generated* quantities, and *reconstructed* quantities. For example, we may simulate an interaction that generates a 3.8 GeV muon. However, due to the resolution of the detector and the complexity of our reconstruction algorithms, when we reconstruct that muon's track, we may estimate it to have an energy of 3.7 GeV. When we reconstruct simulated interactions, we store both the true quantity and the reconstructed quantity in the ntuples. We can then use this simulation to measure how the distributions we reconstruct differ from

Ntuple	Data	Reconstructed MC	Truth	Meta
Reconstructed quantities	x	x		
Shifted reconstructed quantities for evaluating systematic uncer- tainties		x		
True quantities from GENIE		x	x	
Weights for evaluating system- atic uncertainties		x	x	
POT information				x

Table 4.1. Information contained in the different ntuples

what actually occurred in the (simulated) detector, and calculate corrections that can be applied to our data. This procedure is described in detail in chapter 5. The reconstructed Monte Carlo tuples also include additional information that is used to evaluate the effect of systematic uncertainties on our distributions (for example, the effect of changing pion production cross section by 20%, or of underestimating the detector’s energy response to protons by 4%). The procedure for evaluating these is explained in detail in chapter 6.

The ntuple variables used in the analysis are listed in appendix C.

4.2.3. Weighting simulated events

The Monte Carlo is designed to simulate nature as well as possible. To avoid generating more simulation as our knowledge of nature improves, we use a reweighting technique to adjust our simulated distribution in line with the latest external data. When we histogram event distributions in simulation, we weight its contribution to the histogram by a number corresponding to the relative probability of such an event occurring given our latest data constraints, compared to the probability that it would have occurred if nature had been

exactly as predicted when we generated the original simulation. The weight we apply is a product of

- Flux weight, constrained by neutrino-electron scattering
- Non-resonant single pion production weight
- Muon tracking efficiency correction

The simulated events were generated using a simulation of the NuMI beam's energy spectrum. Since these events were generated, MINERvA studied the neutrino-electron scattering cross section [167]. To obtain a **flux weight**, we use alternative flux histograms constrained by the neutrino-electron scattering event rate. As the cross section for this process is well known, we are able to extract a data-constrained flux measurement from our event distribution for the process, as explained in section 3.7.1.3. To apply this constraint, we look up the flux measurement for a simulated event's true neutrino energy as measured with the $\nu - e$ constraint, and divide by the flux for the true neutrino energy in the flux histograms used to generate the simulation to get a weight that we can apply to our simulated event.

A fit to neutrino-deuterium scattering data from bubble chamber experiments [179] constrains the cross section for **non-resonant single pion production** events, which is found to be just $43 \pm 5\%$ of that estimated by GENIE. Thus we apply a weight of 0.43 to all events that correspond to this interaction type. This affects only 3% of our selected sample, as we attempt to remove events that produce pions.

A reanalysis of our **muon tracking efficiency** in MINOS and MINERvA [56] found that our simulation slightly overestimated our ability to reconstruct muons. This study was done by looking at the fraction of the tracks exiting the back of MINERvA and

heading towards MINOS’s fiducial volume that are successfully matched to a MINOS track. We correct for this by weighting our simulation to account for this inefficiency. For the ‘minerva5’ dataset used in this analysis, MINERvA’s tracking efficiency is corrected by a factor of 99.5%, while MINOS tracking efficiency is corrected by a factor of 97.5% for muons with momenta below 3 GeV and 99.5% for those with higher momenta.

4.3. Event selection

4.3.1. Selecting reconstructable events

Event reconstruction is performed by the CCQEAntiNuTool analysis tool, which is written in C++ and run using the MINERvA analysis framework, which is built around the the GAUDI software architecture system [51]. In order to identify a reconstructable event, we require good MINERvA data with good MINOS tracks; no DAQ readout or timing errors, and good quality beam data. Additionally, we require a muon “prong” (a MINERvA track candidate), matched to a MINOS muon track, with an interaction vertex in the MINERvA detector within an apothem of 900 mm from the x-y center of MINERvA and with a z coordinate between 4500 mm and 9050 mm (a larger region than that which will be used for our final analysis cuts).

For Monte Carlo, the MuonIsPlausible() check ensures that the MINOS-matched muon has at least half its hits from Monte Carlo, rather than from the data overlay used to protect against pile-up effects.

This basic selection includes far more events than will be selected by our final cuts. This allows for such possibilities as observing the neutrino background in the anti-neutrino

sample, adjusting our final cuts in order to evaluate systematic effects, and seeing interactions occurring outside the main tracker volume.

4.3.2. Charged-current antineutrino event selection

To identify charged-current antineutrino scattering events, we impose the following reconstruction cuts on our data:

- MINOS-matched μ^+
- Upstream no dead-time cut
- Fiducial volume cut

4.3.2.1. MINOS-matched positive muon. Because the MINOS detector has a magnetic field, it can be used to identify the charge of a muon, based on its curvature in the field. For our anti-neutrino study, we require a reconstructed muon track that can be matched with the track of a positively-charged muon in MINOS (see section 3.8.4 for details of the matching process). Note that this limits our acceptance for events in which the produced muon has a high transverse or low longitudinal momentum as the geometry of the MINOS and MINERvA detectors requires that the muon must be able to travel at least 2 meters beyond the back of MINERvA, and that having done so, it must be incident on the cross sectional area of the MINOS detector (see figure 4.3).

In spite of this limitation, the MINOS matching cut is valuable because it virtually eliminates contamination of the sample with events generated by incident neutrinos (as opposed to anti-neutrinos), which would produce a μ^- instead of a μ^+ . Because the NuMI beam is generated with protons incident on a matter target, it produces more positively than negatively charged mesons, and these decay to give more neutrinos than

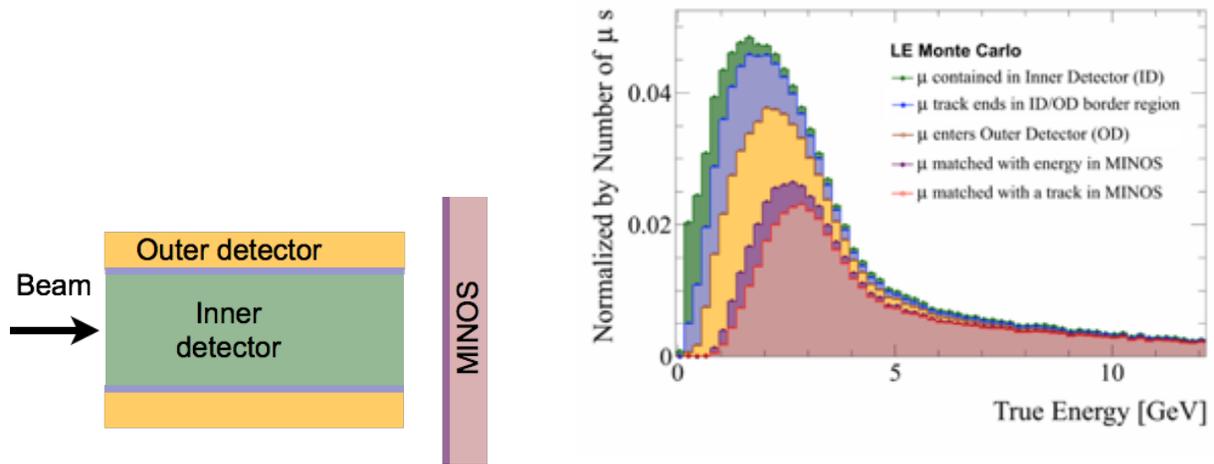


Figure 4.3. Endpoints of muon tracks generated in MINERvA interactions, plot courtesy of J. Ratchford

antineutrinos. Even with the horns set up to preferentially focus the mesons that will produce antineutrinos, there is still a significant contamination of the anti-neutrino beam by neutrinos. The problem is especially severe at higher energies, where the low transverse momentum of the mesons produced at the proton target makes it difficult for the horns to defocus them.

4.3.2.2. Dead-time cut. MINERvA can experience some dead time (see section 3.5.2) after an event has been recorded. A consequence of this can be that the upstream part of a track may not be detected. This can be especially problematic in the case of rock muons - that is, muons produced when neutrinos from the beam interact in the rock upstream of the MINERvA detector, leading to a muon track that enters MINERvA at the first module and continues through the detector. If dead time leads to the upstream part of one of these tracks not being detected, it will appear as if this muon track started part way through the detector, mimicking the signal of a CCQE event.

In order to protect against this, a cut is made on events in which the quantity "tdead", as calculated by the DeadTimeTool is less than two. This quantity is calculated by projecting the muon track upstream by two modules and checking whether the electronics corresponding to each of the strips intersected by the projection or their immediate neighbors were experiencing dead time. If two or more of these strips were in dead time during the event, the event is rejected.

4.3.2.3. Fiducial volume cut. As this analysis is measuring a cross section for scattering on scintillator, we require the reconstructed interaction vertex to be within the fiducial volume of our detector (the scintillator tracker in figure 3.4). In MINERvA coordinates, this means that the vertex must be within a hexagon of apothem 850 mm, and with a z position between 5980 and 8422 mm (corresponding to modules 27 to 80, inclusive).

4.3.3. Quasi-elastic event selection

Once we have identified a selection of good quality, charged-current antineutrino scattering interactions, we must then make cuts to attempt to identify a CCQE-like signal - that is, events with only neutrons and low-energy protons in the hadronic final state. In order to select these events, we look at the recoil energy - the non-muon energy deposited in the detector - as the recoil distribution differs between CCQE and non-CCQE events. Section 3.8.5 details how the recoil energy is reconstructed and calibrated.

The quasi-elastic selection is explained in the following sections, and can be summarized as:

- No non-muon tracks
- Q^2 -dependent recoil energy cut

- 15 GeV maximum longitudinal muon energy

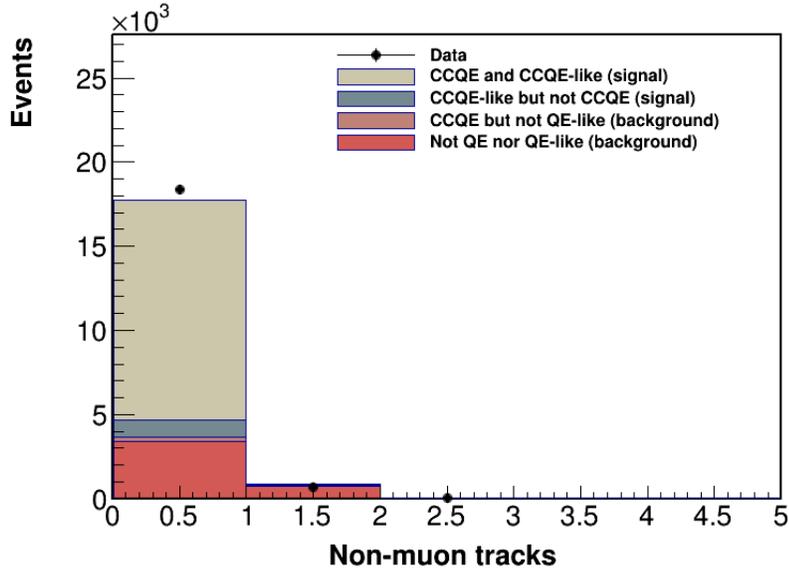


Figure 4.4. Number of non-muon tracks in data (points) and simulation (colors, POT-normalized to data). In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red). All cuts except the track cut are applied.

4.3.3.1. Track cut. As our signal allows only neutrons and low-energy protons (below the tracking threshold of 450 MeV) in the hadronic final state, we require that no tracks (other than the muon track) are present. Figure 4.4 shows the effect this has on the distribution, removing a small number of background events with an extra track, due to either pions or high-energy protons. The effect is modest once the recoil cut has been applied.

4.3.3.2. Q^2 -dependent recoil energy cut. Charged pions and high-energy protons do not always leave reconstructable tracks; they do, however, deposit clusters of energy in the

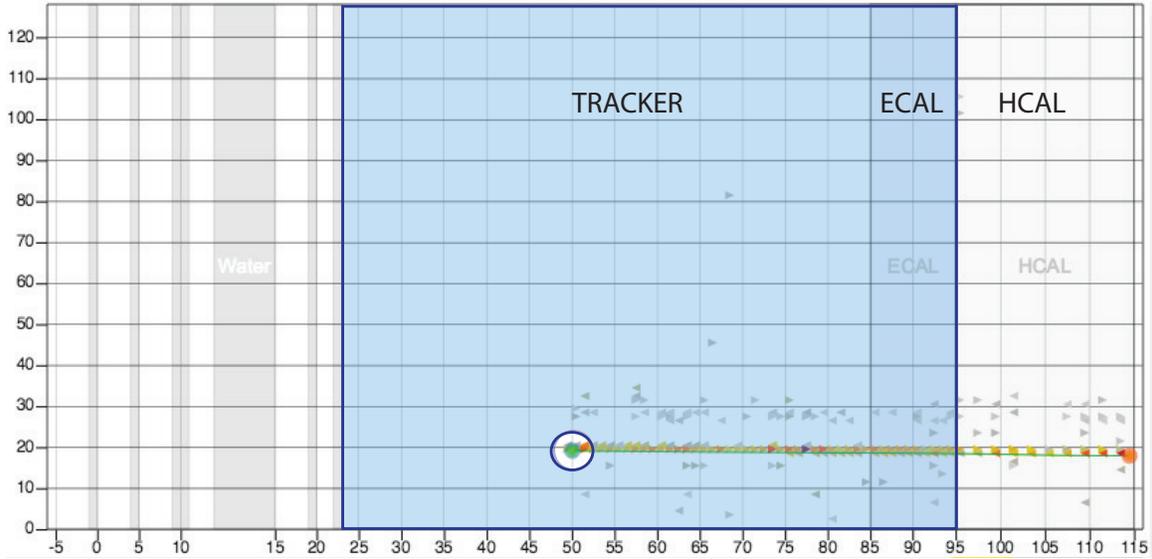


Figure 4.5. Recoil energy region of the MINERvA inner detector (shown in blue).

detector. We thus look at the total calorimetrically measured recoil energy in the detector. The recoil energy is defined as energy from clusters in the tracker or electromagnetic calorimeter regions of the detector (except for low-activity and cross-talk clusters) that are not associated with the muon track, and which have a timestamp between 20 ns before and 35 ns after the timestamp of the muon vertex. When summing the recoil energy, we exclude a sphere of radius 100 mm around the interaction vertex (defined as the start point of the muon track). This is because we know that 2p2h effects, which are not modeled in the current version of our simulation, can produce additional low-energy nucleons; we want to keep these events in our selected sample, as one of our goals is to measure the strength and form of these kinds of interactions. We also know that the efficiency of any cut on energy close to the vertex would have poorly-modeled efficiency, not only because we do not model 2p2h effects, but because final-state interactions, also generate low-energy particles that will deposit energy near the vertex; our simulation does

model these effects, but they are not well understood, and the model may not reproduce true particle distributions.

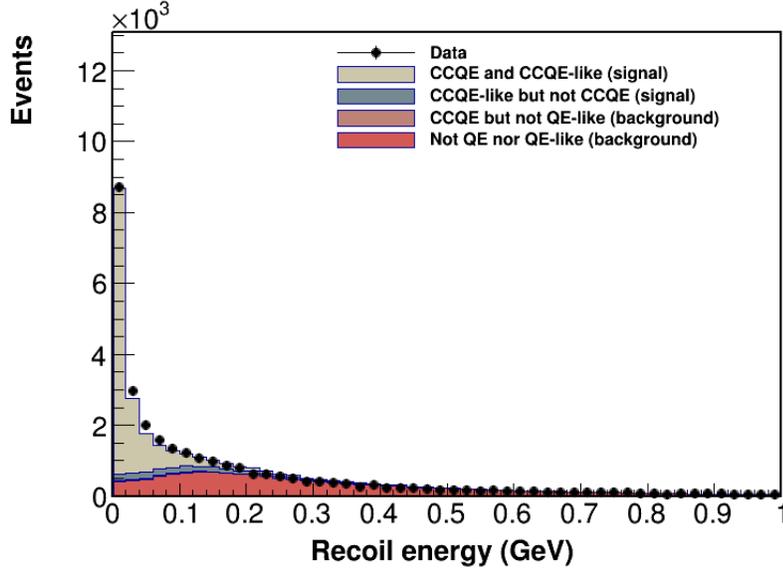


Figure 4.6. Recoil energy for data (points) and simulation (colors, POT-normalized to data). In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red). All cuts except the recoil cut are applied.

To identify CCQE events, we consider the distribution of recoil energy versus the squared four-momentum transfer, Q^2 . As we do not have the ability to measure Q^2 directly in our data, we instead reconstruct it from muon kinematics using the quasi-elastic hypothesis:

$$(4.1) \quad Q_{QE}^2 = 2E_\nu^{QE}(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2$$

$$(4.2) \quad E_{\nu}^{QE} = \frac{m_n^2 - (m_p - E_b)^2 - m_{\mu}^2 + 2(m_p - E_b)E_{\mu}}{2(m_p - E_b - E_{\mu} + p_{\mu} \cos \theta_{\mu})}$$

Where E_{ν} and E_{μ} are the neutrino and muon energy respectively. Muon momentum is represented by p_{μ} , and θ_{μ} represents the angle between the outgoing muon and the incoming neutrino. The proton and muon masses are m_p and m_{μ} respectively and E_b is the binding energy of 30 MeV. Note that this assumes a CCQE interaction on a non-correlated nucleon at rest: for interactions with correlated pairs of nucleons, and for QE-like interactions that were not initially CCQE (such as resonant pion production followed by pion absorption), this hypothesis will not correctly reconstruct the true neutrino energy or Q^2 . In the rest of this thesis, E_{ν}^{QE} or Q_{QE}^2 will refer to the quantities reconstructed with this hypothesis, while E_{ν} and Q^2 will refer to the true (but unmeasurable in our detector) values.

For the previous 2013 antineutrino CCQE analysis at MINERvA [96], we required that recoil energy (in GeV) must be less than $0.03 + 0.3Q_{QE}^2$ (where Q_{QE}^2 is in GeV^2) and that the total recoil energy be less than 0.45 GeV. This cut was chosen to maximize the product of efficiency and purity for the sample. Figure 4.7 shows the distribution of signal and background events in the 2013 analysis, relative to the cut.

For the 2013 MINERvA CCQE analysis, the signal definition was events that were defined by GENIE as having a primary interaction type of CCQE. For this analysis, with the CCQE-like signal definition, we find that the efficiency for CCQE-like events that are not CCQE is poor due to the recoil cut: just 17% for $CC0\pi$ events that are not CCQE (see figure 4.8). (Note: *efficiency* refers to the fraction of signal events that we were able to reconstruct, and is a measure of how well our cuts manage to identify signal events.

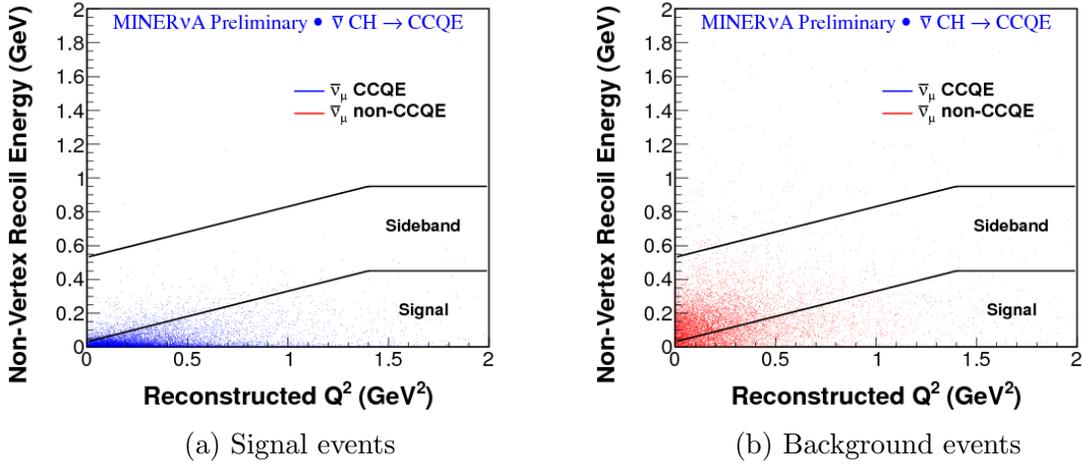
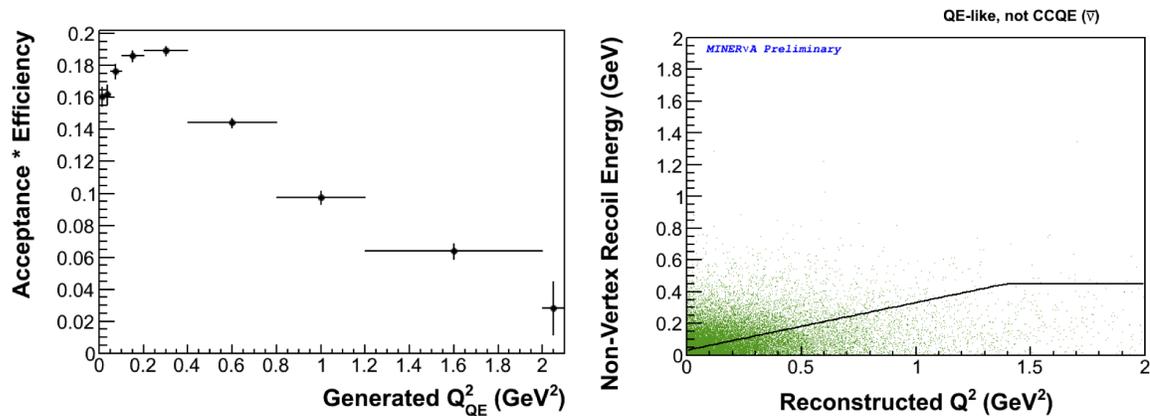


Figure 4.7. Q_{QE}^2 -dependent recoil energy cut used in the previous MINERvA antineutrino CCQE cross section analysis [96]. Blue events are signal, red are background, for the previous analysis' signal definition of true GENIE CCQE.

Acceptance refers to the fraction of signal events that our detector was able to measure - for example, the requirement that muons must be forward enough to hit MINOS limits our angular acceptance. When evaluating these effects, we look at the fraction of signal events that we can reconstruct, compared to the total number of signal events; this ratio gives us the product of efficiency times acceptance. Because we only ever evaluate these two effects together, and cannot separate them, to this product, and will use the terms 'efficiency' and 'acceptance' interchangeably to refer to said product of efficiency and acceptance.)

To improve this efficiency, two adjustments were made:

- Requirement of a maximum proton kinetic energy threshold in the signal definition



(a) Efficiency for non-CCQE $CC0\pi$ events vs. Q^2_{QE} . (b) Recoil distribution of non-CCQE $CC0\pi$ events vs. Q^2_{QE} .

Figure 4.8. Efficiency for non-CCQE events with a $CC0\pi$ signature, and distribution of these events in the recoil energy- Q^2_{QE} phase space

- A “shelf” in the recoil distribution: a minimum recoil energy below which all events are accepted

To select the thresholds for the maximum proton kinetic energy and recoil cut shelf, the purity and efficiency were plotted for various values of these cuts, as shown in figure 4.9.

Figure 4.9a shows the efficiency as the maximum proton kinetic energy and recoil cut shelf are varied. As the cut shelf rises, loosening the recoil cut, efficiency increases as we accept more events into our sample. Introducing lower thresholds on allowed leading proton kinetic energy increases efficiency, as energetic protons that our recoil cut rejects no longer correspond to signal events.

Figure 4.9b shows the purity (the fraction of selected events that correspond to signal, rather than background) as the maximum proton kinetic energy and recoil cut shelf are

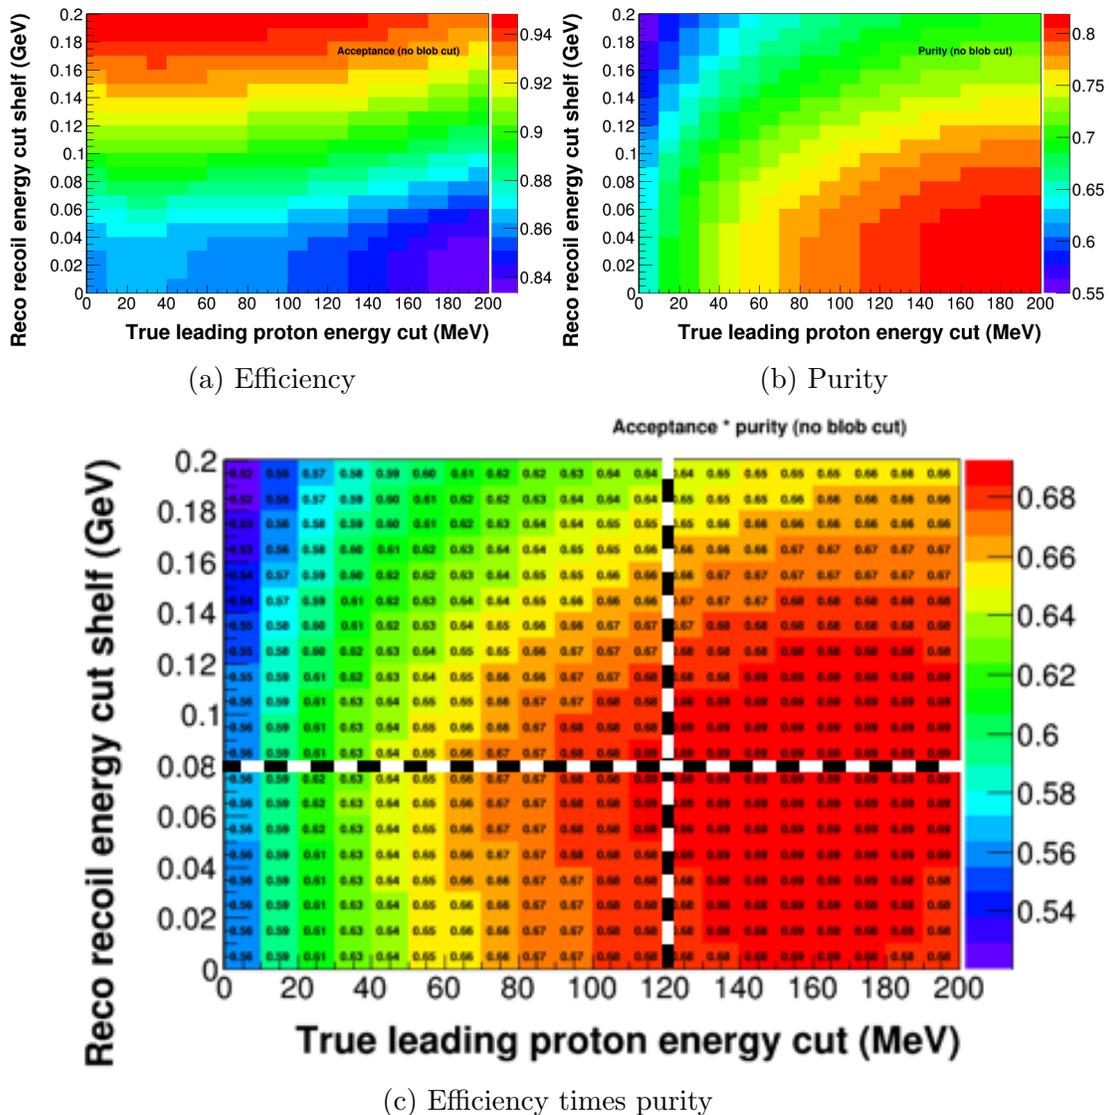


Figure 4.9. Efficiency (top left), purity (top right) and the product of efficiency times purity (bottom) for $CC0\pi$ events with a maximum leading proton kinetic energy threshold as given by the x axis value, where the recoil cut is as shown in figure 4.7, but additionally accepting all events with a reconstructed recoil energy below the y axis value. The black and white lines show the optimized values of a 120 MeV maximum proton kinetic energy in the signal, and a recoil energy shelf of 80 MeV.

varied. As the cut shelf rises, loosening the recoil cut, background events that were previously cut out of our sample are rejected, reducing purity. Introducing lower thresholds on allowed leading proton kinetic energy also decreases purity, as more events with protons in the final state that passed our recoil cut start to be classified as backgrounds.

Figure 4.9c shows the product of the two distributions. The black and white lines show the optimal configuration that maximizes this product: a 120 MeV maximum proton kinetic energy threshold, and a recoil energy shelf of 80 MeV. This gives us a final Q_{QE}^2 -dependent recoil energy cut of

- All events with recoil > 0.45 GeV are rejected
- Events where recoil $> 0.03 + 0.3Q_{QE}^2$ are rejected (recoil in GeV, Q_{QE}^2 in GeV^2),
except:
- All events with recoil < 0.08 GeV are accepted

4.3.3.3. Longitudinal muon momentum cut. We also require our reconstructed muon longitudinal momentum to be less than 15 GeV. This controls muon energy scale uncertainty for very forward muons. In the absence of this cut, some very energetic, forward-going muons have very poor energy reconstruction in MINOS, with the reconstructed energy far higher than the true energy. This is a particular problem when a muon has too much energy to range out in the MINOS tracking calorimeter, and so little transverse momentum that its track curves very little. See section 3.8.4 for more details of how muon momentum is reconstructed in MINOS.

4.3.3.4. Sample purity. After making the cuts, we still have some contamination of background events in our sample. Figure 4.10 shows the purity of the reconstructed sample: the fraction of reconstructed events in each bin that correspond to signal processes.

Figure 4.10a shows how purity varies in the p_T / p_{\parallel} phase space, while 4.10b shows how it varies with reconstructed Q_{QE}^2 and E_{ν}^{QE} . Purity is lowest at high p_T (high Q_{QE}^2) where resonant events are prevalent. The 100% purity in the highest Q_{QE}^2 /lowest E_{ν}^{QE} bin is artificial; it in fact corresponds to a bin in which we have no events at all, due to lack of acceptance.

The overall purity of this sample is 80.9%.

4.3.4. Choice of plotting variables and bins

This analysis is designed to produce double-differential cross sections. By measuring the cross section in two dimensions, we will be able to pick out details hidden by the single-differential cross sections such as those measured in [96] and [107]. This additional information will enable us to differentiate between models that describe very different physics, but which produce similar predictions in the Q^2 space.

Other experiments have chosen to plot double-differential CCQE cross sections vs. the kinematics of the final-state muon. These have the advantage of being directly measurable quantities (as opposed to, for example Q_{QE}^2 , which is reconstructed using the quasi-elastic hypothesis and relying on various assumptions). These measurable quantities are preferred by theorists when generating models' predicted distributions.

MiniBooNE, a short-baseline oscillation experiment at Fermilab, looked at how the cross section for CCQE neutrino and antineutrino scattering varied with muon kinetic energy, T_{μ} , and the cosine of the angle between the muon and the incoming neutrino beam, $\cos \theta_{\mu}$ [29][30]. The MiniBooNE detector's [77] 4π solid angle coverage (it is a spherical mineral oil Cherenkov detector, with photomultiplier tubes in all directions) allows for

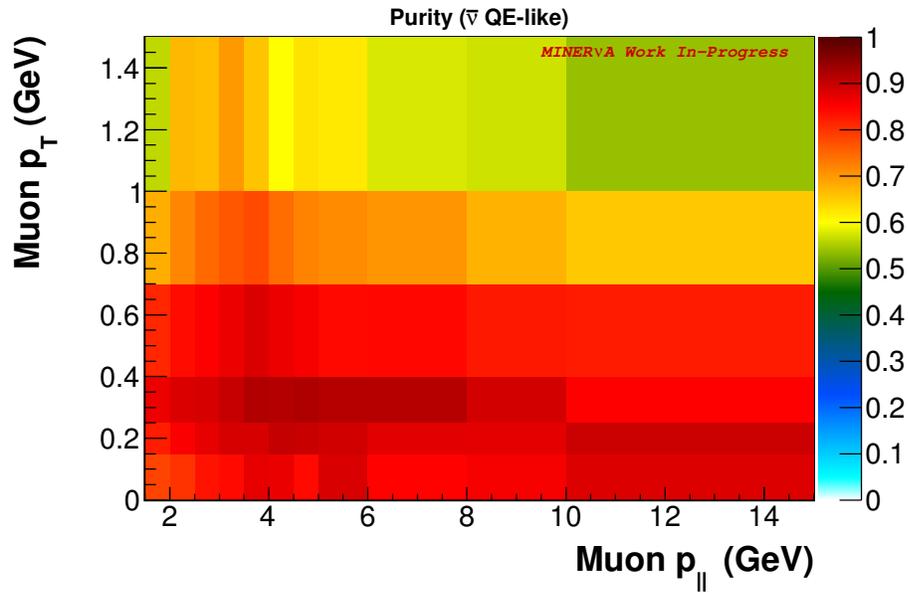
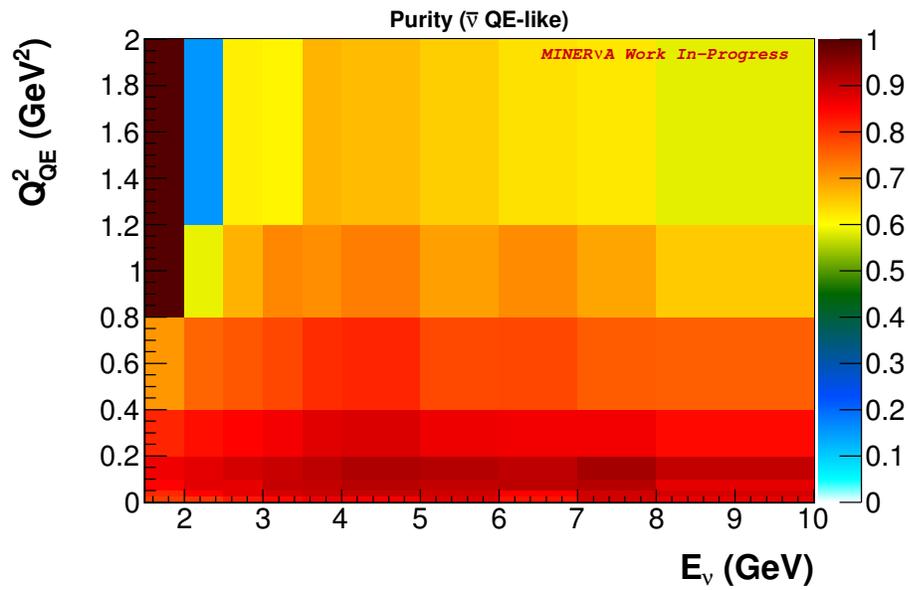
(a) Vs. p_{\parallel} and p_T (b) Vs. E_{ν}^{QE} and Q^2_{QE}

Figure 4.10. Purity of our reconstructed sample

good coverage of this phase space. In MINERvA, however, the limit of approximately $\theta_\mu < 20^\circ$ imposed by the requirement for a MINOS-matched muon, as well as the higher-energy beam that produces more forward-going interaction products, means that we have acceptance only in the range $0.94 < \cos\theta < 1$, rendering this an impractical choice for our detector technology.

T2K, a long-baseline oscillation experiment originating at J-PARC in Japan (and with a far detector at Kamioka, Super-Kamiokande) has published double-differential CCQE neutrino scattering measurements from its near detector ND280, situated 2.5° off of the beam axis. ND280's tracker consists of two fine-grained scintillator-bar targets interspersed with three gas time-projection chambers [11] in a magnetic field, allowing for charge identification as well as particle identification and momentum measurements. As with MiniBooNE, T2K does not have the angular restriction that MINERvA has, and has measured the cross section vs. the muon momentum, p_μ and cosine of the muon angle $\cos\theta_\mu$ with respect to the beam [5][7]. They are able to measure cross sections for $0.6 < \cos\theta_\mu < 1$.

(INGRID, T2K's on-axis near detector [6], has a design that shares some common elements with MINERvA's, with an upstream 'proton module' constructed of scintillator strips, and a downstream region consisting of magnetized iron interspersed with scintillator, similar to the MINOS detector composition. To date, however, INGRID has only measured total cross sections for CCQE, rather than differentials.)

Given our restricted angular acceptance, $\cos\theta_\mu$ is a poor choice for distinguishing the patterns in our distribution. Instead, we choose muon transverse (p_T) and longitudinal momentum (p_{\parallel}), where we have sufficient resolution to be able to reconstruct data in

several bins in each dimension. We also look at the distribution vs. Q_{QE}^2 and E_ν^{QE} , both reconstructed in the quasi-elastic hypothesis from the muon kinematics (see equations 4.1 and 4.2). Note that, as Q_{QE}^2 and E_ν^{QE} are reconstructed from the muon kinematics, they are both functions of p_{\parallel} and p_T . Figure 4.11 shows lines of constant Q_{QE}^2 and E_ν^{QE} , projected onto the p_{\parallel} / p_T phase space. For most of the region considered by this analysis, E_ν^{QE} correlates fairly well with p_{\parallel} , while Q_{QE}^2 is mostly a function of p_T . Note that this simplification breaks down at high p_T and low p_{\parallel} .

This combination has two advantages: it can be projected onto the Q_{QE}^2 axis to give a single-differential distribution that can be compared with our previous $d\sigma/dQ^2$ measurement [96]. Also, as one axis is an approximation of the incoming neutrino energy, it can be used to produce a flux-weighted cross section (we know our flux distribution as a function of neutrino energy).

Note that care must be taken, as E_ν^{QE} does not correspond exactly to true neutrino energy except in the case of a pure quasi-elastic scatter from a stationary nucleon, an assumption which does not apply to all interactions in our signal definition. This gives us a choice: the first option is that when correcting our measured distribution to the truth, we can correct to a true E_ν^{QE} reconstructed from the true muon kinematics (so that we use the formula above, but instead of using our measured muon energy and angle in the calculation, we use the true energy and angle reported by the simulation. This is less model-dependent in terms of the acceptance and unfolding correction, but means that the flux-weighted distribution will be weighted slightly incorrectly, as the flux is a function of true E_ν . Alternatively, when unfolding and acceptance-correcting, we can correct to true E_ν - this may introduce model dependence in the unfolding and acceptance, but allows us

to do the flux weighting correctly. Both methods were used; the differences in cross section were very slight. We will present both cross section results; plots for intermediate stages reflect the unfolding and acceptance correction to true E_ν only, to avoid near-duplication of figures.

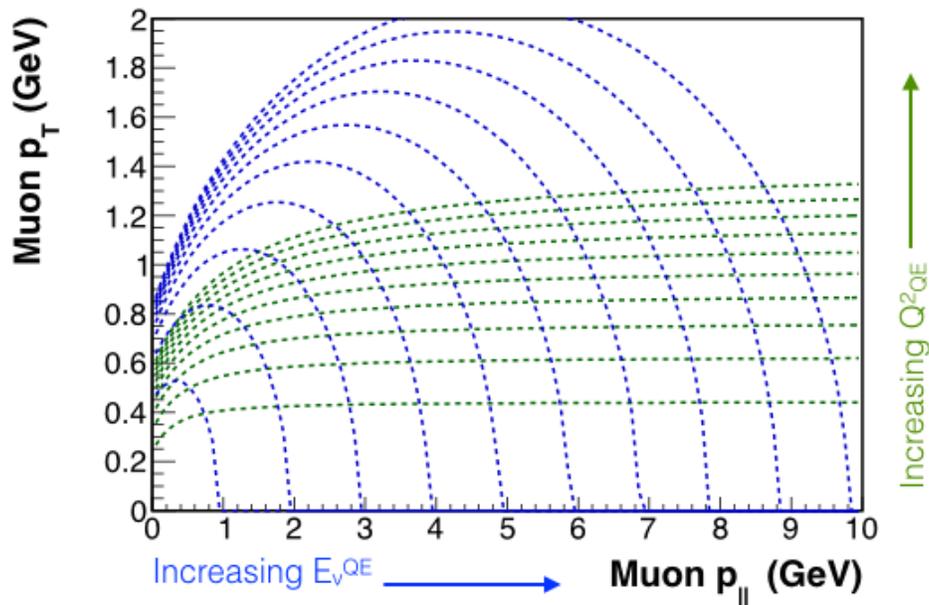


Figure 4.11. Relationship between E_ν^{QE} and Q_{QE}^2 in the quasi-elastic hypothesis, and muon kinematic variables p_T and $p_{||}$. Blue dashed lines show constant values of E_ν^{QE} , increasing in 1 GeV increments from 1 to 10 GeV. Green dashed lines show constant values of Q_{QE}^2 , increasing in 0.2 GeV² increments from 0.2 to 2 GeV².

Ideally, when selecting the binning of our data, we would choose bins such that

- Bins are as narrow as possible to give the maximum detail of the distributions' shapes
- Statistics in each bin are large enough that we are not overwhelmingly dominated by statistical uncertainty

- Each bin is wider than the resolution in the binning quantity at that point, to protect against the majority of events being reconstructed in the wrong bin
- Acceptance should be roughly constant across the bin

For a double-differential cross section, some compromise must be made as the bins form a two-dimensional grid. For example, by choosing narrow p_{\parallel} bins in the low p_T region where statistics are high, we maximize our power to distinguish the distribution's shape, but this leads to low per-bin statistics, and in some cases, to rapidly changing acceptance, in high- p_T bins where the data is more sparse. A binning was selected where the resolution was good enough that we were able to unfold a data-sized warped sub-sample of simulation successfully, reproducing the sample's true bin distribution to within statistical uncertainty. The unfolding study is explained in appendix D. This does, however lead to a small number of bins (6 out of 66, in the high p_T /low p_{\parallel} region where muon angles are above the $\approx 20^\circ$ threshold where the MINOS match requirement leads to poor acceptance) , where the acceptance is so low or so widely-varying that we cannot reliably report a result.

Because of the MINOS-matching restriction, which limits sideways-going muons, we look at data with $p_T < 1.5$ GeV. The MINOS match requirement also means that we will be unable to reconstruct low-energy muons; for this reason, we look at $p_{\parallel} > 1.5$ GeV. The upper limit of $p_{\parallel} < 15$ GeV mimics our reconstruction cut on longitudinal muon momentum.

The choice of bins can be seen in the event distributions in section 4.3.5.

4.3.5. Raw event counts

A total of 18,380 interactions pass our reconstruction cuts for data. 162,328 simulated interactions pass the cuts, of which 128,865 are signal events.

4.3.5.1. Distributions vs. muon transverse and longitudinal momentum. The plots in this section show the distribution of reconstructed events in data (points) and simulation (red line), after all cuts, versus each of our variables. Uncertainties on the data are indicated by error bars; uncertainty on the Monte Carlo is indicated by a pink shaded bar. The data uncertainty is statistical; the Monte Carlo simulation includes all sources of systematic uncertainty. This is due to the standard MINERvA method used to evaluate uncertainties; this method, along with all of the sources of uncertainty evaluated, is described in chapter 6. The systematic uncertainties on the simulation here include uncertainties on the models used to generate the simulation, many of which will be largely irrelevant or have only a small effect on our data, as well as uncertainties on our resolution in various parameters, which will affect our data through the unfolding and efficiency correction.

Figure 4.12 shows the double-differential event distribution in muon transverse and longitudinal momentum, projected onto p_T . Figure 4.13 shows this same distribution projected onto p_{\parallel} . For each variable, the top plot shows the number of events in data and simulation. The simulation, which corresponds to approximately 10 times more protons on target (POT) than the data, has been scaled to the data POT. Additionally, the data in each bin is scaled by the bin width to give a number of events per GeV. The grey hatched area at the bottom of the top plot corresponds to background events in the simulation that nevertheless passed our reconstruction cuts.

The bottom plot in each figure shows a ratio of the number of events in data to the number in simulation, where the simulation, once again, has been POT-normalized to the data. The pink bar shows the fractional uncertainty on the simulation.

With the full two-dimensional distribution, it is harder to show how data relates to simulation, and to show uncertainties. Figure 4.14a shows a two-dimensional distribution of events in data, once again normalized by the bin width in each direction to give a count per GeV per GeV.

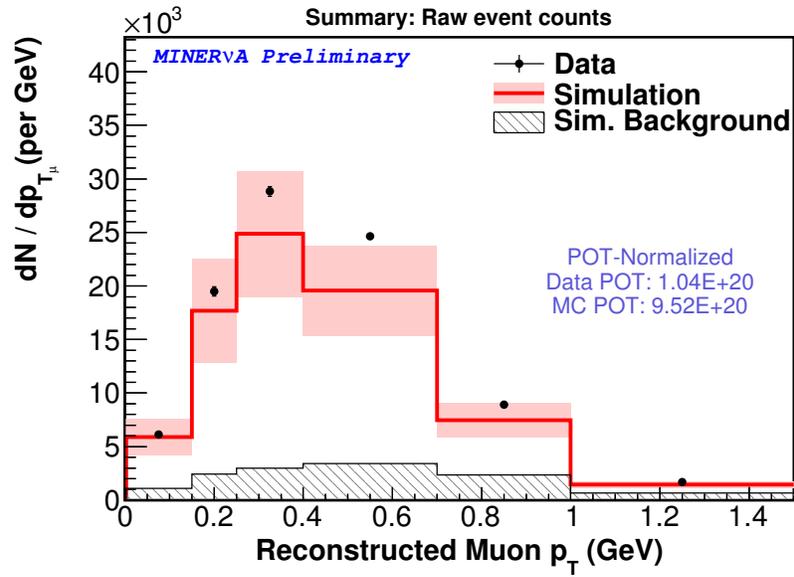
Figure 4.14b shows the pull between data and simulation. This is a function of the number of events N in data and Monte Carlo simulation, and of the uncertainties σ on those measurements:

$$(4.3) \quad \text{pull} = \frac{N_{\text{data}} - N_{\text{MC}}}{\sqrt{\sigma_{\text{data}}^2 + \sigma_{\text{MC}}^2}}$$

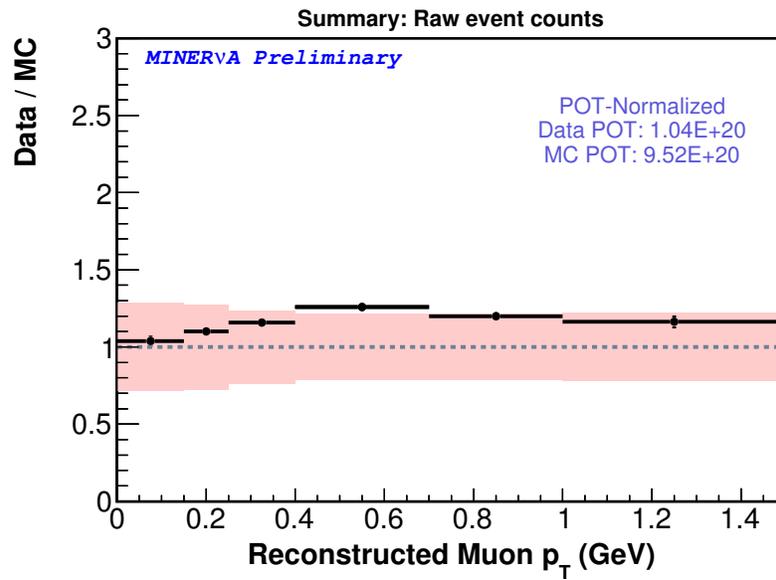
Thus a pull of zero (shown in white) corresponds to data-simulation agreement, a pull of +1 corresponds to a 1σ excess in data and a pull of -1 to a 1σ excess in simulation. Data excesses are colored red, with a stronger color indicating a greater excess, and simulation excesses in blue.

Figure 4.15 shows the data and simulation event distributions versus muon p_T , in bins of p_{\parallel} . Again, the counts are scaled by the bin width in each dimension, to give a distribution per GeV per GeV.

4.3.5.2. Distributions vs. Q_{QE}^2 and E_{ν}^{QE} . This section shows the same distributions as in section 4.3.5.1, but presented vs. Q_{QE}^2 and E_{ν}^{QE} . See section 4.3.5.1 for an explanation of the plots. At this stage, we are only able to reconstruct E_{ν}^{QE} in the quasi-elastic

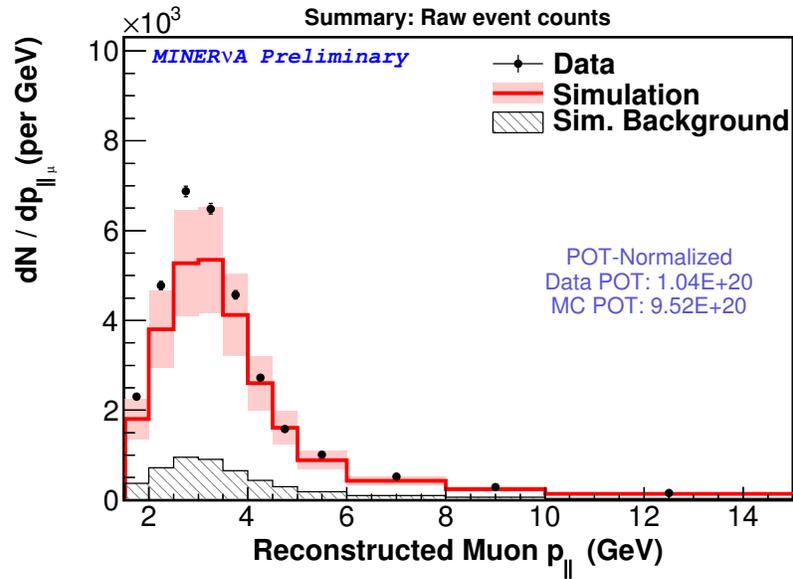


(a) Event counts

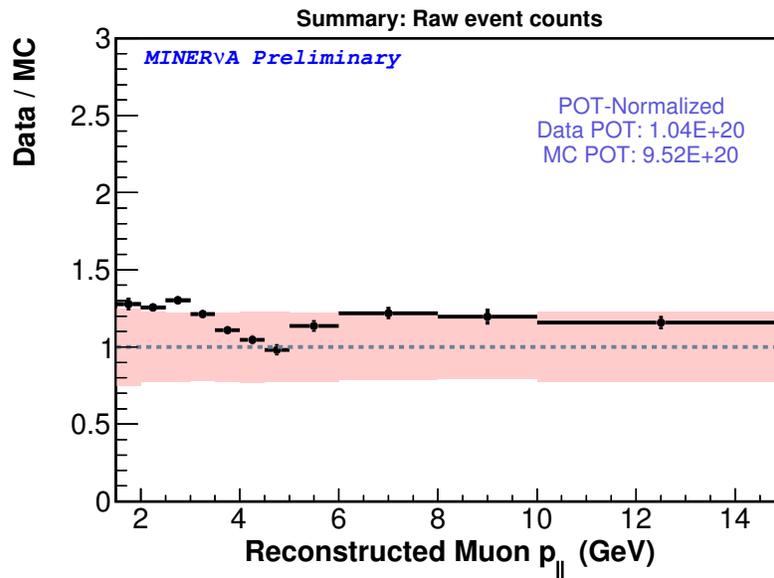


(b) Data/simulation ratio

Figure 4.12. Reconstructed event counts in data and simulation vs. muon p_T (normalized to data POT).

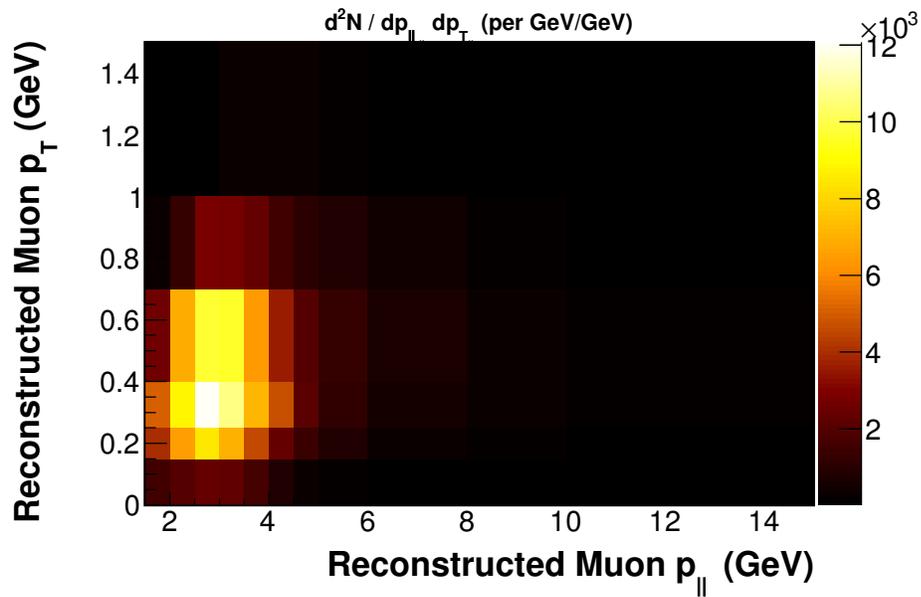


(a) Event counts

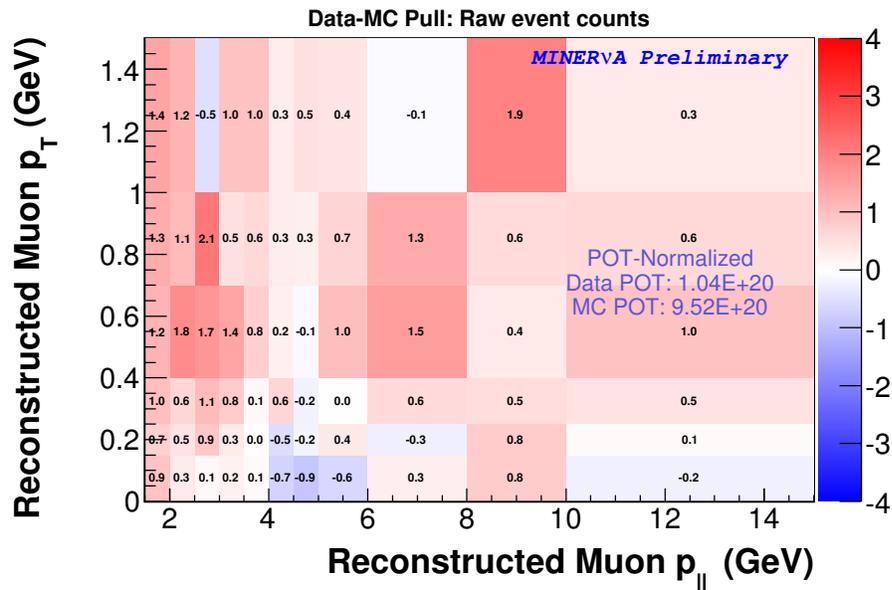


(b) Data/simulation ratio

Figure 4.13. Reconstructed event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).



(a) Event counts in data



(b) Data/simulation pull

Figure 4.14. Reconstructed event counts in data, and pull between data and simulation, vs. muon p_T and p_{\parallel}

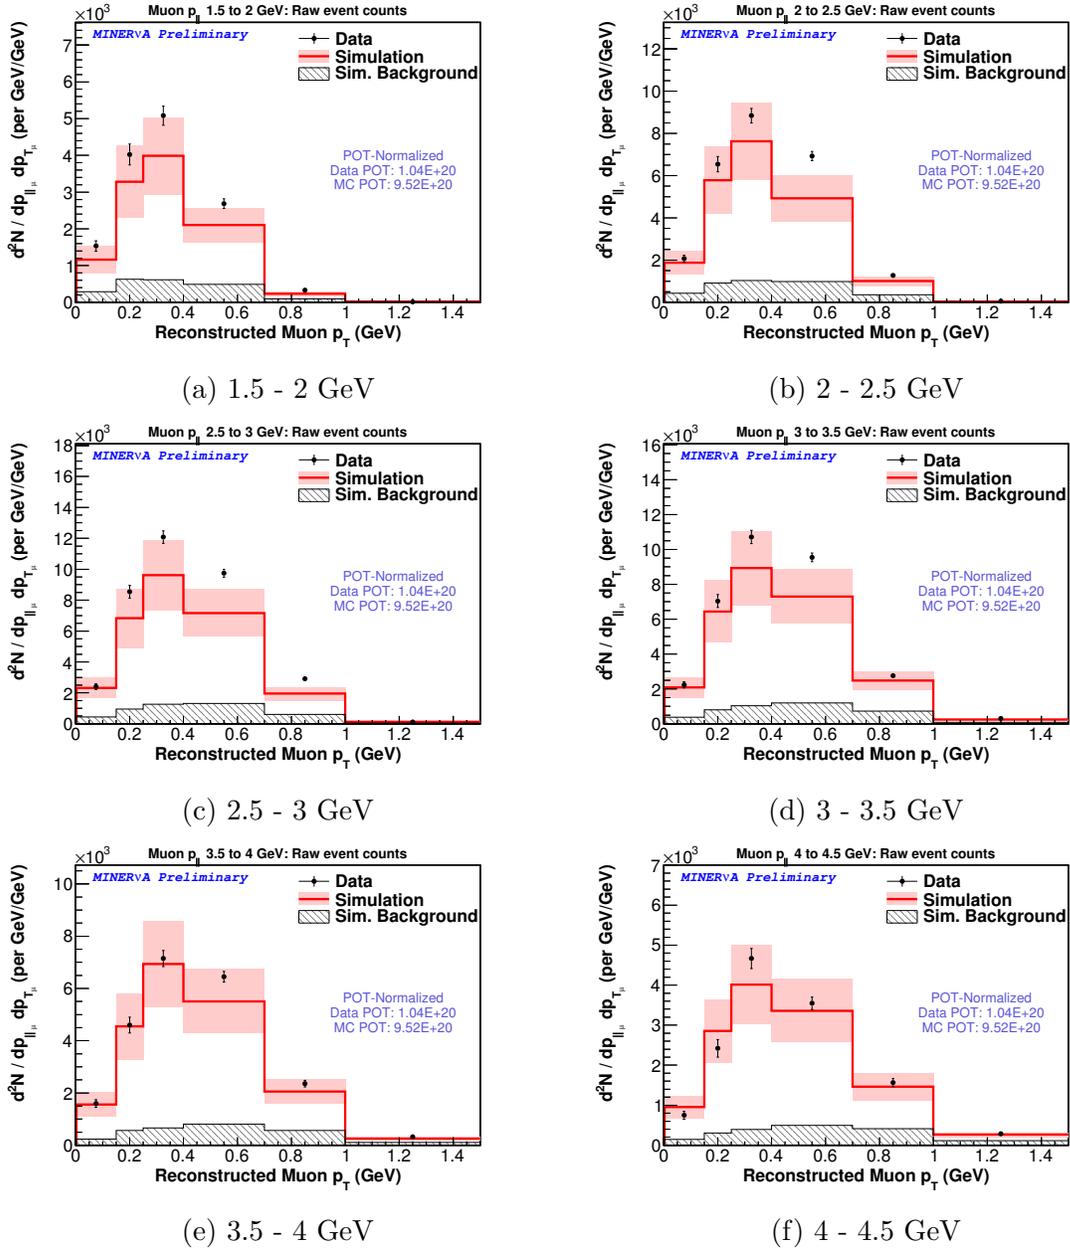
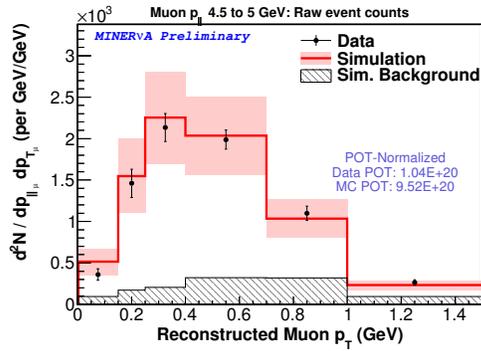
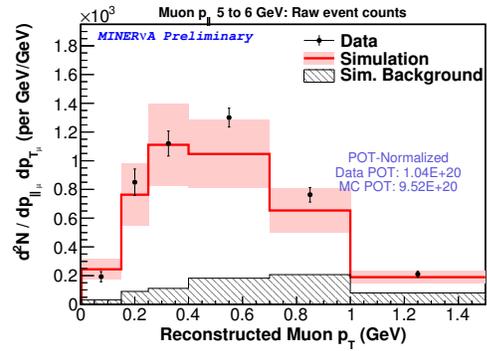


Figure 4.15. Reconstructed event counts vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)

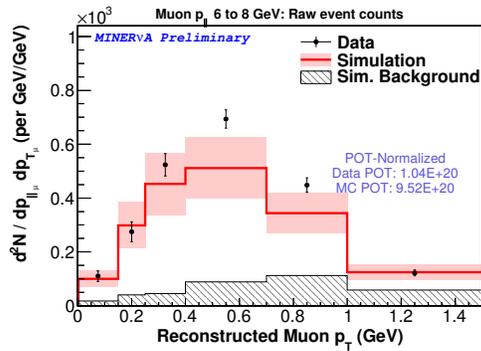
hypothesis from the reconstructed muon variables; we can therefore only show these distributions as a function of reconstructed E_ν^{QE} as this is also our best-effort reconstruction of true E_ν .



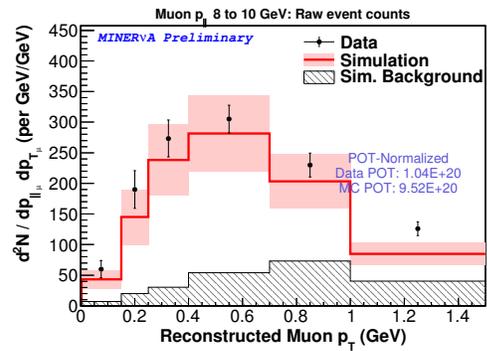
(g) 4.5 - 5 GeV



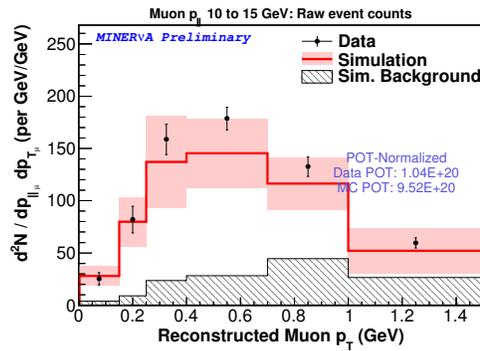
(h) 5 - 6 GeV



(i) 6 - 8 GeV

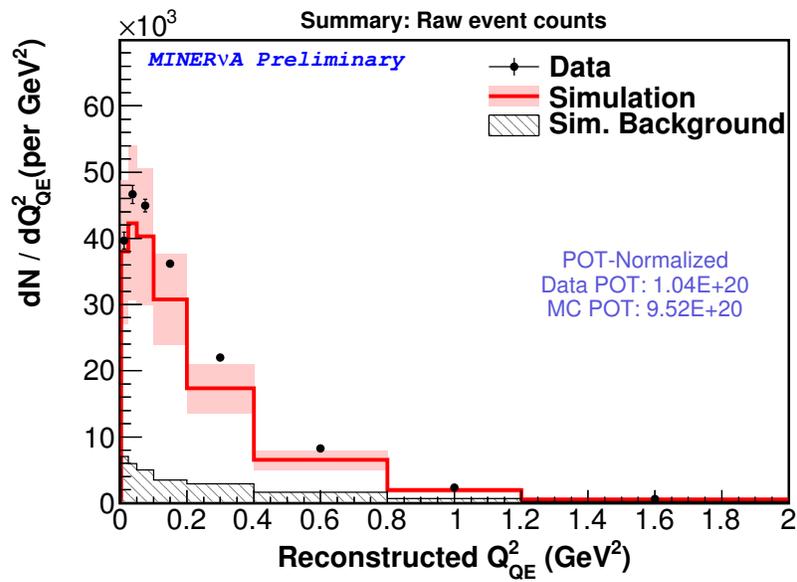


(j) 8 - 10 GeV

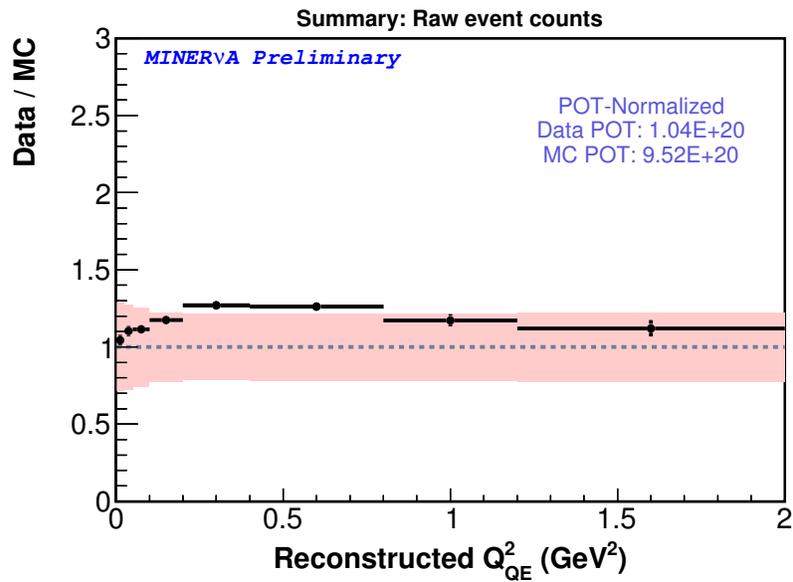


(k) 10 - 15 GeV

Figure 4.15. Reconstructed event counts vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)

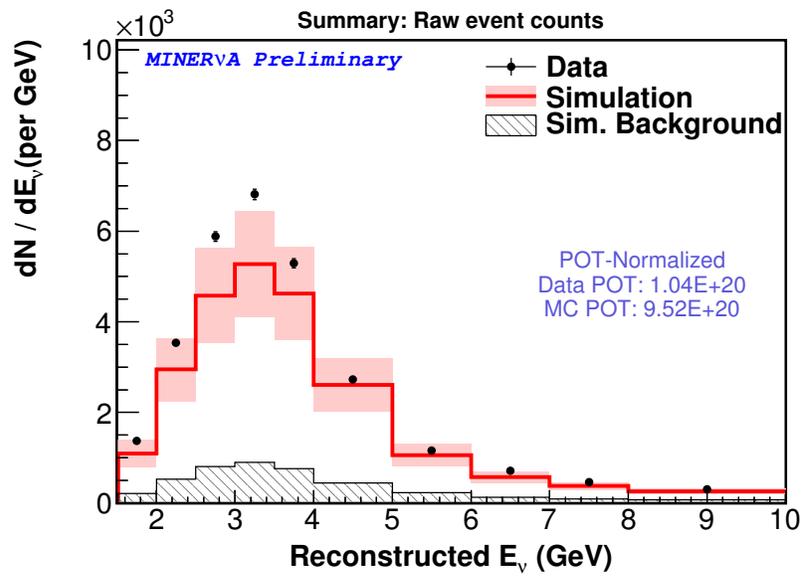


(a) Event counts

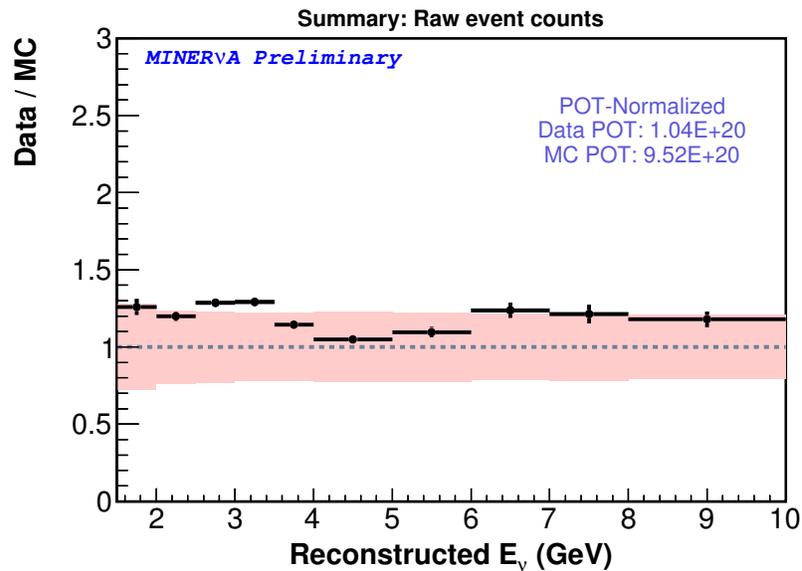


(b) Data/simulation ratio

Figure 4.16. Reconstructed event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).

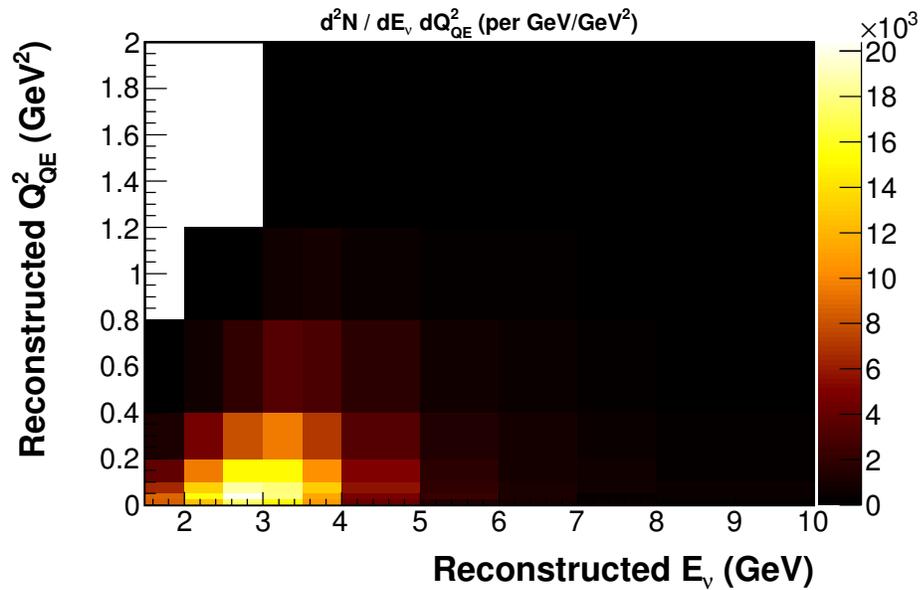


(a) Event counts

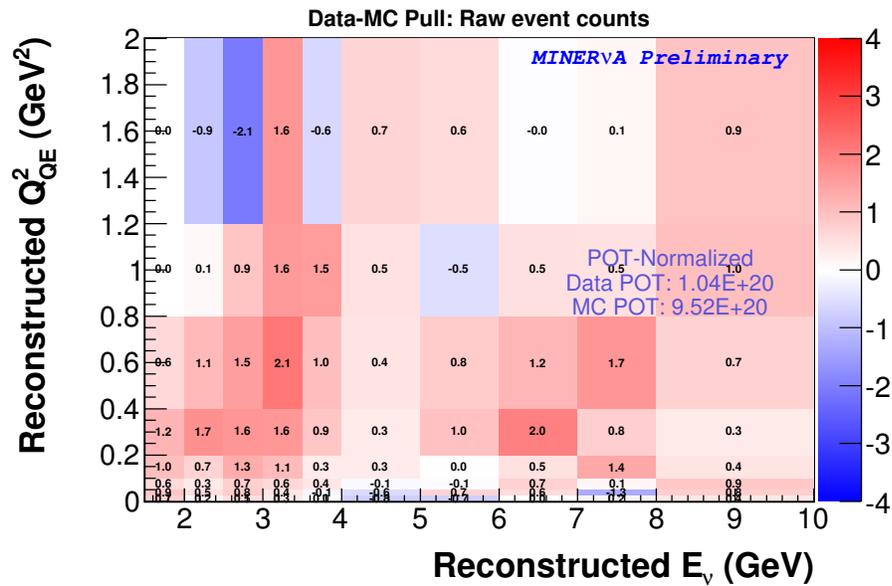


(b) Data/simulation ratio

Figure 4.17. Reconstructed event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).

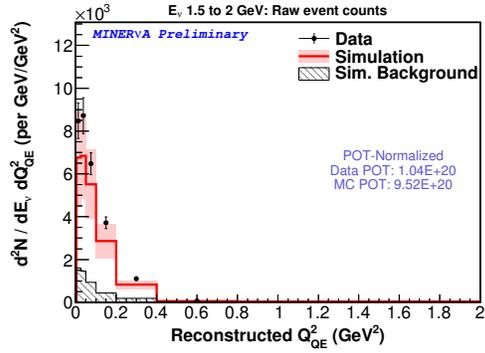


(a) Event counts in data

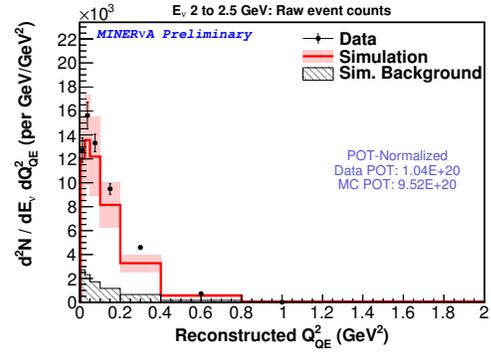


(b) Data/simulation pull

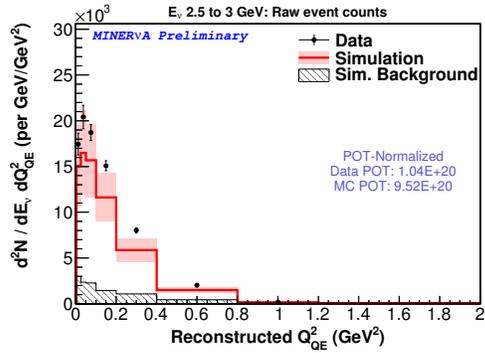
Figure 4.18. Reconstructed event counts in data, and pull between data and simulation, vs. Q_{QE}^2 and E_ν^{QE}



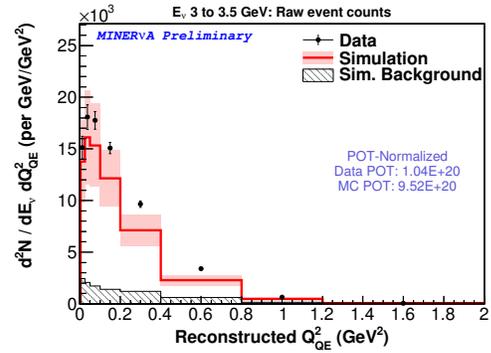
(a) 1.5 - 2 GeV



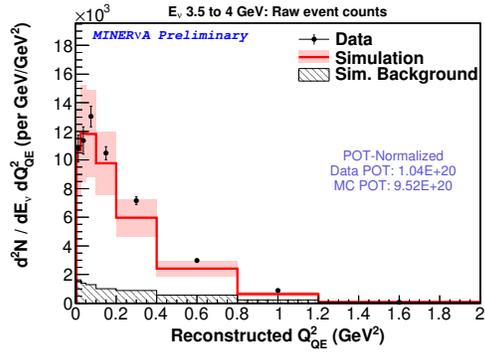
(b) 2 - 2.5 GeV



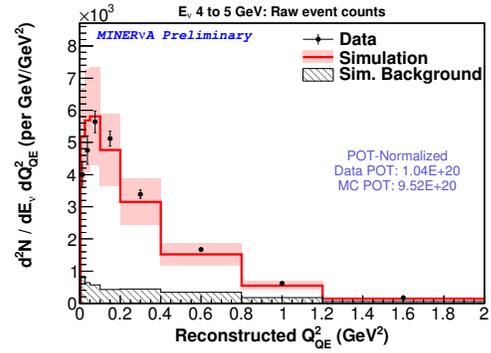
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

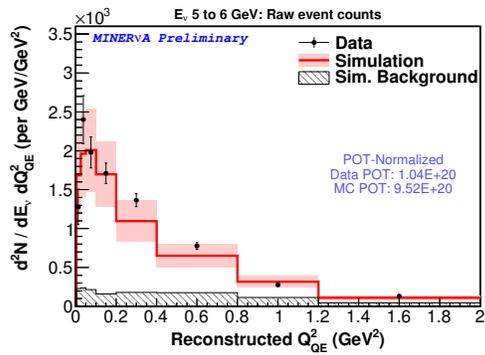


(e) 3.5 - 4 GeV

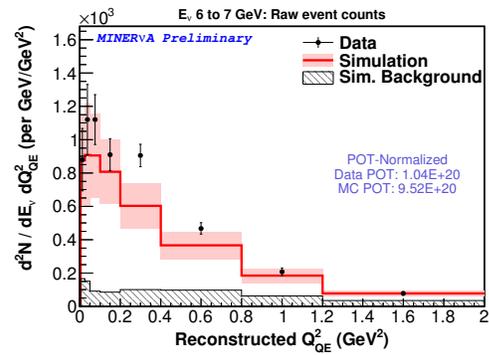


(f) 4 - 4.5 GeV

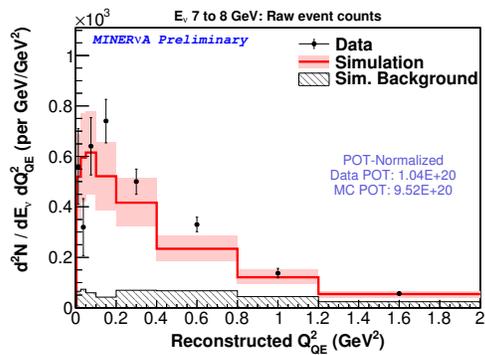
Figure 4.19. Reconstructed event counts vs. Q_{QE}^2 , in bins of E_ν^{QE} (continued in next figure)



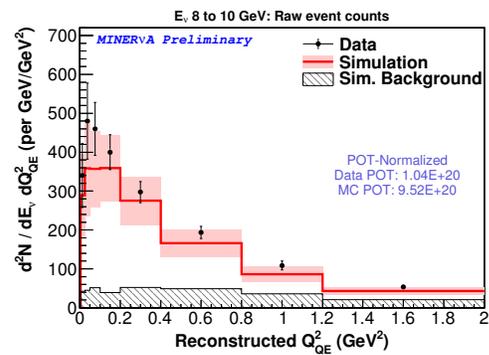
(g) 4.5 - 5 GeV



(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV

Figure 4.19. Reconstructed event counts vs. Q_{QE}^2 , in bins of E_ν^{QE} (continued)

CHAPTER 5

Double-differential cross section calculation**5.1. How the cross section is calculated**

The double-differential cross section versus variables x and y in bin (i, j) is given by:

$$(5.1) \quad \left(\frac{d^2\sigma}{dx dy} \right)_{ij} = \frac{\sum_{\alpha\beta} U_{\alpha\beta ij} (N_{\text{data},\alpha\beta} - N_{\text{data},\alpha\beta}^{bkgd})}{\epsilon_{ij}(\Phi T)(\Delta x_i)(\Delta y_j)}$$

Where

- $N_{\text{data},\alpha\beta}$ is the number of data events reconstructed in bin (α, β)
- $N_{\text{data},\alpha\beta}^{bkgd}$ is the estimated number of background events reconstructed in bin (α, β)
- $U_{\alpha\beta ij}$ is the element of a migration matrix connecting reconstructed bin (α, β) to true bin (i, j)
- ϵ_{ij} is the product of reconstruction efficiency and detector acceptance for events in true bin (i, j)
- Φ is the flux of incoming neutrinos (either integrated or for the given bin - see later discussion)
- T is the number of scattering targets (here, the number of protons)
- Δx_i is the width of bin i
- Δy_j is the width of bin j

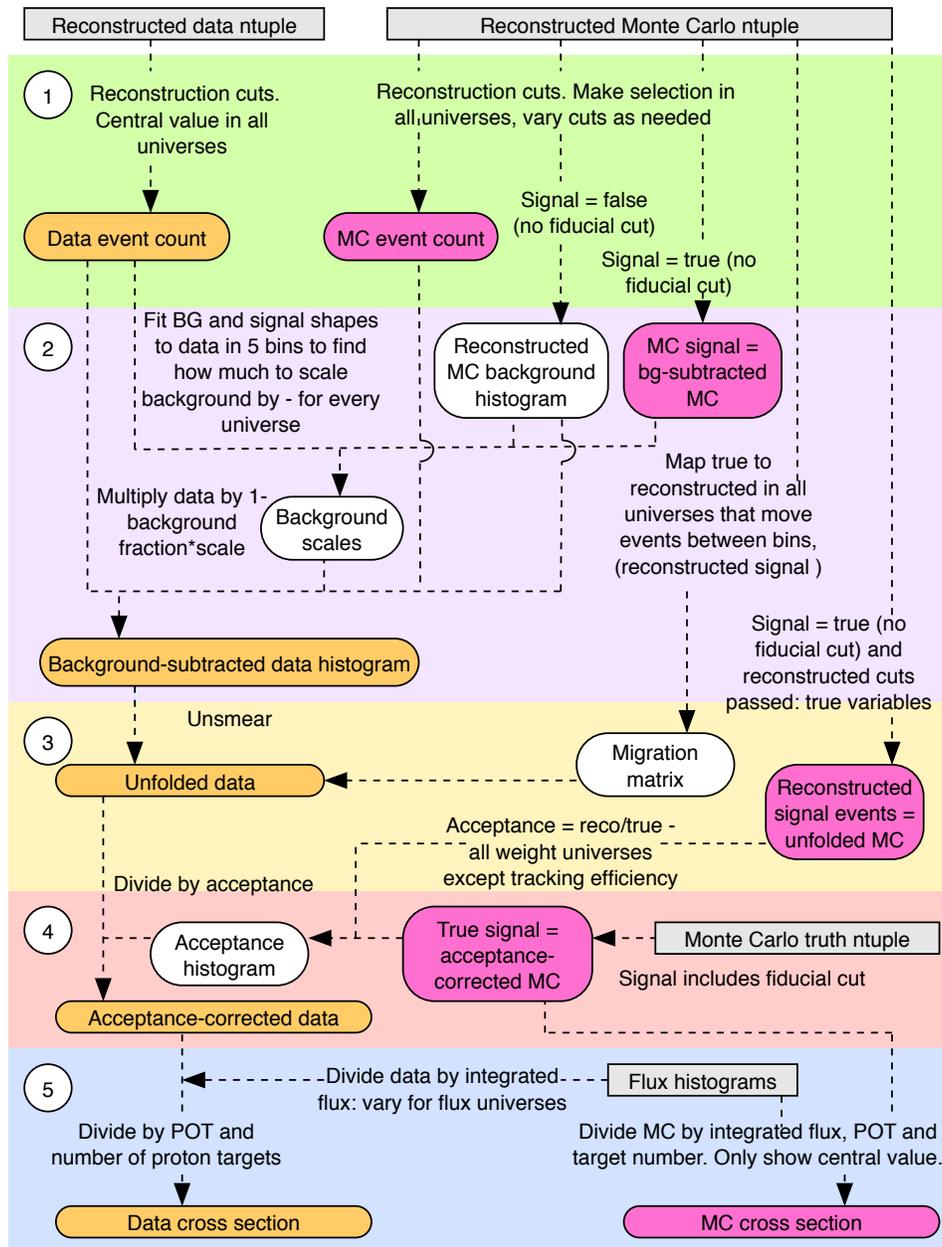


Figure 5.1. Stages of cross section analysis

$N_{\text{data},\alpha\beta}$ is the event distribution shown in section 4.3.5. This chapter will explain how the other components of equation 5.1 are estimated from the simulation, and describe the

stages of the cross section calculation. These stages are summarized as a flow chart in figure 5.1.

5.2. Background subtraction

The term $N_{\text{data},\alpha\beta}^{bkgd}$ in equation 5.1 refers to the estimated number of reconstructed data events that correspond to background processes. By this we mean events that have passed our reconstruction cuts, but which were not actually generated by quasi-elastic-like events, according to our signal definition. Recall that our quasi-elastic-like signal, explained in section 4.1.3, is defined as having a final state containing a μ^+ , any number of neutrons, any number of protons with less than 120 MeV kinetic energy, and no pions, other hadrons, or non-de-excitation photons. Thus, background events in our sample could, for example, correspond to resonant events with pions that did not make a track, and that generated recoil distributions that fell within our cuts. Figure 5.2 shows the distribution of signal and background events vs. muon transverse and longitudinal momentum.

In our simulation, we are able to identify which of our reconstructed events are signal and which are background, by looking at the truth information provided by the GENIE Monte Carlo generator [93], which can tell us the energies and multiplicity of particles in the final state. Thus, a simple way to determine the fraction of events in each reconstructed data bin that correspond to backgrounds would be to assume that this fraction was the same as in the simulation for the same bin:

$$(5.2) \quad N_{\text{data},\alpha\beta}^{bkgd} = \frac{N_{\text{MC},\alpha\beta}^{bkgd}}{N_{\text{MC},\alpha\beta}} \times N_{\text{data},\alpha\beta}$$

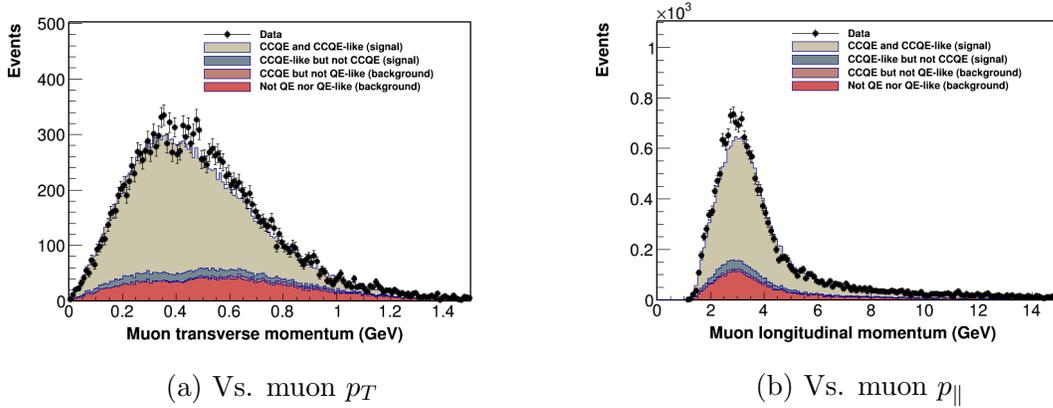


Figure 5.2. Distribution of signal and background events vs. muon transverse and longitudinal momentum. Simulation is POT-normalized to data. In the simulation, beige corresponds to CCQE events with a QE-like signature. Blue events also have a QE-like signature, but were generated by resonant or DIS events where the pion was absorbed by FSI. Both of these are signal. The backgrounds consist of CCQE events with non-QE-like signature (pink), typically where a pion is generated through FSI, and non-CCQE events (resonant, with some DIS) without a QE-like signature (red).

However, in doing this, we are very reliant on the simulation’s ability to correctly predict the strengths of signal and background processes. MINERvA’s charged pion production analysis [90] suggests that GENIE over-predicts the rate of resonant pion production, our most common background. In order to protect against this, we instead use a data-driven fitting procedure to determine the relative fractions of signal and background in our data, by determining the fractions of signal and background processes that would best match our data’s shape.

The result of background subtraction is reconstructed distributions corresponding to only CCQE-like events, plotted vs the reconstructed variables (without correction for any mis-reconstruction due to detector effects, or for loss of efficiency due to the detector’s or the reconstruction algorithm’s limitations).

5.2.1. Background fitting procedure

As the chief factor in discriminating between signal and background (after eliminating events with wrong-sign muons or which have tracks, indicating additional charged particles that would not be produced by a CCQE-like process) is in the recoil energy distribution, it is in this distribution that we look at our signal and backgrounds, to attempt to determine the best-fit fractions of signal and background processes to match our data.

To do this, our Monte Carlo is area-normalized to the data in each of five regions of the p_T/p_{\parallel} phase space, selected for their similarity in signal and background shapes. For each of these bins, the non-vertex recoil energy, after all other cuts, is plotted in data, and in signal and background Monte Carlo. (Events which pass all physics cuts, but whose true vertex falls outside the fiducial volume, are considered signal for this study). The signal and background distributions from simulation are scaled so that the total summed area of the two distributions matches that of the data, allowing us to compare the distributions' shapes. At this stage of the procedure, we define events who fit our signal definition of quasi-elastic-like scattering, but which have a true interaction vertex outside the fiducial volume, as signal; these come from the same physics processes as other quasi-elastic-like interactions, and will thus generate signal-like recoil distributions. The `TFractionFitter` tool, part of the ROOT framework [64], is used to perform a fractional fit of the simulation to data, in which the relative normalizations of the signal and background parts of the simulation are allowed to float, until the best match to data is achieved.

Figures 5.3-5.7 show the recoil distributions in data and (area-normalized) simulation in the five regions of p_T/p_{\parallel} , before and after tuning the signal and background fractions. In each case, a scale is extracted corresponding to the factor by which the background

was scaled to give the best fit. The most likely estimate of background fraction within each bin the data distribution corresponds to the background fraction of the Monte Carlo in that bin, multiplied by this scale factor. Thus, we get our final result:

$$(5.3) \quad D_{si} = D_i \times \left(1 - S_i \frac{M_{bi}}{M_i} \right)$$

where S_i is the scale factor corresponding to bin i . D_i corresponds to the number of data events in bin i , while D_{si} is the estimated number of signal events in data in bin i . M_i is the number of Monte Carlo events in bin i and M_{bi} the number of Monte Carlo background events in bin i . Note that by multiplying the data distribution by (1-scaled background), we protect against double-counting statistical uncertainties, as would occur if we subtracted off the scaled background (adding the background uncertainty in quadrature to the original data uncertainty).

The scales for our p_T vs. $p_{||}$ backgrounds, and for our Q_{QE}^2 vs. E_ν^{QE} backgrounds, are shown in figure 5.8 and in table 5.1. In each case, as suggested by [90], the simulation is found to predict too high a fraction of background events; this is corrected for by scaling the backgrounds down by a fraction varying from 8 to 19% depending on the bin.

The background-subtracted Monte Carlo distribution is generated by selecting only reconstructed events that pass the true CCQE-like signal cuts.

5.2.2. Background-subtracted event distributions

The plots in this section show the distribution of events in data (points) and simulation (red line), after tuned backgrounds have been subtracted. Uncertainties on the data are indicated by error bars; uncertainty on the Monte Carlo simulation is indicated by a

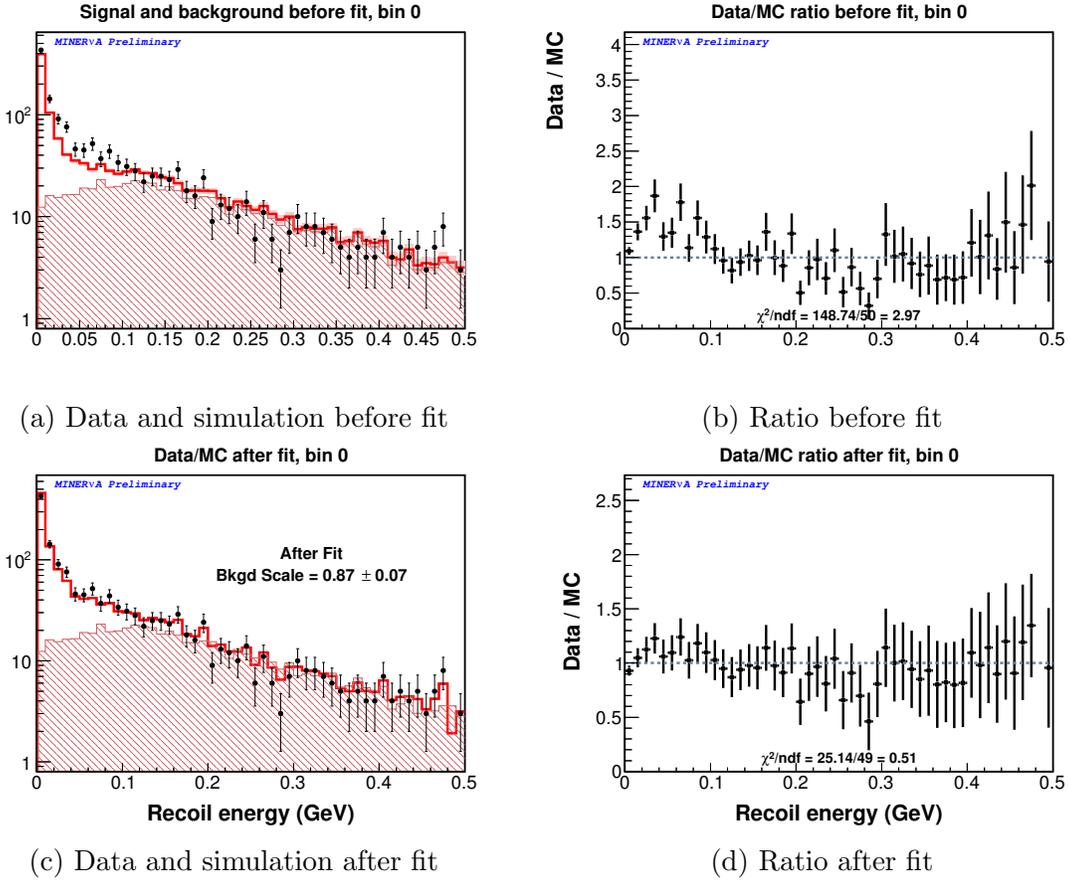


Figure 5.3. Background tuning: $p_T < 0.15$ GeV

pink shaded bar. The background subtraction procedure causes uncertainties (previously all evaluated on the simulation) to move to the data. This is because the fraction of backgrounds subtracted carried systematic uncertainty; using this uncertain distribution in the subtraction calculation means that the uncertainty will be propagated to the result of the subtraction process. The systematic effects on the background-subtracted data distribution are dominated by background model uncertainties. They are explained in detail in section 6.8.2.

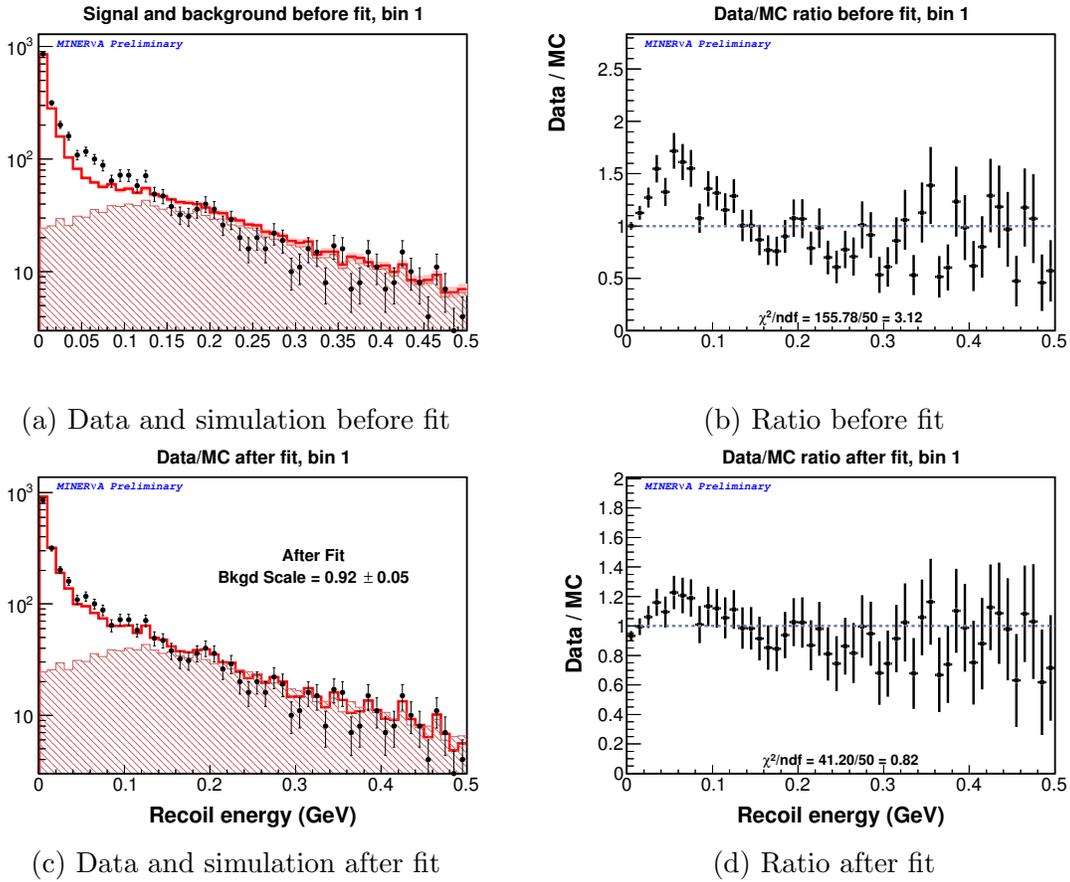


Figure 5.4. Background tuning: $0.15 < p_T < 0.25$ GeV

Figure 5.9 shows the double-differential background-subtracted event distribution in muon transverse and longitudinal momentum, projected onto p_T . Figure 5.10 shows this same distribution projected onto $p_{||}$. For each variable, the top plot shows the number of events in data and simulation. The simulation, which corresponds to approximately 10 times more protons on target (POT) than the data, has been scaled to the data POT. Additionally, the data in each bin is scaled by the bin width to give a number of events per GeV. The bottom plot in each figure shows a ratio of the number of events in data to

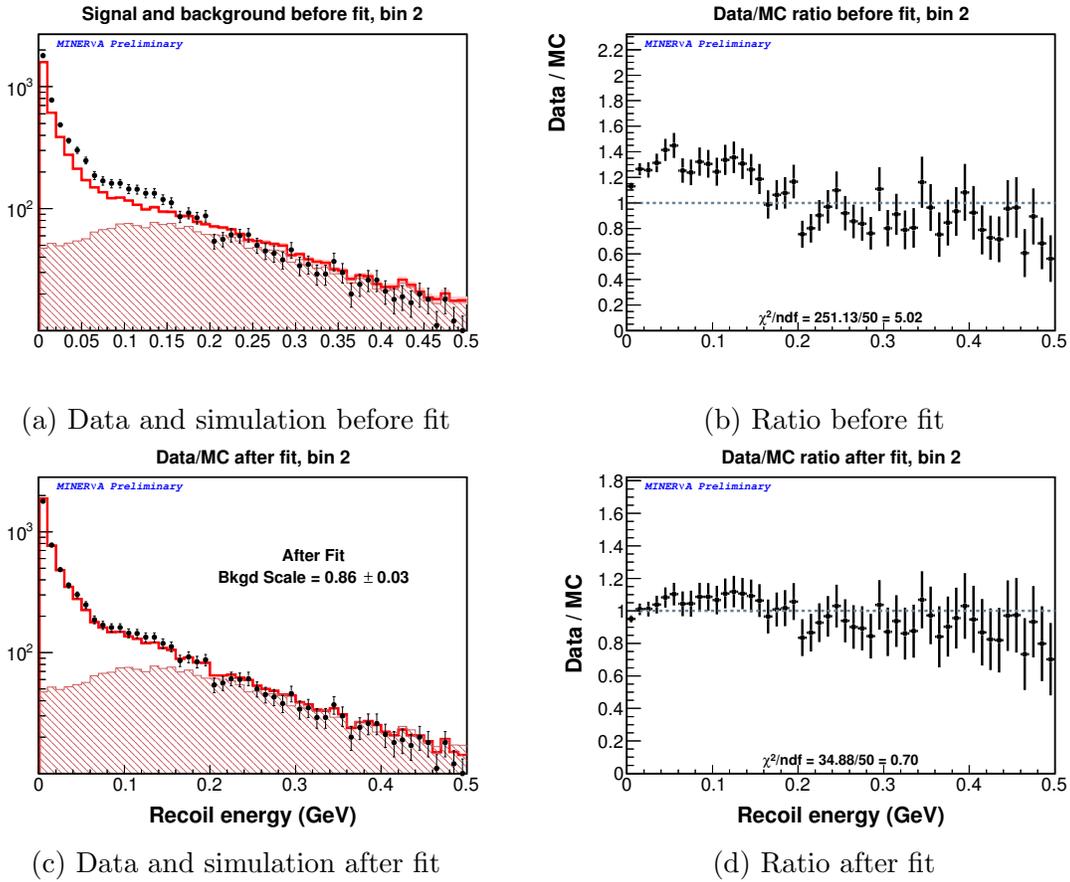


Figure 5.5. Background tuning: $0.25 < p_T < 0.4$ GeV

the number in simulation, where the simulation, once again, has been POT-normalized to the data. The pink bar shows the fractional uncertainty on the simulation.

After background subtraction, 15,265 events remain in our data distribution.

5.3. Unfolding

5.3.1. Detector smearing

The limitations of any particle detector mean that no quantity can be reconstructed with exact precision. For example, in MINERvA, our position resolution is limited by the

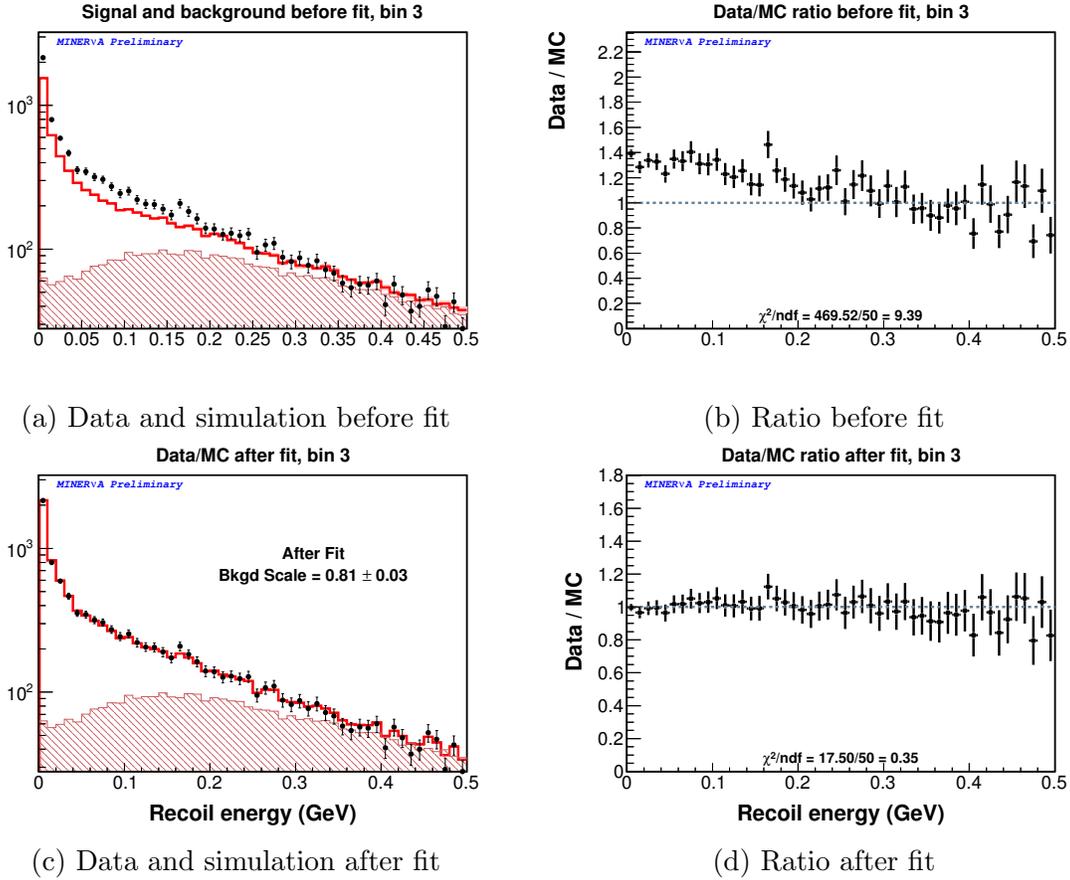


Figure 5.6. Background tuning: $p_T > 0.4$ GeV; $p_{\parallel} < 4$ GeV

size of the scintillator strips; the energy resolution by the detector response to charged particles, calibration etc. Thus a reconstructed quantity (such as the muon energy) may be measured somewhat higher or lower than its true value, and may therefore be reconstructed (“migrate”) into an adjacent, or even more distant, bin. The effect of this is to “smear” a distribution - sharp peaks in a true distribution will tend to broaden in the reconstructed distribution, as demonstrated by the cartoon in figure 5.11.

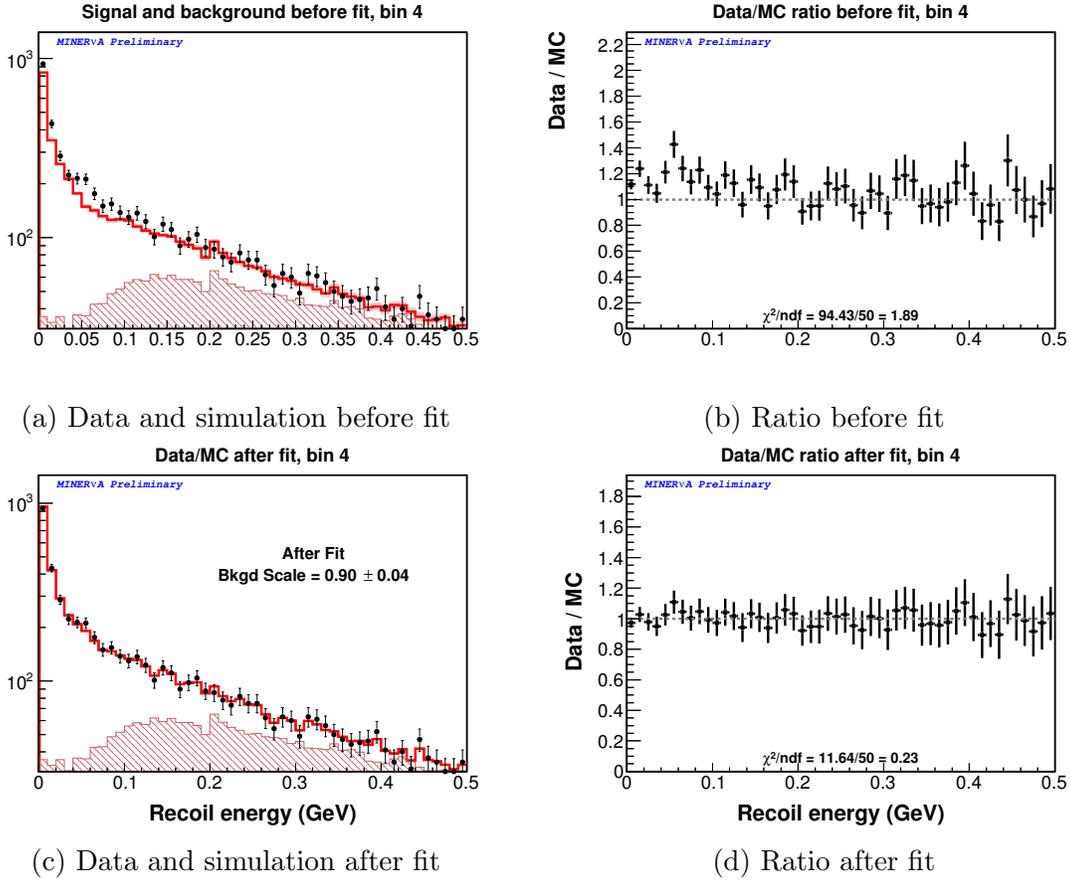


Figure 5.7. Background tuning: $p_T > 0.4$ GeV; $p_{\parallel} > 4$ GeV

5.4. Unsmearing the data

To correct for this effect, we look at Monte Carlo simulation, in which we have access to both generated and reconstructed quantities. Figure 5.12 shows the distribution of signal events in our reconstructed Monte Carlo sample, plotted versus the reconstructed (left) and true generated (right) values of muon longitudinal and transverse momentum. Note that, for this stage of the procedure, we do not enforce the requirement of an interaction vertex in the fiducial volume when we define our ‘signal’ in the reconstructed

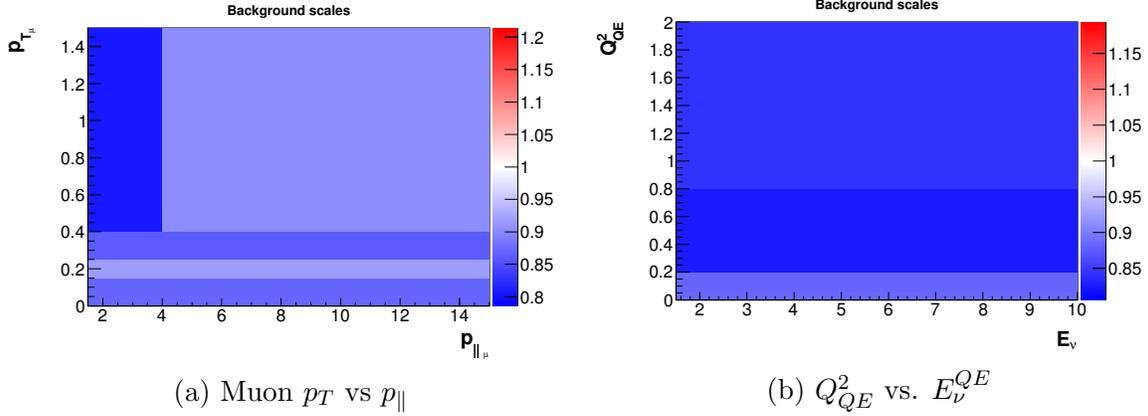
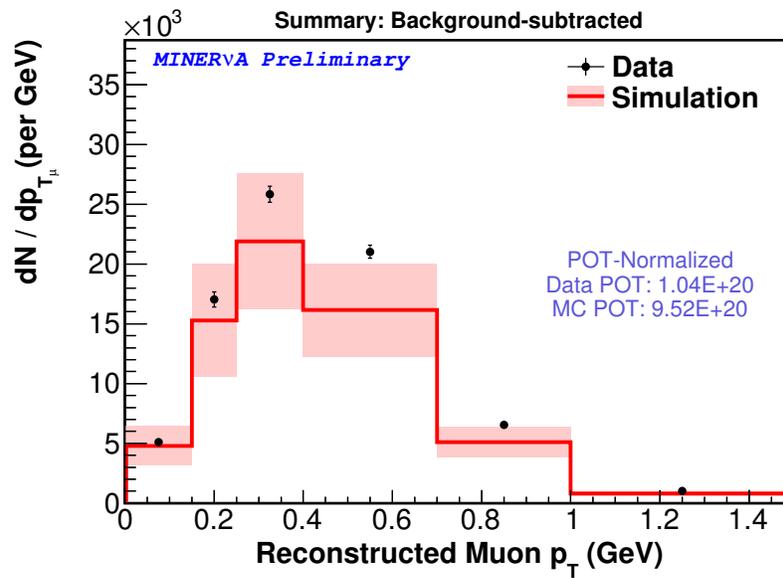


Figure 5.8. Scales by which the background fractions should be multiplied, in order to achieve the best data/simulation shape fit in the recoil distribution

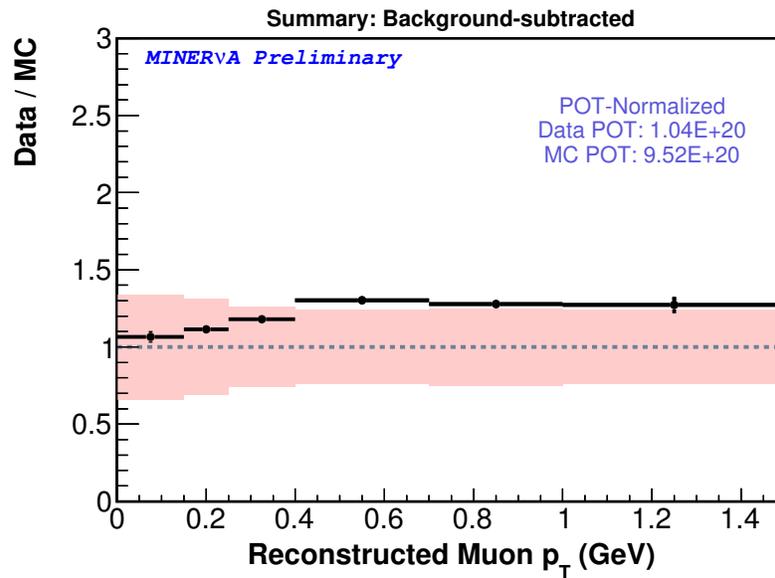
Bin	p_T range (GeV)	$p_{ }$ range (GeV)	Scale	χ^2/DOF before	χ^2/DOF after
0	0-0.15 GeV	1.5-15GeV	0.87 ± 0.07	2.97	0.51
1	0.15-0.25 GeV	1.5-15GeV	0.92 ± 0.05	3.12	0.82
2	0.25-0.4 GeV	1.5-15GeV	0.86 ± 0.03	5.02	0.70
3	0.4-1.5 GeV	1.5-4GeV	0.81 ± 0.03	9.39	0.35
4	0.4-1.5 GeV	4-15GeV	0.90 ± 0.04	1.89	0.23
Bin	Q_{QE}^2 range (GeV ²)	E_{ν}^{QE} range (GeV)	Scale	χ^2 before	χ^2 after
0	0-0.05 GeV ²	1.5-10GeV	0.89 ± 0.05	3.83	0.92
1	0.05-0.2 GeV ²	1.5-10GeV	0.88 ± 0.03	5.23	0.76
2	0.2-0.8 GeV ²	1.5-10GeV	0.83 ± 0.03	8.57	0.30
3	0.8-2 GeV ²	1.5-10GeV	0.84 ± 0.08	1.66	0.28

Table 5.1. Summary of the scales by which the background fractions should be multiplied, in order to achieve the best data/simulation shape fit in the recoil distribution, along with uncertainty in the scale and χ^2 between data and simulation before and after. TFractionFitter introduces additional degrees of freedom to the fit to account for statistical fluctuation in the simulation template shapes, causing the χ^2/DOF after the fit to appear lower than one would expect.

distribution. This is to protect against labeling events as backgrounds when they are in fact from quasi-elastic-like processes.

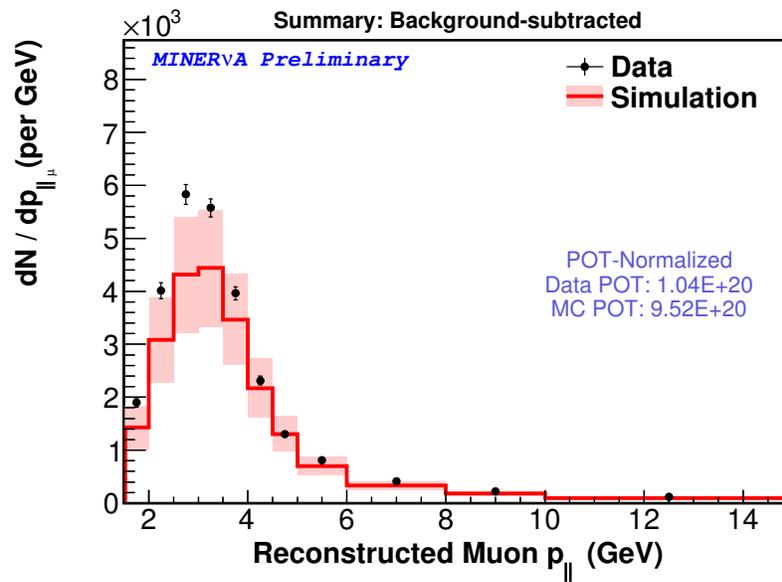


(a) Event counts

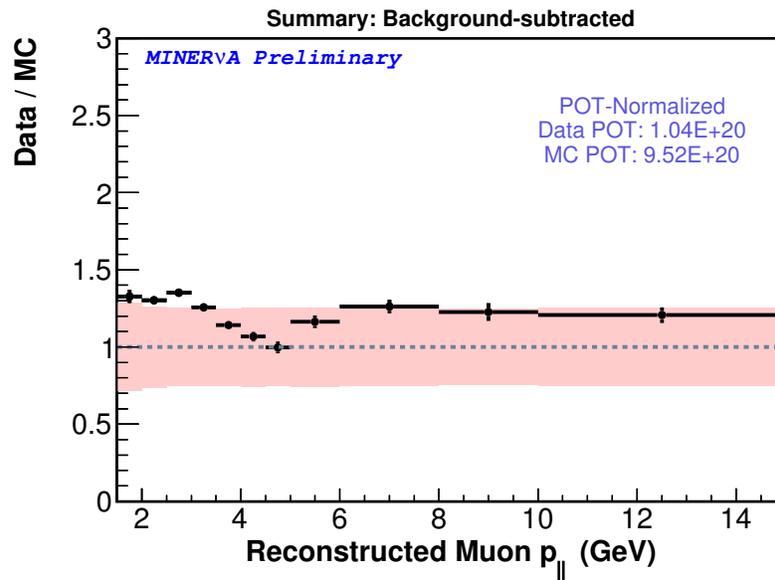


(b) Data/simulation ratio

Figure 5.9. Background-subtracted event counts in data and simulation vs. muon p_T (normalized to data POT).



(a) Event counts



(b) Data/simulation ratio

Figure 5.10. Background-subtracted event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).

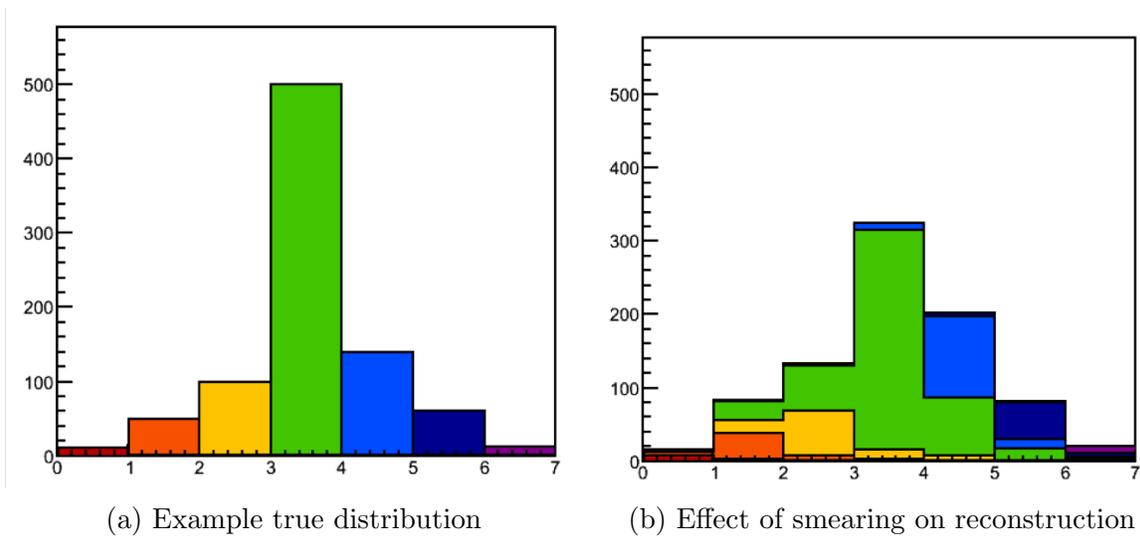


Figure 5.11. Cartoon to demonstrate how a small fractional migration from each bin leads to a smearing effect on the distribution’s shape

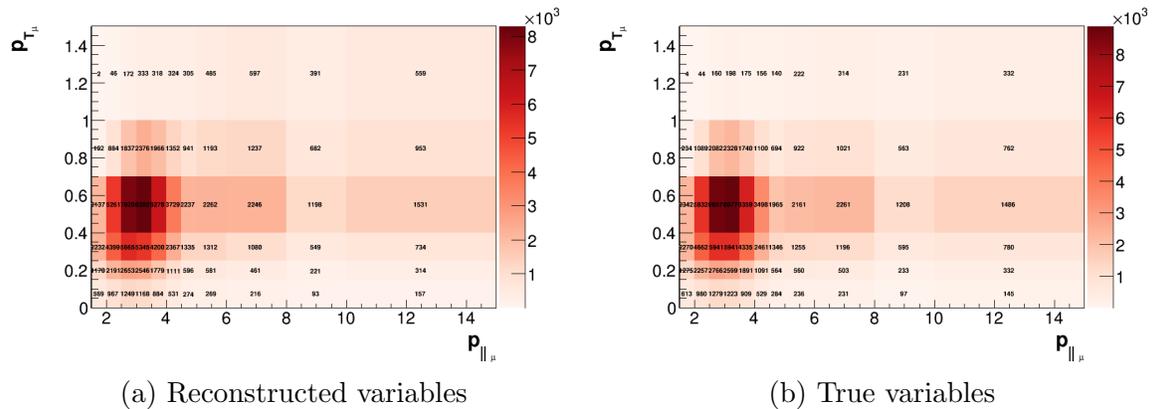


Figure 5.12. Monte Carlo events distribution versus reconstructed and true muon kinematics

The effect of this smearing can be characterized by generating a “migration matrix”, $U_{ij\alpha\beta}$ a histogram in which the x axis corresponds to the bin in which a given event was reconstructed, and the y axis to the bin in which that event was generated.

Note that in this case, a bin refers to a unique combination of p_T and p_{\parallel} , and that, furthermore, we must include overflow and underflow bins, in which either or both of the

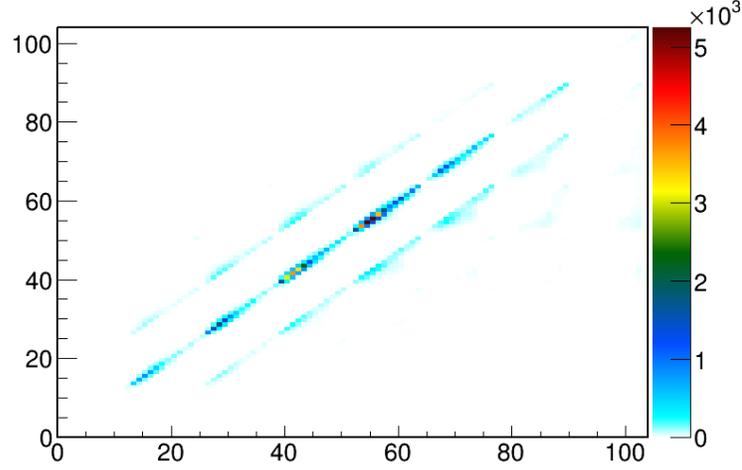


Figure 5.13. Migration matrix for the p_{\parallel} / p_T distribution. The x axis corresponds to reconstructed bins, the y to true.

p_T and p_{\parallel} values fall outside of our measurement range. Thus the total matrix is of size $N \times N$ bins, where the number of bins on each axis is related to the number of p_T and p_{\parallel} bins N_T and N_{\parallel} respectively by $N = (N_T + 2)(N_{\parallel} + 2)$.

The migration matrix for our simulated reconstructed signal distribution is shown in figure 5.13. The x axis indicates bins in the reconstructed variables, where the bins of p_{\parallel} are repeated for each bin of p_T , causing an appearance of $(N_T + 2) \times (N_T + 2)$ subplots. The y axis indicates bins in the true variables, arranged in the same way. Thus any events on the diagonal were reconstructed in the the correct bin of both p_{\parallel} and p_T . An event reconstructed in the wrong bin of p_{\parallel} (but the right p_T bin) will be displayed in another bin in the same subplot; one reconstructed in the wrong p_T bin (but the right p_{\parallel} bin) will appear in a different subplot, but in the same relative position within that subplot.

If each row of this matrix is normalized to 1, it represents the probability that an event generated in a given bin of p_{\parallel} and p_T will be reconstructed in each of the possible

p_{\parallel} / p_T bins. This smearing matrix $U_{ij\alpha\beta}$ can be represented by:

$$(5.4) \quad U_{ij\alpha\beta} = \frac{N_{\text{generated}, ij}^{\text{reconstructed}, \alpha\beta}}{N_{\text{generated}, ij}^{\text{reconstructed}}}$$

Theoretically, inverting this matrix would give an *unsmearing* (also known as *unfolding* - I will use the two terms interchangeably) matrix, corresponding to the probability that a Monte Carlo event *reconstructed* in a given bin of p_{\parallel} and p_T would have been *generated* in a given bin. We can postulate that applying this same unsmearing matrix to a data distribution would correct for the smearing effect on data:

$$(5.5) \quad N_{\text{data}, \alpha} = \sum_i U_{i\alpha}^{-1} N_{\text{data}}^{\text{reconstructed}, i}$$

This is a very large matrix to invert (for our 6 p_T and 11 p_{\parallel} bins, the migration matrix has $(13 \times 8)^2 = 10816$ entries. Additionally, we must consider that, as the unsmearing matrix is derived from the Monte Carlo simulation, the unfolding is dependent on the models used to generate this Monte Carlo, as well as the statistics of the simulated data. Using the matrix inversion method to unfold inflates the statistical uncertainty on the unfolded spectrum. Several alternative unfolding methods are available, which attempt to alleviate these issues. As with all previous MINERvA analyses, we use an iterative Bayesian method [81], implemented in MINERvA's MnvUnfold package, which in turn uses the ROOT package RooUnfold. [25]. This study uses four iterations of Bayesian unfolding.

The unfolding procedure was validated using an ensemble test, in which ten data-sized subsamples of the simulation were selected and warped by an adjustment to the axial mass simulation. These samples were then unfolded using the migration matrix

generated from the full un-warped simulation. We then looked at the mean of the pull in each bin between the unfolded and true distributions, for each of the ten samples, for different numbers of iterations of unfolding. As the number of iterations increases, the pulls tend to a stable value. However, using too many iterations is inadvisable as each successive iteration introduces statistical uncertainty. Thus four iterations was selected as a minimum number of iterations that would stabilize the unfolding. More details of the unfolding validation study are provided in appendix D.

The result of unfolding is reconstructed distributions corresponding to only CCQE-like events, plotted vs the true variables, corrected for smearing between bins due to detector effects, but not for loss of efficiency due to the detector's or the reconstruction algorithm's limitations).

5.4.1. Unfolded event distributions

Event distributions after unfolding, along with data/simulation ratios are shown in figures 5.14 (projected onto p_T), 5.15 (projected onto p_{\parallel}), 5.16 (projected onto Q_{QE}^2) and 5.17 (projected onto E_{ν}^{QE}). The effect of unfolding is to move events between bins (and possibly in and out of the sample, into the not-shown overflow and underflow). The unfolded Monte Carlo simulation corresponds to all signal events that pass our reconstruction cuts, plotted versus the true generated quantities. Uncertainties on the data are indicated by error bars; uncertainty on the Monte Carlo simulation is indicated by a pink shaded bar.

Systematic uncertainties on the data increase due to uncertainties that can move events between bins. For example, an uncertainty on the muon energy scale could cause p_{\parallel} and p_T to be measured slightly higher or lower than their true values, and events that

fall near bin boundaries to be reconstructed into the wrong bin. The effects are explained in detail in section 6.8.3.

At this point comes a choice: whether to unfold to true E_ν or to E_ν^{QE} calculated from the true muon kinematics. Whereas our reconstructed distributions have been the same for both, we must make a decision which variable to correct to. Both were studied, and the differences between them were slight; the plot here shows reconstruction to true E_ν .

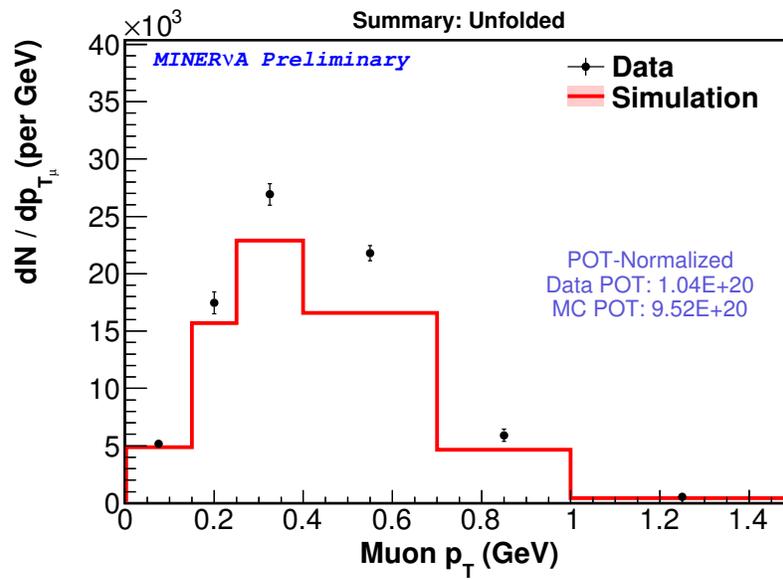
5.5. Efficiency and acceptance correction

5.5.1. Detector acceptance and reconstruction efficiency

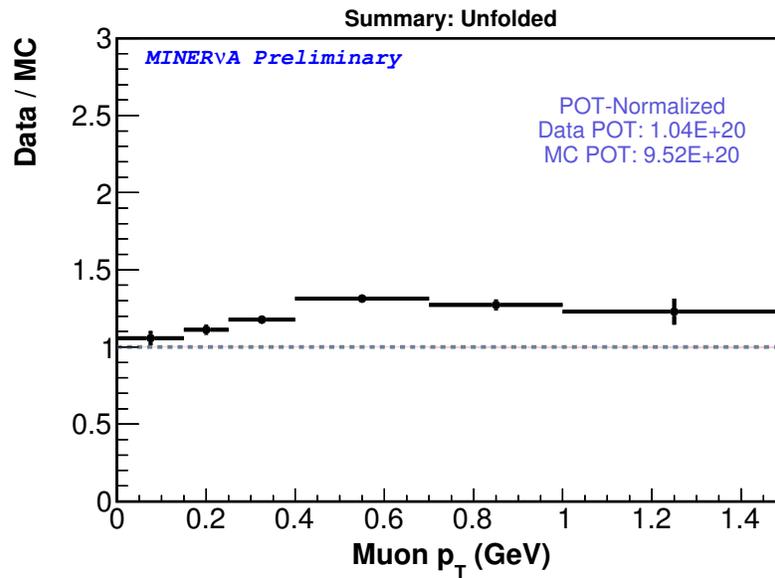
Ideally, to make our cross section measurement, we would detect and reconstruct every signal event during our run period. We are hampered from doing this by the limitations of both our detector's *acceptance*, and the *efficiency* of our reconstruction algorithms.

The detector acceptance refers to the fraction of signal events that we are able to see, based upon the limitations of the detector's technology. The largest effect on our acceptance in this analysis is from the fact that we require final-state muons to be matched with a track in the MINOS detector 2 m downstream, limiting the muon's angle with respect to the beam line to a maximum of around 20° . The MINOS-match requirement also limits our ability to accept muons with low longitudinal momentum, which will stop before they reach MINOS. The temporary loss of detector channels due to dead time also affects the detector acceptance.

The reconstruction efficiency refers to the fraction of signal events that we are able to reconstruct with our cuts. The largest loss of efficiency is due to the Q_{QE}^2 -dependent recoil cut, as some quasi-elastic-like events can have identical recoil signatures to background

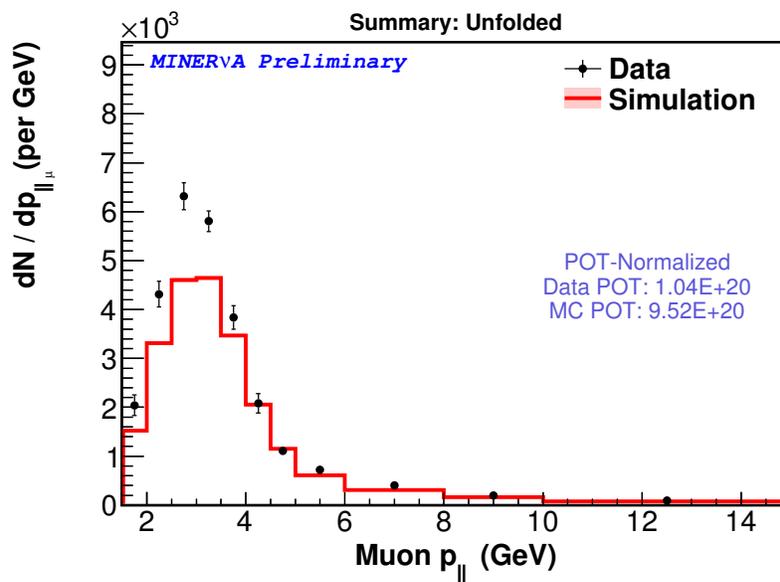


(a) Event counts

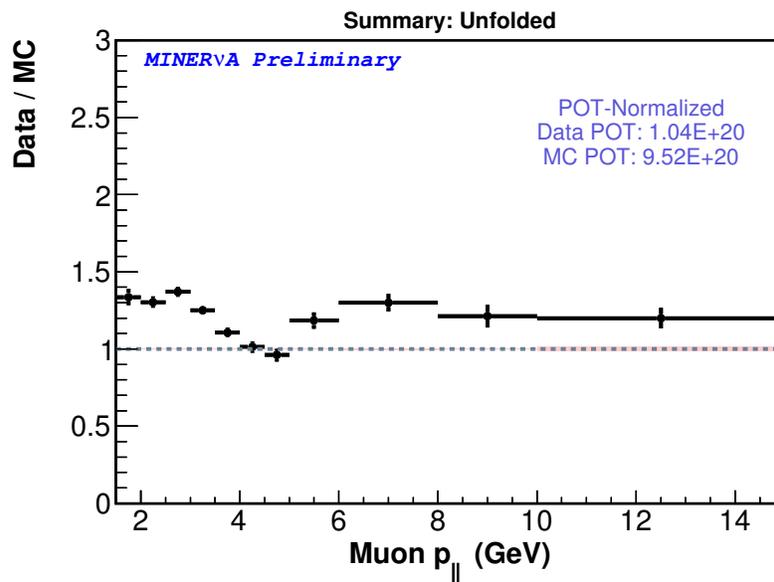


(b) Data/simulation ratio

Figure 5.14. Unfolded event counts in data and simulation vs. muon p_T (normalized to data POT).

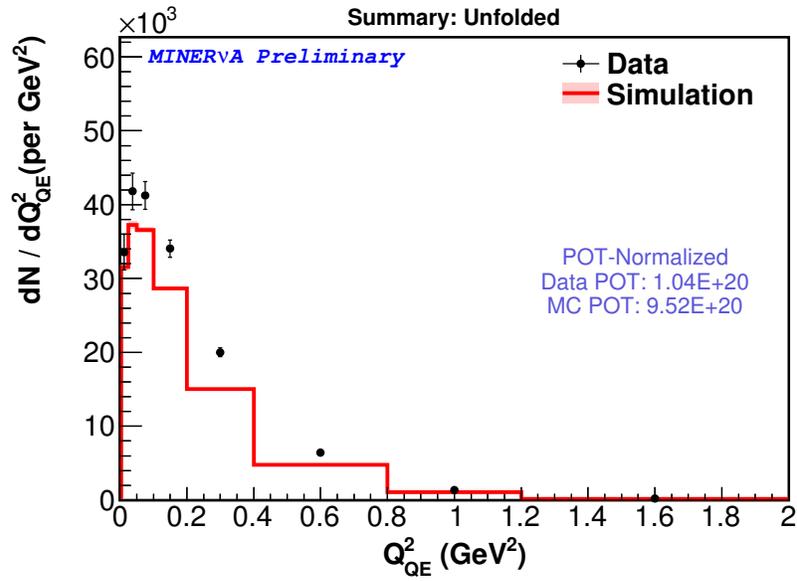


(a) Event counts

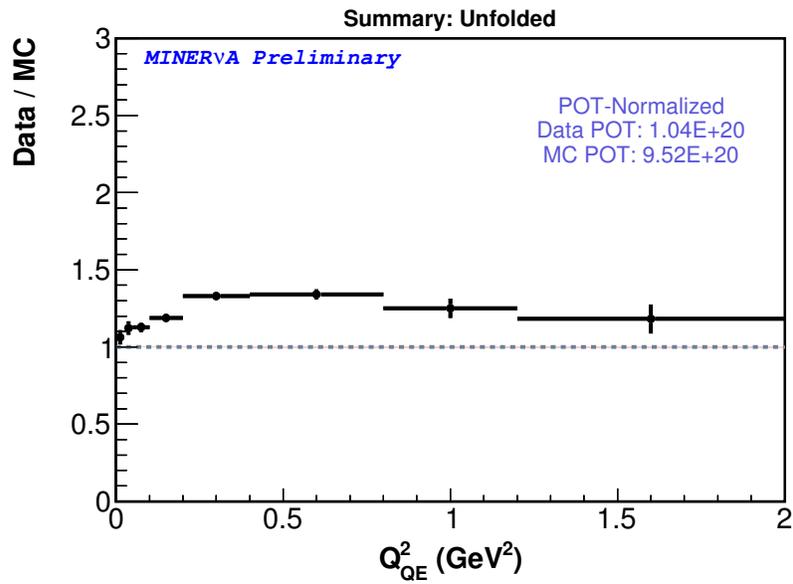


(b) Data/simulation ratio

Figure 5.15. Unfolded event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).

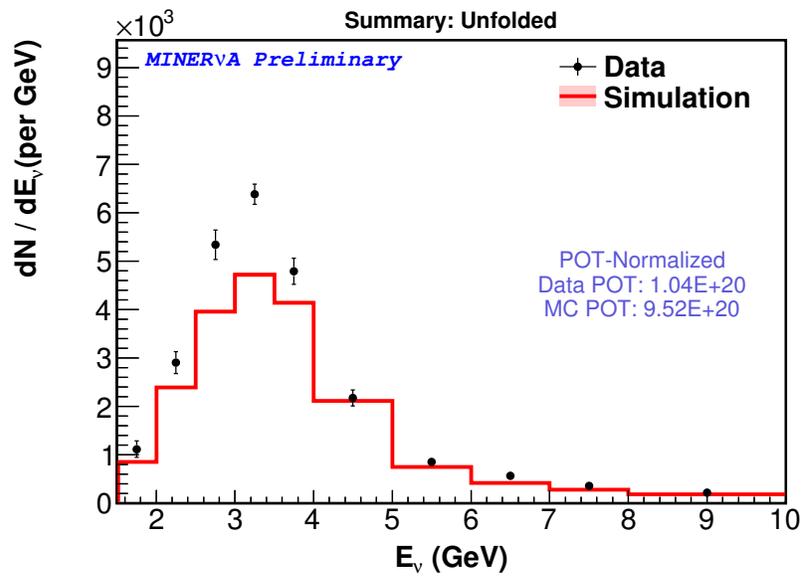


(a) Event counts

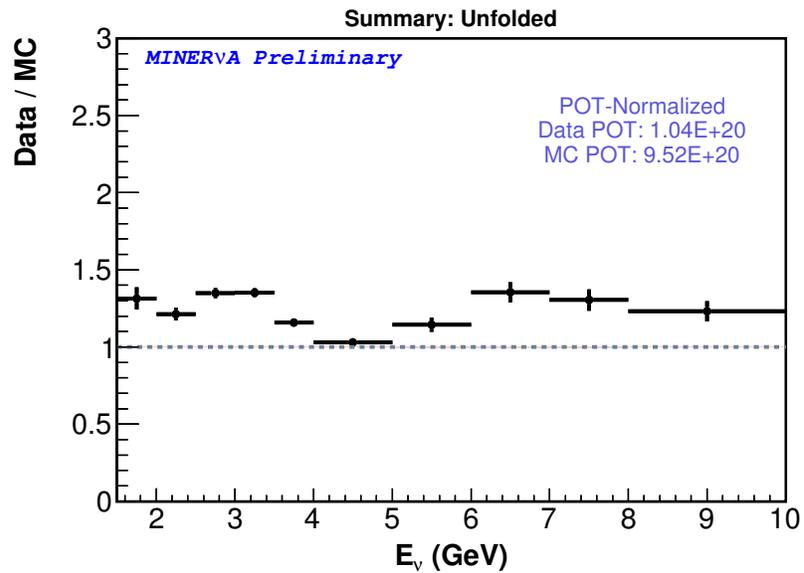


(b) Data/simulation ratio

Figure 5.16. Unfolded event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).



(a) Event counts



(b) Data/simulation ratio

Figure 5.17. Unfolded event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).

events. While loosening this cut would increase efficiency, it would be at the cost of purity, as background events moved into the selected sample. The cut has been chosen to maximize the product of efficiency and purity, while ensuring that the sample includes some quasi-elastic-like events that do not result from an initial quasi-elastic interaction (see the discussion in section 4.3.3.2. Reconstruction efficiency can also be lost when we reconstruct a track due to an allowed sub-120 MeV proton, where we fail to reconstruct the muon’s track in MINOS (which can happen in the case of MINOS event pile-up, when two tracks particles pass through MINOS at the same time). It is also possible to incorrectly reconstruct an interaction’s vertex outside of the fiducial volume, or to misreconstruct a muon’s forward momentum as being above 15 GeV (for very high-energy, straight-going muons that do not range out in MINOS, the curvature may be so slight that we are unable to make a good-quality momentum measurement).

As both the reconstruction efficiency and detector acceptance lead to the same effect — reducing the number of signal events that make it into our reconstructed sample — we instead calculate the product of acceptance and efficiency, corresponding to the fraction of signal events that we manage to include in our sample after reconstruction cuts. In the rest of this section, the words ‘efficiency’ and ‘acceptance’ will be used interchangeably to refer to this product. It is this product that is referred to by the term ϵ_{ij} in equation 5.1.

5.5.2. Correction for efficiency and acceptance

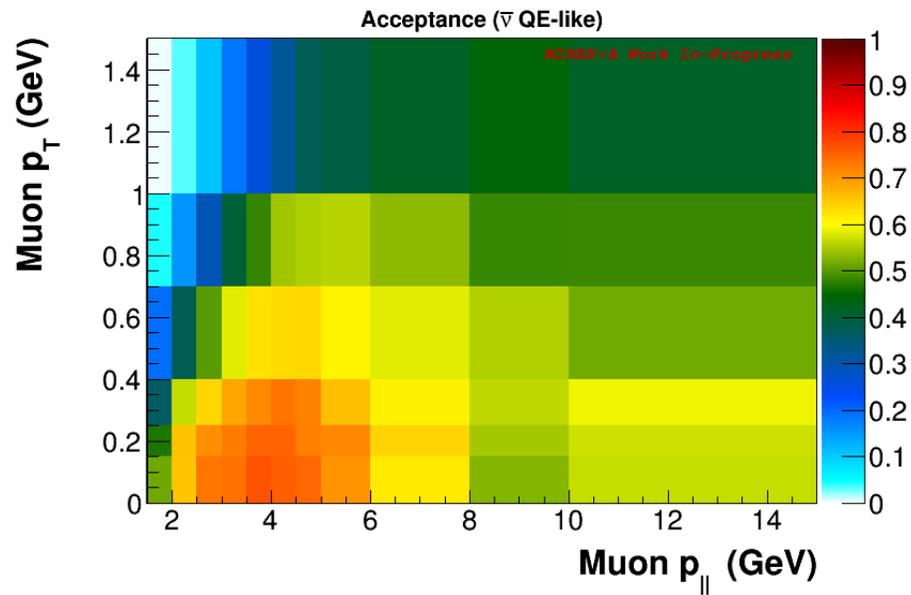
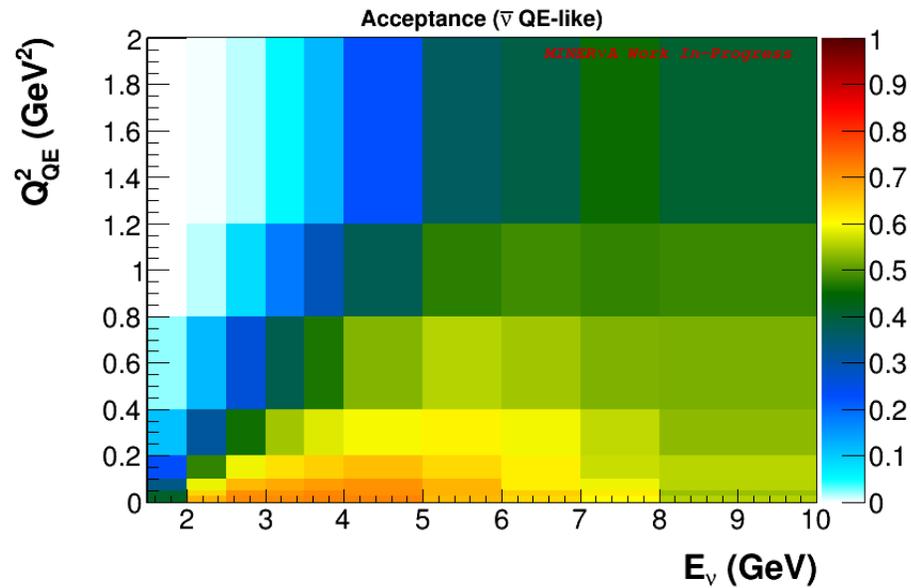
To estimate the efficiency, we again turn to the simulation. In the generated bin i, j , the efficiency is given by

$$(5.6) \quad \epsilon_{ij} = \frac{N_{ij}^{\text{generated and reconstructed}}}{N_{ij}^{\text{generated}}}$$

To determine the denominator of equation 5.6, the number of generated events, $N_{ij}^{\text{generated}}$, we look at our true simulated distribution (stored in a separate Truth ntuple, as explained in 4.2.2, which includes generator-level information about the all simulated interactions, including those that were not reconstructable, typically because of no MINOS match). We apply our full signal cut (including the fiducial volume cut) and histogram the event count vs. the true variables. When constructing this distribution, we do not include the tracking efficiency weight (a reconstruction effect) in the event weight.

To get the numerator of equation 5.6, $N_{ij}^{\text{generated and reconstructed}}$, we select all events in the reconstructed Monte Carlo ntuple that pass both the signal cuts and the reconstruction cuts. In this case, when making the signal cut, we do not require the true vertex to be within the fiducial volume. This selection corresponds to the background-subtracted Monte Carlo. We histogram the event count vs. the true variables, giving a distribution that corresponds to the background-subtracted, unfolded Monte Carlo.

Figure 5.18 shows the product of efficiency and acceptance vs. p_{\parallel} and p_T (figure 5.18a) and vs. E_{ν} and Q_{QE}^2 (figure 5.18b). The extremely poor acceptance at high p_T / low p_{\parallel} (which corresponds approximately with the high Q_{QE}^2 / low E_{ν}^{QE} region; see figure 4.11) is due to the MINOS match requirement; these are high-angle muons that will bypass MINOS. Efficiency also decreases at higher energies, where interactions are more inclined to generate large amounts of recoil energy, which may be vetoed by our Q_{QE}^2 -dependent recoil cut as they mimic the recoil distributions of non-quasi-elastic-like events. Again,

(a) Vs. p_{\parallel} and p_T (b) Vs. E_{ν} and Q_{QE}^2 Figure 5.18. Efficiency \times acceptance of our reconstructed sample

the acceptance was calculated versus both E_{ν} and E_{ν}^{QE} , with only slight differences; E_{ν} is shown.

The overall efficiency \times acceptance of the sample is 50.6%.

5.5.3. Efficiency-corrected distributions

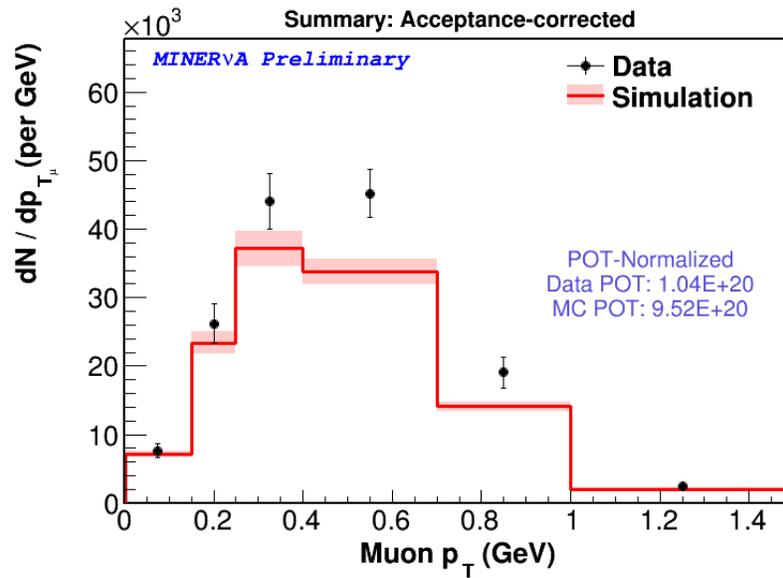
Event distributions after efficiency correction, along with data/simulation ratios, are shown in figures 5.19 (projected onto p_T), 5.20 (projected onto p_{\parallel}), 5.21 (projected onto Q_{QE}^2) and 5.22 (projected onto E_{ν}). The effect of acceptance correction is to add back in signal interactions that we predict to have taken place, but to have not been reconstructed. The efficiency-corrected Monte Carlo simulation corresponds to the true distribution produced by the event generator. Uncertainties on the data are indicated by error bars; uncertainty on the Monte Carlo simulation is indicated by a pink shaded band. Section 6.8.4 explains how efficiency correction affects uncertainties - its most significant effect is to increase statistical uncertainty in bins with very poor acceptance.

5.6. Cross section calculation from acceptance-corrected distribution

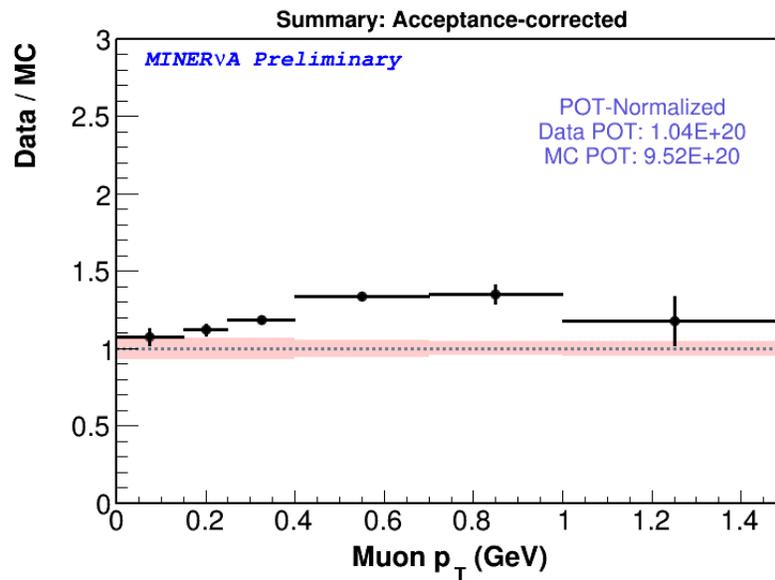
To convert an acceptance-corrected distribution to a cross section, we must divide by the number of targets (protons, in this case), the total number of protons on target (POT), and the flux per POT. These are summarized in table 5.2 and explained in this section.

Quantity	Value
Protons on target (data)	1.041×10^{20}
Protons on target (simulation)	9.515×10^{20}
Number of targets	1.76474×10^{30} protons
Integrated flux	$2.36008 \times 10^{-8} \bar{\nu}_{\mu}/\text{cm}^2 / \text{POT}$

Table 5.2. Quantities we divide by to convert our acceptance-corrected distributions to flux-integrated cross sections

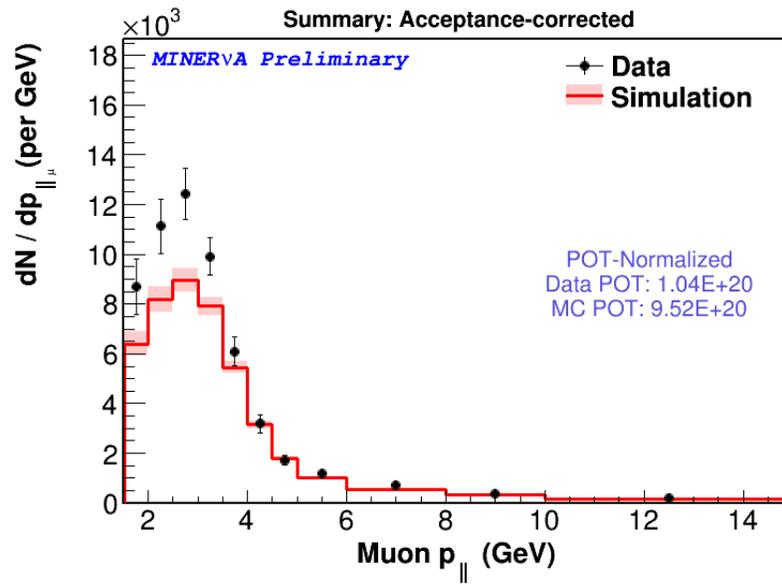


(a) Event counts

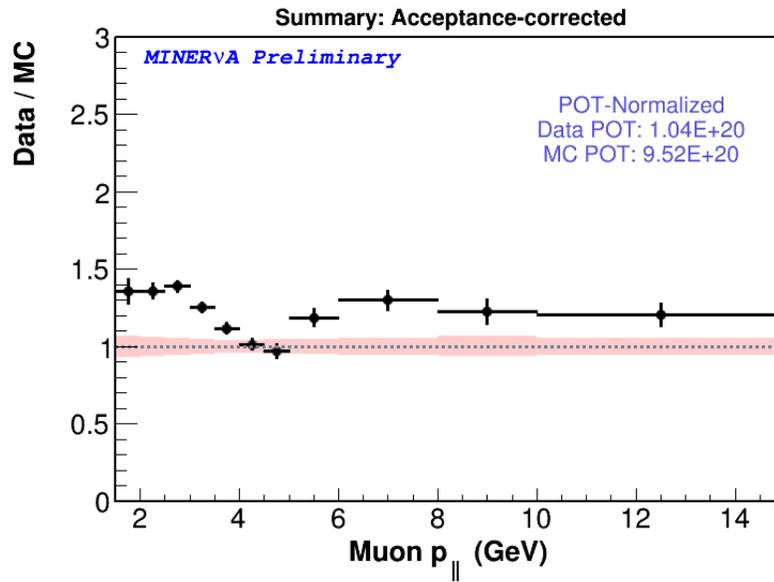


(b) Data/simulation ratio

Figure 5.19. Acceptance- and efficiency-corrected event counts in data and simulation vs. muon p_T (normalized to data POT).

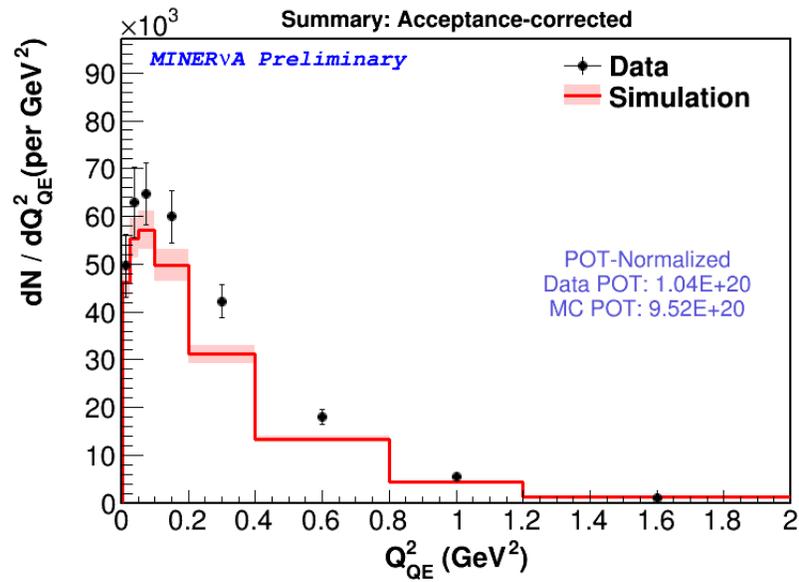


(a) Event counts

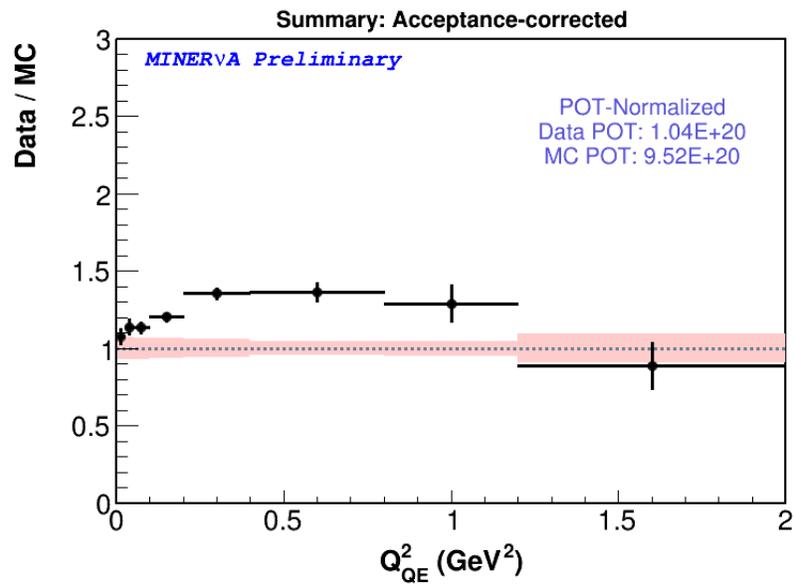


(b) Data/simulation ratio

Figure 5.20. Acceptance- and efficiency-corrected event counts in data and simulation vs. muon p_{\parallel} (normalized to data POT).

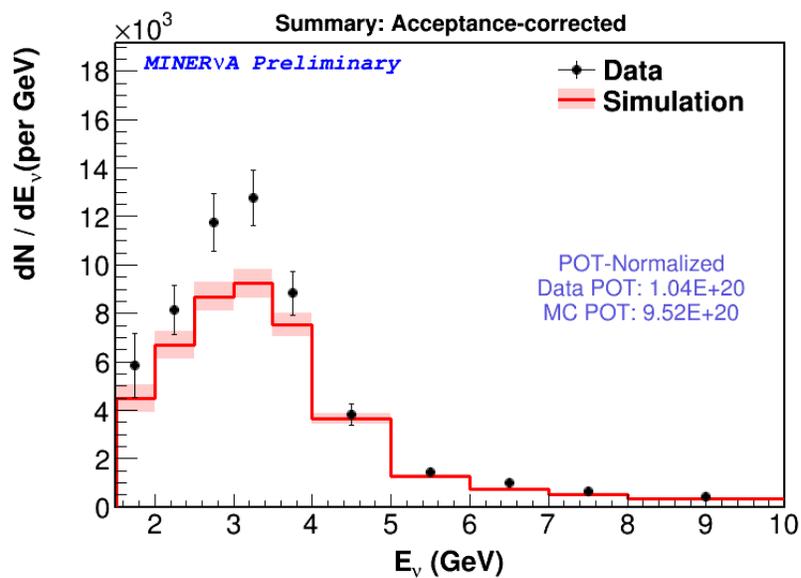


(a) Event counts

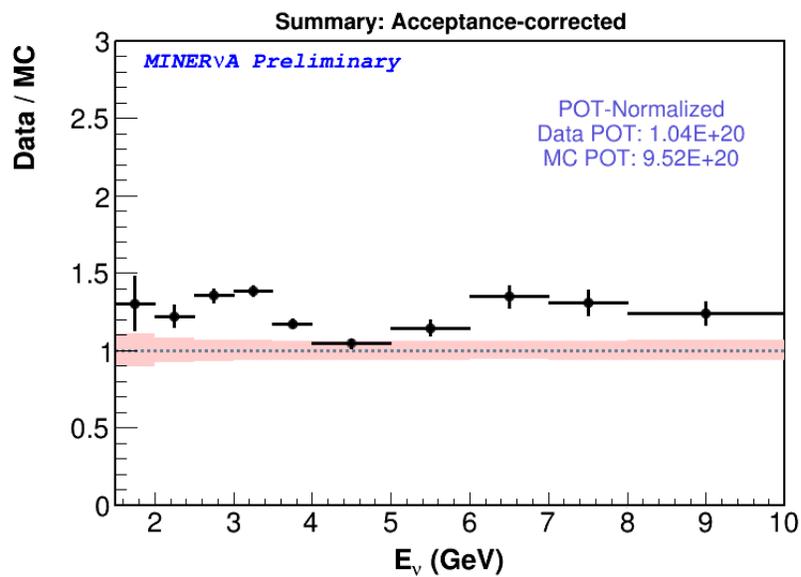


(b) Data/simulation ratio

Figure 5.21. Acceptance- and efficiency-corrected event counts in data and simulation vs. Q_{QE}^2 (normalized to data POT).



(a) Event counts



(b) Data/simulation ratio

Figure 5.22. Acceptance- and efficiency-corrected event counts in data and simulation vs. E_ν^{QE} (normalized to data POT).

5.6.1. The antineutrino flux

The “flux” Φ refers to the energy spectrum of the antineutrino beam. The NuMI beam’s “PPFX” flux prediction is explained in detail in [36], and is summarized in section 3.7.1 of this document. In order to convert a measured (background-subtracted, unfolded and acceptance-corrected) event count, we must divide by the number of incoming neutrinos per unit area that generated the distribution. The flux prediction from PPFX is given per proton on target; we therefore multiply this by the total number of protons on target (1.041×10^{20}) for our sample. Figure 5.23 shows this flux distribution as a function of incoming neutrino energy.

(Note that this refers to true neutrino energy E_ν , which differs from the neutrino energy estimate from muon kinematics in the quasi-elastic hypothesis, E_ν^{QE} . Thus, this flux weighting is strictly correct for the distributions vs E_ν , and slightly incorrect for that vs E_ν^{QE} . The differences will be shown below, and are very subtle.)

There are two alternative ways in which we can divide by the flux to obtain a cross section. Previous MINERvA quasi-elastic measurements [96][107][191] reported a flux-integrated cross section, meaning that they integrated the total flux distribution in some energy range and divided by the resulting number, which corresponded to the number of neutrinos per POT per unit area in that entire range. This was a practical solution as these analyses measure the single-differential cross section $d\sigma/dQ_{QE}^2$, which is not directly correlated with neutrino energy. (They are not, however, entirely unrelated; as these distributions show, the neutrino energy places restrictions on the available four-momentum that can be transferred, and when publishing a flux-integrated cross section

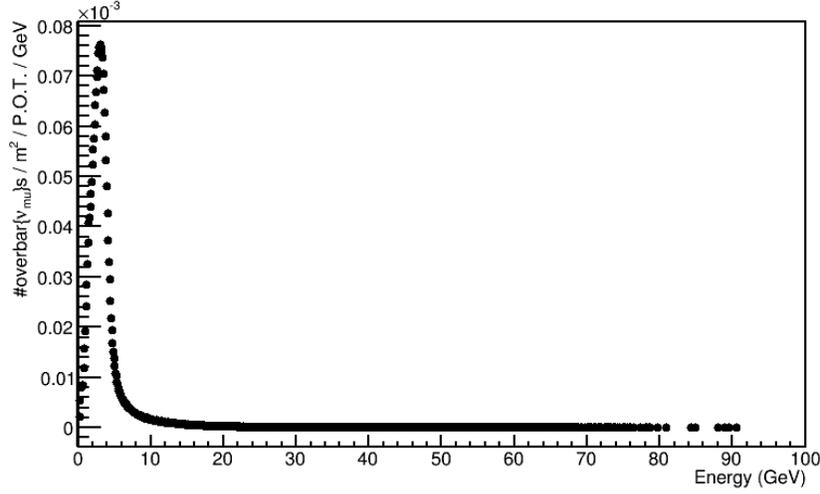


Figure 5.23. The energy spectrum (flux) of antineutrinos used in this analysis. This flux is given in units of the number of antineutrinos per m^2 per GeV per proton on target.

such as these, it is necessary to also publish the flux in order to understand how exactly the cross section was generated.)

For distributions in the p_T / p_{\parallel} phase space, we report flux-integrated cross sections. We do the same for the single-differential cross section $d\sigma/dQ_{QE}^2$. We integrate over the entire available flux range of 0-100 GeV, to get a total integrated flux of $2.36008 \times 10^{-8} \text{ cm}^{-2}$ per proton on target.

For the E_{ν}^{QE} / Q_{QE}^2 distribution, flux integration makes little sense. While E_{ν}^{QE} is not exactly equal to true $E\nu$, they are very highly correlated (see figure 5.24). In this case, the more relevant distribution is a flux-weighted cross section, where the number of events in each bin is scaled down by the flux in that bin (in which case we can produce a total cross section vs. neutrino energy $\sigma(E_{\nu})$ or $\sigma(E_{\nu}^{QE})$ rather than a differential cross section $d\sigma/dE_{\nu}$ or $d\sigma/dE_{\nu}^{QE}$).

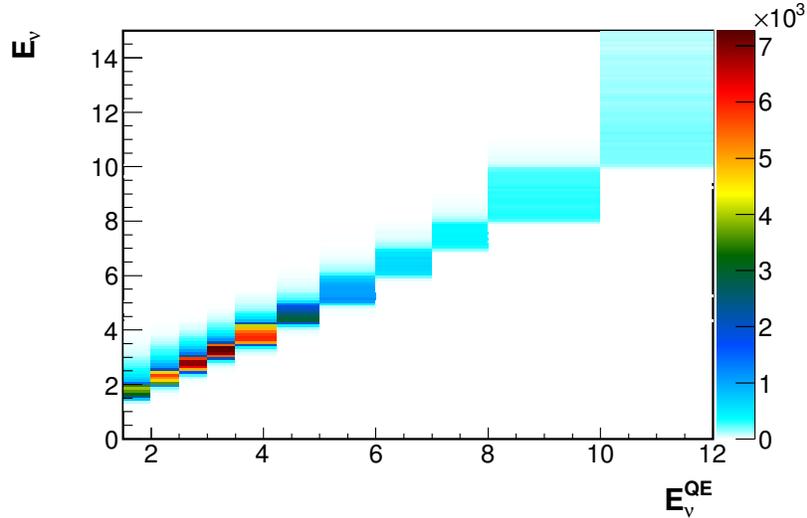


Figure 5.24. Mapping of the true E_ν (in the binning of the flux histogram) to E_ν^{QE} reconstructed in the quasi-elastic hypothesis, from our truth-level simulation ntuple

The flux correction introduces a significant increase in systematic uncertainty due to the flux, while leaving all other fractional uncertainties constant, as explained in section 6.8.5.

The fluxes used for this normalization are central-value weighted fluxes from the PPFX (generation 2, thin target) flux, generated with the MINERvA ‘compute_flux’ script, without the neutrino-electron scattering constraint.

5.6.2. Number of targets

The final stage in the cross section is to divide by the number of targets, T , that are available for a neutrino to interact with. The target for an antineutrino quasi-elastic scatter ($\bar{\nu}_\mu p \rightarrow \mu^+ n$) is a proton. The scintillator tracker (on which this cross section is measured) is not made of a single element, but of a polymer, polystyrene, together with

doping agents, epoxy and light-tight coating (its composition is listed in 3.1). However, the vast majority consists of polystyrene molecules, in which there are an equal number of carbon and hydrogen atoms. The fiducial volume contains 1.76474×10^{30} protons, of which 1.44182×10^{30} are in carbon nuclei, and 2.44088×10^{29} are in hydrogen. The uncertainty in the mass model (1.4%) is taken into account by the systematic uncertainty analysis.

Note that for quasi-elastic-like scattering, it is possible that a scattering process could originate on a neutron (e.g. $\bar{\nu}_\mu n \rightarrow \mu^+ \Delta^-$) where the resonance decays $\Delta^- \rightarrow n\pi^-$ and the pion is absorbed). We may also scatter from a correlated nucleon pair. Thus, while we use the number (1.76474×10^{30}) of protons as our target number, these factors must be included in our simulation.

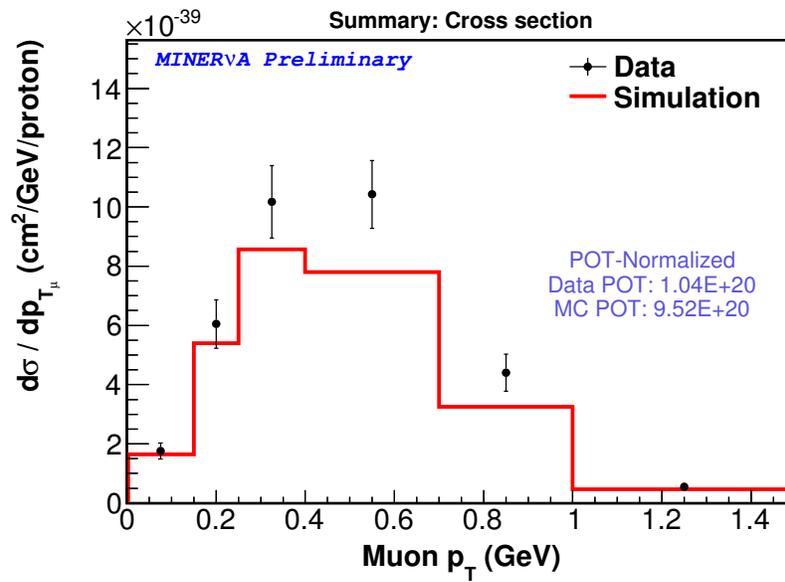
5.7. Cross section distributions

5.7.1. Quasi-elastic-like differential cross sections vs. p_T and p_{\parallel}

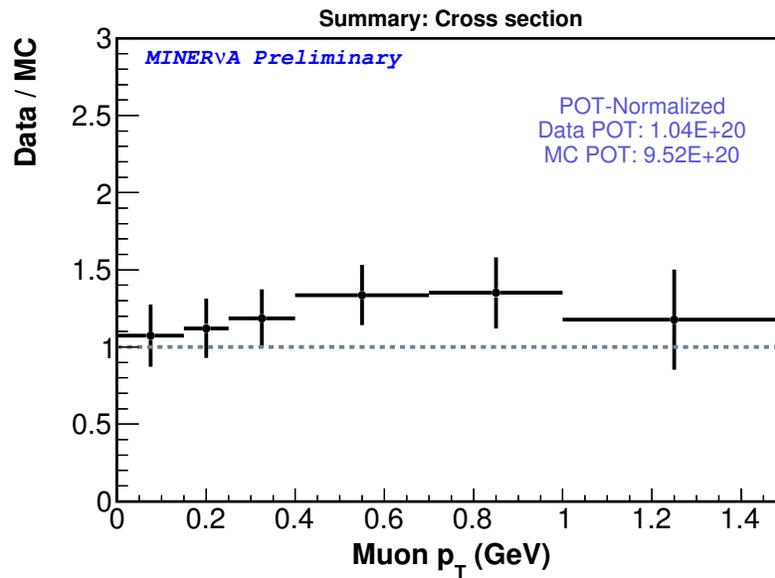
Differential cross sections vs. p_T and p_{\parallel} and data/simulation ratios are shown in figures 5.25 ($d\sigma/dp_T$) and 5.26 ($d\sigma/dp_{\parallel}$). Uncertainties on the data are indicated by error bars. At this stage, we do not report an uncertainty on the simulation; this uncertainty is precisely what we hope to reduce with this analysis.

5.7.2. Quasi-elastic-like cross sections vs. Q_{QE}^2 and E_ν

This section shows cross section distributions presented vs. Q_{QE}^2 and E_ν . Uncertainties on the data are indicated by error bars. No uncertainty is reported on simulation.

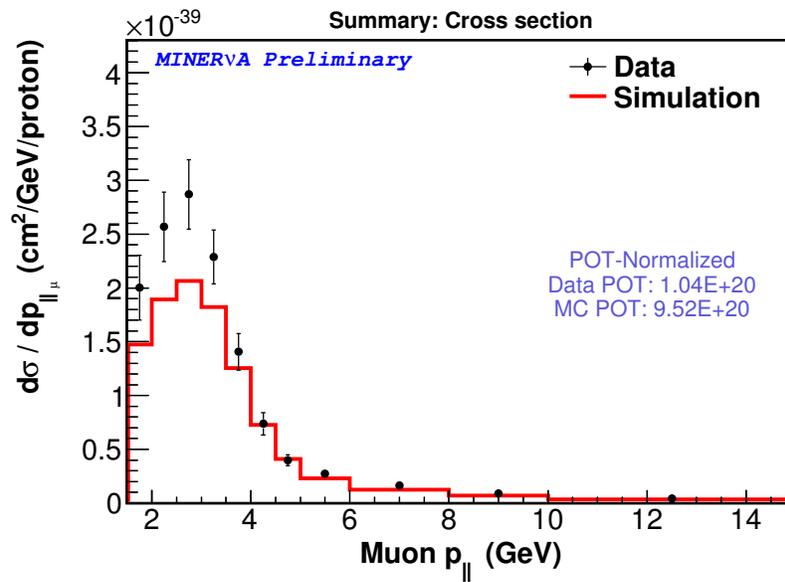


(a) Cross section

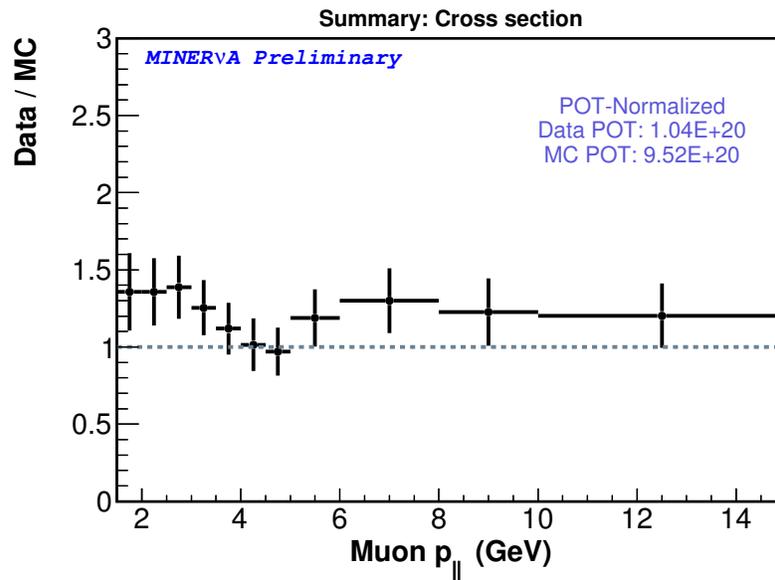


(b) Data/simulation ratio

Figure 5.25. Differential quasi-elastic-like cross section in data and simulation vs. muon p_T (normalized to data POT).



(a) Cross section



(b) Data/simulation ratio

Figure 5.26. Differential quasi-elastic-like cross section in data and simulation vs. muon $p_{||}$ (normalized to data POT).

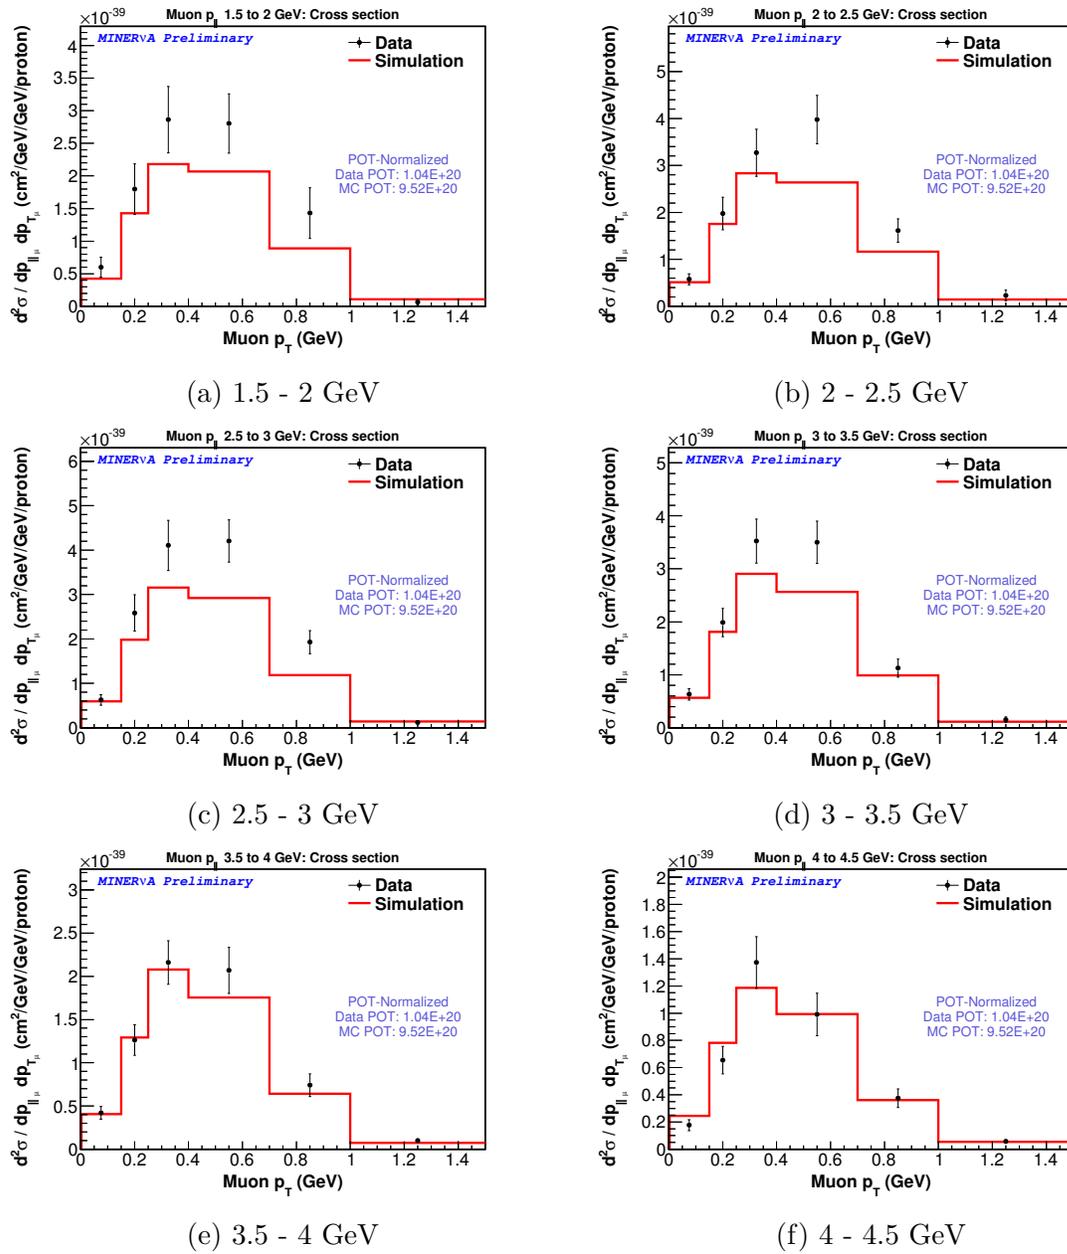
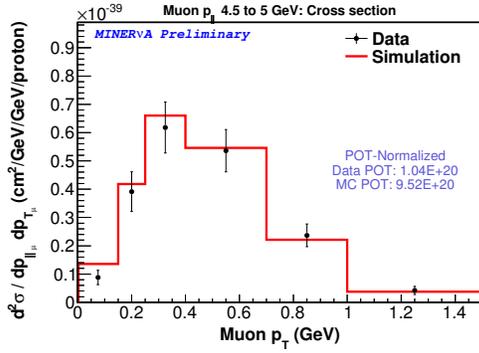
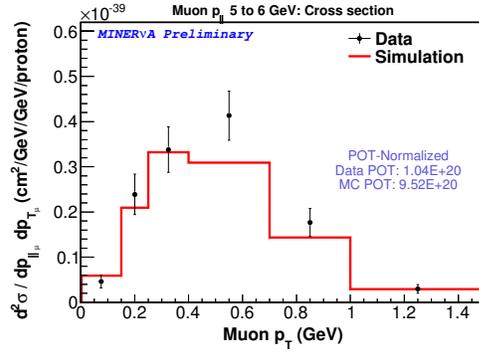


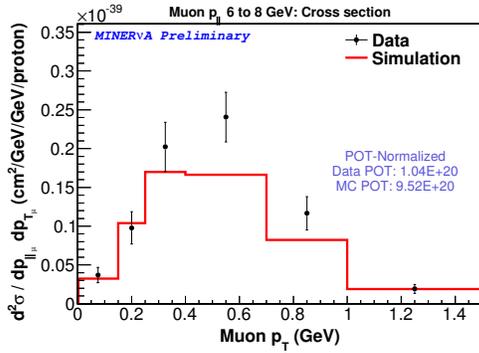
Figure 5.27. Double-differential quasi-elastic-like cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)



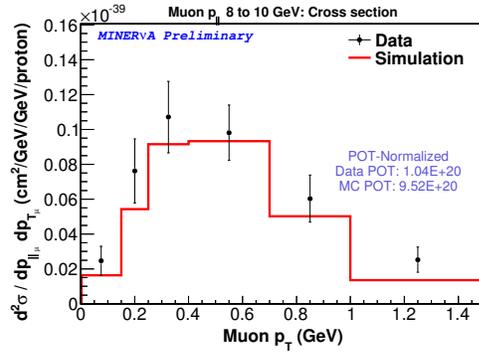
(g) 4.5 - 5 GeV



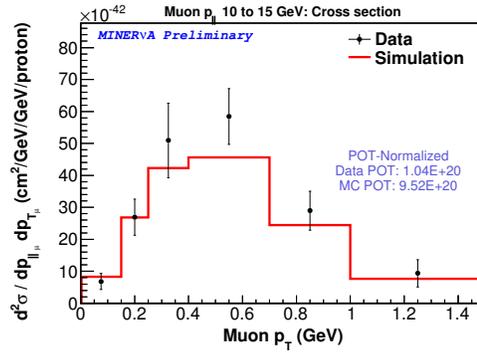
(h) 5 - 6 GeV



(i) 6 - 8 GeV



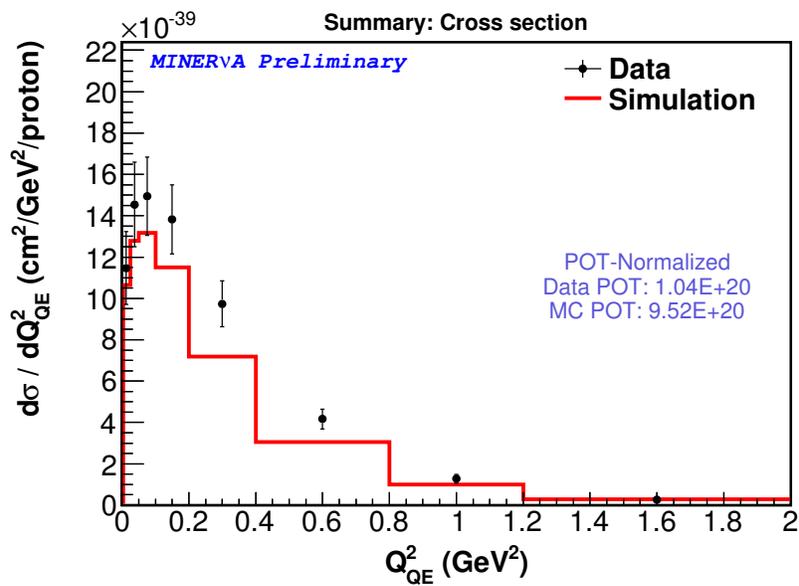
(j) 8 - 10 GeV



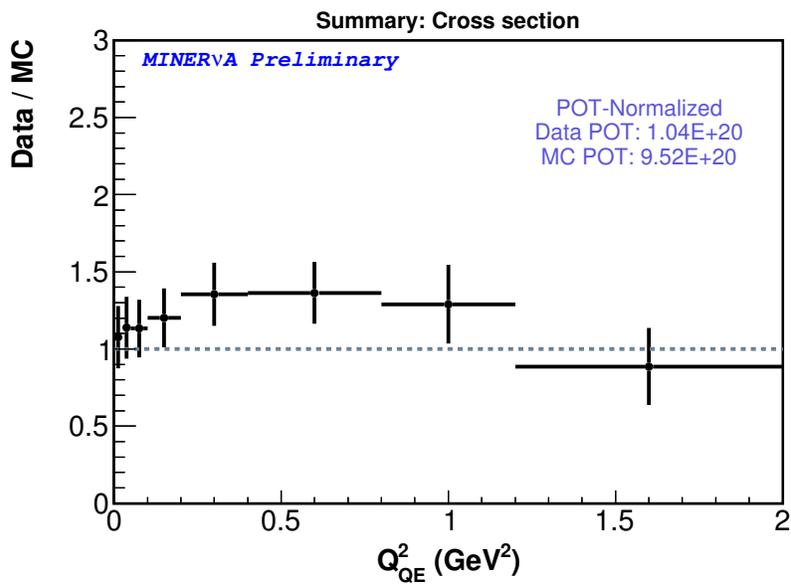
(k) 10 - 15 GeV

Figure 5.27. Double-differential quasi-elastic-like cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)

The single-differential cross section, $d\sigma/dQ_{QE}^2$ is shown in figure 5.28. This is a flux-integrated cross section, similar to those in section 5.7.1, and is generated by projecting a double-differential flux-integrated $d^2\sigma/dQ_{QE}^2 dE_\nu$ distribution onto the Q_{QE}^2 axis.



(a) Cross section



(b) Data/simulation ratio

Figure 5.28. Flux integrated quasi-elastic-like differential cross section $d\sigma/dQ_{QE}^2$ in data and simulation (normalized to data POT).

For the cross section vs the neutrino energy, we present a flux-weighted total cross section $\sigma(E_\nu)$, as explained in section 5.6.1, rather than an integrated differential cross section.

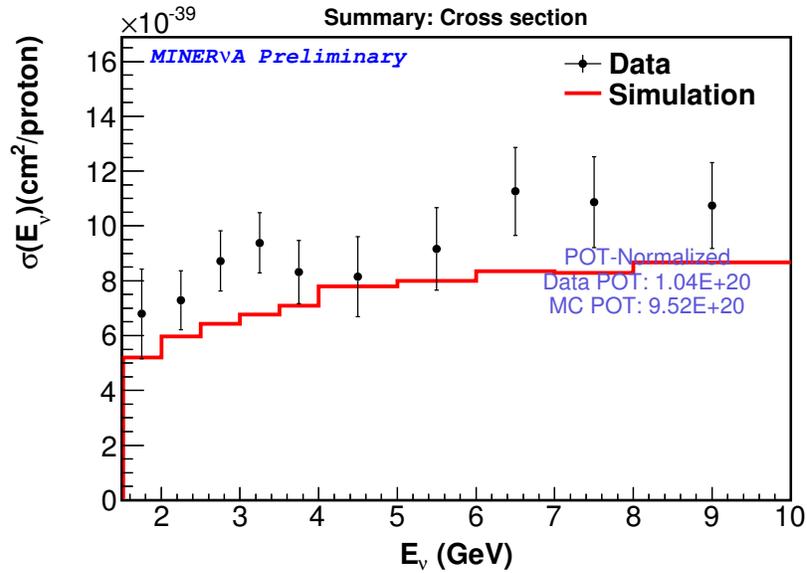
For comparison, we also present the cross section where we unfold and efficiency-correct to E_ν^{QE} rather than to E_ν ; a more valid procedure for the unfolding correction, given that we reconstructed E_ν^{QE} from muon kinematics, but a less valid procedure for the flux-weighting. As can be seen, there is little difference between the $\sigma(E_\nu)$ and $\sigma(E_\nu^{QE})$.

The double differential cross sections are also flux-weighted, displaying $d\sigma(E_\nu)/dQ_{QE}^2$.

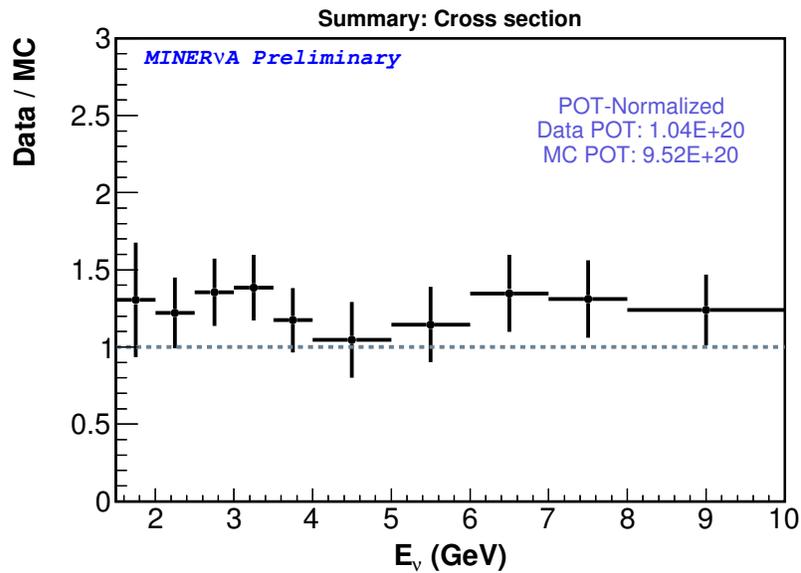
The following chapter will discuss how these cross section measurements compare with those from other experiments.

5.7.3. True CCQE cross sections

The main focus of the analysis was the calculation of a quasi-elastic-like double-differential cross sections shown above, which correspond to our measurement for the signal definition described in section 4.1.3. As an extension to the analysis, however, we also calculated a true CCQE cross section, using the signal definition described in section 4.1.2. Recall that, for the quasi-elastic-like cross section, our signal corresponded to interactions with a quasi-elastic-like final state, even if that final state was generated by a resonant or DIS interaction followed by FSI. For the true CCQE definition, our signal corresponds only to events where the initial interaction was quasi-elastic, even if FSI created final-state particles such as pions that mimicked a non-quasi-elastic interaction. The signal also includes 2p2h events where a CCQE interaction takes place on a correlated pair. These

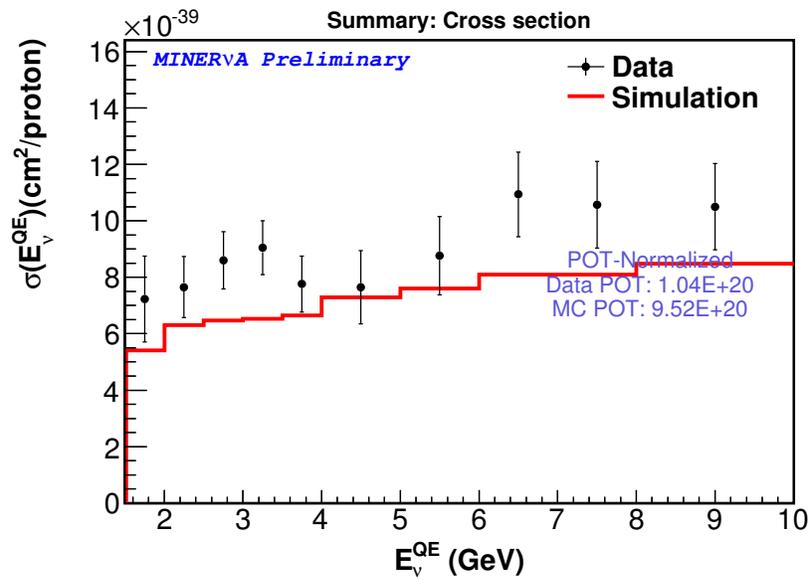


(a) Cross section

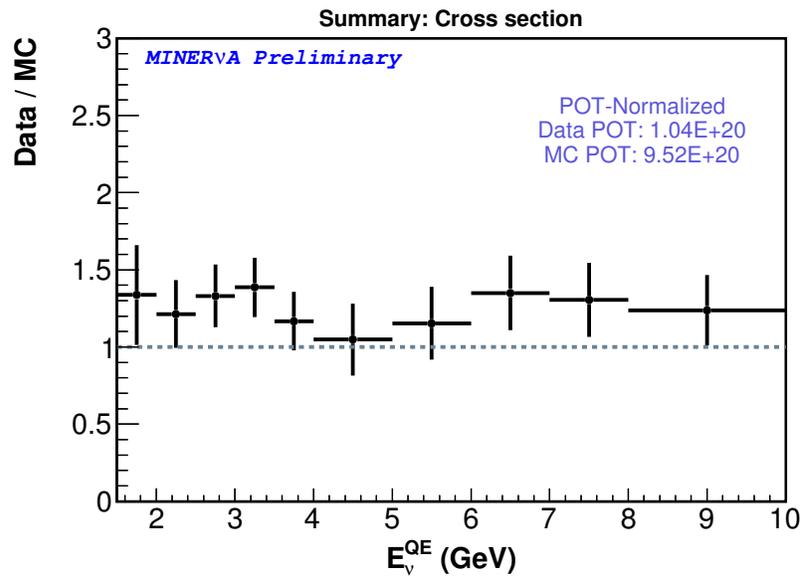


(b) Data/simulation ratio

Figure 5.29. Energy-dependent total quasi-elastic-like cross section $\sigma(E_\nu)$ in data and simulation(normalized to data POT).



(a) Cross section



(b) Data/simulation ratio

Figure 5.30. Energy-dependent total quasi-elastic-like cross section $\sigma(E_\nu^{QE})$ in data and simulation(normalized to data POT).

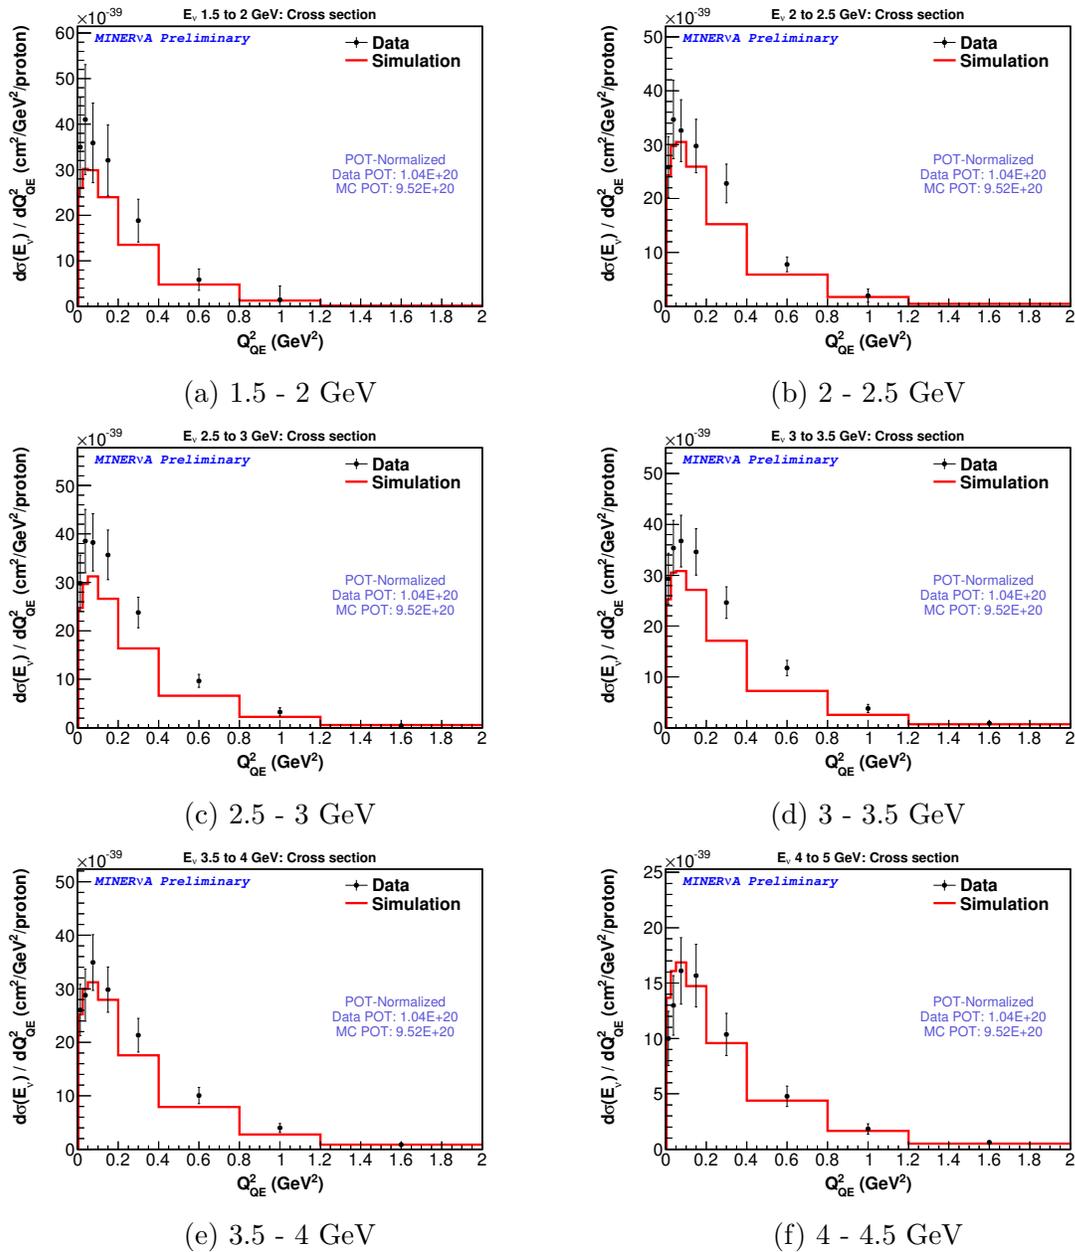


Figure 5.31. Double differential quasi-elastic-like cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued in next figure)

are not included in the current GENIE simulation; a systematic uncertainty is included to take account of the effect this has on backgrounds and acceptance.

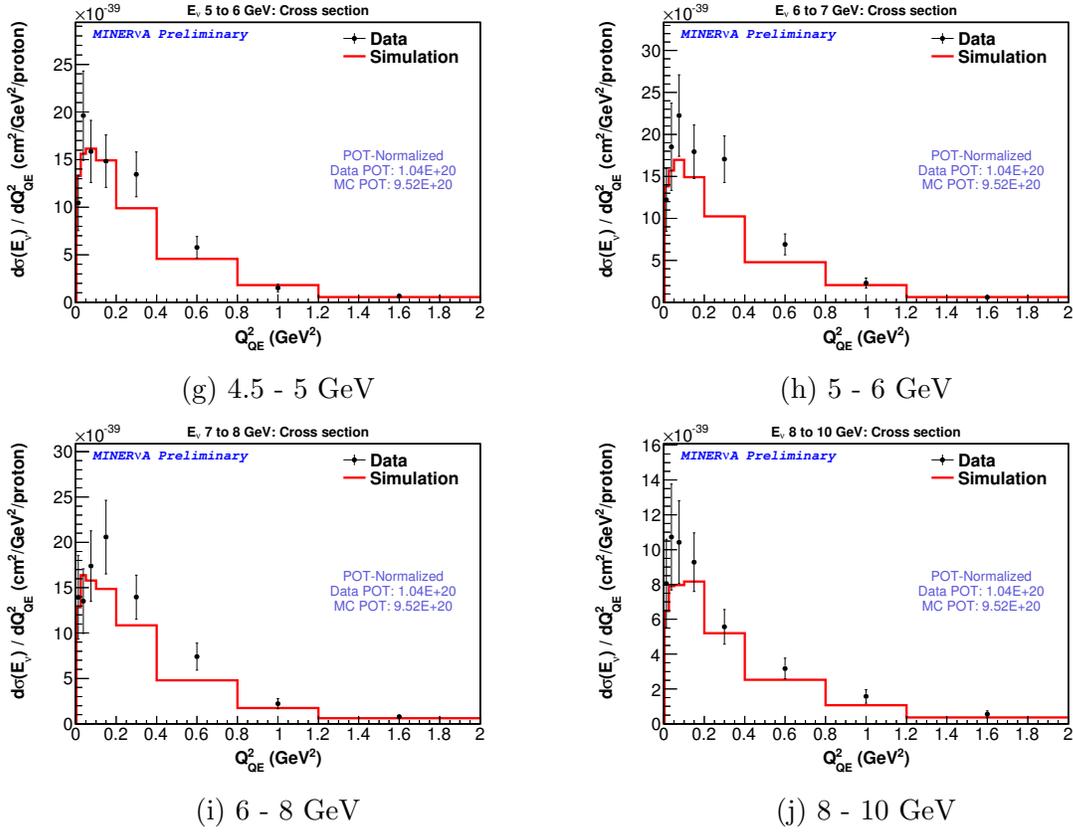


Figure 5.31. Double differential quasi-elastic-like cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued)

The true CCQE single-differential cross section projections are shown in figure 5.32, while double differential cross sections are shown in figure 5.33 ($d_\sigma^2/dp_T dp_\parallel$), and 5.34 ($d\sigma(E_\nu)/dQ_{QE}^2$).

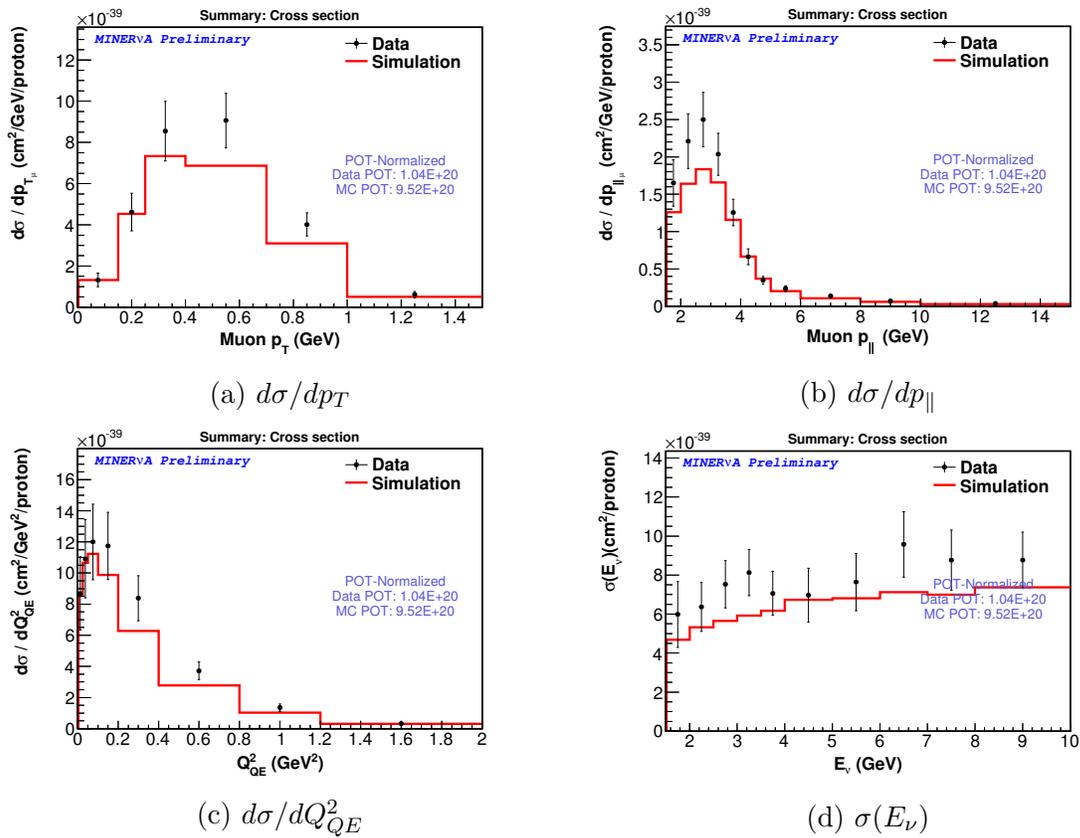


Figure 5.32. True CCQE cross section measurements in data and simulation. As before, the differential cross sections are flux-integrated.

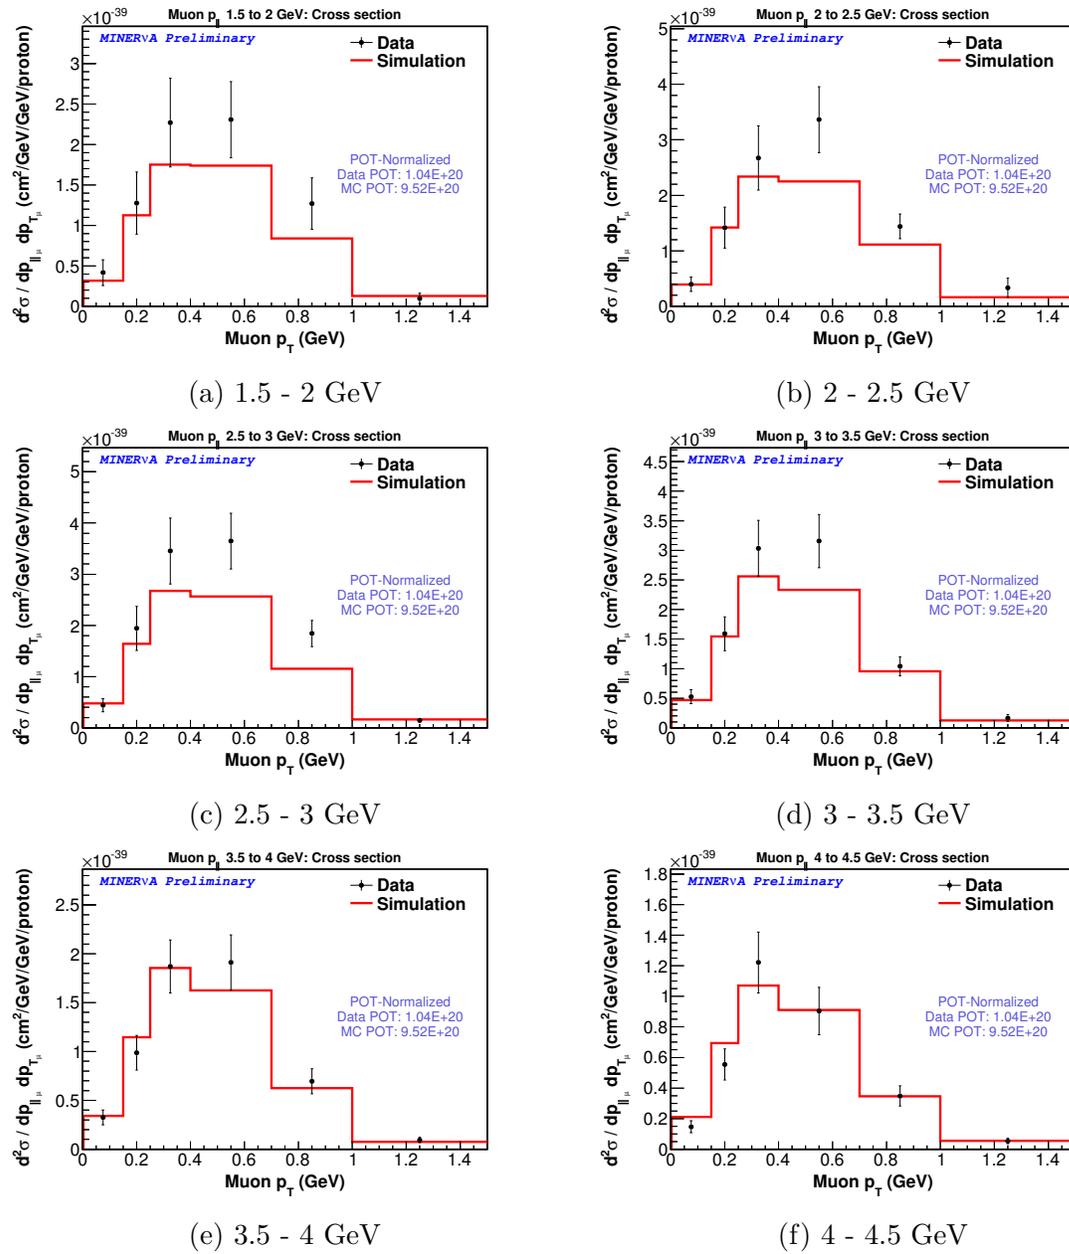
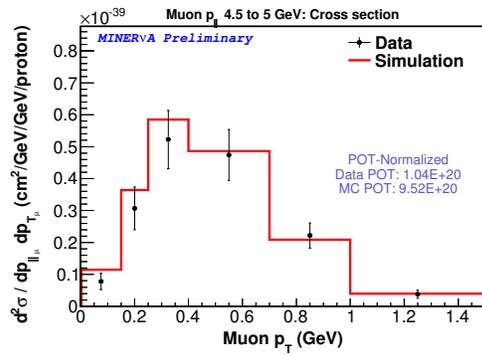
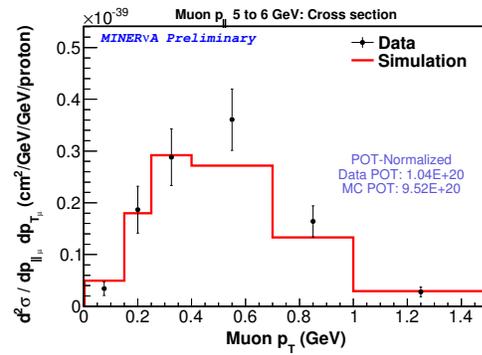


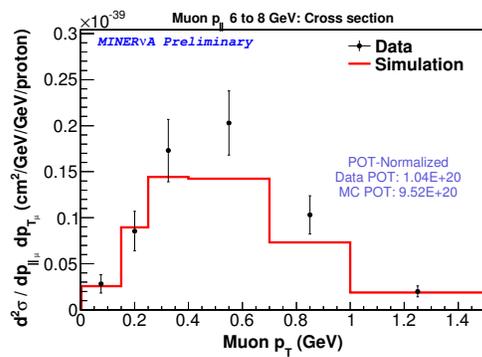
Figure 5.33. Double-differential flux-integrated true CCQE cross section $d^2\sigma/dp_T dp_{\parallel}$ vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)



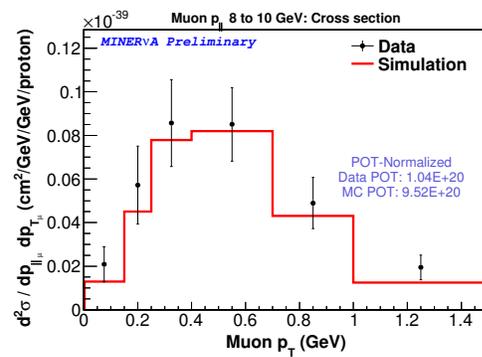
(g) 4.5 - 5 GeV



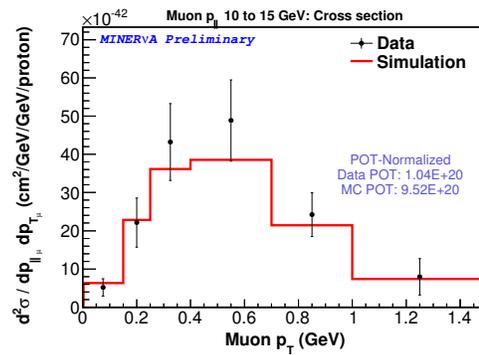
(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV



(k) 10 - 15 GeV

Figure 5.33. Double-differential flux-integrated true CCQE cross section $d^2\sigma/dp_T dp_{||}$ vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)

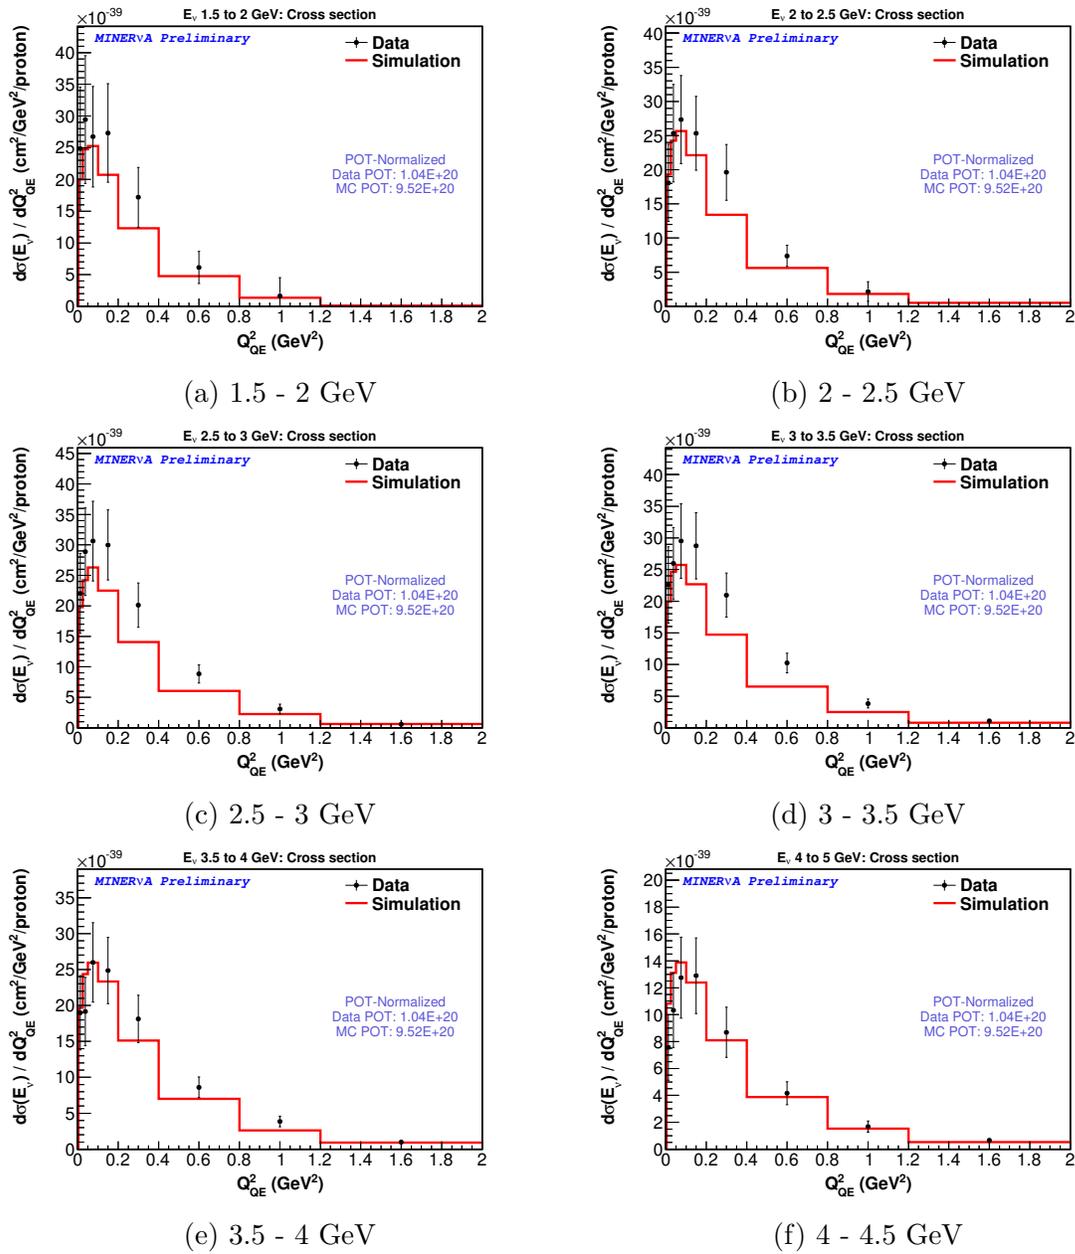
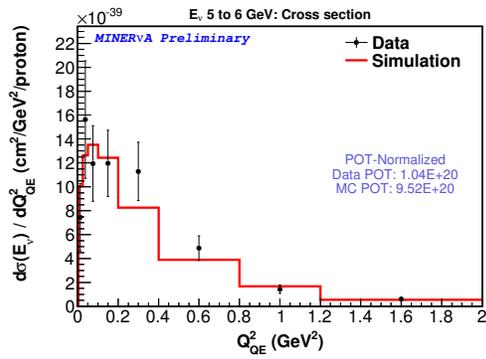
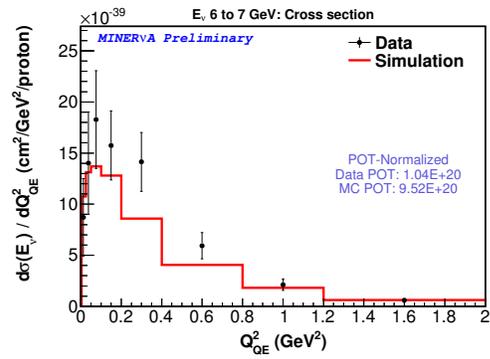


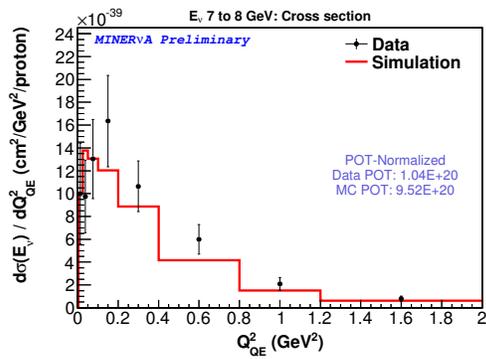
Figure 5.34. True CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued in next figure)



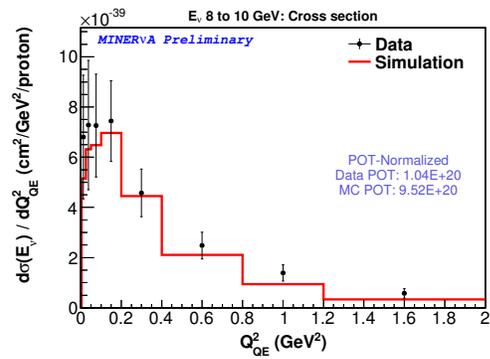
(g) 4.5 - 5 GeV



(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV

Figure 5.34. True CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$, in bins of E_ν (continued)

CHAPTER 6

Systematic uncertainties**6.1. Statistical and systematic uncertainties****6.1.1. Statistical uncertainty**

In a counting measurement such as MINERvA, our results are subject to some level of uncertainty. These uncertainties can be divided into two categories: statistical and systematic. A statistical uncertainty on a measurement is an inevitable consequence of a stochastic process; our event counts in our different energy or momentum bins are based on a random process, and this element of randomness leads to an uncertainty in the precision of our event count. This statistical uncertainty applies to our data histograms, but it is also applicable to our Monte Carlo simulation (referred to hereafter just as Monte Carlo), as this is generated by a procedure that uses random number generators. In order to obtain a total fractional statistical uncertainty on our measurements, we can add the data and Monte Carlo fractional statistical uncertainties in quadrature (as they are uncorrelated, being generated by completely separate random processes). For this study, we have roughly ten times more Monte Carlo events than data events; for this reason, data statistics are the dominant factor in the total statistical uncertainty.

6.1.2. Systematic uncertainty

While the statistical uncertainty corresponds to the level of precision to which we can measure a quantity by taking only a finite number of samples, systematic uncertainties represent the level of precision to which we can measure quantities, given the limitations of our detector, the theoretical models we use in our reconstruction, and our knowledge of the incident neutrino beam. In this chapter, we will focus on sources of systematic uncertainties, how they can be quantified, and how we propagate uncertainties into our cross-section measurements.

6.2. Evaluating systematic uncertainties on our measurements

Systematic uncertainties can be evaluated for many elements of the models and measured variables that we use to reconstruct our result. Each must be separately evaluated and quantified: for example, hypothetically, we may determine that:

- Due to the scintillator-strip structure of our detector, we can determine the position of an interaction vertex to the nearest centimeter
- The cross section model we have for a proton producing a pion in our beam target is accurate to $\pm 10\%$
- Pile-up in $2 \pm 1\%$ of data events leaves us unable to reconstruct the event, and this is not modeled in our Monte Carlo. While we can correct for this bias, we must evaluate the uncertainty on the correction

Once these individual uncertainties have been evaluated, we must determine the effect that the uncertainties on these components will have on our final cross section distribution. We evaluate this using a “many-universe” method. When analyzing the Monte

Carlo, as well as generating our simulated “central-value” event distribution (that is, the distribution generated with our best estimate of simulation parameters), we re-generate our distributions using an altered Monte Carlo (for example, a Monte Carlo where we have produced 10% more pions in our target, or one in which every muon energy measurement is 2% higher). We can imagine that this would be the result of our analysis in an alternate universe where, for example, the physics has slightly different parameters than used in our central-value measurement, or in which our detector has a slightly different response. For that universe, the systematic uncertainty on each bin’s measurement would correspond to the difference between the central-value event count in the bin, and the shifted event count. In the case where we want to quantify the result of varying the same quantity in more than one way (for example increasing the muon energy by 2%, and decreasing it by 2%), we generate a histogram for each, corresponding to multiple shifted universes. When we generate multiple universes for a systematic effect, the magnitude of the systematic uncertainty in each bin due to the effect corresponds to the average difference between each shifted histogram and the central value in that bin. In this analysis, we are able to evaluate some uncertainties (such as lack of precision in measured muon track angles) with only one shifted universe; for others (such as uncertainty in our final-state re-interaction rates) we use two, shifting both up and down within uncertainties in the re-interaction cross sections. For the complicated effects due to our lack of knowledge of the neutrino flux, we evaluate 100 universes, varying many parameters within their uncertainties to generate each.

As the Monte Carlo is used in all stages of the cross section calculation, alterations to the simulation introduce changes to, for example, the acceptance calculation, predicted

background levels, and unsmearing matrix. When shifted versions of these are used to produce a cross section, those differences will propagate to the cross section that is generated. The extent of these differences to the cross section will give us an indication of the effect that the uncertainty in question has on our final result.

In the MINERvA analysis framework, we facilitate the many-universe calculation by converting our central-value histogram from a ROOT TH2D (2-dimensional histogram) object to an MnvH2D object. This is part of MINERvA's custom PlotUtils package, which is built to hold the shifted histograms corresponding to the different systematic uncertainties' universes. The MnvH2D implements operations that must be performed on all universes, simplifying the cross section calculation procedure, in which each step (background subtraction, unfolding, efficiency correction etc) must be performed for every individual universe. In order to do this, not only our raw Monte Carlo distributions, but also all other objects generated from the Monte Carlo (efficiency histograms, migration matrix, background scales) must be generated in all of our shifted universes.

There are several ways in which we can generate alternate universes, appropriate for different situations:

- Smearing and scaling of observables
- Re-weighting techniques
- Alternative simulations

These will be discussed below.

6.2.1. Smearing and scaling of observables

Some quantities, if varied, will affect which events will pass our reconstruction cuts, entering the sample. For example, an inaccuracy in measuring the position of an event vertex in the detector will affect which events pass the requirement that the interaction must have taken place in the detector’s fiducial volume. In this case, when generating Monte Carlo, we add to our tuple one or more alternative versions of the quantity (or quantities: for example, the x, y and z positions). These could involve simply adjusting the quantity by a fixed scale for each event, corresponding to a potential bias (this scale may be a function of the variable, for example a 1% increase, or of other variables (for example, seeing the effect on reconstructed neutrino energy by making a 1% increase to final-state muon energy)). Alternatively, we may “smear” the quantity, by adding a random amount selected from a Gaussian distribution with a width corresponding to the 1σ uncertainty on the quantity, corresponding to a perceived lack of precision. In some situations, we may want to see the effect of more than one adjustment to the quantity - for example, we may generate one variable corresponding to an energy 2% higher than our central reconstructed value, and another corresponding to an energy 2% lower.

The set of histograms corresponding to each universe for a given systematic (e.g. the histograms for shifting muon energy up and shifting it down) is stored in an object called an `MnvLatErrorBand`, corresponding to a “lateral” error band, in which events may move from one bin to another, or may move in and out of the sample. (For example, if we are increasing muon energy measurements by 10%, an event with a 0.95 GeV muon that would have once fallen into the < 1 GeV bin will now fall into the > 1 GeV bin). The `MnvLatErrorBand` class not only enables storage of the individual histograms for each

universe, but also includes methods to calculate the systematic uncertainty due to the combined effect of those universes, to perform operations on all universes at once, such as multiplying every universe’s histogram by a constant. The MnvPlotter utility, also part of the PlotUtils package, can then be used to provide visualizations of these effects.

Examples of uncertainties that are treated in this way are our ability to measure muon energy and angle (which affect the values of p_{\parallel} and p_T that we record, moving events from bin to bin or in and out of the samples space, and our detector’s response to different particles, which would affect our measured recoil energy, changing which events pass the recoil cut.

6.2.2. Re-weighting

While some uncertainties affect which events will pass our reconstruction cuts, and which bins they fall into, others affect the probability that a given Monte Carlo event is likely to exist at all - for example, the probability of a final-state interaction producing a pion may be 10% higher or lower than we assumed when making our central value tuples. In this case, when generating Monte Carlo, an extra field is included to store a weight corresponding to the event’s relative likelihood when our quantity (say the pion production rate) is shifted. For a given uncertainty, we may store a single weight, or we may store multiple ones (for example, one corresponding to an increased pion production rate, and another corresponding to a decreased rate). Each of these adjusted weights can be thought of as corresponding to a “universe” in which the parameter (such as the pion production cross section) had a different value. For example, increasing the pion production cross

section would give an increased weight, or probability, to all events with a pion in the final state, while having no effect (a weight of 1) on events with no pion.

After plotting the central-value histogram, we use our same cuts and plot another histogram for each of our universes, weighting each event by its stored weight value corresponding to that universe. As with the cut-modification systematics, the value of the uncertainty in each bin is the difference between the central-value event count in the bin, and the shifted one. With two universes, it is the average difference between each shifted histogram and the central value.

Whereas, for the smearing and scaling systematics, we built a lateral error band, in this case we have a “vertical” error band, represented by the `MnvVertErrorBand`. This is because in this case, each event stays in its original bin, but with a different weight. A vertical error band can be populated by specifying, for each event, a vector corresponding to the weight of that event in each universe. Again, the `PlotUtils` functionality enables us to extract absolute and fractional systematic errors from this structure.

More universes are used for more complicated situations wherein many parameters combine to have a total effect on the produced cross-section uncertainty. The flux systematics are especially complex, as there are many components that combine to have an effect on the overall neutrino cross section. For these more complicated systematics, we use a “many universe” technique. In this case, we consider each parameter that has an effect on the final result, and evaluate the systematic uncertainty on that parameter. We then create a large number of “universes” by sampling the parameter space and then calculating each event’s weight in each one. A sampling universe corresponds to a situation in which each of the contributing parameters is separately smeared by a random number

from a gaussian distribution around its central value, with a width of its uncertainty. For example, if we have a central horn current of 200kA and a central POT of 9×10^{19} , one universe might correspond to the case where the horn current was 201kA and POT 8.9×10^{19} , and another to the case where current was 202kA and the POT 9.1×10^{19} . The set of shifted parameters is used to calculate a weight for each event in each universe. For each bin, the event counts for that bin in all the universes are histogrammed, and the width of the resulting distribution gives the systematic uncertainty in that bin.

We use the reweighting method to evaluate the effect of changing the central value probabilities used by the Monte Carlo when calculating the cross sections for initial interactions and for different final-state interaction processes. In this case we typically use 2 universes. For flux uncertainties, we use 100.

6.2.3. Alternate simulations

Some variations to the Monte Carlo cannot be simulated simply by applying weights. In this case, we generate a complete new set of Monte Carlo, changing all of these non-reweightable parameters at once. By performing a central-value analysis on an alternate sample, we get a single shifted universe corresponding to the combined effect of these changes. We add this shifted sample as a lateral error band with 1 universe. We use one alternate sample that includes changes to the models for the formation zone and effective nuclear radius, as well as incorporating an alternative hadronization model and a varied detector response to cross talk. The alternate Monte Carlo sample uses a total of 9.337×10^{20} POT.

Additionally, an additional sample of Monte Carlo is used to simulate the effect due to 2-particle-2-hole meson exchange current (MEC) events, as modeled by Nieves et al. [162]. While this effect is modeled in the newest version of GENIE, it is not available in version 2.8.4, which is used for this analysis. Thus a sample generated of MEC interactions generated with a more recent version is used to evaluate the effect on our distribution of adding the effect to the simulation. This sample uses 2.485×10^{21} POT (more POT are needed to generate sufficient statistics for this sample than for our standard samples, as the 2p2h processes are rarer than 1p1h CCQE).

6.2.4. How each uncertainty is evaluated

Table 6.1 details which uncertainties are analyzed by which of the methods listed above. For convenience, the uncertainties are grouped by the type of physics that they are measuring. In this, GENIE [93] refers to our Monte Carlo simulation package. FSI stands for final-state interactions. All of these uncertainties will be described in detail later in this chapter.

6.2.5. The covariance matrix

In all cases, the systematic uncertainty is are evaluated by forming a covariance matrix. This matrix contains information about the value of a given of systematic uncertainty in a given bin, as well as information about the extent of correlation between bins for that uncertainty. For example, consider how we create our many universes by changing some set of parameters, and recalculating our event distributions. If universes in which the event count in bin i increases also lead to an increased event count in bin j , this is

Uncertainty	Category
Meson exchange current model	Alternate simulation
Other GENIE signal uncertainties	Re-weighting
GENIE background uncertainties	Re-weighting
Flux uncertainties	Re-weighting
FSI uncertainties	Re-weighting
Detector particle responses	Smearing/scaling
Neutron path length weight	Re-weighting
Muon energy/angle uncertainties	Smearing/scaling
GENIE variation uncertainties	Alternate simulation
Mass model	Re-weighting
Tracking efficiency	Re-weighting

Table 6.1. Summary of techniques used to evaluate the individual systematic uncertainties

indication of positive correlation between bins i and j . If universes in which the event count in bin i increases give a decreased event count in bin j , this is indication of a negative correlation between bins i and j .

The covariance matrix is an $N \times N$ matrix, where N is the total number of bins in our two-dimensional distribution: for example, in the p_T vs p_{\parallel} case, N is the number of p_T bins (including over- and under-flow) multiplied by the number of p_{\parallel} bins (including over- and under-flow). For any two bins, i and j , (where each bin corresponds to a single combination of p_T vs p_{\parallel}), the value of the covariance matrix element is given by

$$(6.1) \quad M_{ij} = M_{ji} = \frac{\sum_{k=1}^n (x_{i_k} - \bar{x}_i)(x_{j_k} - \bar{x}_j)w_k}{\sum_{k=1}^n w_k}$$

Where:

- M_{ij} is the matrix element corresponding to the covariance between bins i and j .

Note that the covariance matrix is symmetric: $M_{ij} = M_{ji}$

- n is the total number of universes
- w_k is a weight to be applied to universe k (for unweighted universes, w_k is always 1)
- x_{i_k} refers to the event count in bin i , in universe k
- \bar{x}_i is the mean event count in bin i , averaged over all universes: $\bar{x}_i = \frac{\sum_{k=1}^n x_{i_k} w_k}{\sum_{k=1}^n w_k}$

The magnitude of the uncertainty in bin i is then given by the square root of the matrix element $[i, i]$:

$$(6.2) \quad \sigma_i = \sqrt{\frac{\sum_{k=1}^n (x_{i_k} - \bar{x}_i)^2 w_k}{\sum_{k=1}^n w_k}}$$

or, for unweighted universes, by the familiar standard deviation formula:

$$(6.3) \quad \sigma_i = \frac{\sqrt{\sum_{k=1}^n (x_{i_k} - \bar{x}_i)^2}}{\sqrt{n}}$$

While the diagonal elements of the covariance matrix are used to calculate the magnitude of the uncertainty in each bin, the off-diagonal elements are used in chi-squared calculations, to determine the goodness of fit between distributions (for example between data and Monte Carlo distributions, or between measured results and different theoretical models. The formula for the chi-squared is given by:

$$(6.4) \quad \chi^2 = \sum_{i=1}^N \sum_{j=1}^N (x_i - y_i) M_{ij}^{-1} (x_j - y_j)$$

where:

- N is the number of bins (for calculation of χ^2 , we would tend to ignore overflow and underflow bins)
- M_{ij}^{-1} is the $[i, j]$ matrix element of the inverse of the covariance matrix
- x_i is the value in the i th bin of the quantity for which we have calculated the covariance matrix
- y_i is the value in the i th bin of the model against which we are comparing

If a model describes the data perfectly, we would expect a chi-square per degree of freedom around 1. The number of degrees of freedom is typically equal to the number of bins used in the calculation. Note that taking a shape-only comparison removes one degree of freedom (once all but one of the bins' values are known, there is only one possible allowed value for the final bin that will meet the area-normalization condition).

6.2.5.1. Correlation matrix. The correlation matrix provides a measure of how correlated an uncertainty is between bins. An element X_{ij} of the correlation matrix is related to the corresponding element M_{ij} of the covariance matrix by:

$$(6.5) \quad X_{ij} = \frac{M_{ij}}{\sigma_i \sigma_j}$$

The possible values in the correlation matrix range from -1 (completely anti-correlated; if our systematic effect causes the value to increase in bin i , it will cause it to decrease in bin j , and vice versa) to +1 (completely correlated; if our systematic effect causes the value to increase in bin i , it will increase correspondingly in bin j).

An example of an effect with complete positive correlation would be something that changed the normalization of the event count, such as our tracking efficiency effect (see

section 6.4.2). Note that, by definition, the diagonals of the correlation matrix, which correspond to the systematic uncertainty in a single bin, will always have a value of +1.

Negative correlation will typically be observed in shape-only distributions, where we impose a requirement that the normalization of each universe must be the same. For this reason, an effect that increases the event count in one bin must, by necessity, decrease it in others, in order to keep the total event count constant.

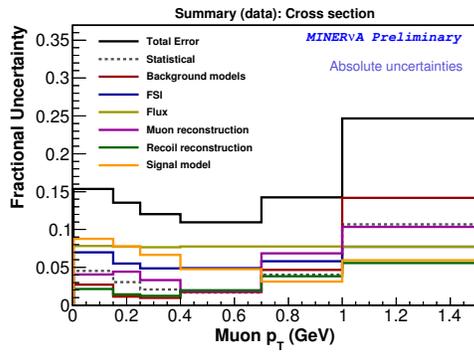
A value of zero in the correlation matrix indicates that the effects of the given systematic on the two bins in question are completely uncorrelated.

6.2.6. Summary of uncertainties on raw data distributions

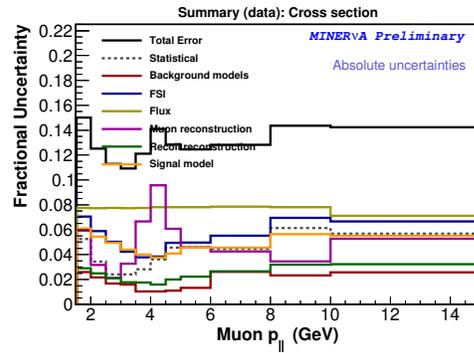
As there are many systematic uncertainties to consider, we categorize them in groups as follows:

- Flux uncertainties
- Muon reconstruction uncertainties
- Primary interaction uncertainties
- Hadron interaction uncertainties (FSI)
- Recoil uncertainties

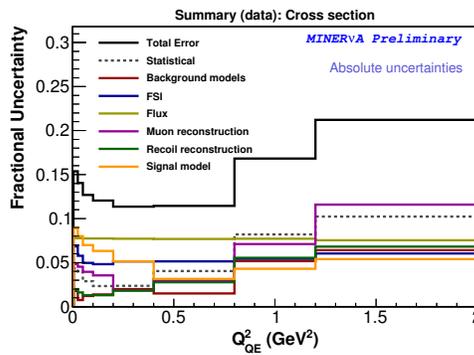
Figure 6.1 summarizes the contributions of these groups to the total fractional uncertainty on our raw distributions, while figure 6.2 shows their effect on the shape of the distribution (as opposed to the total normalization). These uncertainties are broken down into more detailed, two-dimensional bins, in section 6.9.1. More detail of the uncertainties that make up these groups will be given later in this document.



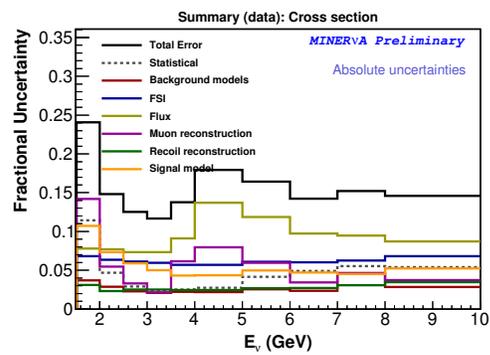
(a) Versus p_T



(b) Versus p_{\parallel}



(c) Versus Q_{QE}^2



(d) Versus E_{ν}

Figure 6.1. Summary of fractional uncertainties on the cross section

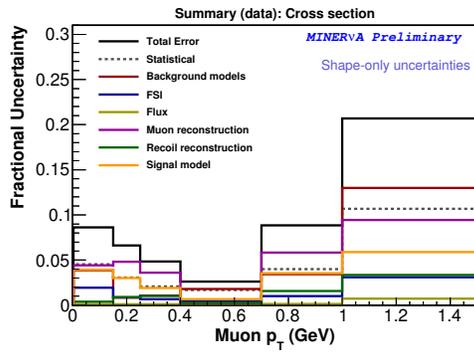
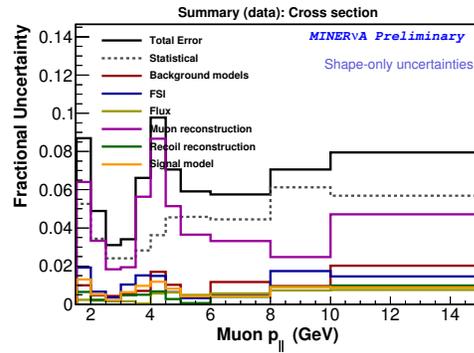
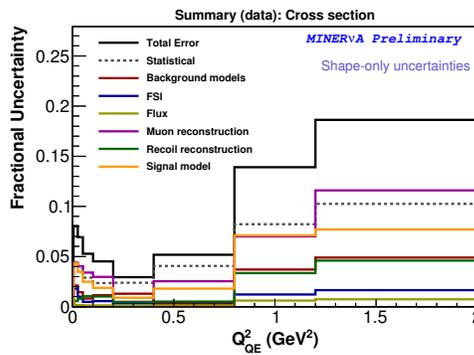
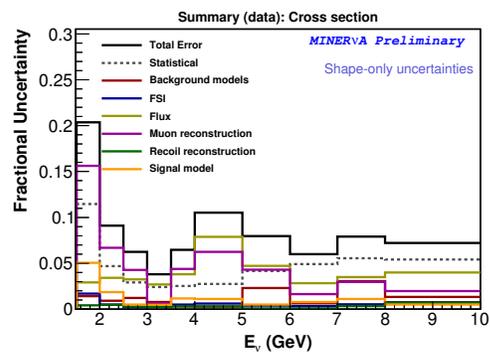
(a) Versus p_T (b) Versus p_{\parallel} (c) Versus Q_{QE}^2 (d) Versus E_{ν}

Figure 6.2. Summary of fractional shape uncertainties in the cross section

6.3. Flux uncertainties

The neutrino “flux” refers to the rate and energy distribution of incoming neutrinos. Our beam consists of antineutrinos with a broad spectrum of energies, and because the flux and energy distribution of these antineutrinos cannot be directly measured, it is necessary to calculate an expected flux distribution based upon what we know about how the beam is produced.

The neutrino beam is generated when protons from Fermilab’s NuMI beam-line hit a graphite target, producing various particles, including pions and kaons. These are then focused using a pair of magnetic focusing horns, whose current can be adjusted to favor either positive mesons (which will decay to produce neutrinos) or negative mesons (which will produce antineutrinos). Adjusting the horn current can also favor different energy profiles for the resulting neutrino beam. The focused mesons decay according to mechanisms such as

$$(6.6) \quad \pi^- \rightarrow \mu^- + \bar{\nu}$$

Muons are filtered out as the beam traverses 240 meters of rock, resulting in a neutrino beam. The NuMI beam is described in detail in [98].

The calculation of the NuMI flux is complex, and is described in more detail in [38].

The sources of uncertainty can be categorized as being due to:

- Hadron production rates for the NuMI beam, as measured by Fermilab’s PPFX (Package to Predict the Flux) method
- Focusing of the hadron beam by the focusing horns

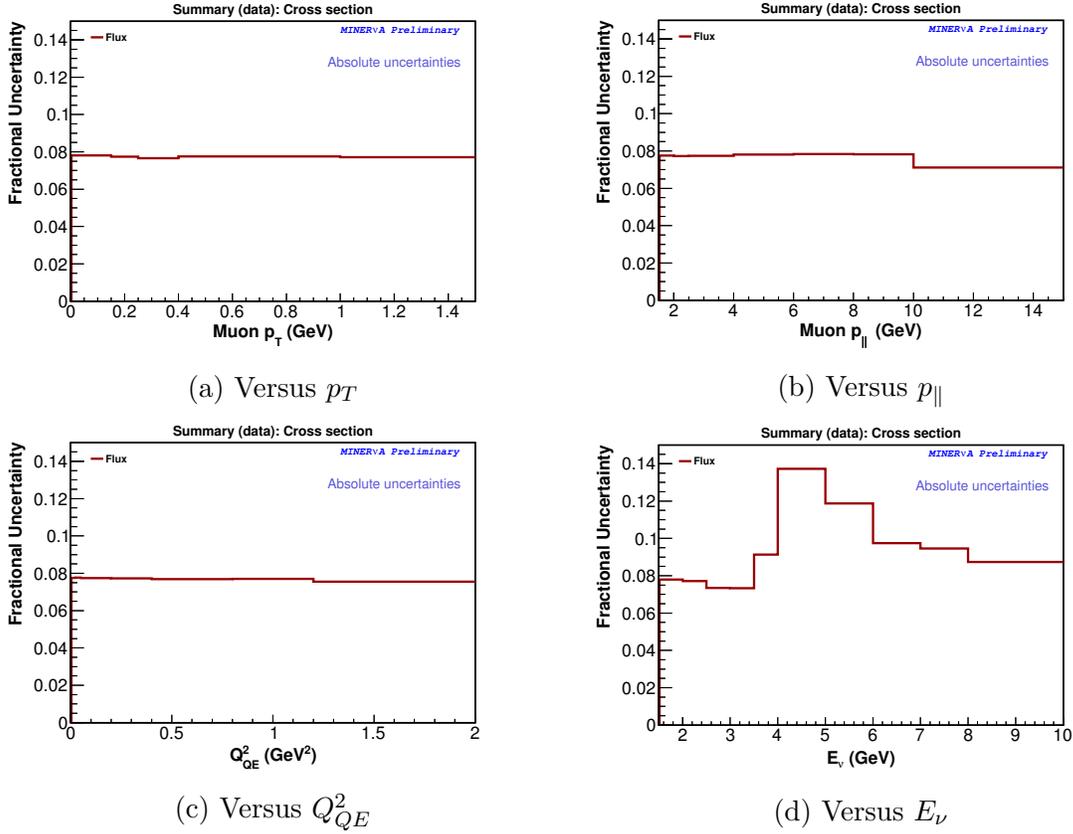


Figure 6.3. Summary of fractional uncertainties on the final cross section due to flux

The relative magnitude of the flux uncertainties is summarized in figure 6.3.

In each case, flux uncertainties are treated at the Monte Carlo generation level. Various parameters are smeared by throwing a random number from a Gaussian distribution with a width of 1 standard deviation of that parameter. Using the smeared parameters, a weight is calculated for each event, based on the probability of that event being generated given the new parameter set. This process is repeated 100 times, so that for each event, there is an array of 100 different weights, corresponding to the 100 different parameter sets that were randomly selected.

For each of these 100 sets of weights, we generated a histogram, using our standard cuts. These 100 histograms correspond to how our Monte Carlo event distribution would look, given those 100 different sets of smeared parameters. For each bin of our distribution, the 100 values corresponding to the event counts in the 100 smeared universes are plotted. Additionally, we apply the neutrino-electron scattering constraint, explained in section 3.7.1.3 and in [167]. To do this, for each simulated event, we look up the relevant flux value based on the $\nu - e$ constrained flux histogram for the universe, and divide this value by the unweighted central flux value used to generate the simulation, to get a corrected weight. A detailed description of the sources of flux uncertainty is given in section 3.7.1

6.4. Muon reconstruction uncertainties

There is some level of imprecision in how well we can measure the properties of a muon moving through our detector. As this analysis calculates cross sections based on the muon kinematics, we must quantify the extent to which we are able to measure these properties.

The muon energy scale, included in this section, is our most significant individual uncertainty over much of our phase space. This is unsurprising, as our analysis is so dependent on muon reconstruction and kinematics.

The relative magnitude of the muon reconstruction uncertainties is summarized in table 6.2 and figure 6.4.

Parameter	Variation	Effect on cross section
Muon energy scale	11 MeV(assay), 30 MeV (dE/dx), 2-3.2% (MINOS)	1 – 15%
Tracking efficiency	3-4%	1%
Muon angle smearing	1 mr	4% (high p_T / p_{\parallel} only)
Muon angle bias	1 mr	< 1%
Vertex smearing	x :0.91 mm, y : 1.25 mm, z : 1 cm	< 0.5%
Mass model	1.4%	1.4%

Table 6.2. Summary of muon reconstruction uncertainties

The muon reconstruction calculations use only 1 or 2 universes and quantify various aspects of our uncertainty in measuring the properties of a muon moving through the detector.

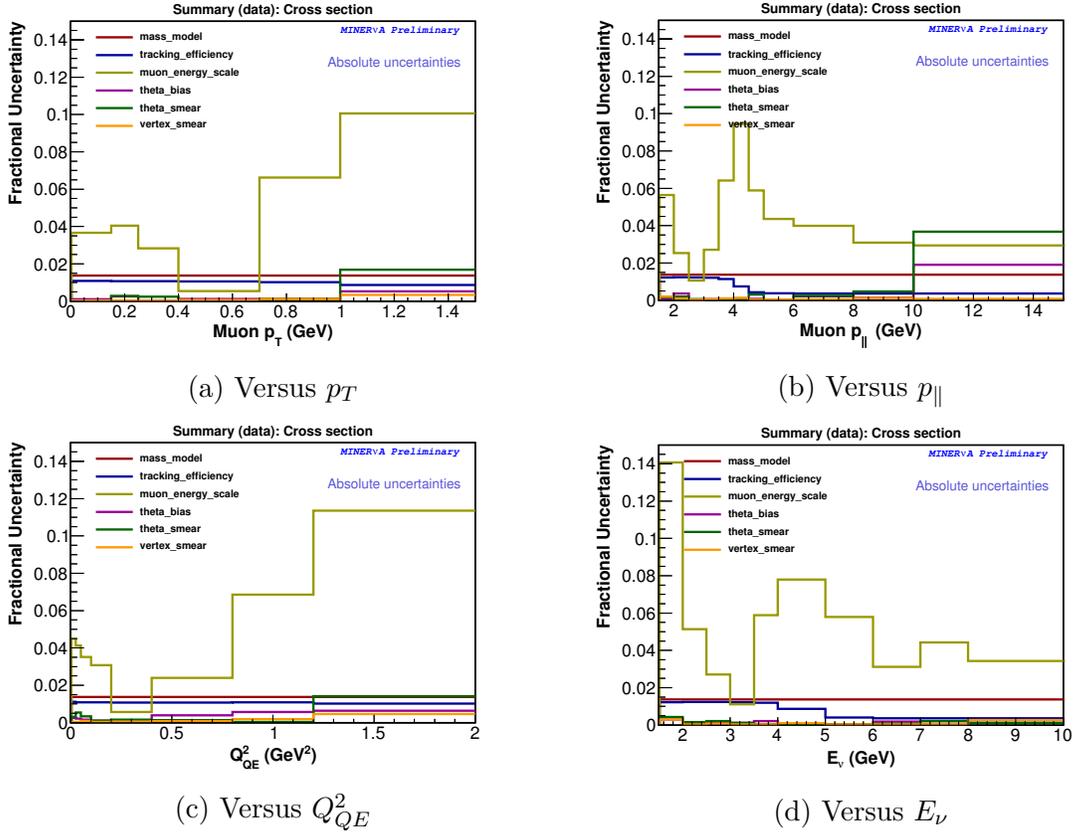


Figure 6.4. Summary of fractional uncertainties on the cross section due to muon reconstruction

6.4.1. Muon energy scale

For this analysis, we require a MINOS-matched muon. We reconstruct the muon energy using a combination of the momentum measured in MINOS (both by range and curvature) and by the energy deposition rate in the MINERvA detector.

In order to calculate this systematic, we generated histograms in two universes, corresponding to a reconstructed muon energy shifted up by 1σ and down by 1σ . These are used to create two shifted histograms in a lateral error band. The momentum shift corresponds to changes in three different properties:

- An 11MeV shift corresponding to our uncertainty in MINERvA’s material assay, and
- A 30 MeV shift corresponding to our uncertainty in the measurement of energy deposition rate in MINERvA, $\frac{dE}{dx}$, and
- A MINOS momentum-dependent shift [37] corresponding to
 - 2% for muons whose momentum is measured by range, added in quadrature to
 - 0.6% for muons with momentum is measured by curvature to be above 1GeV, or
 - 2.5% for muons with momentum is measured by curvature to be below 1GeV

These shifts are calculated using the `calculateMomentumCorrection` method of `MuonUtils`. The total muon energy scale uncertainty is calculated by adding the components in quadrature. Note that the MINOS range uncertainty is applied to all MINOS-matched muons; the additional curvature uncertainty is only non-zero if it was not possible to measure the muon’s momentum by range, and it instead had to be measured by curvature (that is, if the muon did not stop in the calorimeter region of MINOS).

Because our cuts (specifically the Q_{QE}^2 -dependent recoil energy cut, as we reconstruct Q_{QE}^2 from muon kinematics) depend on the muon momentum, this means that, when generating the shifted histograms, we must make our cuts using these shifted values, changing the total number of events that pass the cuts.

Our event distributions are plotted versus the muon longitudinal and transverse momentum ($p_{\parallel\mu}$ and $p_{T\mu}$) or versus the neutrino energy and 4-momentum transfer (E_ν and Q_{QE}^2) calculated in the quasi-elastic hypothesis from muon kinematics. Because all of

these variables depend on the muon momentum, the shifted universe plots are plotted with the axes corresponding to the shifted variables (that is, the variables $p_{T\mu}$, Q_{QE}^2 etc. are calculated as functions of the shifted muon momentum).

6.4.2. Tracking efficiency

This uncertainty corresponds to the combined uncertainty in three sources of event loss. These are for MINERvA tracking efficiency (such as that due to events lost in the dead time between triggers), for MINOS tracking efficiency, and for matching between MINERvA and MINOS tracks . With each of these factors, an uncertainty is associated.

Our Monte Carlo simulates events one at a time. However, in the actual MINERvA detector, when a channel has experienced activity, there is a reset period or “dead time” of around 100ns following the 150ns window in which charge is counted[37]. This, it is possible that a second neutrino event can take place in the dead time. We attempt to simulate this in the Monte Carlo using a data overlay method [108].

The tracking efficiency in the MINERvA detector is measured by taking muon tracks that have been matched in MINOS, extrapolating them back, and seeing what percentage are also matched as muon tracks in the MINERvA detector. The study is detailed in [183]. MINERvA tracking efficiency is 99.5%. In the same document, the muon tracking efficiency in MINOS is also discussed. Our analysis depends on our identification of a MINOS-matched muon. However, it has been observed that muon tracks in MINOS may be misidentified, particularly if two or more MINOS tracks overlap in time. While this phenomenon may occur in data, we do not simulate these overlapping events in Monte

Carlo. As it is a function of pile-up, it is both intensity- and momentum-dependent, the effect being more significant below 3 GeV (97.5% efficiency) than above 3 GeV (99.5%).

The total correction and the uncertainty on this correction are calculated using the `MnvNormalizer` class in the `MINERvA PlotUtils` package. This utility provides a total weight (by multiplying the individual weights) and a total uncertainty on the correction (by summing the uncertainties on each in quadrature). The only component of this uncertainty that is not constant over our whole phase space is that due to the MINOS reconstruction, which is more significant at lower energies; apart from this, the uncertainty is flat.

6.4.3. Muon angle smearing

The study in [200] looks at how well we can measure the angle of muon tracks. This is done by cutting a track in half, and attempting to reconstruct the angle of each half of the track, and seeing how well these agree (MINERvA has no magnet, so tracks should be straight). This method was used on both data and Monte Carlo tracks, with no difference seen above the milliradian scale.

To quantify the effect of possible differences in resolution below this scale, we “smear” the angle θ_μ between the muon and the beam in Monte Carlo. We do this by adding a small shift to the muon angle in the $x - z$ plane, based on a random number taken from a Gaussian distribution with a width of 1 milliradians. This shift (which may be positive or negative) is then added to the original angle in this plane, θ_x , to produce a shifted angle. An analogous procedure is performed for θ_y , and the shifted angles are combined.

As the shifts will move some angles up and some down, we need only 1 shifted histogram (1 universe). Some of our cuts (specifically the Q_{QE}^2 -dependent recoil energy cut) depend on the muon angle. When generating the shifted histogram, we must cut on these shifted values, changing the total number of events that pass the cuts.

Our plotting variables, the muon longitudinal and transverse momentum ($p_{\parallel\mu}$ and $p_{T\mu}$) and the neutrino energy and 4-momentum transfer (E_ν and Q_{QE}^2) calculated in the quasi-elastic hypothesis from muon kinematics, also depend on the muon angle. Thus the shifted universe plot is plotted with the axes corresponding to the shifted variables (that is, the variables $p_{T\mu}$, Q_{QE}^2 etc. are calculated as functions of the shifted muon angle).

The angular smearing is most significant for events where the muon's trajectory is at a large angle to the beam direction, and thus contributes at high p_T and low p_{\parallel} .

6.4.4. Muon angle bias

The NuMI beam points downwards by 0.05887 radians, in order to point at the MINOS far detector, in the Soudan mine in Minnesota. (The coordinate system of the detector is such that the z direction is along the axis of the detector, with the beam in the y - z plane). MINERvA studies [168] [165] using very forward-going electrons from neutrino-electron scattering in MINERvA detect a small bias in the beam angle. To correct for this, when correcting from the MINERvA coordinate system to the beam coordinate system, we include an additional correction in data only of -3mr in the $y - z$ plane and -1mr in the $x - z$ plane. As there is some uncertainty on this bias (quantified in [85]), we account for this by generating two universes, one in which bias is increased by 1mr, and one in

which it is decreased by 1mr. As the uncertainty is biggest in the $x - z$ plane, it is in this direction that we make these shifts.

As our cuts (specifically the Q_{QE}^2 -dependent recoil energy cut) depend on the muon angle, the number of events that pass the cuts will change when using these shifted values. We also plot against variables (E_ν , Q_{QE}^2 , $p_{\parallel\mu}$ and $p_{T\mu}$) calculated with this shifted angle. For this systematic, we use two universes, corresponding to an over-and under-estimate of the bias. Its effect is modest over most of the phase space, and contributes little when added in quadrature to the other muon reconstruction uncertainties.

6.4.5. Vertex smearing

The position of the interaction vertex (in this case, the starting vertex of the muon track) is used to determine whether an event takes place within the fiducial tracker region.

The study in [199] looks at differences in vertex resolution between data and Monte Carlo, calculating an uncertainty of 0.91mm for the x coordinate and of 1.25mm for the y coordinate (the y-z plane being that containing the beam and the central axis of the detector). In the x and y directions, we find the resolution by plotting the difference between true and reconstructed positions (the residual) in our Monte Carlo. These form a double-gaussian distribution, which can be fitted, and the width of the central peak taken as the resolution. As our detector consists of planes of scintillator transverse to the z direction, our resolution in the z direction is less easy to calculate. The residual distribution no longer forms a Gaussian, instead forming a much more complicated distribution that is heavily dependent on the vertex's z position within its strip. We thus consider a

conservative value of 1 cm as our possible resolution difference in the z direction. To create a shifted value for the vertex position, we smear our position simultaneously by 0.91 mm in x , 1.25 mm in y , and 1 cm in z . (To “smear” a quantity by some amount means that we add a value to the quantity, taken at random from a Gaussian distribution with a width corresponding to the amount). We create a single shifted universe by applying our cuts to the shifted events. As the fiducial volume cut depends on vertex position, a different set of events will be selected after the shift. However the effect is very minor, and barely contributes when the uncertainties are added in quadrature.

6.4.6. Mass model

The MINERvA tracker consists of scintillator strips coated with a reflective coating and attached with glue. Thus, while it mostly consists of the well-defined scintillator, there is also a small contribution to the total mass from the glue and coating, the exact mass of which is known to a precision of $\pm 1.4\%$. The mass and its uncertainty are documented in [177]. The mass model uncertainty is tracked in the shared MINERvA module MnvNormalizer. While this is not directly related to the muon measurement, we include it in this category as it affects the muon energy scale.

6.5. Interaction model uncertainties

The models used by our Monte Carlo generator, GENIE 2.8.4, [93] include various parameters which can only be measured by experiment. As such, these parameters carry uncertainties, which can affect the cross-section distributions produced by GENIE. Uncertainties are evaluated for the following parameters used in the primary interaction

model; that is, in the model for the initial neutrino-nucleon interaction probability distribution, prior to final-state effects or transit of the reaction products through the detector. The GENIE uncertainties are summarized in [113], and are explained in detail in the remainder of this section.

Most GENIE uncertainties are evaluated using the re-weighting method. In general, we use two universes, one corresponding to an increase in some parameter, and one to a decrease. The method for evaluating each of these uncertainties is the same: the function `fillGenieWeightBranches` from the MINERvA utility `AnaUtils/GenieWeightExtensions.cc` is used when generating tuples. For each event, this populates an array of 7 weight values for each of the listed uncertainties; these weights correspond to the probability of the event occurring if the parameter in question were shifted by -3σ , -2σ , -1σ , no shift, $+1\sigma$, $+2\sigma$ and $+3\sigma$ respectively. For each uncertainty, we use the weights corresponding to the -1σ and $+1\sigma$ shifts to create two universes. The uncertainty due to the parameter are then calculated using the reweighting technique described in section 6.2.2.

We can divide the re-weightable GENIE interaction model uncertainties into two groups - those on the signal model (charged-current quasi-elastic interactions) and those on background models (resonant and deep inelastic scattering interactions). Background event model uncertainties contribute to cross section uncertainty as they affect the fraction of events that we subtract in the background subtraction process. The signal model also makes a small contribution as we fit the signal and background shapes in order to perform our background fits, and to adjust the background rates in each universe. Additionally, we must note that for a QE-like signal definition, where our signal is defined by the final-state signature, some signal events will actually result from resonant or DIS events where

final-state interactions lead to a QE-like final state, while some background events will in fact be primary CCQE interactions in which a pion is later generated. This means that the background subtraction will introduce a small amount of uncertainty from the CCQE “signal” model due to the subtraction of these CCQE (but not QE-like) backgrounds.

GENIE uncertainties that change particle fates (such as changing the hadronization model) cannot be modeled using the re-weighting method. In this case, we generate an alternative simulated sample in which these models have been adjusted.

6.5.1. Background model uncertainties

Table 6.3 lists a summary of background model parameters, including their central values and uncertainties.

6.5.1.1. M_A (Elastic scattering) (MaNCEL). Unlike in quasi-elastic scattering, elastic scattering is a neutral-current effect, involving no charge exchange: a neutrino scatters off a nucleon, exchanging a Z boson. As this is a weak interaction, it also contains an axial component and, as with the quasi-elastic version, involves an axial form factor and the corresponding axial mass M_A . A $\pm 25\%$ shift in the elastic-scattering axial mass is simulated in the MaNCEL error band. As we require a MINOS-matched muon, which is easily identifiable (and which would not be produced by neutral-current event), we have very little elastic background, and this uncertainty is at less than 0.5% across all our bins.

6.5.1.2. η (Elastic scattering) (EtaNCEL). GENIE uses the Ahrens[34] model of elastic scattering. This includes a form factor

$$(6.7) \quad G_A(Q^2) = \frac{1}{2} \frac{G_A(0)}{\left(1 + \frac{Q^2}{M_A^2}\right)^2} (1 + \eta)$$

Parameter	Abbreviation	Central value	Variation	Effect on cross section
Axial mass for neutral-current elastic scattering	MaNCEL	$M_A = 0.99$ GeV	$\pm 25\%$	$< 0.5\%$
Strange axial form factor for neutral-current elastic scattering	EtaNCEL	$\eta = 0.12$	$\pm 30\%$	$< 0.5\%$
Normalization of neutral-current resonance production	NormNCRES	1	$\pm 20\%$	$< 0.5\%$
Axial mass for resonance production	MaRES	$M_A^{RES} = 1.12$ GeV	$\pm 20\%$	3-6%
Vector mass for resonance production	MvRES	$M_V^{RES} = 0.84$ GeV	$\pm 3\%$	$< 1\%$
Deep inelastic scattering normalization	NormDISCC	1		$< 0.5\%$
Non-resonant 1-pion production strength ($\nu : n$ or $\bar{\nu} : p$)	Rvn1pi	1	$\pm 5\%$	$< 0.5\%$
Non-resonant 2-pion production strength ($\nu : n$ or $\bar{\nu} : p$)	Rvn2pi	1	$\pm 50\%$	$< 0.5\%$
Non-resonant 1-pion production strength ($\nu : p$ or $\bar{\nu} : n$)	Rvp1pi	1	$\pm 50\%$	$< 0.5\%$
Non-resonant 2-pion production strength ($\nu : p$ or $\bar{\nu} : n$)	Rvp2pi	1	$\pm 50\%$	$< 0.5\%$
GENIE variations	genie_variations			2%

Table 6.3. Summary of background uncertainties

In this formula, η is a numeric parameter corresponding to the strange axial form factor. The GENIE default value is $\eta = 0.12$. The EtaNCEL uncertainty models shifts to this value. Again, this refers to elastic scattering, an extremely small background for this

analysis. An $\pm 30\%$ shift in the value of η still leads to an EtaNCEL error band of less than 0.5% across all our bins.

6.5.1.3. NC Resonance normalization (NormNCRES). Neutral-current resonant interactions are similar to charged-current resonant interactions, except for a Z boson is exchanged instead of a W^\pm . Again, we expect to see a nucleon and pion produced, with the pion possibly being affected by final-state interactions to produce an elastic-like signal. However, as no muon is produced in elastic interactions, it is easy for us to reject this background by requiring a MINOS-matched muon track, and a 1σ shift in the total event count due to NormNCRES gives less than a 0.5% effect to our distribution.

6.5.1.4. M_A and M_V (resonance production) (MaRES and MvRES). Charged-current resonance production refers to a charged-current interaction in which the struck nucleon is raised to an excited state (most commonly one of the Δ_{1232} resonances, spin $\frac{3}{2}$ states comprising u and d quarks, which decay to a nucleon and a pion). While the products of this interaction (muon, nucleon and pion) differ from the CCQE signature (muon and nucleon only), the pion produced may undergo final state interactions within the nucleus, meaning that it may be absorbed and not exit the nucleus, leading to a quasi-elastic-like final state of just a muon and nucleon. Alternatively, the pion may exit the nucleus but not travel far enough to produce an identifiable track. For these reasons, resonant events form a significant proportion of our background for a CCQE signal definition, and comprise part of our signal for a CCQE-like definition where our signal is based on final-state particles rather than on primary interaction.

GENIE models resonant interactions are using the Rein Sehgal model [178]. Like the Llewellyn-Smith CCQE model, this includes nucleon form factors, in the form of a

“transition form factor”:

$$(6.8) \quad G^{V,A}(q^2) = \left(1 - \frac{q^2}{4m_{\mathcal{N}}^2}\right)^{1/2-n} \left(\frac{1}{q^2 - M_{V,A}^2}\right)$$

Where $M_V = 0.84$ GeV and $M_A = 1.12$ GeV and the pre-factor depends on the number n of oscillator quanta in the resonance. GENIE accounts for uncertainty in the values of these form factors, with a $\pm 20\%$ shift in M_A^{RES} , and a $\pm 10\%$ shift in M_V^{RES} . However, fits to electroproduction data [141] constrain M_V to a smaller uncertainty, of the order of 3%, as explained in [157]. Thus the MvRES weights are scaled for each event so that the difference between the shifted and central values for each event is 0.3 times the GENIE shift.

These uncertainties mainly affect our background rate and shape; the relatively large uncertainty on M_A causes MaRES to be the largest background interaction uncertainty. It is particularly important at higher Q^2 , above 1GeV^2 (or correspondingly, at p_T above 0.5 GeV), where the resonant production cross section is higher.

6.5.1.5. CC DIS normalization (NormDISCC). At high interaction energies, it is possible that, instead of scattering off an intact nucleon, a neutrino can scatter off one of the quarks within the nucleon. This effect is known as deep inelastic scattering, and is typically characterized by the ejection of several hadrons from the nucleus. For a true CCQE signal, this represents a small part of the background, as DIS events tend to have higher recoil energies, and thus fail our reconstruction cuts. However, charged-current DIS events that produce many nucleons can be CCQE-like in that they produce only nucleons and a muon in the final state. Additionally, DIS events tend to occur at high Q^2

and are this not prevalent in this study's phase space. The $\pm 1\sigma$ shift represented by the NormDISCC error band produces an effect less than 0.5% over our entire phase space.

6.5.1.6. Pion production from non-resonant interactions (Rvn1pi, Rvn2pi, Rvp1pi, Rvp2pi). While the resonant interactions described earlier account for the majority of pion production, other inelastic processes, as described by Bodek-Yang [60], also take place. These four uncertainties describe non-resonant pion production. Rvn1pi and Rvn2pi correspond to 1- and 2-pion production respectively from $\nu : n$ or $\bar{\nu} : p$ interactions; Rvp1pi and Rvp2pi correspond to 1- and 2-pion production respectively from $\nu : p$ or $\bar{\nu} : n$. In each of these cases, a 1σ change corresponds to a $\pm 50\%$ shift in the strength of the interaction type. GENIE includes values for all of these shifts; the Rvn1pi is additionally constrained by a fit to pion-production data on deuterium from bubble-chamber experiments at Argonne and Brookhaven National Laboratories [179]. This showed that GENIE's estimate of the strength of this effect was overestimated; the actual effect was $43 \pm 5\%$ of that estimated by GENIE. Thus we reweight both our central value to reduce the probability of events like these, as well as reducing the uncertainty to just a 5% shift in the effect. All of the non-resonant pion production processes have a small effect on our sample as we attempt to reject events with pions.

6.5.1.7. Alternative GENIE Sample. Other effects within the GENIE simulation cannot be modeled by reweighting, as they change particle fates. In this case, a separate sample of Monte Carlo must be generated, and the analysis run on this sample. The central values from performing each step of the cross-section calculation on this sample generate a single universe of shifted distributions. The difference between this and the central value in each bin gives a measurement of the level of uncertainty due to the

combined effect of the processes involved. (Note that, in order for this to be valid, the distributions generated with the shifted Monte Carlo must be normalized to the same number of protons on target (POT) as the central-value Monte Carlo, to ensure a like-for-like comparison).

We use a single alternative GENIE sample, which includes the effects of several processes, listed below.

Effective nuclear radius correction. The effective nuclear radius also affects the distance within which a particle is susceptible to final-state interactions. If the effective nuclear radius is increased, final-state particles have a longer journey before they are free of the nucleus, increasing the probability that they will undergo FSI. The central value effective nuclear radius used by GENIE for carbon-12 is 3.2 fm.

Formation zone. When a deep inelastic scattering interaction occurs, the quarks produced can propagate some distance through the nucleus, with low interaction probability, before forming hadrons. GENIE[93] uses a formation time of 0.342 fm/c, corresponding to the SKAT model [49]. Increasing the radius of this formation zone thus reduces the amount of nuclear matter through which the eventual hadrons must propagate to escape the nucleus, and thus the amount of time for which they are susceptible to final-state interactions.

Alternative hadronization model. GENIE [93] uses the AGKY[197] hadronization model for deep inelastic scattering processes, with a phenomenological description of the hadron multiplicity based on KNO scaling at low invariant masses ($< 2.3\text{GeV}/c^2$), with a gradual change to the PYTHIA/JETSET hadronization model for high invariant masses ($< 3\text{GeV}/c^2$). The KNO-based system includes a p_T -based rejection scheme which is

designed to reproduce experimental data. Turning this off gives an alternative model , which produces a more isotropic momentum distribution for the produced hadrons. [188]

Cross talk. Cross talk [37] is an electronics effect, by which light incident on a given photomultiplier tube channel induces a small amount of signal in a neighboring channel. For example, light on the channel corresponding to the muon track, in which a lot of energy is deposited, could induce a signal in the channel next to it, either within the PMT, or at the optical fiber–PMT junction. Several safeguards are in place to minimize this effect. For example, the cables are “woven” such that neighboring channels in the detector are far apart in electronics space, and vice versa - thus cross talk hits will appear to have come from a part of the detector far from the track or energy blob that “caused” them, and will not be mistakenly included as part of that track. Although MINERvA attempts to filter cross talk, there is still some uncertainty as to what fraction of the remaining recoil energy comes from this source. Based on the recommendations in [195], where muon measurements of cross talk are compared with bench test measurements, this uncertainty is $\pm 20\%$. The alternative GENIE sample increases the rate of cross talk by 1σ .

Note that cross talk is not really a GENIE model effect - it affects the recoil energy reconstruction. However, as a non-reweightable effect that can be simulated, it is convenient to include it in this alternate sample, to save the processing load of generating an entire separate sample that just varies the cross talk model.

6.5.1.8. Summary of background model uncertainties. The relative magnitude of the background interaction uncertainties is summarized in figure 6.5. For ease of reading, uncertainties of less than 0.5% in all bins are not drawn.

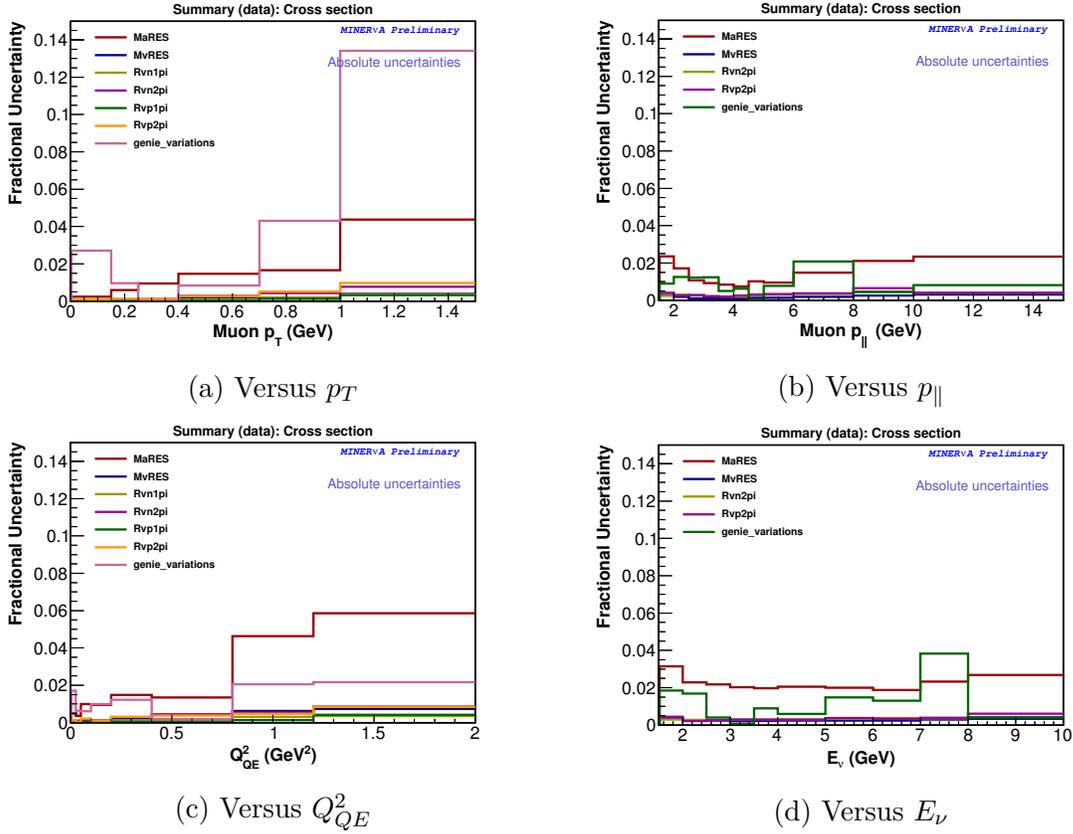


Figure 6.5. Summary of fractional uncertainties on the final cross section due to background interaction model uncertainties

6.5.2. Signal model uncertainties

Table 6.4 lists a summary of these parameters, including their central values and uncertainties.

6.5.2.1. M_A (CCQE scattering, shape only) (MACCQEShape). The Llewellyn-Smith model, used by GENIE to model the quasi-elastic scattering cross section, includes an axial form factor F_A , which must be measured by experiment. GENIE assumes a dipole model for this - that is, it assumes that the axial potential takes an exponential form proportional to, $\rho(r) = \rho_0 e^{-Mr}$, whose Fourier transform gives a form factor of the form

Parameter	Abbreviation	Central value	Variation	Effect on cross section
CCQE axial mass (fixed normalization)	MACCQEshape	$M_A = 0.99 \text{ GeV}$	$\pm 15\%$	$< 2\%$
CCQE normalization	NormCCQE		+20% - 15%	2-4%
Vector form factor model	VecFFCCQEshape	BBBA05	Dipole	$< 1\%$
Pauli suppression	CCQEPauliSupViaKF	$\kappa = 1.007$	30%	$< 2\%$
Meson exchange currents	MEC	None	26% of CCQE	5-10% (max at low p_{\parallel} / p_T)

Table 6.4. Summary of CCQE signal model uncertainties

$F_A(Q^2) = \frac{F_A(0)}{(1 + \frac{Q^2}{M_A^2})^2}$. The parameter $F_A(0)$ is well-known from neutron beta decay; the axial mass, M_A , however, must be measured in neutrino scattering experiments. Based on the results of bubble-chamber measurements on deuterium, GENIE uses a value of $M_A = 0.99 \text{ GeV}$. However, other experiments, on heavier nuclei such as carbon, have used shape fits to data to extract higher values; for example, MiniBooNE [29] found $M_A^{eff} = 1.35 \pm 0.17 \text{ GeV}$.

A change in the value of M_A changes both the shape and normalization of the predicted cross section distribution. A 1σ shift in MACCQEshape uncertainty parameter corresponds to making a $\pm 15\%$ shift in M_A , while holding normalization constant. As these variations in the CCQE shape are largely due to 2-particle-2-hole processes, whose effect we model separately in the meson exchange current model systematic uncertainty, this is a conservative estimate.

6.5.2.2. CCQE normalization (NormCCQE). Uncertainty in the overall normalization of the predicted quasi-elastic event rate is the most significant factor in the absolute primary interaction uncertainty. However, as this affects only the normalization (and not the shape) of the CCQE events, it provides only a small contribution to the shape-only distribution. A $+1\sigma$ shift in the NormCCQE parameter corresponds to a scaling of $+20\%$ while a -1σ shift is -15% . As these variations in the CCQE normalization are largely due to 2-particle-2-hole processes, whose effect we model separately in the meson exchange current model systematic uncertainty, this is a conservative estimate.

6.5.2.3. CCQE vector form factor model (VecFFCCQEshape). The Llewellyn-Smith CCQE cross section formula includes two vector form factors, which are related to the electric and magnetic form factors of the proton and the neutron. While initial models assumed a basic dipole form for the electromagnetic potential, this has been found to be a poor model, especially at high Q^2 where interactions are influenced by the quark substructure, and the nucleon cannot be considered as a homogenous particle. GENIE uses the BBBA05 model [63] for this vector form factor; however there are several different, and more recent models, based on parameterization of electron scattering results. To account for the effect on the quasi-elastic cross-section shape of varying the calculated form factor, the VecFFCCQEshape parameter changes the vector form factor model from BBBA05 to a dipole [91]. The total normalization is held constant, so that this change affects only the shape.

6.5.2.4. Pauli blocking in CCQE (CCQEPauliSupViaKF). Because nucleons are fermions, the Pauli exclusion principle states that no two neutrons or protons can occupy the same energy state. [66] This means that the final-state nucleon produced must have

an energy above the Fermi energy E_F , as states up to this energy are filled, according to the Fermi-Dirac distribution. This limits the cross section at low Q^2 , where nucleons with a low initial energy may not receive enough additional energy to break free of the nucleus. For a given energy transfer w and binding energy E_B , the initial state nucleon must have an energy $E_N > \kappa(E_F - w + E_B)$. The strength of the Pauli blocking is controlled by the parameter, κ . While this would be expected to be equal to 1, the MiniBooNE experiment found a better fit to their CCQE neutrino scattering data using a model with $\kappa = 1.007$ [29]. To take this into account, we introduce a 35% uncertainty on κ and use this to generate the CCQEPauliSupViaKF error band. Even this rather conservative uncertainty on κ leads to only a minimal influence on our overall event counts, with less than a 0.5% effect in all bins.

6.5.2.5. Meson exchange current model. The version of GENIE used in this analysis (2.8.4) does not model 2-particle-2-hole meson exchange current interactions, in which the neutrino can scatter from a correlated pair of nucleons within the nucleus, rather than from an individual nucleon. However, later versions of GENIE do include this model. A sample of 2.48×10^{21} POT of Monte Carlo was generated with MEC interactions, following the model of Nieves et al. [162][126]. (The large amount of POT was necessary to generate the statistics necessary for a good estimate of the cross section, as the cross section for 2p2h interactions is around 1/4 that of the free-nucleon CCQE cross section[100]). The magnitude of the MEC event count distribution, when POT-normalized to the Monte Carlo, serves as a single error band estimating the magnitude of this uncertainty on the signal model.

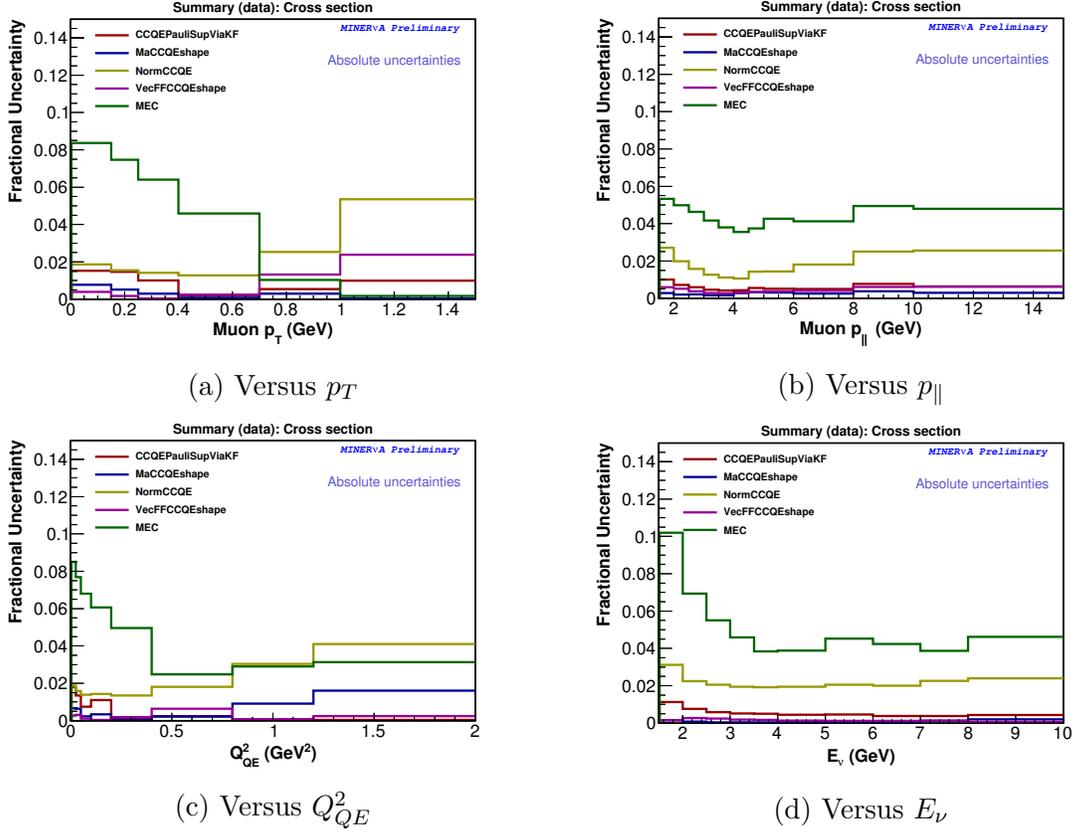


Figure 6.6. Summary of fractional uncertainties on the final cross section due to CCQE model uncertainties

6.5.2.6. Summary of CCQE signal model uncertainties. The relative magnitude of the CCQE signal model uncertainties is summarized in figure 6.6. For ease of reading, uncertainties of less than 0.5% in all bins are not drawn.

6.6. Final-state interaction uncertainties

In heavy nuclei, the primary interaction does not necessarily determine the final state. Instead, it is possible that the particles produced by this initial interaction can re-interact during their passage through the nuclear medium. This could for example, lead to a quasi-elastic primary event with a pion in the final state, as the neutron produced initially

interacts with another proton or neutron in the nucleus. Conversely, a resonant pion production event may have a quasi-elastic-like signature, if the pion is re-absorbed and does not exit the nucleus. We refer to these as “final-state interactions”. GENIE [93] models the probability of them occurring using the INTRANUKE-hA intranucleon hadron cascade package [89].

Uncertainty in the hadron interaction model affects our total event rate uncertainty much less significantly than does our uncertainty in the primary interaction, flux, or muon reconstruction. Its overall effect is around 2% across most of our phase space.

As with the primary interaction uncertainties, these are calculated with two universes reweighted to $-\pm 1\sigma$ as generated by the `GenieWeightExtensions` function `fillGenieWeightBranches`. For details of how this works, see section 6.5. The sources of the individual uncertainties are explained below, while table 6.5 summarizes them.

6.6.1. Mean free paths (MFP_N and MFP_pi)

A 1σ variation of the mean free path for nucleons (MFP_N) and for pions (MFP_pi) corresponds to a change of $\pm 20\%$. Note that this parameter corresponds to the mean free path of produced hadrons *within* the nucleus; the shorter the mean free path, the greater the probability that FSI will occur before the particle escapes the nucleus. According to [93], the mean free path for a nucleon produced at MINERvA energies is a few femtometers within the nucleus.

The mean free path of the pion has a fairly significant effect on our total hadronic interaction uncertainty (about a 0.5% uncertainty overall), as it affects whether a pion produced in a non-CCQE interaction will exit the nucleus (a background event that is

Parameter	Abbreviation	Variation	Effect on cross section
Neutron mean free path	MFP_N	$\pm 20\%$	$< 1\%$
Pion mean free path	MFP_Pi	$\pm 20\%$	$< 1\%$
Nucleon elastic scattering cross section	FrElas_N	$\pm 30\%$	$< 1\%$
Pion elastic scattering cross section	FrElas_Pi	$\pm 10\%$	$< 1\%$
Nucleon inelastic scattering cross section	FrInel_N	$\pm 40\%$	$< 1\%$
Pion inelastic scattering cross section	FrInel_Pi	$\pm 40\%$	3-5%
Nucleon charge exchange cross section	FrCEx_N	$\pm 50\%$	$< 1\%$
Pion charge exchange cross section	FrCEx_Pi	$\pm 50\%$	$< 1\%$
Nucleon absorption cross section	FrAbs_N	$\pm 20\%$	$< 2\%$
Pion absorption cross section	FrAbs_Pi	$\pm 20\%$	3-5%
Nucleon pion production cross section	FrPiProd_N	$\pm 20\%$	$< 1\%$
Pion pion production cross section	FrPiProd_Pi	$\pm 20\%$	$< 1\%$
DIS hadronization model adjustment	AGKYxF1pi	$\pm 20\%$	$< 1\%$
Pion angle distribution (resonant events)	Theta_Delta2Npi	Isotropic Rein-Sehgal	$\rightarrow < 0.5\%$
Resonant decay photon branching ratio	RDecBR1gamma	$\pm 50\%$	$< 0.5\%$

Table 6.5. Summary of hadron interaction uncertainties

likely to fail our reconstruction cuts), or whether it will re-interact, potentially producing a CCQE-like signature (and probably passing our reconstruction cuts). The mean free path for nucleons has a smaller (less than 0.5%) effect. While CCQE events always produce a neutron (which would be affected by this), our reconstruction cuts do not require that it

is detected, so re-interaction would have a small effect, unless it produces a proton that is subsequently tracked. There may also be some effect due to background events (e.g. DIS) that produce protons, as this may affect whether that proton is tracked (the event fails our reconstruction cuts) or not (it may pass).

6.6.2. Nucleon and pion fates

GENIE includes the capacity to determine uncertainties due to the following particle “fates”:

- Elastic (FrElas_N, FrElas_pi)
- Inelastic (FrInel_N, FrInel_pi)
- Charge exchange (FrCEX_N, FrCEX_pi)
- Absorption (FrAbs_N, FrAbs_pi)
- Pion production (FrPiProd_N, FrPiProd_pi)

These refer to the probability of a hadron (produced by our primary interaction) undergoing different forms of final state interaction. In *elastic* scattering, the nucleus remains in the ground state, and the hadron retains its same charge. Our 1σ uncertainty on FrElas_N (nucleon elastic scattering) is $\pm 30\%$, while for pions (FrElas_pi), it is $\pm 10\%$. It should be noted that FrElas_N correlates completely with changes to the mean free path (MFP_N) for the nucleon. While events containing only neutrons, our CCQE signal particles, are unlikely to be rejected by our reconstruction cuts regardless of whether they have re-scattered in the nucleus, this could affect whether we mistakenly accept events with protons (if their subsequent scattering leads to them having such low energy as to not

be identified and rejected by our track and recoil energy cuts). Because of this, FrElas_N is one of our most significant hadron interaction uncertainties, at around 1%.

An *inelastic* scatter can break up the nucleus. To produce our uncertainty due to this, we vary the cross sections for inelastic rescattering (FrInel_N) and (FrInel_pi) by $\pm 40\%$.

Charge exchange occurs where the final state hadron has a different charge from that originally produced (for example, a CCQE-produced neutron undergoes FSI, causing a proton to be ejected). For both nucleons (FrCEX_N) and pions (FrCEX_pi), we vary the cross section for this process by $\pm 50\%$. The nucleon situation could be important for the true CCQE sample as a neutron produced in a CCQE antineutrino interaction could exit the nucleus as a proton, making it more likely to be rejected by our recoil and track cuts.

When a hadron is *absorbed* in the nucleus, it does not appear in the final state. This can be a cause of non-CCQE primary interactions with a CCQE-like signature, as pions produced by mechanisms such as resonance production are subsequently absorbed. Thus a background event for the true-CCQE analysis could be misidentified as signal, while these events will be legitimate signal in a CCQE-like analysis. As the final state matches the signal we are selecting for with our reconstructed cuts, we select a large number of these events. Because of this, the pion absorption cross-section is the largest component of the hadron interaction uncertainties for the CCQE analysis, with around a 1% fractional uncertainty. The 1σ variation that produces this uncertainty for (FrAbs_pi) is a $\pm 30\%$ variation in the pion absorption cross section, while that for nucleon absorption (FrAbs_N) is varied by $\pm 20\%$. While the nucleon absorption cross section has a smaller effect than that of pion absorption, it also contributes significantly to the total hadron interaction uncertainty, particularly at higher Q^2 . Events affected by this are those in which protons

are produced (e.g. by DIS and hadronization) but the protons are absorbed in the nucleus, leading to a neutron-only final hadronic state. While these events are not true CCQE, they would have a CCQE-like signal and would tend to be selected by our reconstruction cuts.

Pion production in final-state interactions (where the nucleon has enough energy to produce an additional particle) can cause a primary CCQE interaction to mimic a resonant interaction. An event like this (FrPiProd_N) would be signal for a CCQE signal definition, but background for a CCQE-like, due to the presence of a pion in the final state. Additionally, the pion would make the event likely to fail recoil energy and track cuts. Additionally, a pion from an initial interaction could produce another if it re-interacts (FrPiProd_pi). The 1σ variation for each of these cases is $\pm 20\%$.

6.6.3. Hadronization model (AGKYxF1pi)

In deep inelastic scattering, the incoming neutrino scatters off a parton (quark or gluon) within the nucleon, breaking up the nucleon. As the strong force prevents quarks and gluons from existing outside of hadrons, interactions of this type form hadronic showers, in a process known as “fragmentation” or “hadronization”. GENIE simulates this using the AGKY hadronization model [197], which is described in [93]. The AGKYxF1pi parameter varies the distribution of the pion Feynman parameter, x_F , in 1-pion states of this model.

DIS events tend to have a high recoil energy due to the hadron shower, and are thus only a small background to our sample. For that reason, this has only a small effect on our event count (less than 0.5% in almost the entire sample), this being most noticeable

at high energies and Q^2 (high muon momentum, both transverse and longitudinal), where DIS events are more prevalent.

6.6.4. Delta decay angle (Theta_Delta2Npi)

The most common pion production mechanism in neutrino scattering is resonant scattering, in which the struck nucleon is excited to a Δ resonance state, which then decays to a nucleon and a pion. In GENIE, the resulting pions are modeled with an isotropic distribution. The Theta_Delta2Npi uncertainty models the effect of introducing a more realistic angular distribution for these pions, using the Rein-Sehgal model [178]. For the CCQE study, which aims to reject events with pions, this has a minimal ($< 0.1\%$) effect on our event counts.

6.6.5. Resonant decay photon branching ratio (RDecBR1gamma)

The RDecBR1gamma uncertainty varies the branching ratio to photons in the decay of resonant states. A 1σ variation to this branching ratio corresponds to $\pm 50\%$. For the CCQE study, which aims to reject events resonant events and those with photons, this has a negligible ($< 0.1\%$) effect on our event counts.

6.6.6. Summary of final-state interaction uncertainties

The relative magnitude of the final-state interaction uncertainties is summarized in figure 6.7. For ease of reading, uncertainties of less than 0.5% in all bins are not drawn.

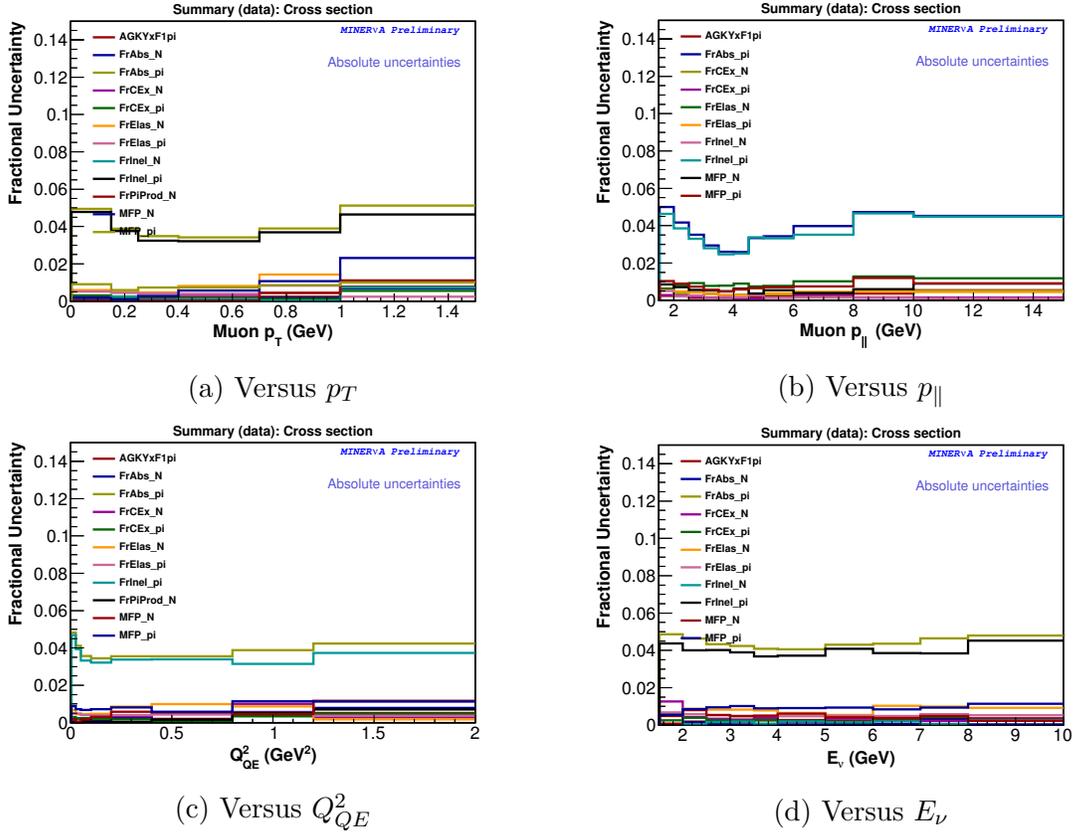


Figure 6.7. Summary of fractional uncertainties on the final cross section due to final-state interaction uncertainties

6.7. Recoil reconstruction uncertainties

As explained in 4.3 of our selection criteria involves a Q_{QE}^2 -dependent cut on the non-vertex recoil energy — that is, the energy deposited in the detector that forms isolated energy “blobs”, plus the dispersed energy deposits that are not included in a blob or track.

The measured recoil energy can be caused by charged particles in the detector, or by a neutron interacting with a detector nucleus. We model the detector’s calorimetric energy response to the various particles using GEANT4 [27]. To model our uncertainty in

the detector simulation, we use the method described in section 6.2.1, where we generate shifted universe histograms by adjusting the isolated and dispersed recoil energy by some amount, and regenerating our event distribution. Typically, we generate two shifted universes, corresponding to $\pm 1\sigma$ shifts in the parameter in question.

The individual recoil reconstruction uncertainties are listed in 6.6 and discussed below.

Parameter	Variation	Effect on cross section
Neutron path length	Energy-dependent	2 – 6%
Proton response	3-4%	< 0.5%
Pion response	4%(400-1900 MeV) or 5%	< 1%
Other (E-M) response	$\pm 3\%$	< 0.5%

Table 6.6. Summary of recoil reconstruction uncertainties

6.7.1. Neutron path length

CCQE antineutrino events have a hadronic final state consisting of a neutron. In order to reject events that include other final-state particles, while retaining those with a neutron, we need to know how neutrons interact in the MINERvA tracker. GEANT4 is used to simulate this. In order to evaluate a systematic on the neutron interaction model, we use a reweighting technique, corresponding to a variation in the mean free path of neutrons in the detector, based on comparisons between GEANT and various path length measurements [190].

The technique is described in [154] and can be summarized as follows: Assuming a neutron scatters only once, the probability that it will have scattered after a distance x is given by

$$(6.9) \quad p(x, \lambda) = \frac{e^{-x/\lambda}}{\lambda}$$

Where λ is the interaction length in the detector material. To get the cumulative probability that an interaction would have occurred within a radius r (proportional to cross section), one would need to integrate:

$$(6.10) \quad P(r) = \int_0^r p(x, \lambda) dx.$$

To generate weights, we need to take a ratio of a shifted to a default cross section, when a path length is shifted.

$$(6.11) \quad \text{weight} = \frac{P(r, \lambda(1 + \delta))}{P(r, \lambda)}$$

For small shifts $\lambda \rightarrow \lambda(1 + \delta)$, this can be expanded in a power series, giving a weight of

$$(6.12) \quad \text{weight} = 1 - \frac{r\delta}{\lambda(e^{r/\lambda} - 1)} + \mathcal{O}(\delta^2)$$

for neutrons that interact in the detector and a weight of

$$(6.13) \quad \text{weight} = 1 + \frac{r\delta}{\lambda} + \mathcal{O}(\delta^2)$$

for neutrons that exit the detector without interacting.

The default mean free path λ is taken from the GEANT simulation and, for the scintillator tracker, is related to the neutron's momentum by

$$(6.14) \quad \lambda = \frac{1}{0.0010528 - \frac{0.3269}{\min[p_n, 1000]} + \frac{159.55}{\min[p_n, 1000]^2}}$$

To make our shifts, we substitute this formula for λ , as well as r , the distance that the neutron travels through the detector, and insert a shift fraction δ corresponding to the fraction by which GEANT differs from measured path lengths in other studies. This value is dependent on the neutron kinetic energy (here in MeV), and corresponds to 25% for $T_n < 25$ MeV; $10\% + 0.6\%(50 - T_n)$ for $25 < T_n < 50$ MeV; 10% for $50 < T_n < 150$ MeV, $10\% + 0.6\%(50 - T_n) + \frac{T_n - 150}{1500}$ for $150 < T_n < 300$ MeV and 20% for $T_n > 300$ MeV. This corresponds to the measured data-Monte Carlo discrepancies in the < 25 , $50 - 150$ and > 300 MeV ranges, with linear interpolation between these regions to give a continuous distribution.

For this uncertainty, we use two universes — one in which the path length is shifted up by δ and one in which it is shifted down. As CCQE events always have a neutron in the final state, this is an important recoil energy systematic, with the highest significance of all our recoil uncertainties.

6.7.2. Charged particle response (proton, pion and others)

The recoil energy in each generated event can be divided into constituents that are generated by neutrons, protons, pions and other particles (mostly electrons/positrons and photons from electromagnetic showers, although there is also a small kaon contribution). To generate uncertainties for the detector's response to the different particles, for each event, we shift the fraction of both the isolated and dispersed energy contributions to the recoil energy for a given particle type by a percentage. To make two universes, we make a shift up and a shift down.

Protons are divided into three energy groupings. For low kinetic energy protons below 50 MeV, we shift energies by 4% - that is, the dispersed energy is shifted (dispersed energy $\times(1 \pm 0.04 \times$ proton fraction), and correspondingly for the isolated energy. For protons from 50-100 MeV, the shift is 3.5%; for above 100 MeV, we use 3%. These were determined by the MINERvA test beam, as documented in [84].

The proton response affects our event rate measurement $< 1\%$ across our whole phase space, as the track cut removes many protons, and only a small amount of those that remain pass into our selected sample by making this shift.

The pion calorimetric response has been constrained by recent test beam studies [125] to an accuracy of 4%, for pions with a kinetic energy between 400 and 1900 MeV. We thus separate our pions into two categories - “constrained” within this energy range, and “unconstrained” outside of it. Pions within the constrained range have their energy fraction varied by $\pm 4\%$, while others have it varied by $\pm 5\%$, as detailed in [37]. The pion response has only a minor effect ($< 1\%$ across our whole phase space) on our cross section, as the shift makes only a modest difference to the number of pions that will mistakenly be reconstructed.

For the other (basically electromagnetic, though it may include small amounts of kaons) response, we vary the fraction of the recoil due to other particles by $\pm 3\%$, as recommended in [37]. This uncertainty was derived by observing the energy response for Michel electrons (electrons from muon decay), which have a well-known energy spectrum. This change mainly affects the Q^2 ($p_{T\mu}$) shape, and contributes its maximum of around 1% uncertainty at low Q^2 .

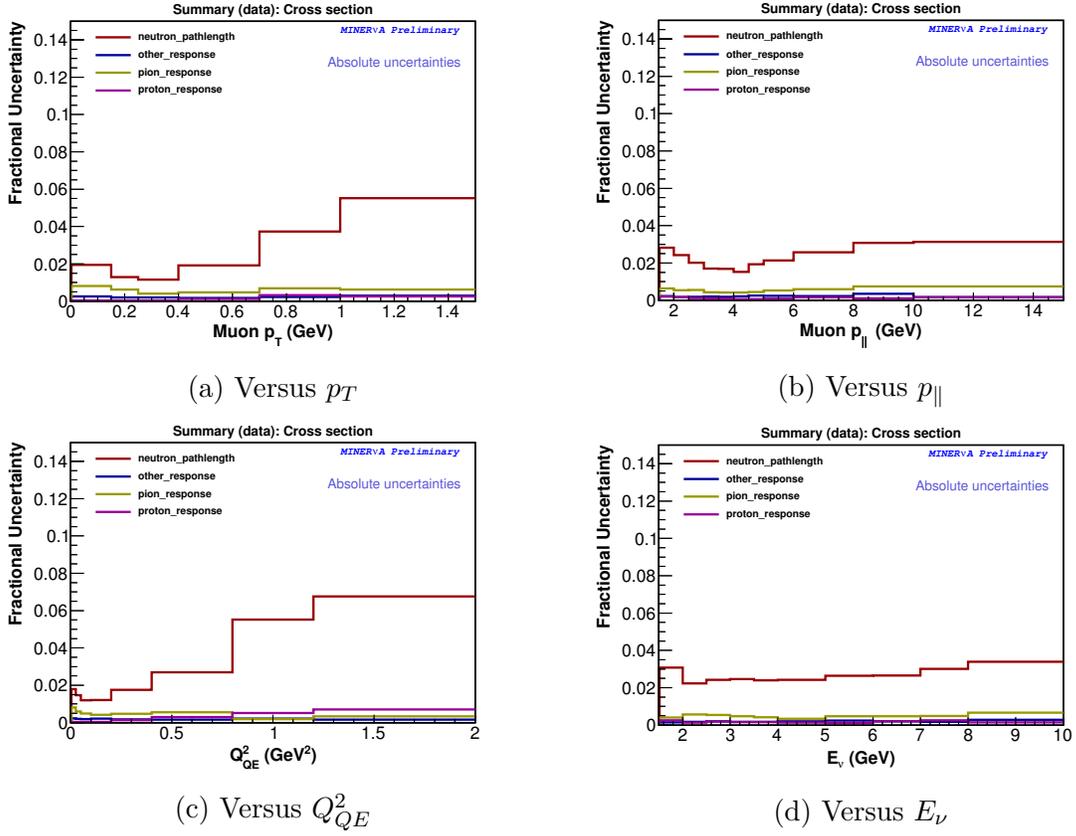


Figure 6.8. Summary of fractional uncertainties on the final cross section due to recoil energy reconstruction

6.7.3. Summary of recoil reconstruction uncertainties

The relative magnitude of the recoil reconstruction uncertainties is summarized in figure 6.8. For ease of reading, uncertainties of less than 0.5% in all bins are not drawn.

6.8. Propagation of uncertainty through the analysis stages

The cross section calculation consists of several stages. Systematic effects must be considered at each stage. As the Monte Carlo is used to evaluate the magnitude of the effects, the uncertainty universes are initially added to Monte Carlo-based distributions. As these distributions are used with the raw data distribution in various stages of the cross section calculation, the uncertainties transfer from the Monte Carlo to the data histograms, as will be demonstrated in this section.

The objects generated from Monte Carlo include:

- Raw Monte Carlo distributions
- Background scales (used in background subtraction)
- Migration matrix (used for unsmearing)
- Truth histogram (used for acceptance-correction)
- Flux histograms (integrated to apply flux correction)

Table 6.7 summarizes at which stages the individual categories of systematics come into play. For example, separate migration matrices must be made for the systematic universes that move muon transverse and longitudinal energies between bins. Conversely, for systematic universes where a quantity is weighted, that weight must be applied in the true histogram used for the acceptance-correction phase. In situations where a specific quantity is not evaluated for a given uncertainty's universes, the central value is used. In this table, the muon reconstruction uncertainties refer to muon energy scale, angle bias and smearing, and vertex smearing. Particle responses refer to the proton, pion and other particle responses, in the recoil systematic category.

Uncertainty	Raw	Bkgd scales	Migration matrix	Truth	Flux
GENIE signal uncertainties	x	x		x	
GENIE background uncertainties	x	x		x	
Flux uncertainties	x	x		x	x
FSI uncertainties	x	x		x	
Particle responses	x	x	x		
Neutron path length weight	x	x		x	
Muon reconstruction	x	x	x		
Genie variation uncertainties	x	x		x	
Mass model	x	x		x	
Tracking efficiency	x	x		x	

Table 6.7. Summary of which uncertainties are separately evaluated when generating different analysis histograms from Monte Carlo data

6.8.1. Raw Monte Carlo distributions

Figure 6.9 shows the distribution of uncertainties on the raw Monte Carlo distributions. For ease of viewing, the individual uncertainties are combined into summary groups, and the distributions are projected onto p_T and $p_{||}$. Note that, following the standard MINERvA systematic uncertainty evaluation procedure, the custom is to initially evaluate the effect of all uncertainties on the reconstructed simulation. A large portion of this is model uncertainty which will have little to no effect on our final cross section. Some sources of uncertainty, however, will affect the data when the uncertain simulation is used to estimate the amount of background to subtract, and to unfold and efficiency-correct our data.

As expected, the dominant uncertainty on our Monte Carlo is that of the CCQE signal model - precisely the thing that we are hoping to improve with this measurement. The effect of this uncertainty on our final cross sections is much smaller than the uncertainty

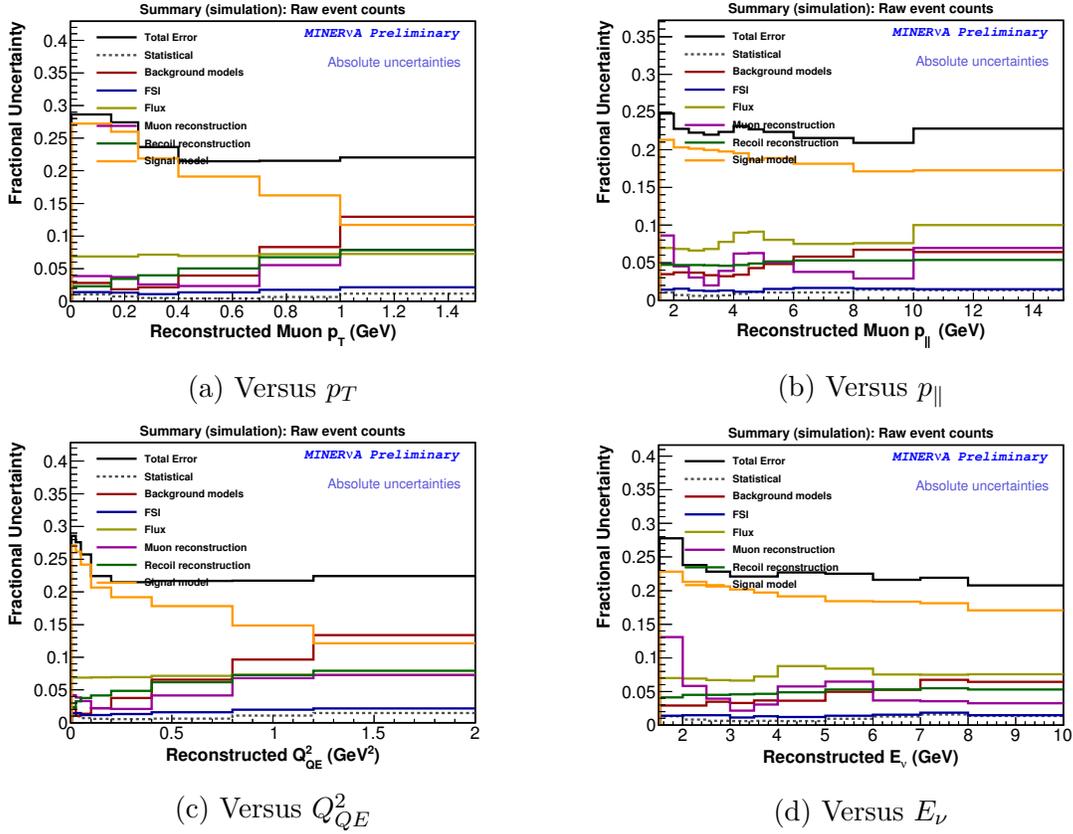


Figure 6.9. Summary of fractional uncertainties on raw Monte Carlo distributions. Many of these will be irrelevant to the final cross section result; some will affect background subtraction, unfolding, and efficiency correction.

it indicates on the simulation. With results such as ours, future versions of generators will be able to decrease the CCQE model uncertainties.

Figure 6.10 shows the uncertainties on the subsample of the Monte Carlo that corresponds to background events. In this case there is little contribution from the CCQE model, with a greater fractional contribution from uncertainties on the resonant and DIS process that constitute the background. The small residual contribution from the CCQE model is due to the small proportion of CCQE events that subsequently produce a pion via FSI. As these do not have a CCQE-like final state, they become part of our background.

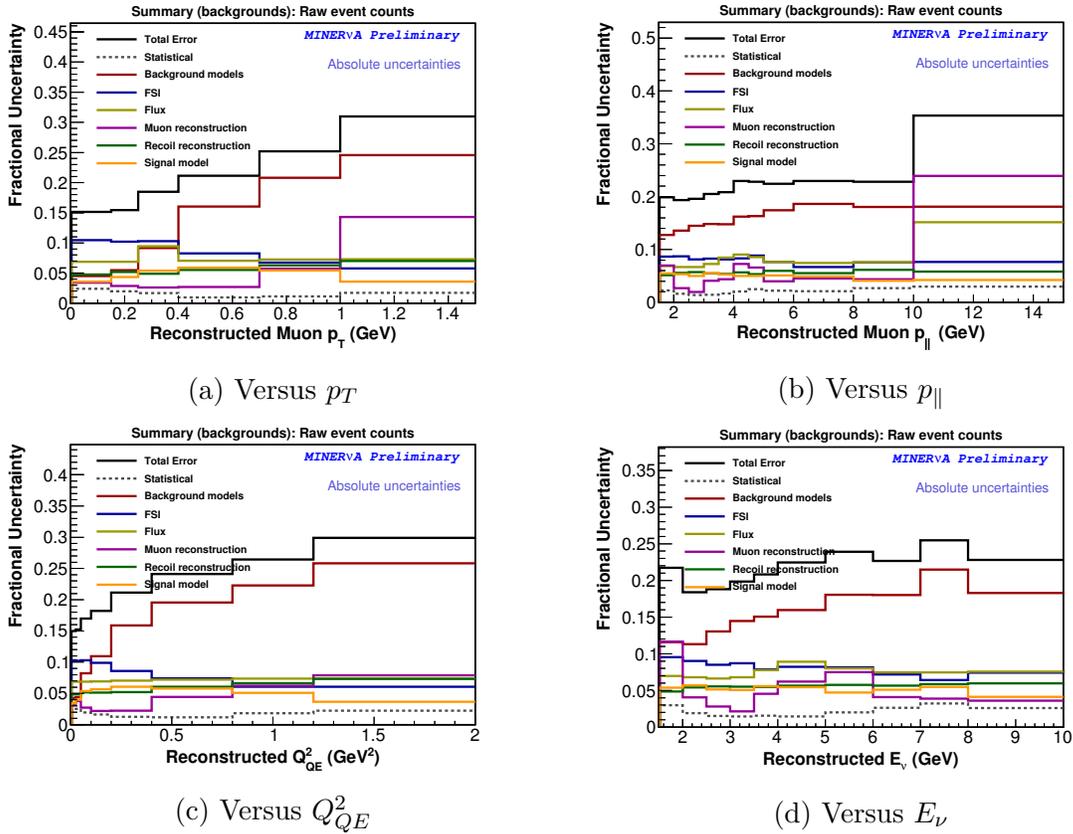


Figure 6.10. Summary of fractional uncertainties on raw Monte Carlo background distributions, as a fraction of the total background event count in simulation. The data-driven background scaling means that the effect of these on the cross section will be reduced.

If a “true GENIE CCQE” signal definition were used, there would be no contribution at all from CCQE events.

6.8.2. Background-subtracted data distributions

Figure 6.11 shows the distribution of uncertainties on the data after subtracting the scaled background fraction extracted from the Monte Carlo. Because of the data-driven scaling,

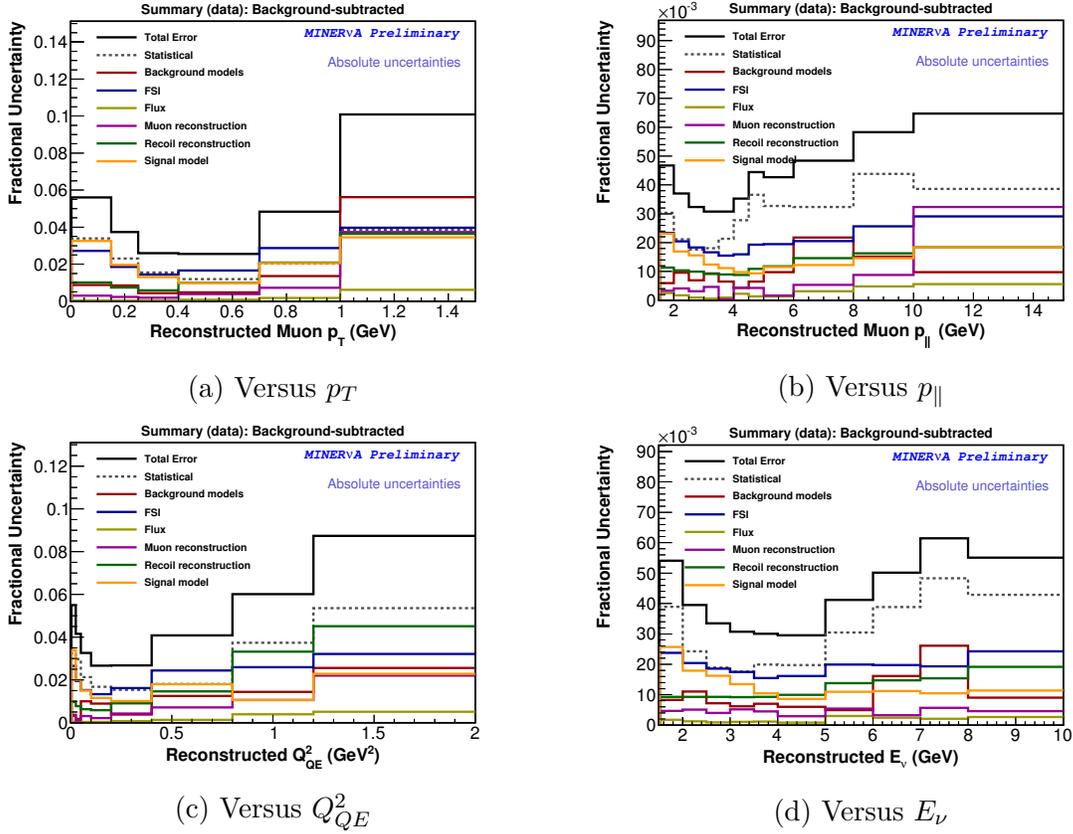


Figure 6.11. Summary of fractional uncertainties reconstructed data event count distributions, after tuned backgrounds have been subtracted

the effect on the cross section uncertainty from the background subtraction is modest. It is dominated by background model, FSI and recoil model uncertainties.

6.8.3. Unfolded data distributions

Figure 6.12 shows the distribution of uncertainties on the data after unsmearing using the migration matrix extracted from the Monte Carlo, which moves events from their reconstructed to their true bins. This generates a significant contribution from muon

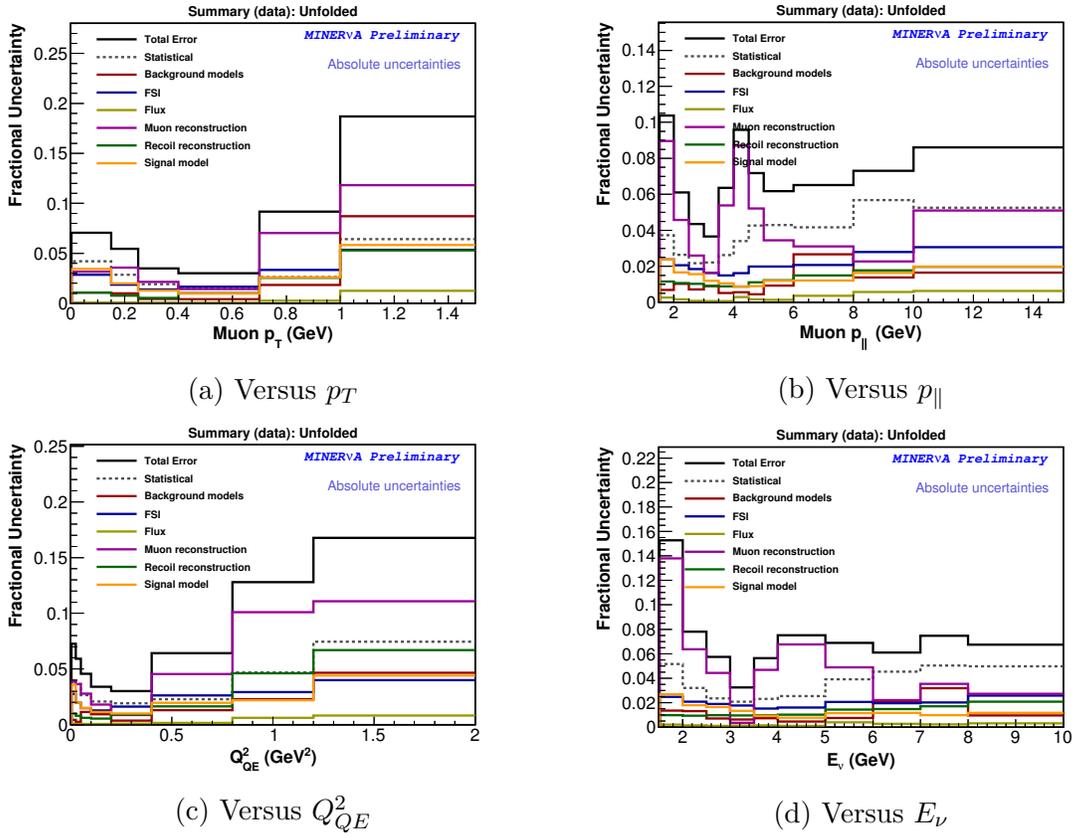


Figure 6.12. Summary of fractional uncertainties on data distributions, after backgrounds have been subtracted and the distribution has been unfolded

reconstruction uncertainties, which can cause events to move between bins. The muon energy scale dominates, as shown in figure 6.13

6.8.4. Acceptance-corrected data distributions

Figure 6.14 shows the distribution of uncertainties on the data after acceptance correction to the true Monte Carlo distribution. The main effect of this is to introduce large statistical uncertainties in bins with very low acceptance, particularly at high p_T , where

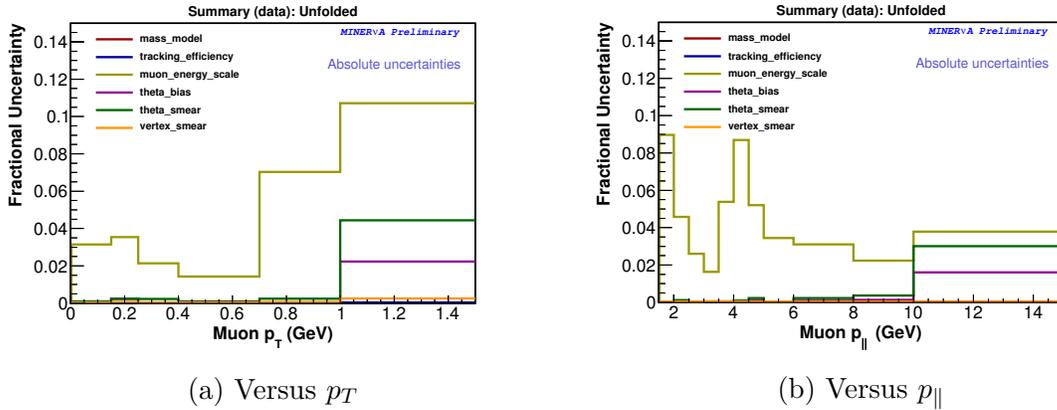


Figure 6.13. Summary of muon reconstruction uncertainties on data distributions, after backgrounds have been subtracted and the distribution has been unfolded

muons' high-angle trajectories make them very unlikely to hit the front face of MINOS, and at low $p_{||}$ where they lack the energy to reach the MINOS detector.

6.8.5. Cross section uncertainties

Figure 6.15 shows the distribution of uncertainties on the final cross sections. This introduces uncertainty on the flux measurement; to convert from an acceptance-corrected event count to a cross section, we must divide by the integrated flux. Our flux histograms have 100 universes each corresponding to the PPFX and beam-focusing uncertainties. This causes each flux universe to be scaled differently, introducing uncertainty. All other uncertainties are unchanged, as we divide by the central value.

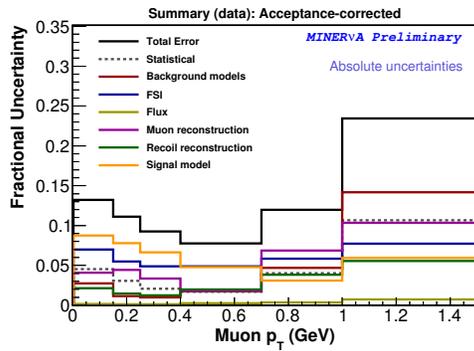
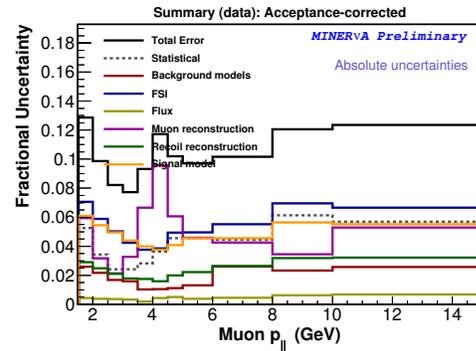
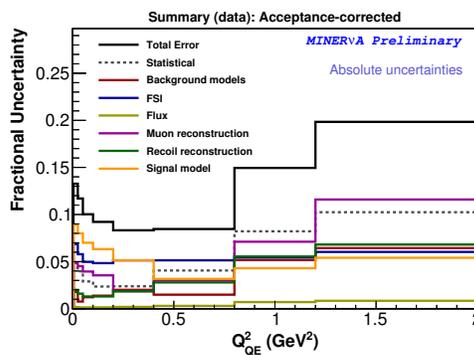
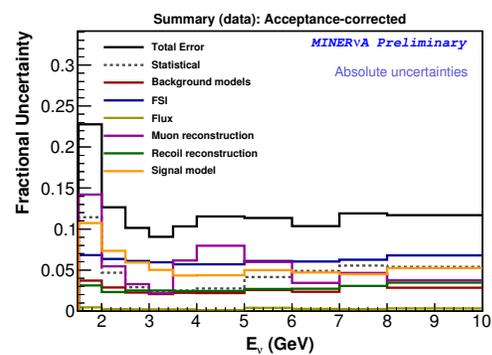
(a) Versus p_T (b) Versus p_{\parallel} (c) Versus Q_{QE}^2 (d) Versus E_{ν}

Figure 6.14. Summary of fractional uncertainties on acceptance-corrected data distributions

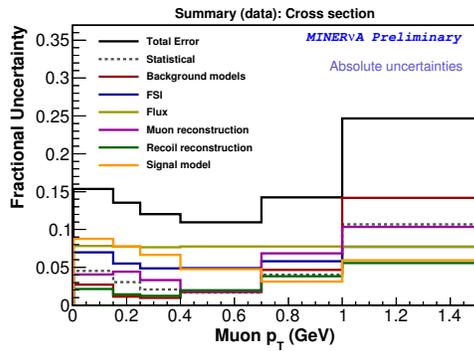
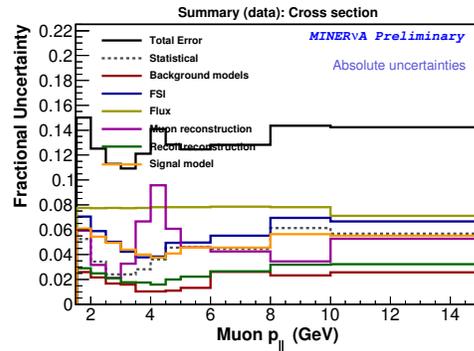
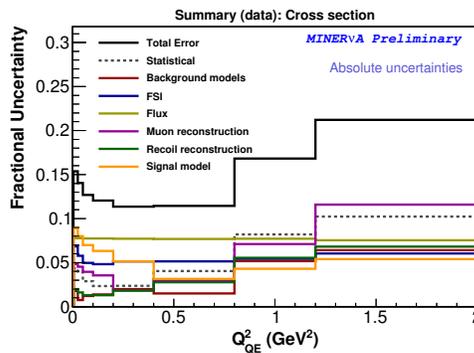
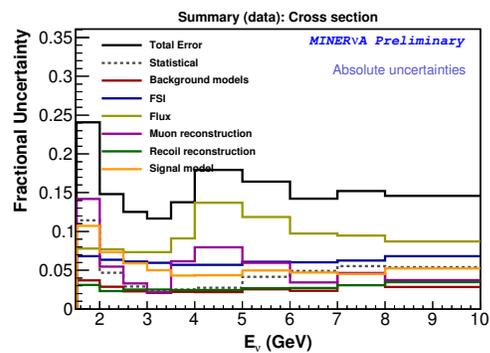
(a) Versus p_T (b) Versus p_{\parallel} (c) Versus Q_{QE}^2 (d) Versus E_{ν}

Figure 6.15. Summary of fractional uncertainties on the final data cross section distributions

6.9. Detailed breakdown of uncertainties vs. p_{\parallel} and p_T

As there are so many systematic uncertainties to consider, we categorize them in groups as follows:

- Flux uncertainties
- Muon reconstruction uncertainties
- Signal model uncertainties
- Background model uncertainties
- Final-state interaction uncertainties
- Recoil reconstruction uncertainties

To obtain a combined fractional uncertainty for a group, the individual group fractional uncertainties are added in quadrature:

$$(6.15) \quad \sigma_{\text{total}} = \sqrt{\sigma_1^2 + \sigma_2^2 + \dots}$$

By a “fractional” uncertainty, we mean the uncertainty in a given bin as a fraction of the total event count in that bin. This approach assumes that the uncertainties are uncorrelated. Note that in some cases an individual uncertainty line can in fact incorporate a few correlated uncertainties (for example, see the muon energy scale uncertainty discussed in section 6.4.1, which includes uncertainties in the MINERvA energy deposition rate and in the MINOS reconstruction of the muon energy).

Note also that the uncertainties shown here are only the values from the diagonal of a covariance matrix that shows how uncertainties are correlated from bin to bin. While the

off-diagonal terms of this matrix are not included in these plots, they will become important when making χ^2 comparisons to simulation and other theoretical models, where the way that the bins vary together is just as important as the magnitudes of the uncertainties in individual bins.

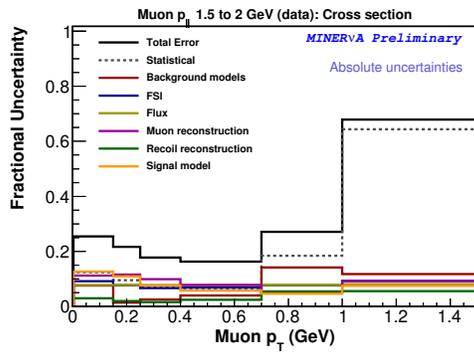
We present fractional uncertainties on the two-dimensional distributions, presented in bins of muon longitudinal and transverse momentum ($p_{\parallel\mu}$ and $p_{T\mu}$) and the neutrino energy and four-momentum transfer calculated from the quasi-elastic hypothesis (E_ν and Q_{QE}^2).

For each of these plots, the systematic uncertainty groups are represented by the colored lines. The grey dashed line represents the statistical uncertainty on the data.

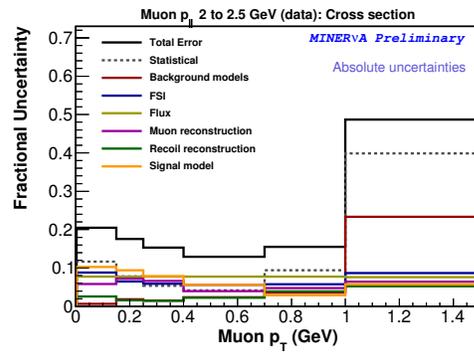
6.9.1. Fractional systematics vs. $p_{\parallel\mu}$ and $p_{T\mu}$

Figure 6.16 shows the grouped systematic uncertainties plotted vs. the muon transverse momentum, with one plot per bin of muon longitudinal momentum.

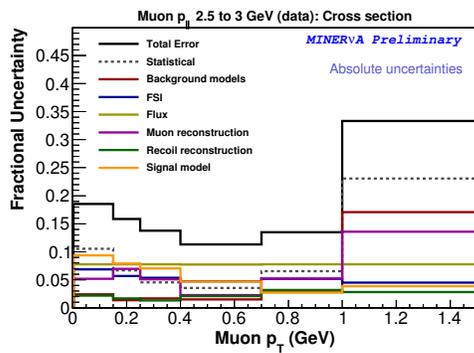
For these absolute uncertainties, the flux dominates, except for in regimes (low p_T , and high p_T / low p_{\parallel}) where there is little data, and we become statistics-dominated.



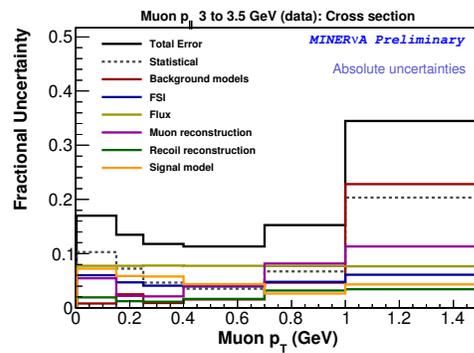
(a) 1.5 - 2 GeV



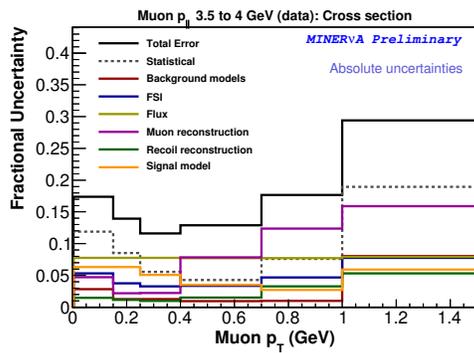
(b) 2 - 2.5 GeV



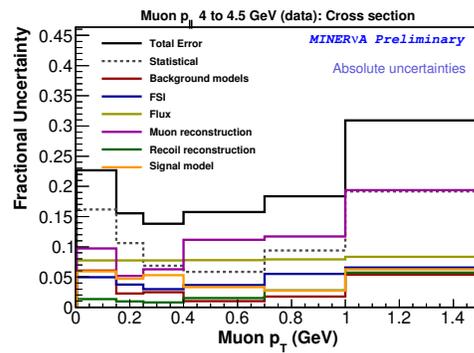
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

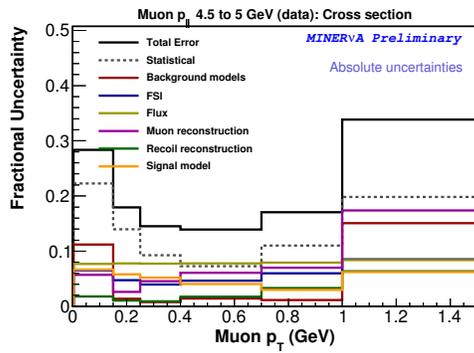


(e) 3.5 - 4 GeV

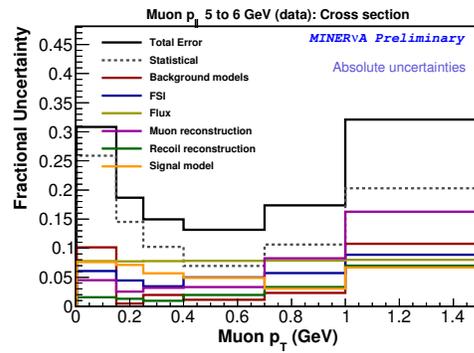


(f) 4 - 4.5 GeV

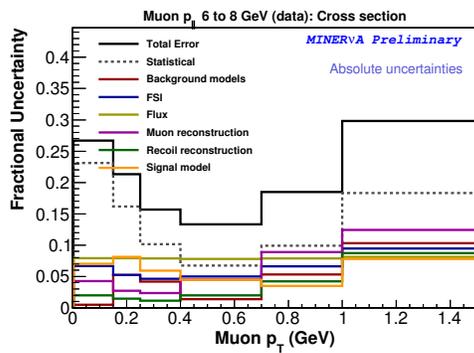
Figure 6.16. Absolute fractional uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued in next figure)



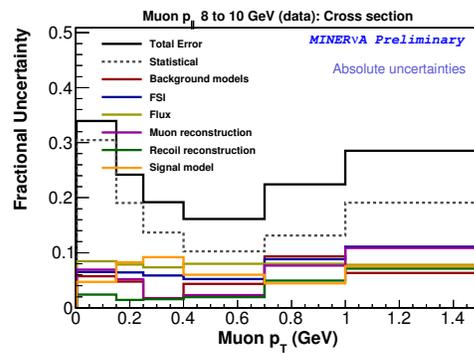
(g) 4.5 - 5 GeV



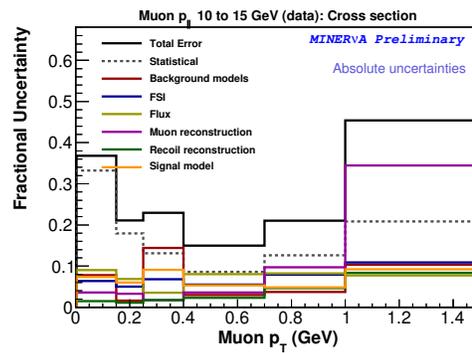
(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV

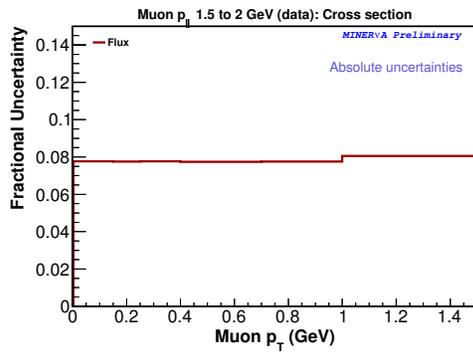


(k) 10 - 15 GeV

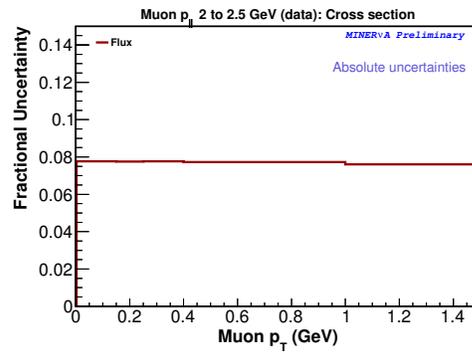
Figure 6.16. Absolute fractional uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued)

6.9.1.1. Flux systematics vs. $p_{\parallel\mu}$ and $p_{T\mu}$. The flux lines in figures 6.16 are formed by adding two separate flux systematics in quadrature. Figure 6.17 shows a breakdown of these into the individual components: beam focusing (see section 3.7.1.2), and hadron production (see section 3.7.1.1). For all muon momenta, the hadron production model dominates the flux uncertainty, although both sources contribute similarly to the shape uncertainty as the hadron production varies little with p_T and p_{\parallel} .

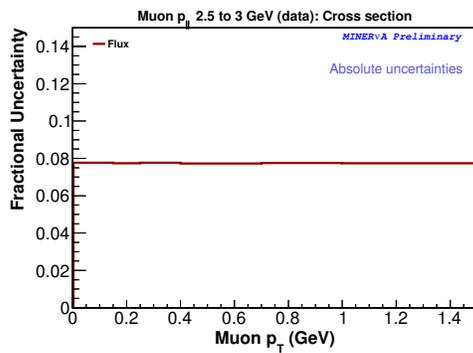
The data is shown in tabular form in table E.1.



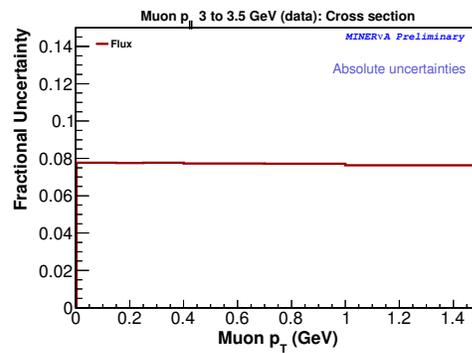
(a) 1.5 - 2 GeV



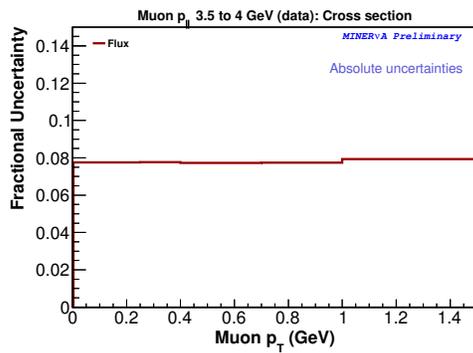
(b) 2 - 2.5 GeV



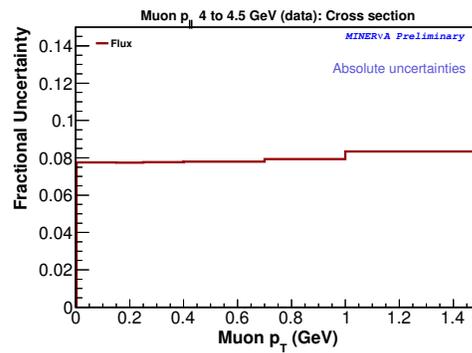
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

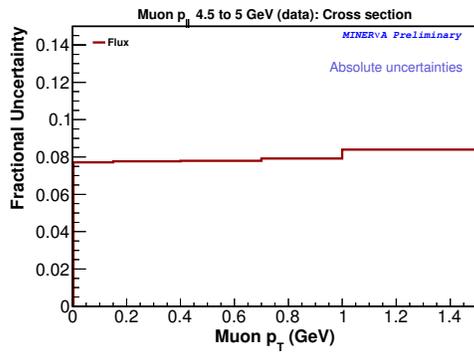


(e) 3.5 - 4 GeV

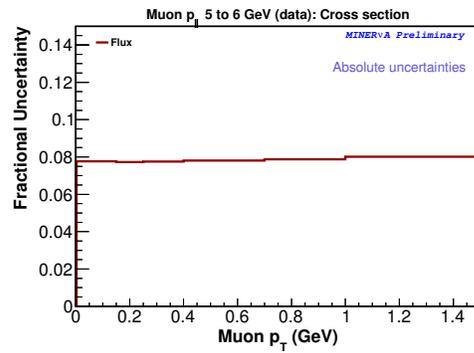


(f) 4 - 4.5 GeV

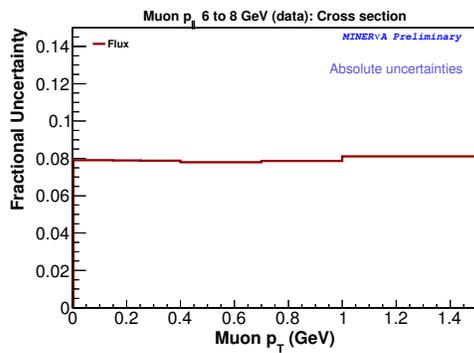
Figure 6.17. Flux uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



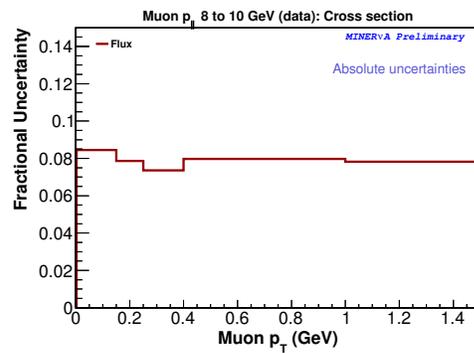
(g) 4.5 - 5 GeV



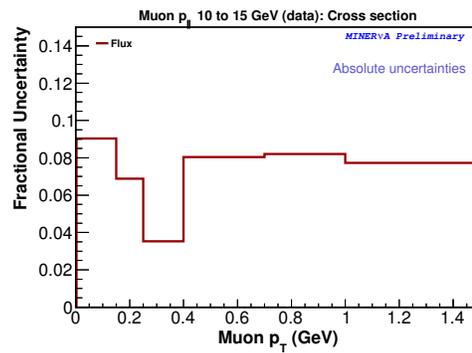
(h) 5 - 6 GeV



(i) 6 - 8 GeV



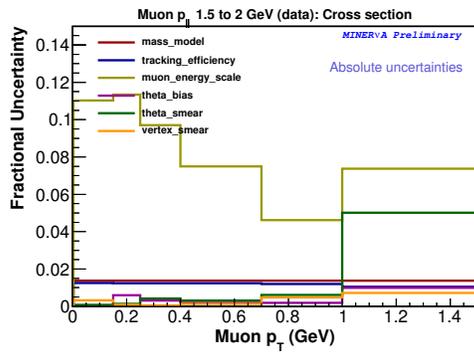
(j) 8 - 10 GeV



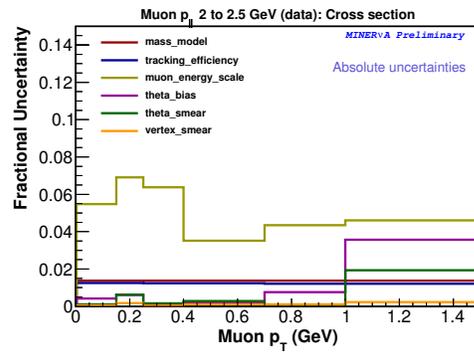
(k) 10 - 15 GeV

Figure 6.17. Flux uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

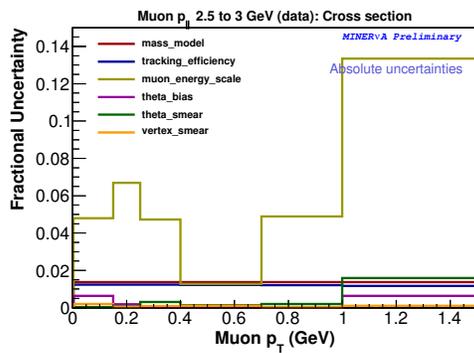
6.9.1.2. Muon reconstruction systematics vs. $p_{\parallel\mu}$ and $p_{T\mu}$. The muon reconstruction lines in figure 6.16 are formed by adding the individual muon reconstruction systematics in quadrature. Figure 6.18 and table E.2 show a breakdown of these into the individual components: muon energy scale (see section 6.4.1), tracking efficiency (section 6.4.2), muon angle smearing (section 6.4.3), muon angle bias 6.4.4) and vertex smearing 6.4.5). Over most of our momentum range, muon energy scale dominates the muon reconstruction systematics, although at very forward-going muons, uncertainty on the muon's angle becomes dominant.



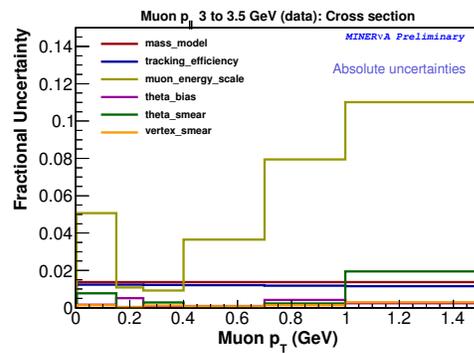
(a) 1.5 - 2 GeV



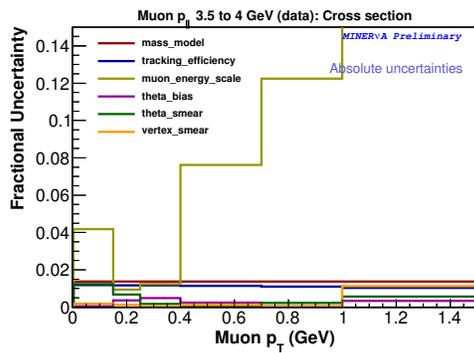
(b) 2 - 2.5 GeV



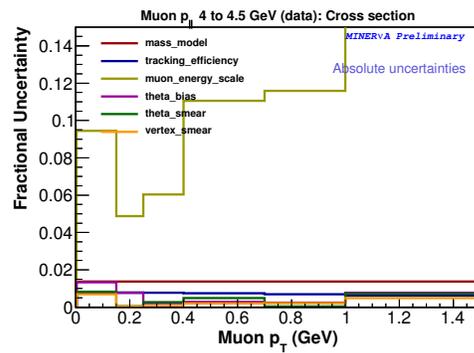
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

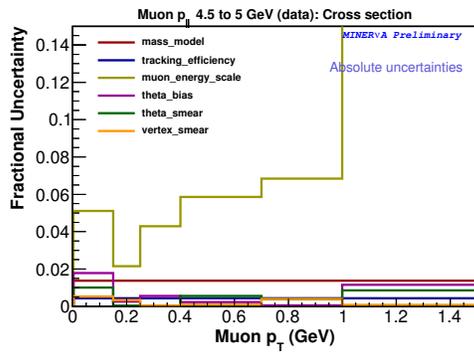


(e) 3.5 - 4 GeV

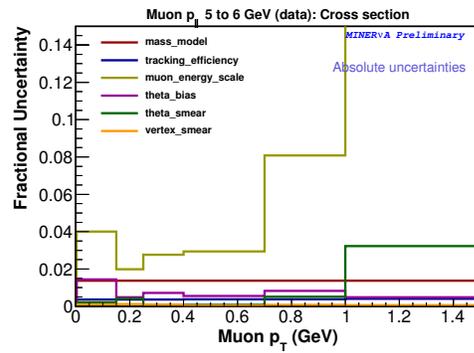


(f) 4 - 4.5 GeV

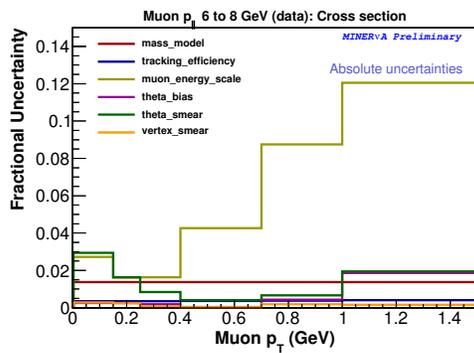
Figure 6.18. Muon reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



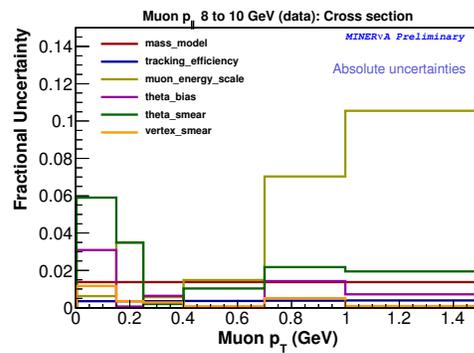
(g) 4.5 - 5 GeV



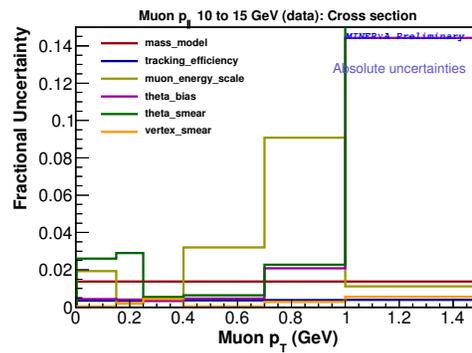
(h) 5 - 6 GeV



(i) 6 - 8 GeV



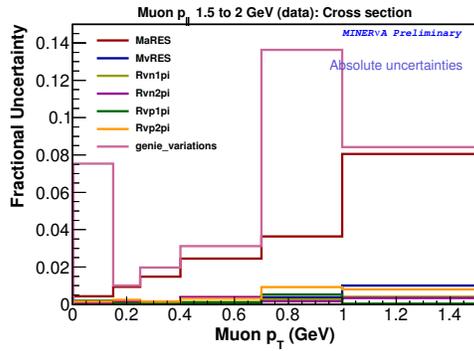
(j) 8 - 10 GeV



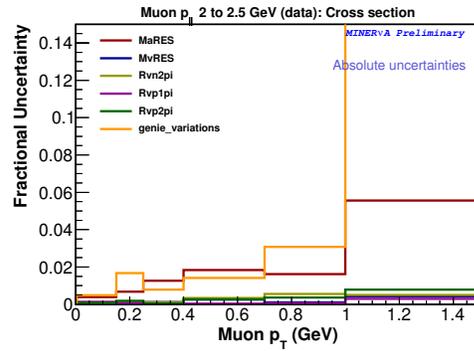
(k) 10 - 15 GeV

Figure 6.18. Muon reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

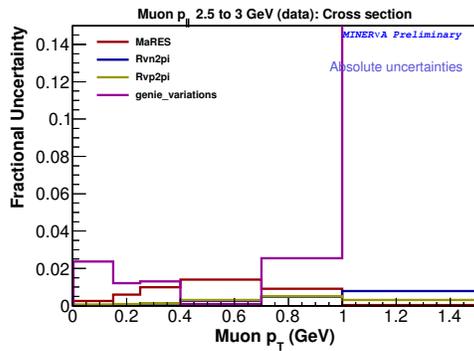
6.9.1.3. Background interaction model systematics vs. p_T and p_{\parallel} . The background interaction model lines in figure 6.16 are formed by adding the individual background interaction model systematics in quadrature. Figure 6.19 and table E.3 show a breakdown of these into the individual components.



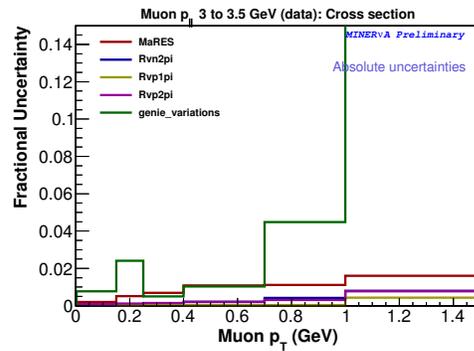
(a) 1.5 - 2 GeV



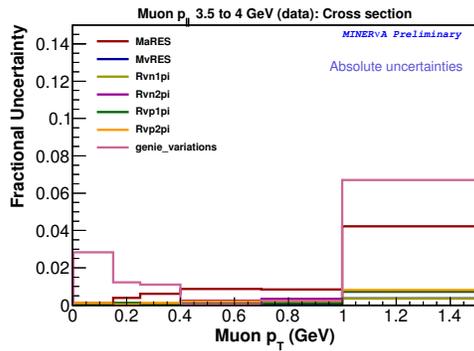
(b) 2 - 2.5 GeV



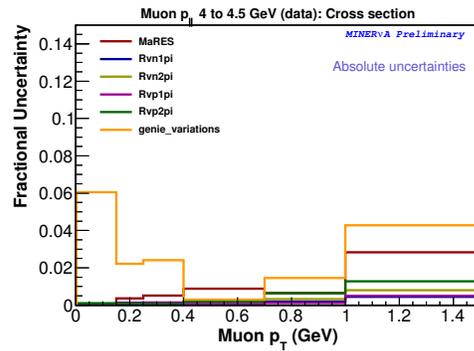
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

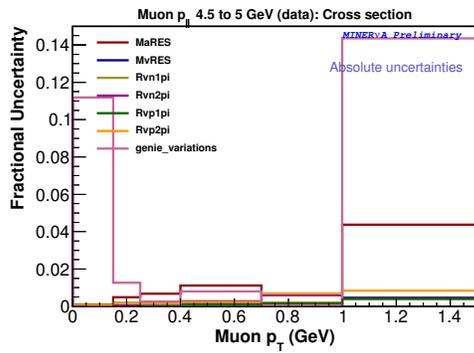


(e) 3.5 - 4 GeV

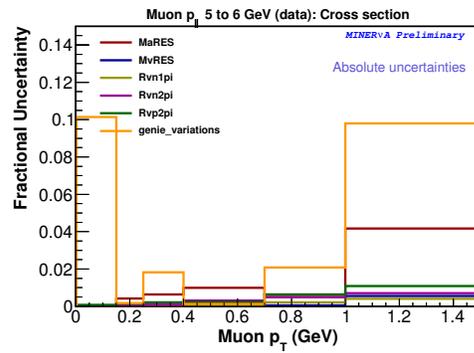


(f) 4 - 4.5 GeV

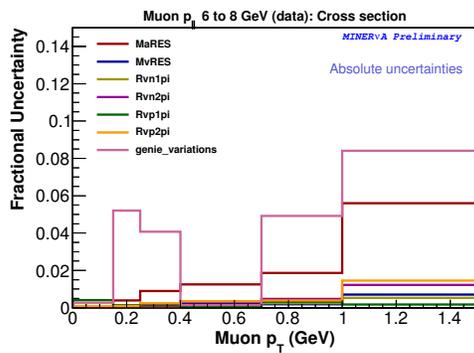
Figure 6.19. Background interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



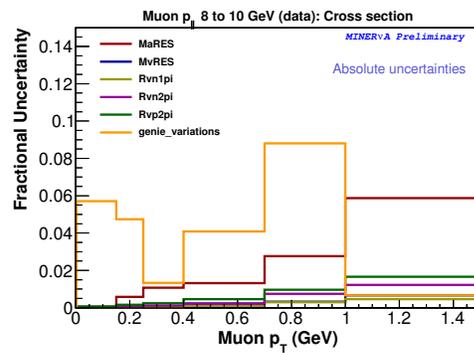
(g) 4.5 - 5 GeV



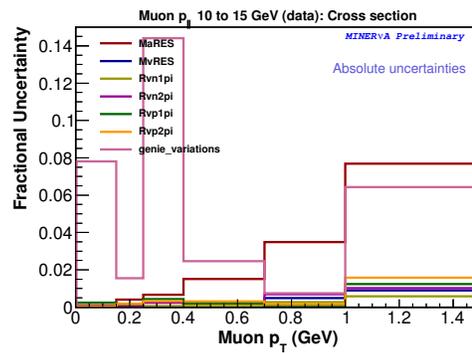
(h) 5 - 6 GeV



(i) 6 - 8 GeV



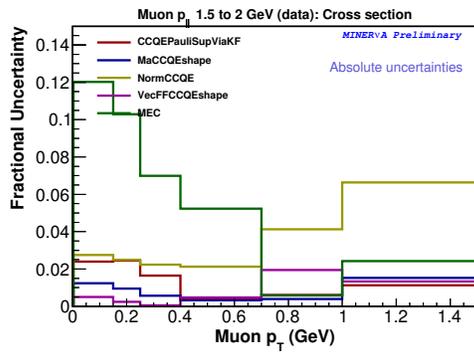
(j) 8 - 10 GeV



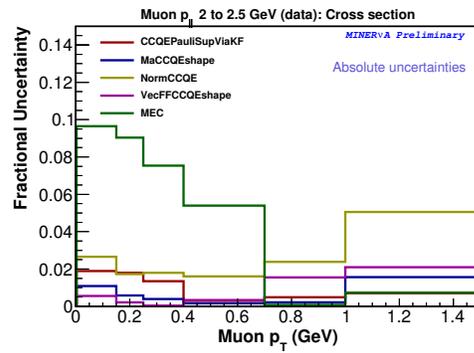
(k) 10 - 15 GeV

Figure 6.19. Background interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

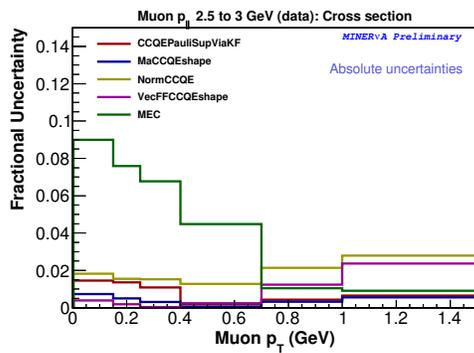
6.9.1.4. CCQE signal model systematics vs. p_T and p_{\parallel} . The CCQE signal interaction model lines in figure 6.16 are formed by adding the individual signal model systematics in quadrature. Figure 6.20 and table E.5 show a breakdown of these into the individual components.



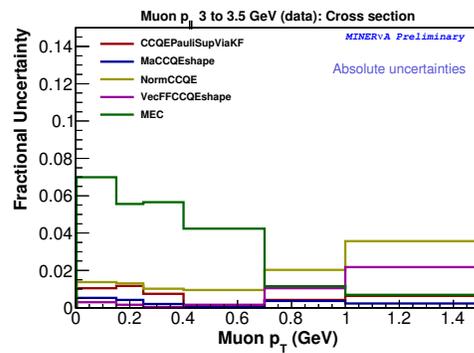
(a) 1.5 - 2 GeV



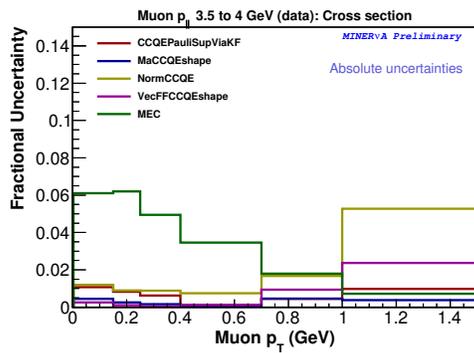
(b) 2 - 2.5 GeV



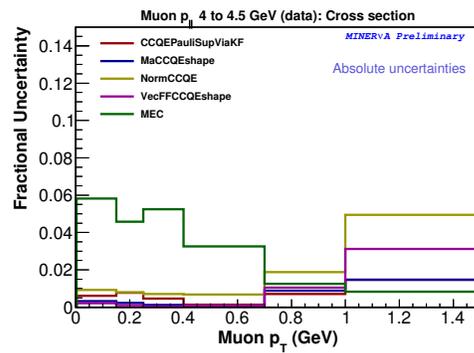
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

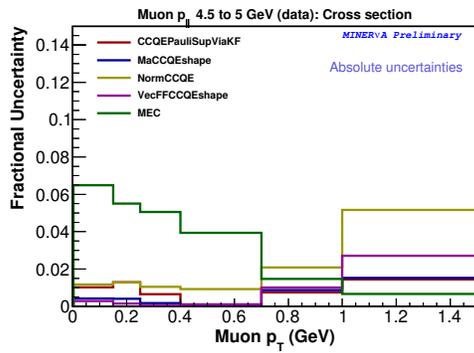


(e) 3.5 - 4 GeV

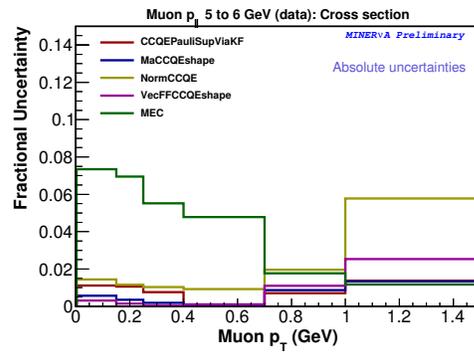


(f) 4 - 4.5 GeV

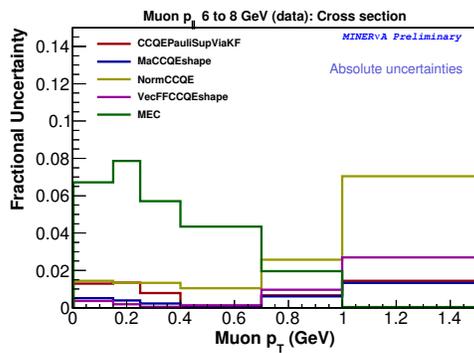
Figure 6.20. CCQE signal model uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



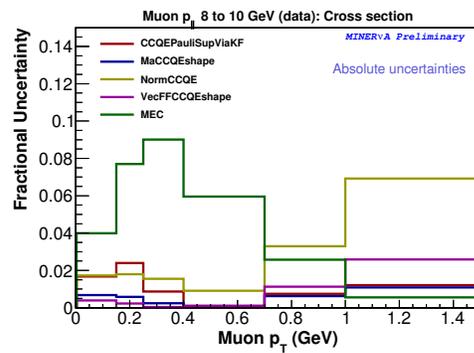
(g) 4.5 - 5 GeV



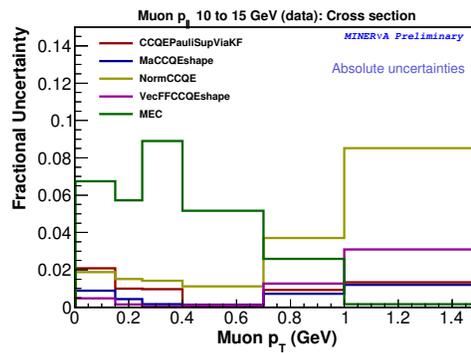
(h) 5 - 6 GeV



(i) 6 - 8 GeV



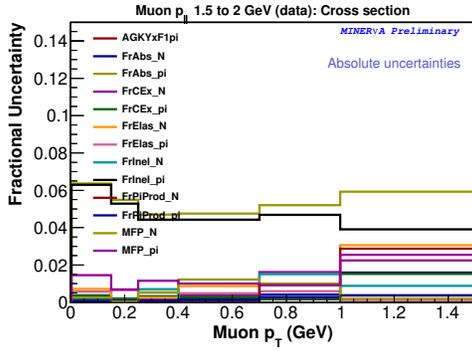
(j) 8 - 10 GeV



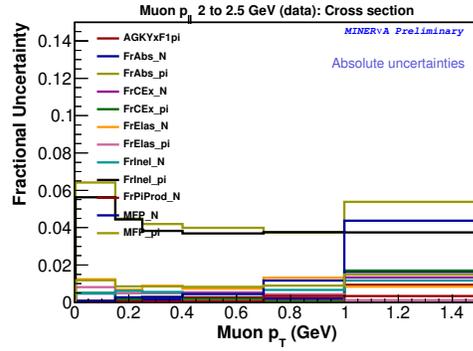
(k) 10 - 15 GeV

Figure 6.20. CCQE signal model uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

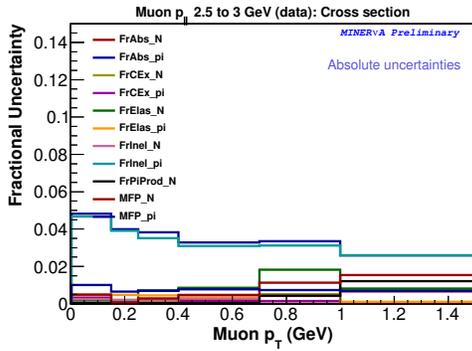
6.9.1.5. Final-state hadron interaction systematics vs. $p_{\parallel\mu}$ and $p_{T\mu}$. The hadron interaction lines in figure 6.16 are formed by adding the individual final-state hadron interaction systematics in quadrature. Figure 6.21 and table E.6, show a breakdown of these into the individual components.



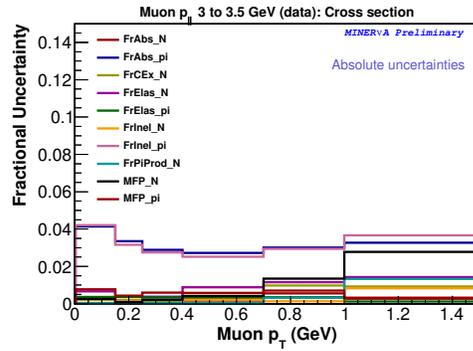
(a) 1.5 - 2 GeV



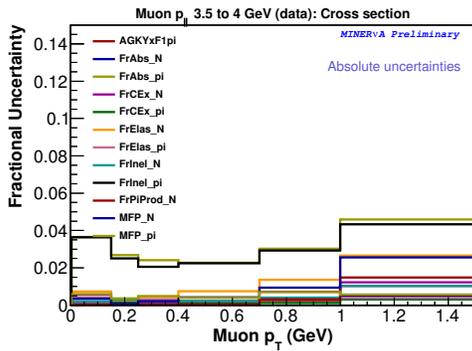
(b) 2 - 2.5 GeV



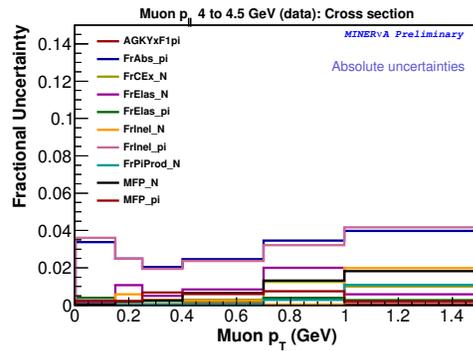
(c) 2.5 - 3 GeV



(d) 3 - 3.5 GeV

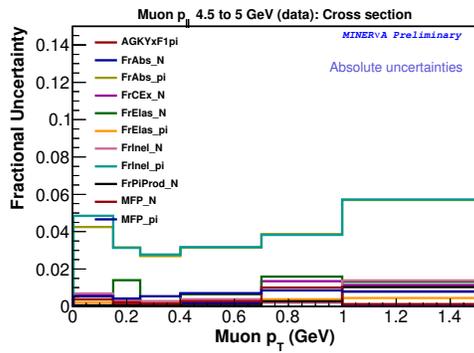


(e) 3.5 - 4 GeV

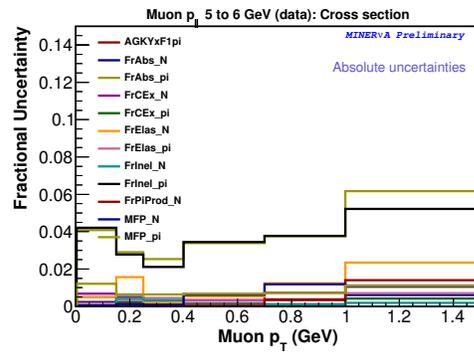


(f) 4 - 4.5 GeV

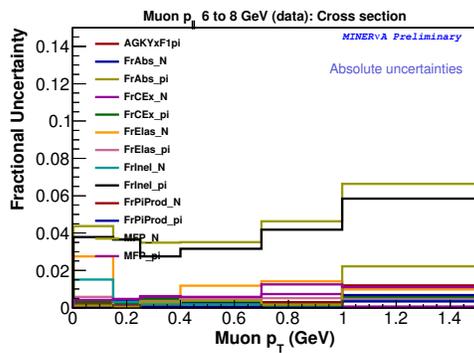
Figure 6.21. Final-state interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



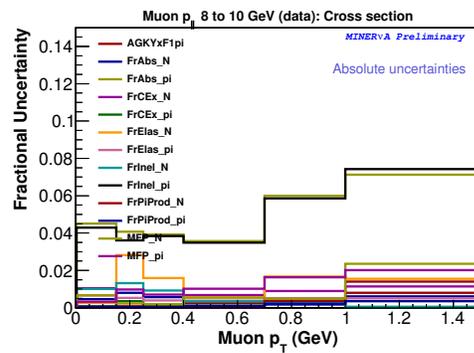
(g) 4.5 - 5 GeV



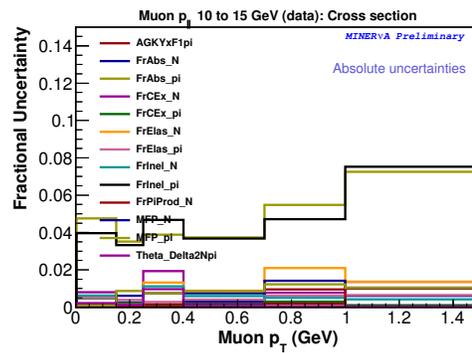
(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV



(k) 10 - 15 GeV

Figure 6.21. Final-state interaction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

6.9.1.6. Recoil reconstruction systematics vs. p_T and p_{\parallel} . The recoil reconstruction lines in figure 6.16 are formed by adding the individual recoil systematics in quadrature. Figure 6.22 and table E.9 show a breakdown of these into the individual components.

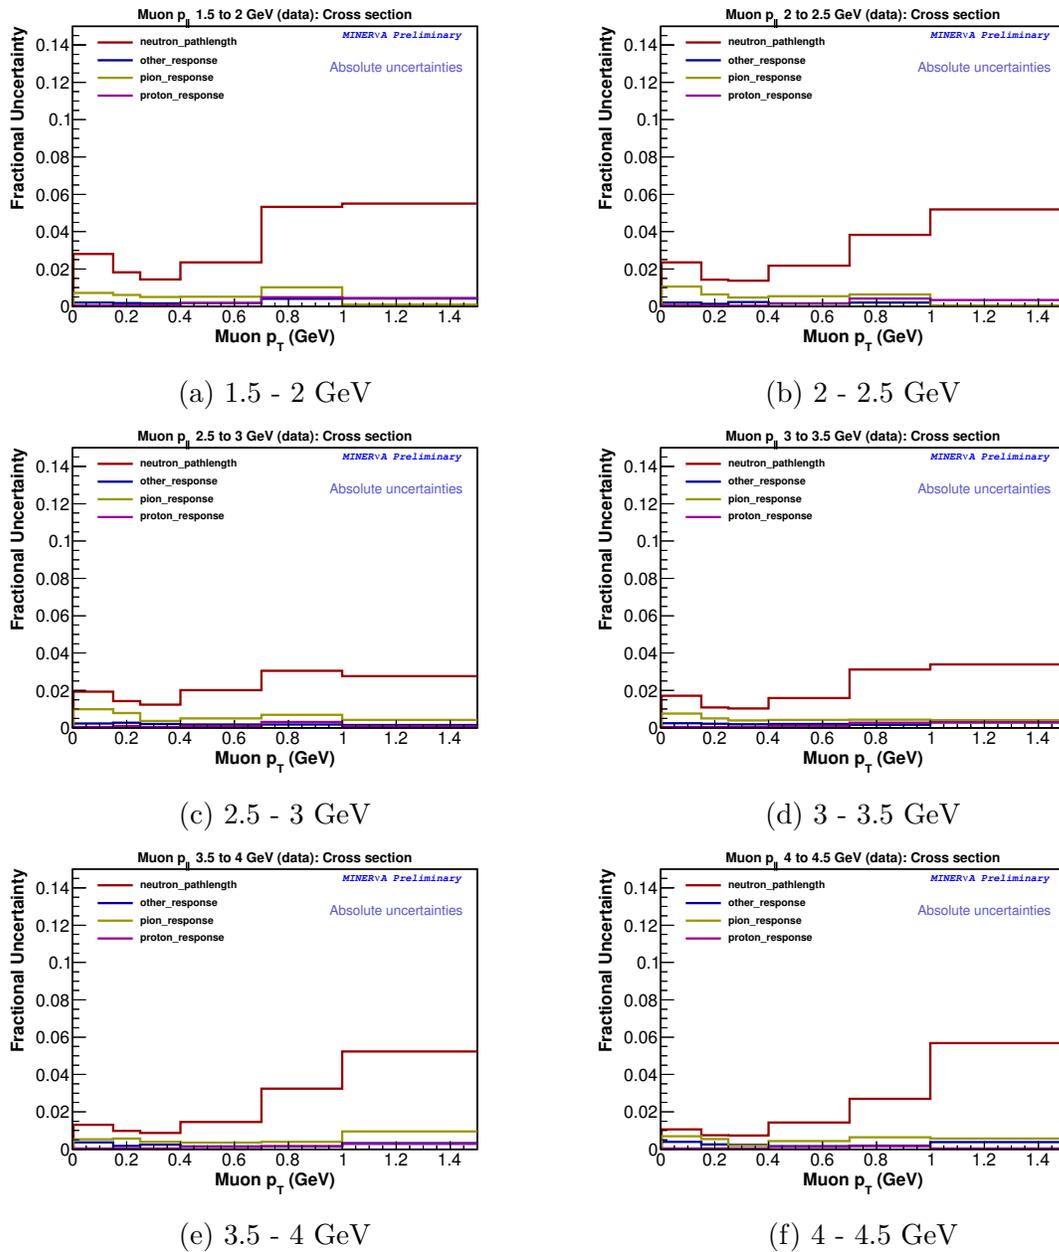
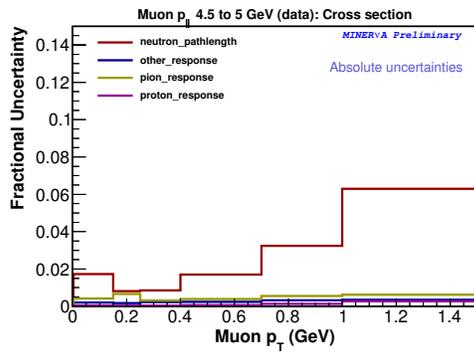
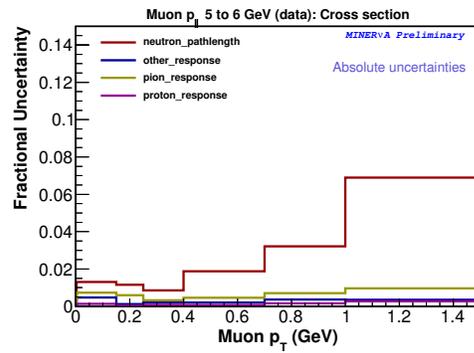


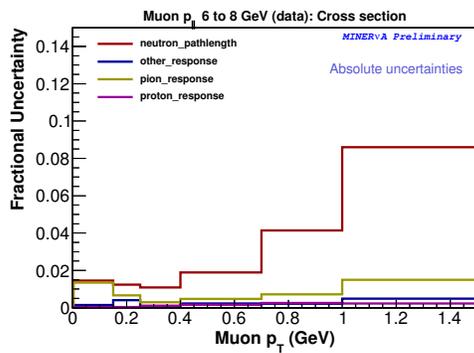
Figure 6.22. Recoil reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued on next page).



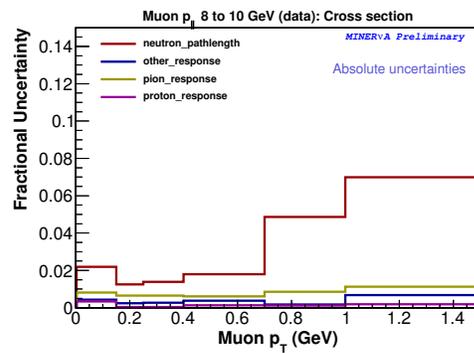
(g) 4.5 - 5 GeV



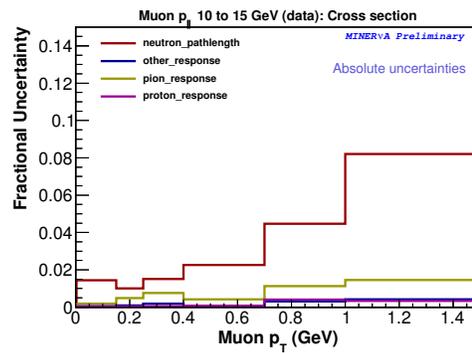
(h) 5 - 6 GeV



(i) 6 - 8 GeV



(j) 8 - 10 GeV



(k) 10 - 15 GeV

Figure 6.22. Recoil reconstruction uncertainties on the final cross section vs. muon transverse momentum, in bins of muon longitudinal momentum (continued).

CHAPTER 7

Discussion of results**7.1. Alternative nuclear models**

As explained in chapter 2, the nucleus is a complex environment, and the interactions between nucleons affect neutrino-nucleus scattering distributions. Several different theories attempt to explain the effects. The NuWro event generator [121] includes several of these models; the details of their implementations are described in [119]. All of the effects are described in chapter 2; this section includes a summary of the models available and how they can be combined. We will then compare the predictions of the different models to our cross section measurement.

NuWro's models consist of the following components:

- **Basic nuclear model:** Global or local Fermi gas, or spectral function
- **Random Phase Approximation:** Turned on or off
- **2p2h multi-nucleon effects:** Nieves and Transverse Enhancement models

The two Fermi gas models can be combined with the RPA effect. In each case, and in the case of spectral functions, an additional component from one or other of the 2p2h models can be added. A brief summary of the different models is given below. In addition, GENIE 2.10 allows one to generate samples of 2p2h events in the Nieves model; we will also compare our data to a GENIE sample with this additional effect.

7.1.1. Basic nuclear models

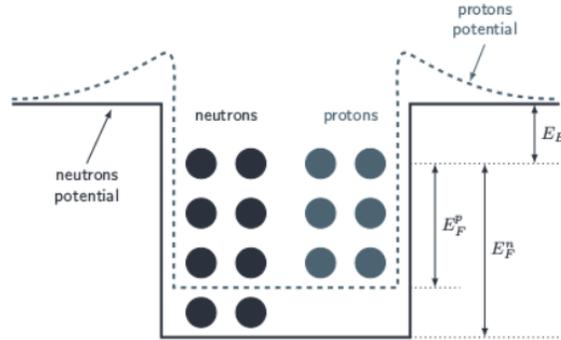


Figure 7.1. Cartoon of Fermi potential, from [119]

The **global relativistic Fermi gas model** (RFG or GFG) is the nuclear model used by GENIE [93], the simulation program we used for the unfolding, background estimation and acceptance correction. The model is described in section 2.3.1. It treats nucleons as independent particles moving in a fixed Fermi potential. GENIE takes the Fermi momentum k_F to be 221 MeV for carbon, with a binding energy E_B of 30 MeV for protons and 34 MeV for neutrons. It also includes Pauli blocking [144] and a Bodek-Ritchie high-momentum tail [59]. NuWro uses the same model as GENIE for quasi-elastic scattering. However NuWro's treatment of final-state interactions is rather different from GENIE's [119].

The **local relativistic Fermi gas model** (LFG), described in section 2.3.1.2, replaces the constant Fermi momentum with a distribution that depends on a particle's position in the nucleus. Cross section predictions using the LFG model are higher than those using the global Fermi gas model [119], although the effects are minor at our energies; figure 7.2 indicates that the effect on the total energy-dependent cross section is most significant

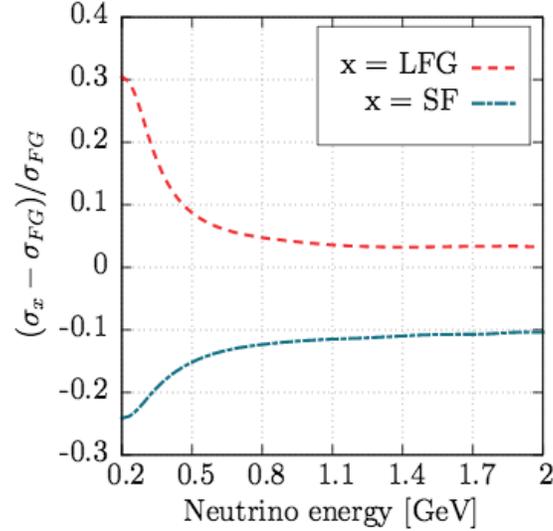


Figure 7.2. Fractional difference between NuWro’s cross sections for local Fermi gas and spectral function models, compared to the global Fermi gas. As seen in these comparisons, the LFG model enhances the cross section; spectral functions reduce it. Note that these effects are most pronounced at neutrino energies lower than the minimum (around 1.5 GeV) analyzed in this study. Plot reproduced from [119]

at low energies, flattening to a near-constant enhancement of less than 5% to the total cross section above in our energy range of above 1.5 GeV.

NuWro also includes a spectral function (SF) model, as explained in section 2.3.1.3. This models the nucleus with a shell model of energy levels, including a component corresponding to short-range correlations between pairs of nucleons. Cross sections calculated using spectral functions are smaller than those calculated with the Fermi gas model [119]. Again, in our energy range, figure 7.2 indicates an almost-flat reduction at the 10% level in the energy-dependent cross section.

All of these models use the Llewellyn Smith [145] model for the basic free-nucleon cross section. This model, as explained in section 2.2.2.3, has a free parameter, the axial

mass M_A . GENIE takes this to be 0.99 GeV, close to the world average. MiniBooNE, however, extracted a value of $M_A = 1.35$ GeV from their quasi-elastic fits [29]. We therefore also include these models with $M_A=1.35$ GeV in our comparisons.

7.1.2. RPA

The **Random-Phase Approximation** described in section 2.4, refers to long-range correlations that have a polarization effect, screening the weak charge of the W boson. The effect of RPA is to reduce cross sections at low Q^2 . This is a 1-particle-1-hole effect — that is, it does not cause multi-nucleon knockout. RPA is not part of the 2p2h models, and may occur in addition to those effects. When modeling distributions, we can turn on this effect to modify the Fermi gas distributions.

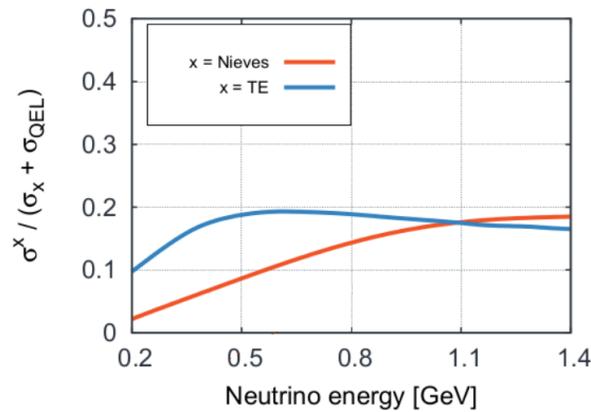


Figure 7.3. Additional fractional contribution of the Nieves and TEM multi-nucleon effects, compared to the quasi-elastic cross section, as predicted by NuWro. Reprinted from [119].

7.1.3. 2p2h multi-nucleon effects

Figure 7.3 shows the additional contributions to the cross section from the 2-particle-2-hole (2p2h) effects caused by correlations and meson exchange currents. These effects are believed to occur in addition to the 1p1h quasi-elastic cross section, and thus always represent an enhancement to the cross section. NuWro has provided us with two alternative models; these both attempt to simulate similar physics processes, but using different approaches. The **Nieves** model, described in section 2.5.3 is a prediction generated by evaluating the Feynman diagrams that lead to meson exchange currents (MEC), where pions and other particles are exchanged between pairs of nucleons, leading to the possibility of scattering from these correlated pairs, and to possible multi-nucleon knockout. The **Transverse Enhancement Model** or TEM, explained in section 2.5.2, attempts to parameterize an enhancement seen in electron scattering that is believed to have been caused by meson exchange current processes. These multi-nucleon enhancements can be added to any of NuWro's basic nuclear models, with or without RPA.

GENIE also provides an implementation of the Nieves model. None of the other nuclear models were available in GENIE at the time of writing, though there are plans to add them to future versions.

7.2. Comparison of quasi-elastic-like cross sections vs. p_T and p_{\parallel} with models

In order to compare our cross section with the many models from NuWro (and the one model from GENIE), we divide them up based on the basic nuclear model, and whether or not we include the RPA effect. Thus, we will show a set of double-differential cross section plots corresponding to these six configurations:

- NuWro Global Relativistic Fermi gas, without RPA effects
- NuWro Local Relativistic Fermi gas, without RPA effects
- NuWro Global Relativistic Fermi gas, with RPA effects
- NuWro Local Relativistic Fermi gas, with RPA effects
- NuWro Spectral Functions, without RPA effects
- GENIE Global Relativistic Fermi gas, without RPA effects

On each of the NuWro plots, we will show the data, along with lines representing the following NuWro models:

- $M_A = 0.99$ GeV, with no 2p2h effects
- $M_A = 1.35$ GeV, with no 2p2h effects
- $M_A = 0.99$ GeV, with the Nieves model for 2p2h contributions
- $M_A = 0.99$ GeV, with the transverse enhancement model for 2p2h contributions

We will also include a line corresponding to GENIE's prediction, which is always the global Fermi gas with $M_A = 0.99$ GeV, without RPA or 2p2h effects.

The sole GENIE plot will compare data to GENIE's GFG prediction with and without a 2p2h contribution from the Nieves model.

7.2.1. Global Relativistic Fermi gas, without RPA effects

These plots contain variations on the global Fermi gas model, with no contribution from the RPA effect. By comparing the blue Nieves and purple transverse-enhancement lines with the light green line corresponding to the NuWro with no MEC, we can see the contribution from 2p2h effects in the p_T / p_{\parallel} phase space. For all values of p_{\parallel} , Nieves predicts a greatest enhancement at low p_T , while TEM predicts enhancement at higher

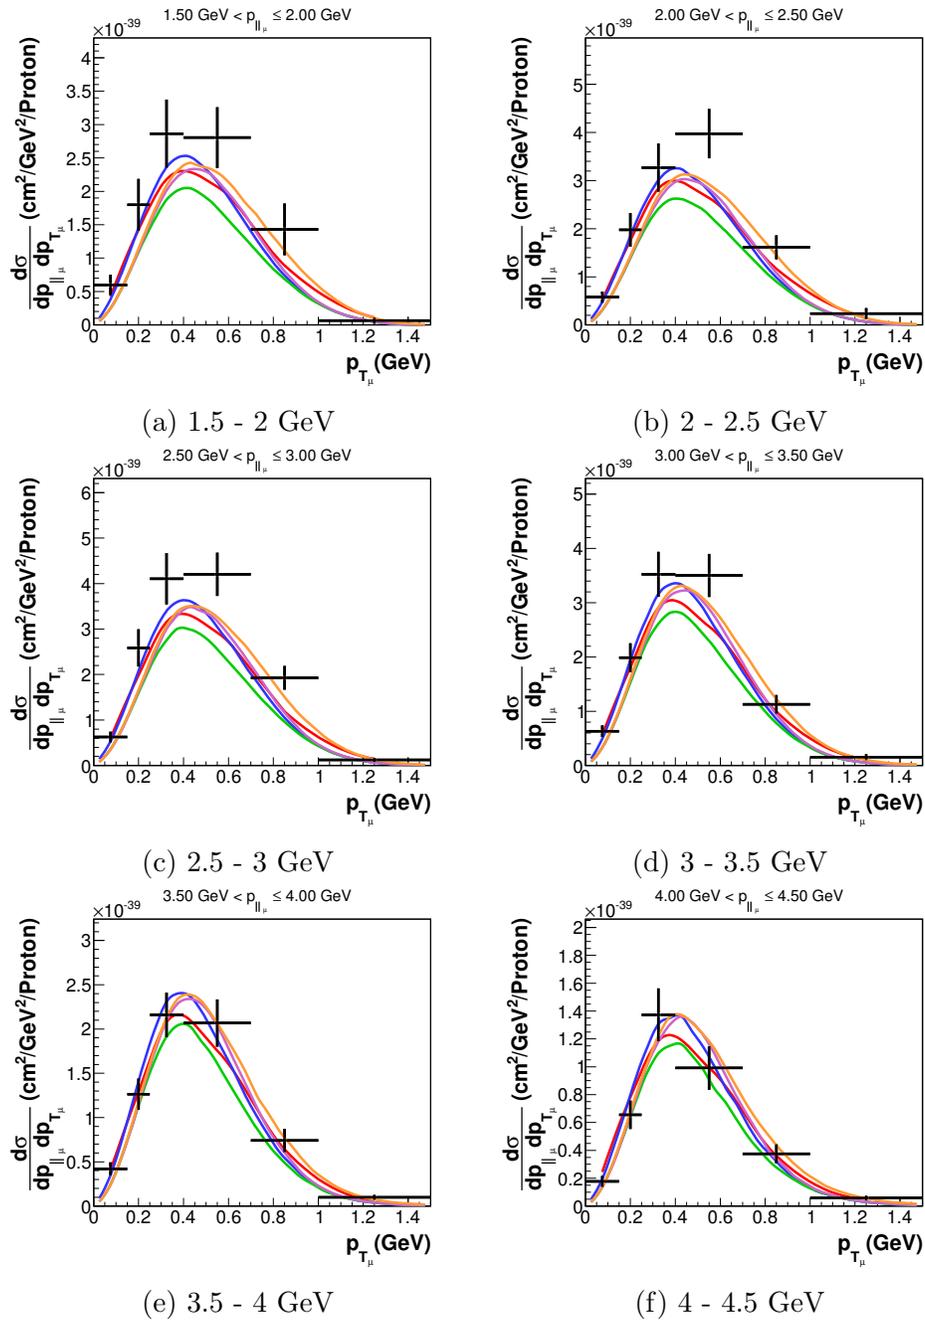


Figure 7.4. Comparison of measured quasi-elastic-like cross section with the **Global Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects (continued in next figure)

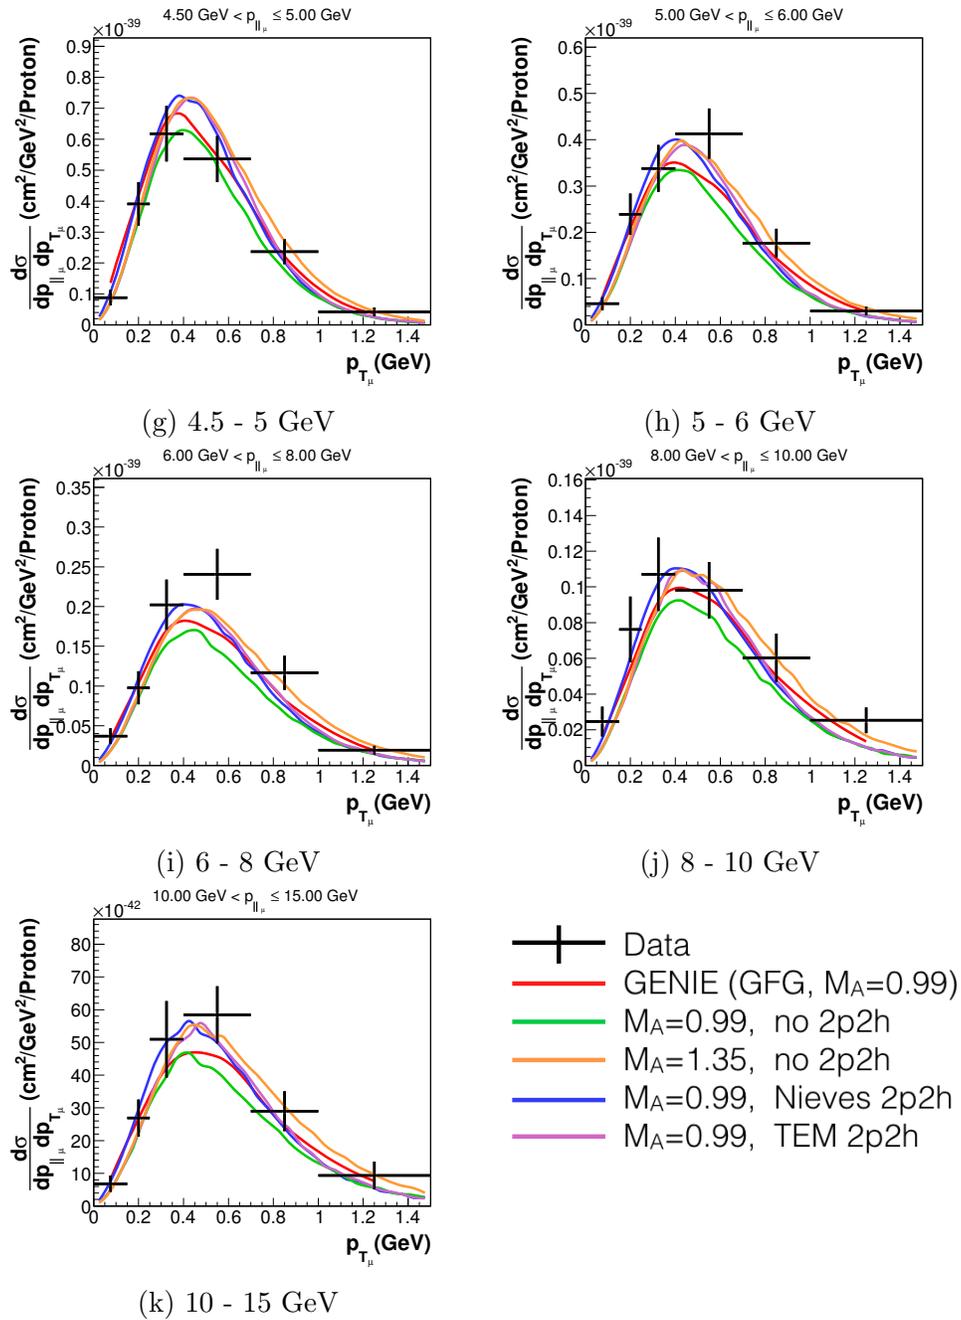


Figure 7.4. Comparison of measured quasi-elastic-like cross section with the **Global Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects (continued)

p_T , around 0.6 GeV. At high p_T , both of these converge to join the non-enhanced NuWro line.

A particular feature of these plots is that the pale green line represents NuWro's modeling of the non-enhanced global Fermi gas, with no RPA or 2p2h effects; in other words, this is modeling the same effects as GENIE. However, the NuWro line is markedly different from the GENIE (red) line; NuWro's prediction is approximately 80% of GENIE's, with a particularly low cross section at low p_T and low p_{\parallel} . A possible explanation for this could be that NuWro and GENIE use rather different models for final-state interactions, which contribute to the quasi-elastic-like cross section.

7.2.2. Local Relativistic Fermi gas, without RPA effects

Compared to the global Fermi gas, the local Fermi gas model generates an enhancement to the cross section, especially at low p_T and low p_{\parallel} . Its effect does not appear sufficient to explain the data excess at low p_T and p_{\parallel} . The GENIE model in the plots remains as global Fermi gas. A local Fermi gas model will be added to GENIE in an upcoming release, but is not available in GENIE 2.8.4; it will be an interesting exercise to compare the GENIE LFG model with this data.

7.2.3. Global Relativistic Fermi gas, with RPA effects

The effect of RPA is to suppress the cross section at low Q^2 (compare figure 7.6, which includes RPA, to figure 7.4, which does not). This mostly translates to a suppression at low p_T , across the whole p_{\parallel} range. As the NuWro cross sections are already lower than our measurement, the RPA effect makes the models appear less consistent with our

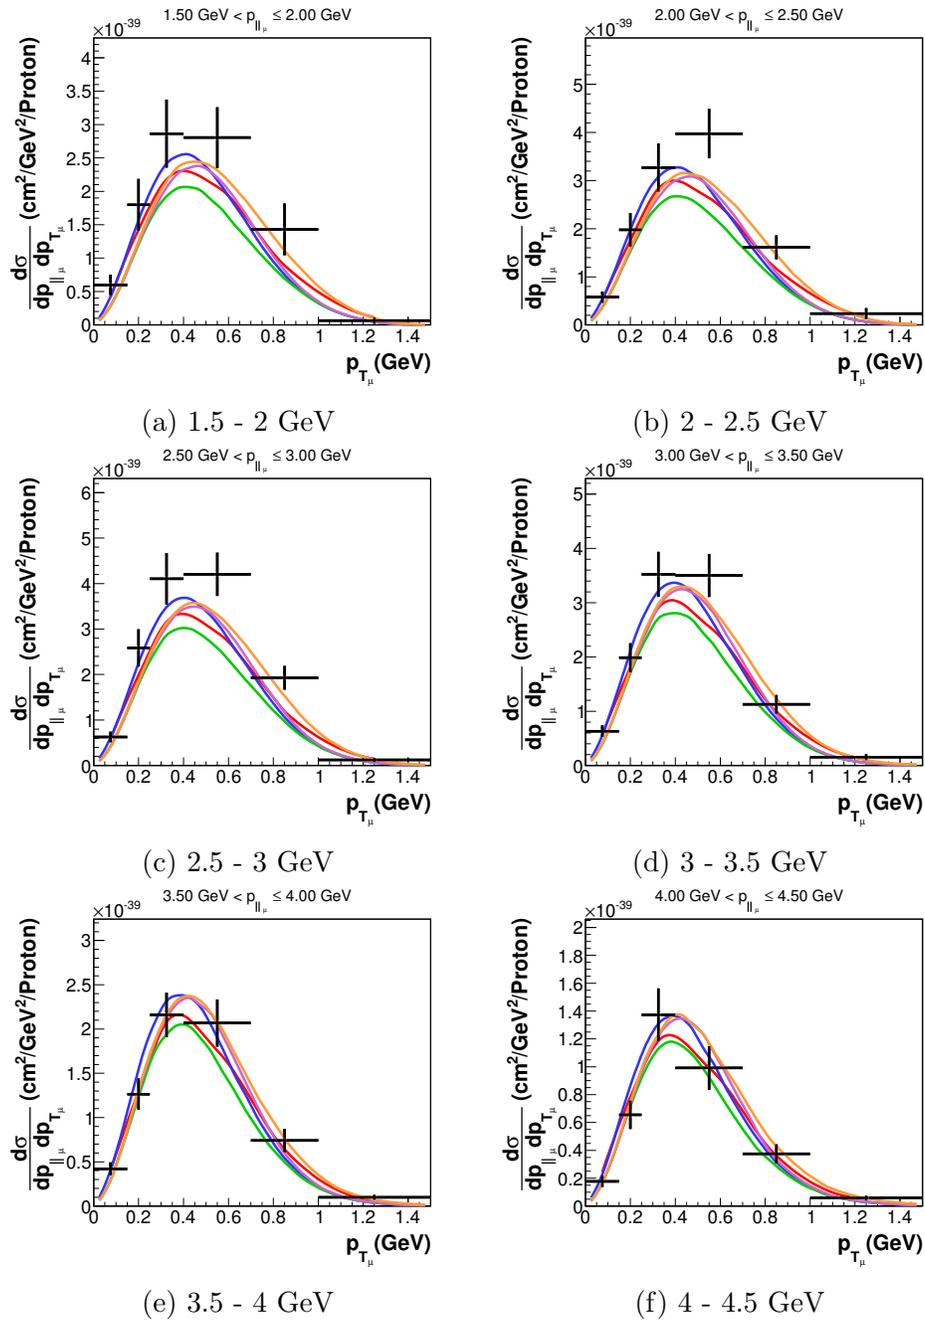


Figure 7.5. Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects, to GENIE (continued in next figure)

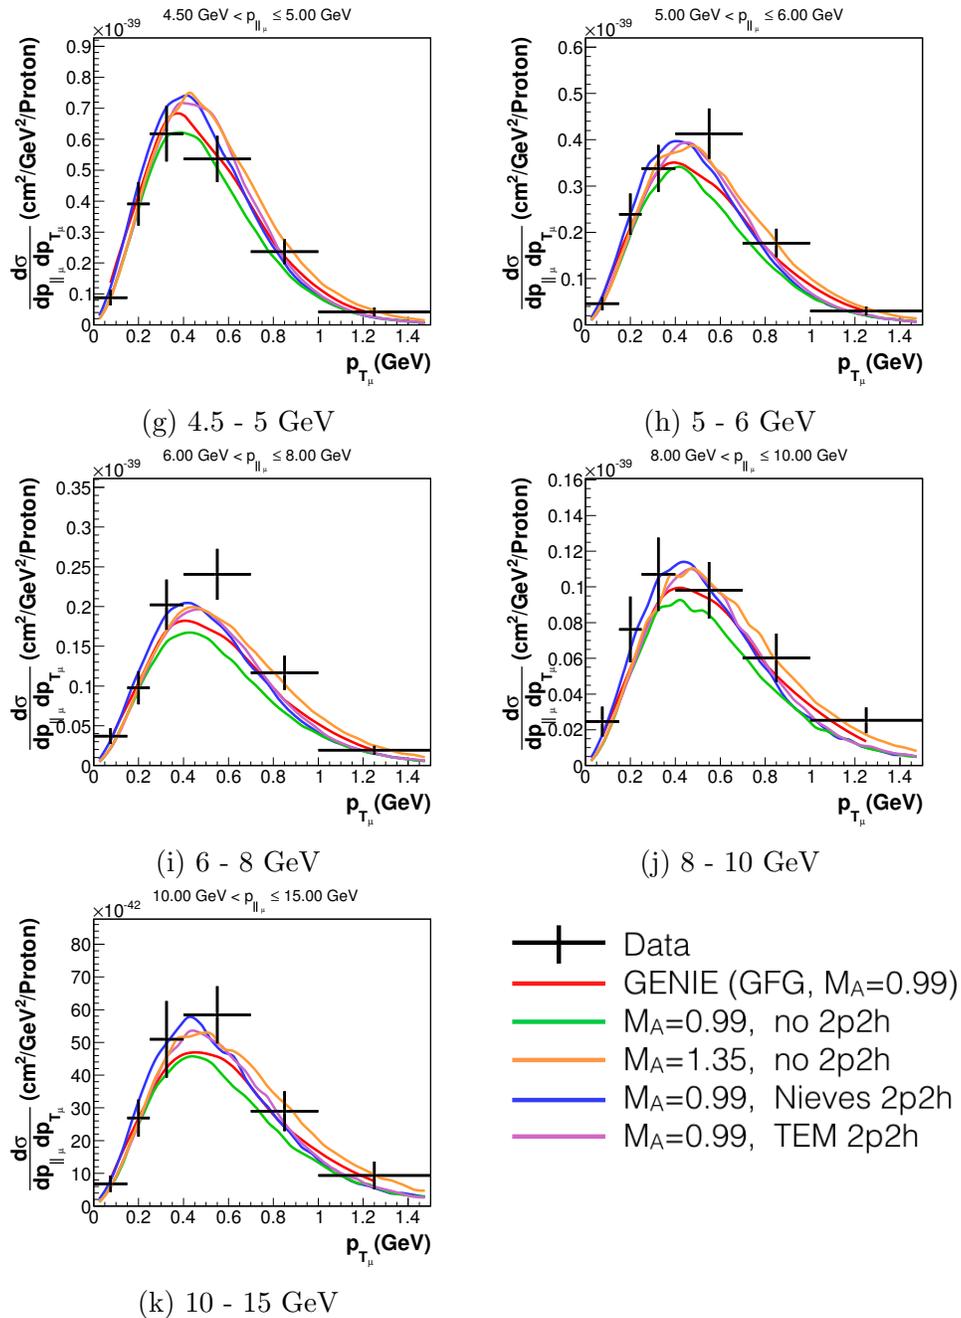


Figure 7.5. Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, **without RPA**, with and without 2p2h effects, to GENIE (continued)

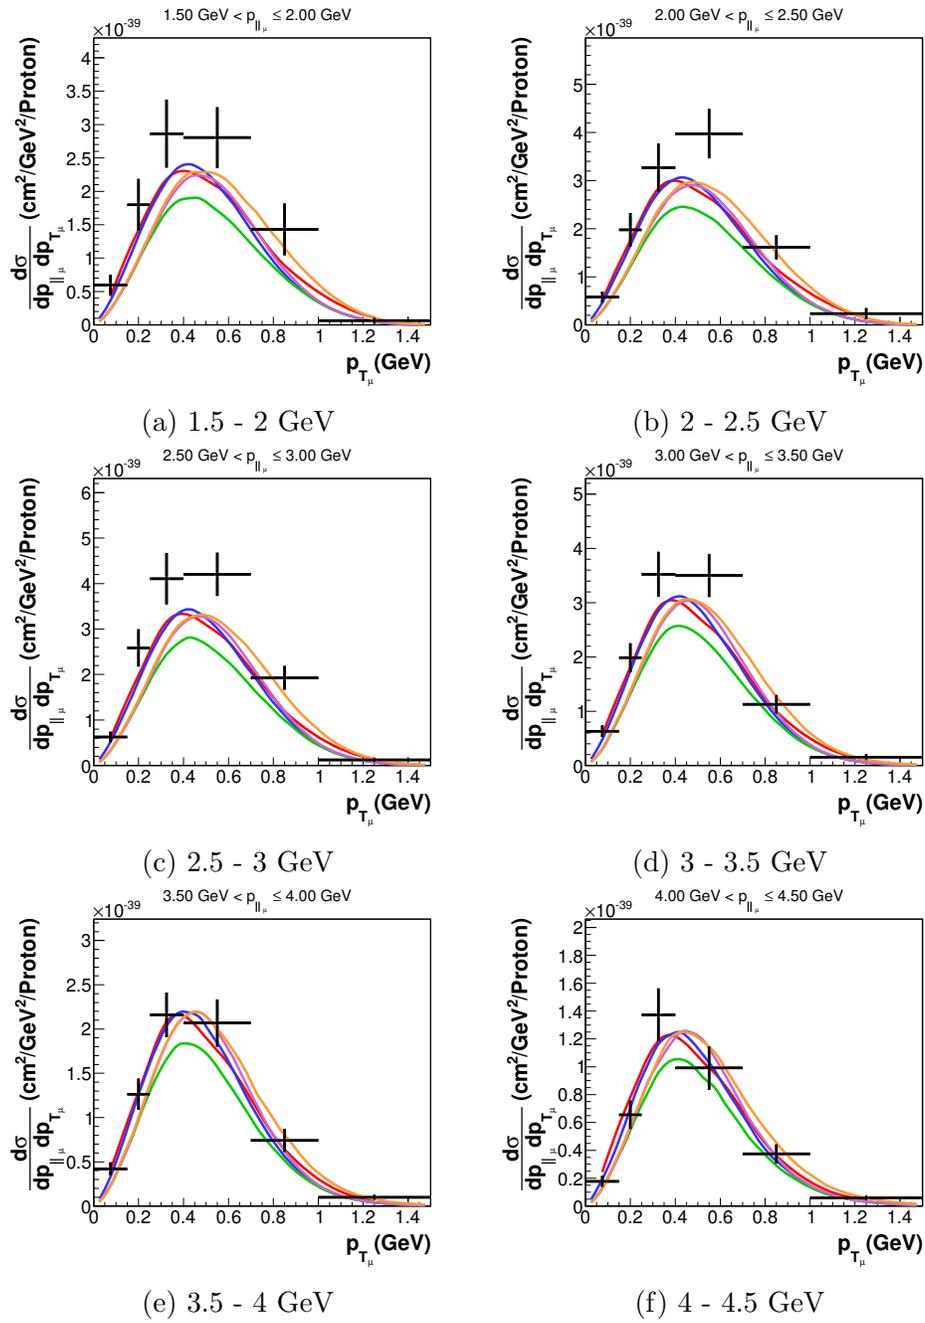


Figure 7.6. Comparison of measured quasi-elastic-like cross section and **Global Relativistic Fermi Gas** models, with **RPA**, with and without 2p2h effects, to GENIE (continued in next figure)

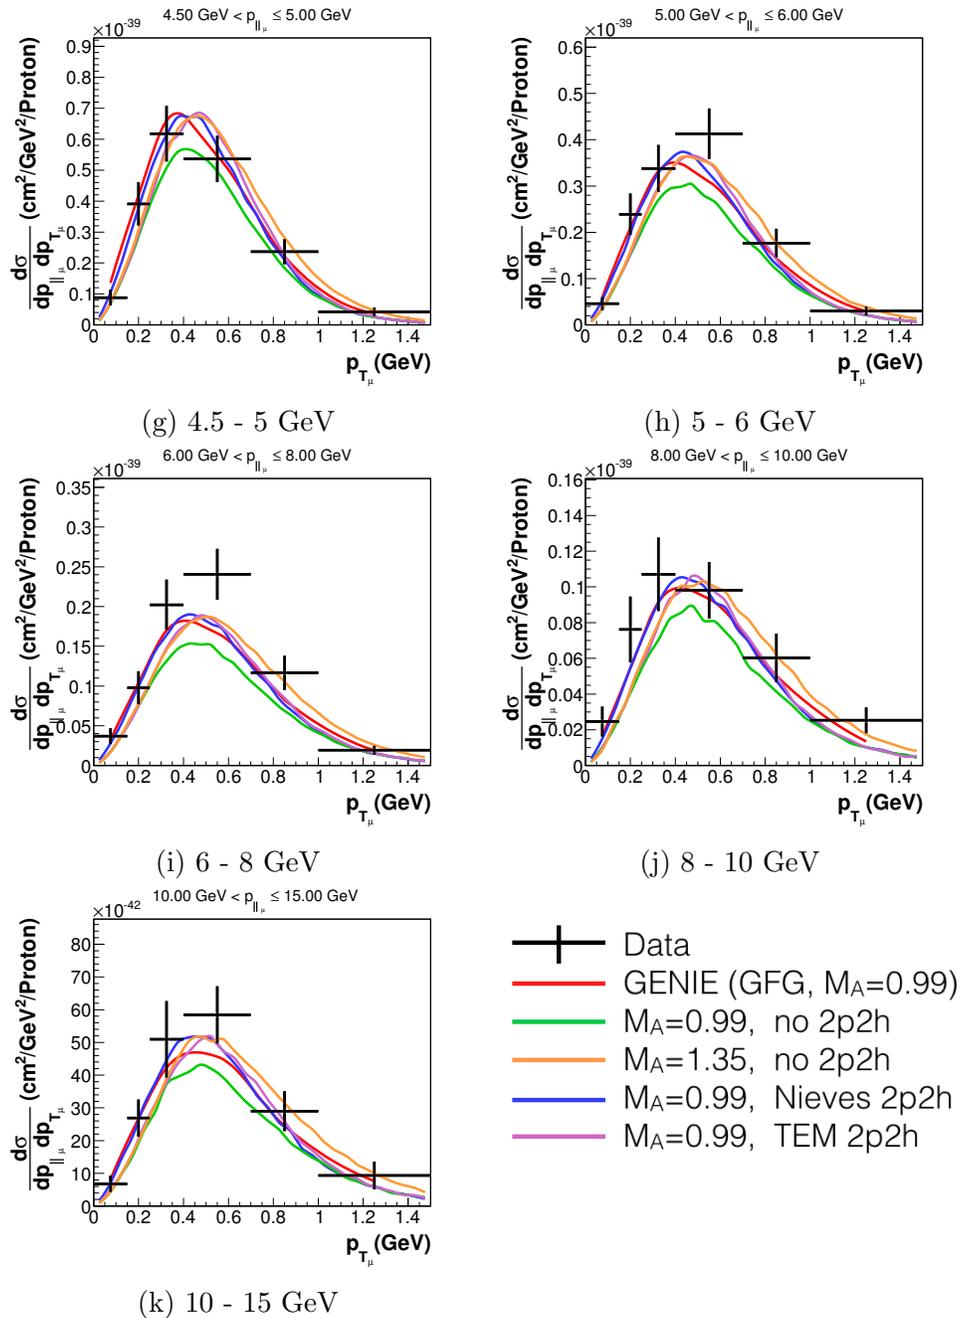


Figure 7.6. Comparison of measured quasi-elastic-like cross section and **Global Relativistic Fermi Gas** models, with RPA, with and without 2p2h effects, to GENIE (continued)

distribution. A slight enhancement due to RPA at high p_T may provide better agreement with data, but our uncertainty in this area is large. There is no RPA model currently available in GENIE.

7.2.4. Local Relativistic Fermi gas, with RPA effects

The combined effect of RPA's low p_T suppression of RPA and the increased cross section due to the LFG leads to a lower cross section than for the basic RFG models at low p_T , with an enhancement at higher p_T .

7.2.5. Spectral functions

Spectral functions produce a lower cross section than does the global Fermi gas model, as shown in figure 7.2. As NuWro already under-predicts the data with its Fermi Gas prediction, spectral functions exacerbate this issue.

7.2.6. 2p2h effects modeled in GENIE

GENIE's implementation of the Nieves model was used to generate a sample of 2p2h effects (also used to calculate the MEC systematic error band, as explained in section 6.2.3). When POT-normalized to our standard simulation, the 2p2h sample contains 26% as many 2p2h events as the standard simulation contains CCQE events (there are 18% as many QE-like 2p2h events in the POT-normalized Nieves sample as there are 1p1h QE-like events in the standard sample). Figure 7.9 shows the effect of adding this 2p2h effect to GENIE's cross section prediction. The model with 2p2h added appears to show better agreement with the data than the model without 2p2h.

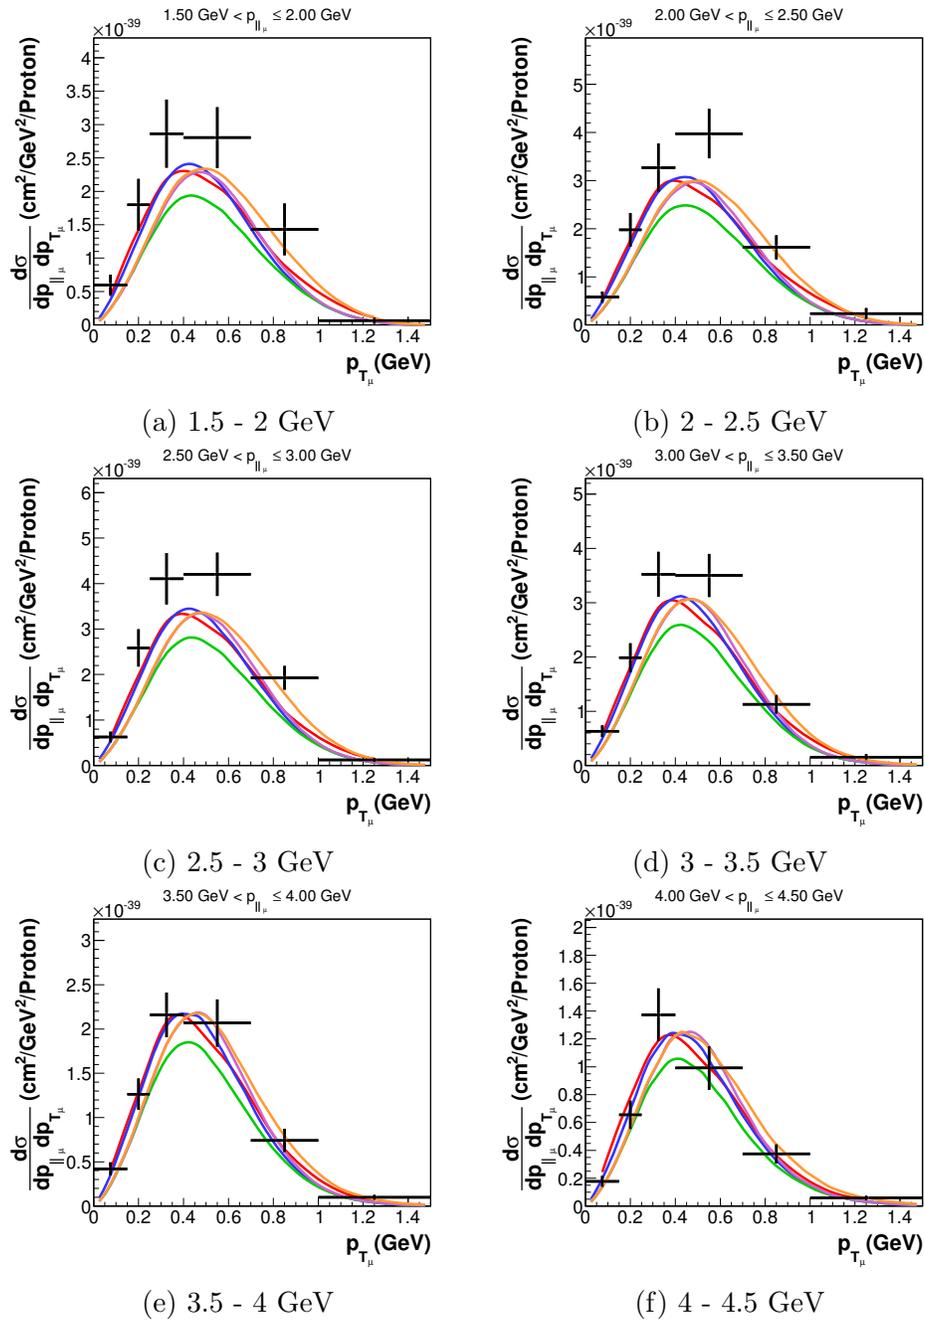


Figure 7.7. Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, with **RPA**, with and without 2p2h effects, to GENIE (continued in next figure)

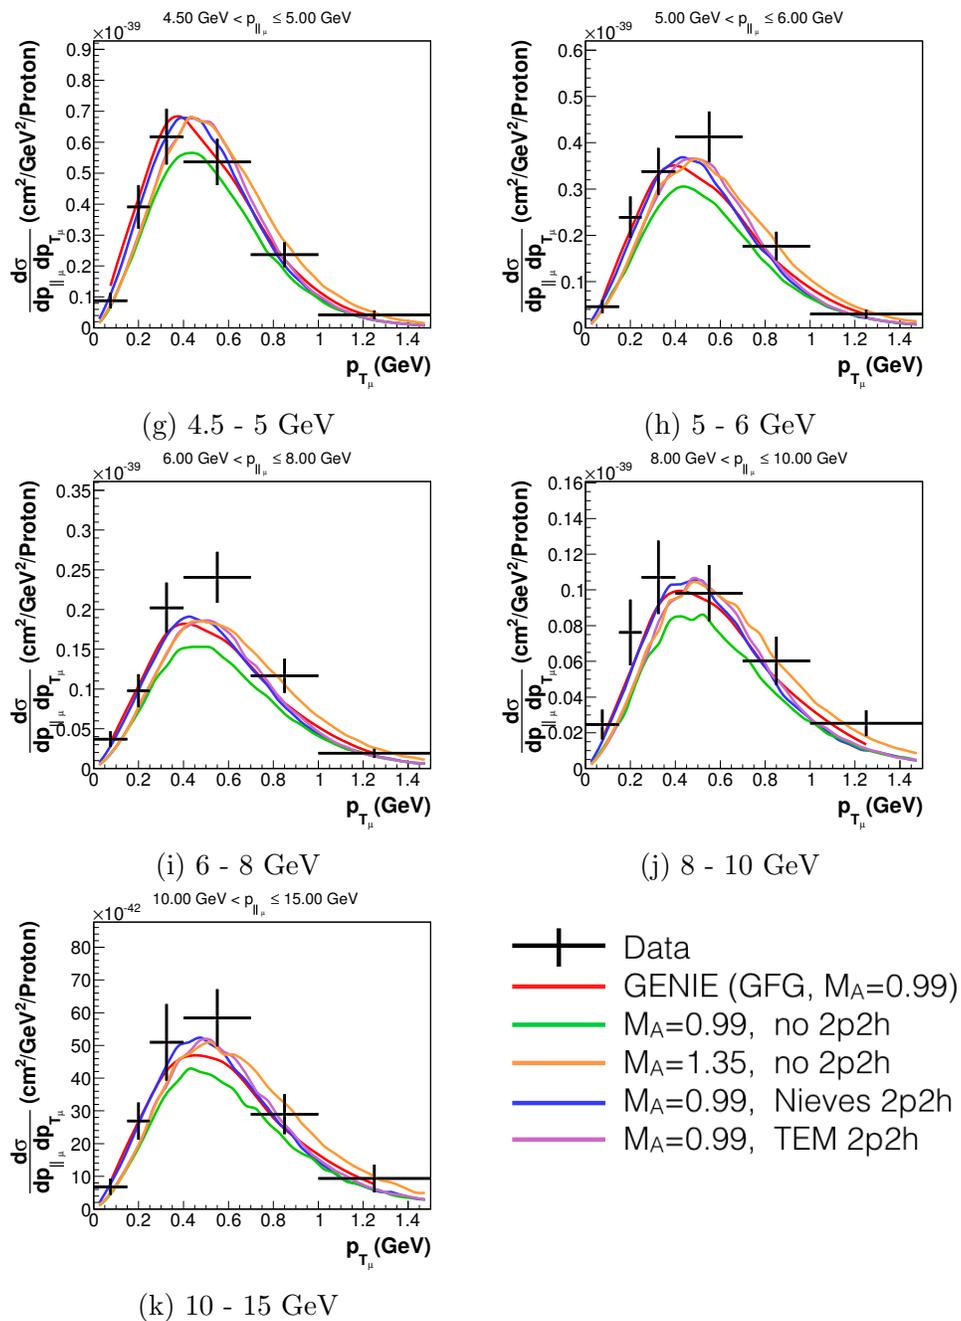


Figure 7.7. Comparison of measured quasi-elastic-like cross section with the **Local Relativistic Fermi Gas** model, with **RPA**, with and without 2p2h effects, to GENIE (continued)

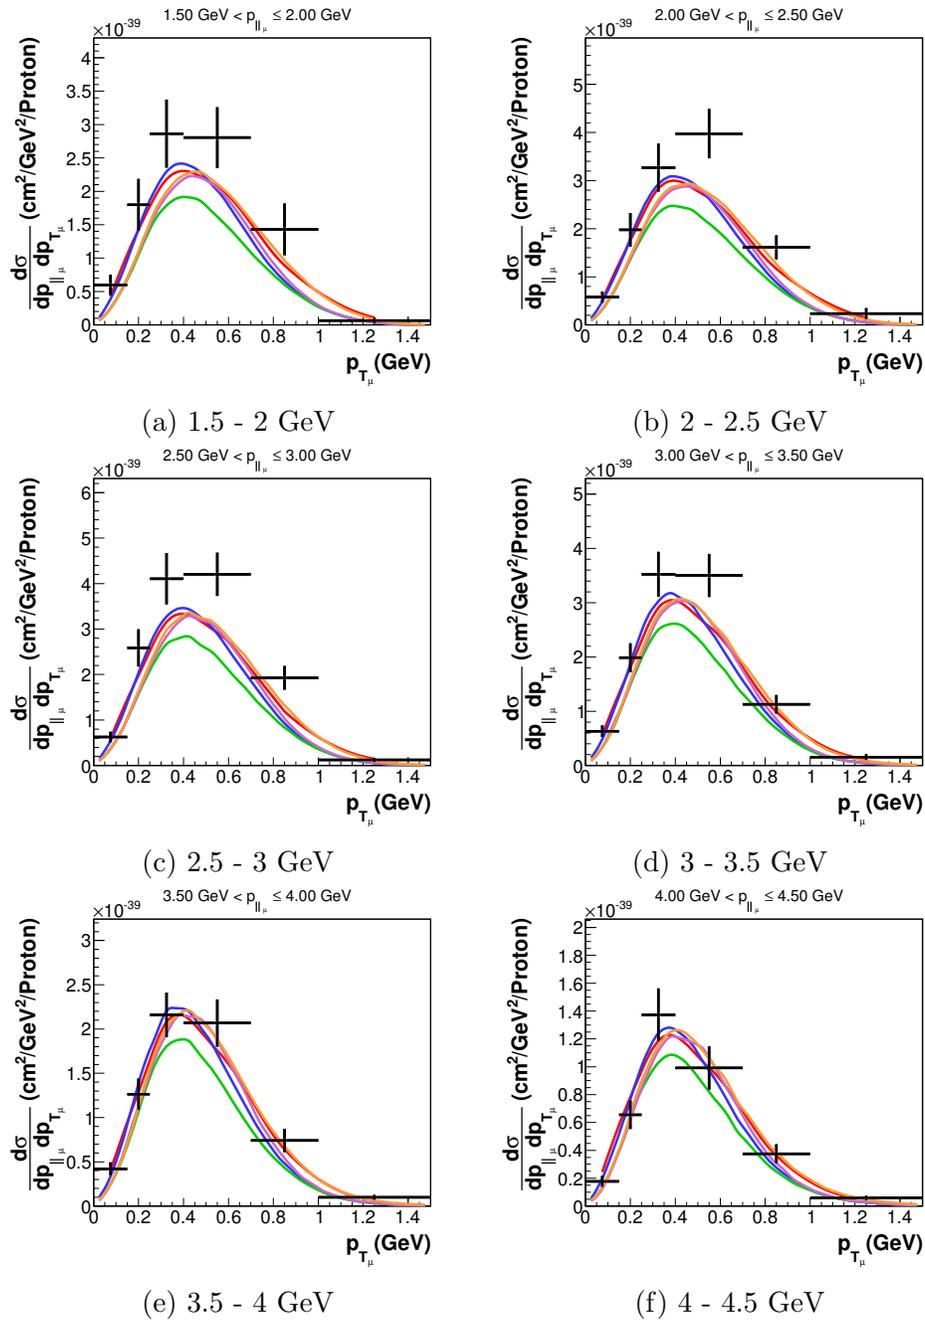


Figure 7.8. Comparison of measured quasi-elastic-like cross section with the **Spectral function** model, with and without 2p2h effects, to GENIE (continued in next figure)

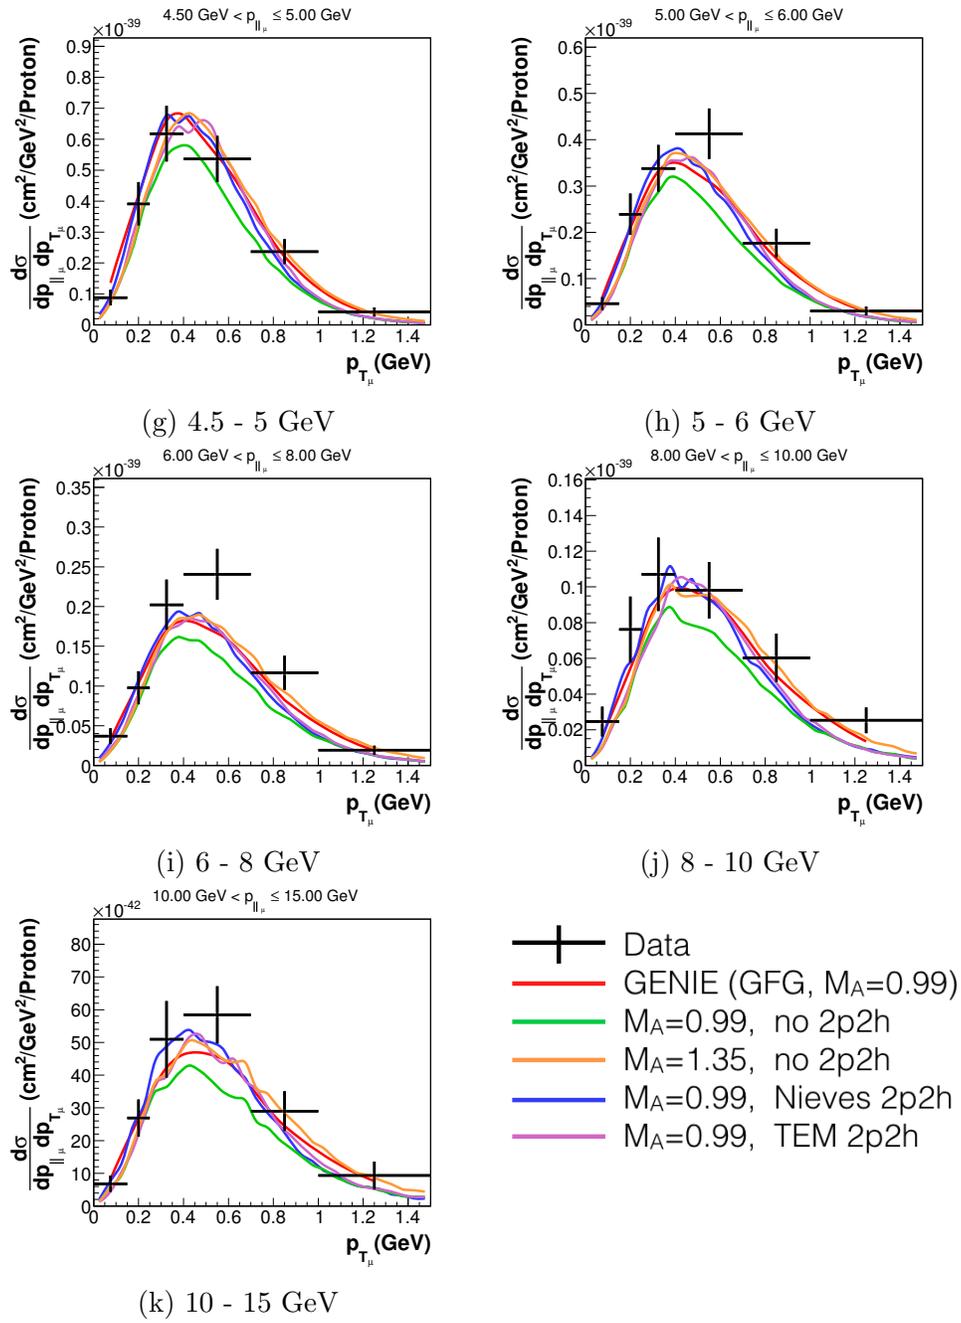


Figure 7.8. Comparison of measured quasi-elastic-like cross section with the spectral function model, with and without 2p2h effects, to GENIE (continued)

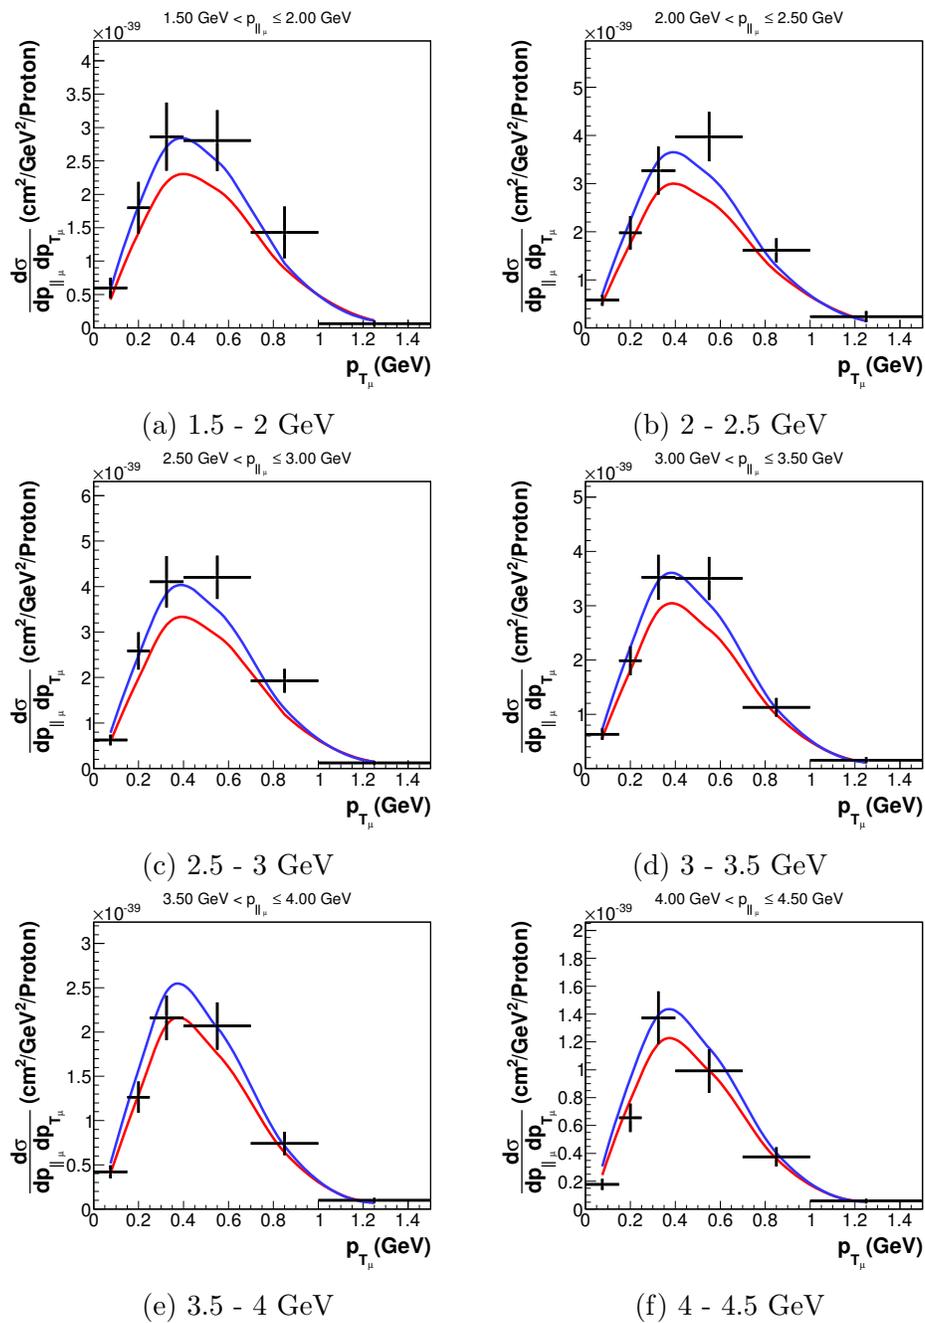


Figure 7.9. Comparison of measured quasi-elastic-like cross section and GE-NIE with 26% 2p2h effect, to GE-NIE with no 2p2h (continued in next figure)

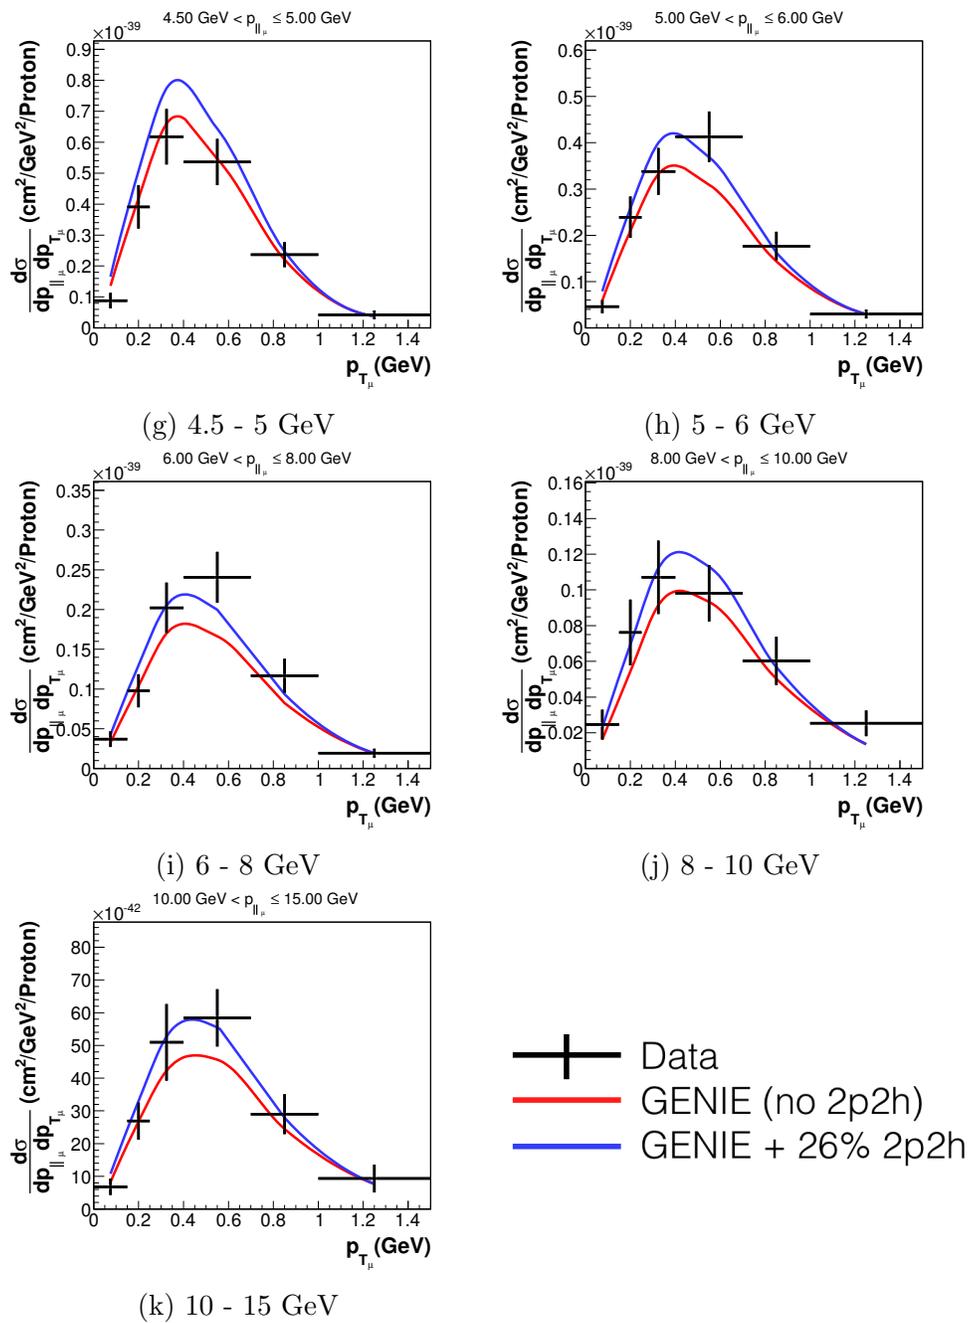


Figure 7.9. Comparison of measured quasi-elastic-like cross section and GENIE with 26% 2p2h effect, to GENIE with no 2p2h (continued)

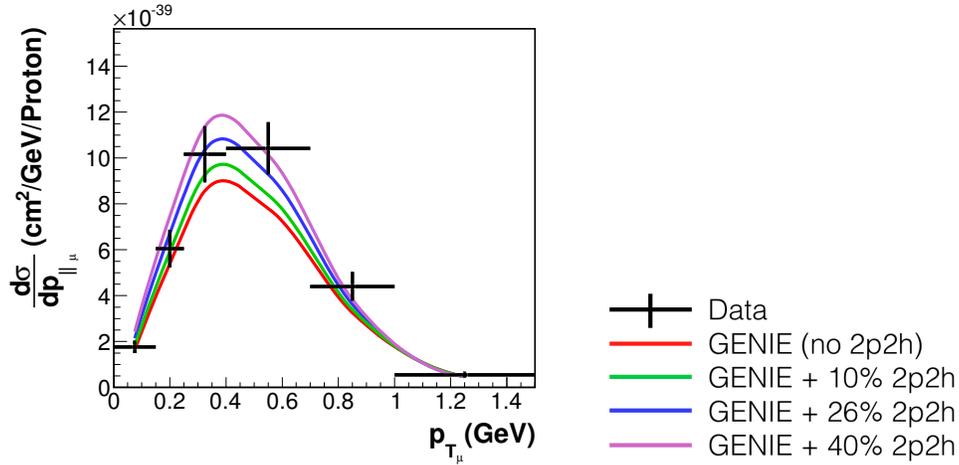


Figure 7.10. Comparison of our measured quasi-elastic-like cross section $d\sigma/dp_T$ with GENIE without 2p2h, and with 10, 26 (the standard amount) and 40% 2p2h per CCQE interaction, simulated by the Nieves MEC model.

A future enhancement could be to fit for the amount of 2p2h that best matches our data. Figure 7.10 shows our data in comparison to three different simulated MEC strengths, generated using GENIE's implementation of the Nieves meson-exchange current model. To perform a fit to these templates, we would need to minimize a goodness of fit quantity, such as the two-dimensional χ^2 . Note that we would not currently be able to include other factors such as RPA and the possibility of switching between the global and local Fermi has models in this fit, as those are not yet implemented in GENIE. A new version of GENIE, to be released later in 2016, will include these features.

7.2.7. Projection of data and models onto p_T

An interesting summary can be found by projecting the data and model comparisons separately onto muon $p_{||}$ and p_T , to obtain single-differential cross sections. Figure 7.11 shows the projection of all the models to give $d\sigma/dp_T$. An excess of data over simulation

at intermediate p_T is not accounted for by any of the models. However, a corresponding increase in cross section is shown in that same intermediate p_T range by the transverse enhancement model (TEM). While that model predicts a consistently smaller overall cross section than our data, we should remember that this model is an enhancement to the NuWro Fermi gas models, which predict a smaller cross section than GENIE due to the difference in FSI modeling. It is therefore possible that GENIE with the TEM enhancement would better reproduce our data distribution. This model is to be added to a new release of GENIE.

7.2.8. Projection of data and models onto p_{\parallel}

We can also project our data onto muon p_{\parallel} . Here, we see that the models show far less variation in shape, with most of the differences between them coming from an overall normalization of the cross section. None of the models replicate our shape in muon p_{\parallel} ; neither do any of them match the normalization.

7.3. Comparison of true CCQE distributions with models

A comparison of the true CCQE (see section 4.1.2) cross section measurements to models allows us to look at distributions that do not depend on the final-state interaction models. (Our standard quasi-elastic-like definition, as explained in section 4.1.3, includes contributions from resonant interactions that include FSI.) This allows us to distinguish whether differences between data and models for the quasi-elastic-like distributions are more likely to be due to FSI or resonant cross section modeling, or whether they are more likely to be due to other effects such as correlations.

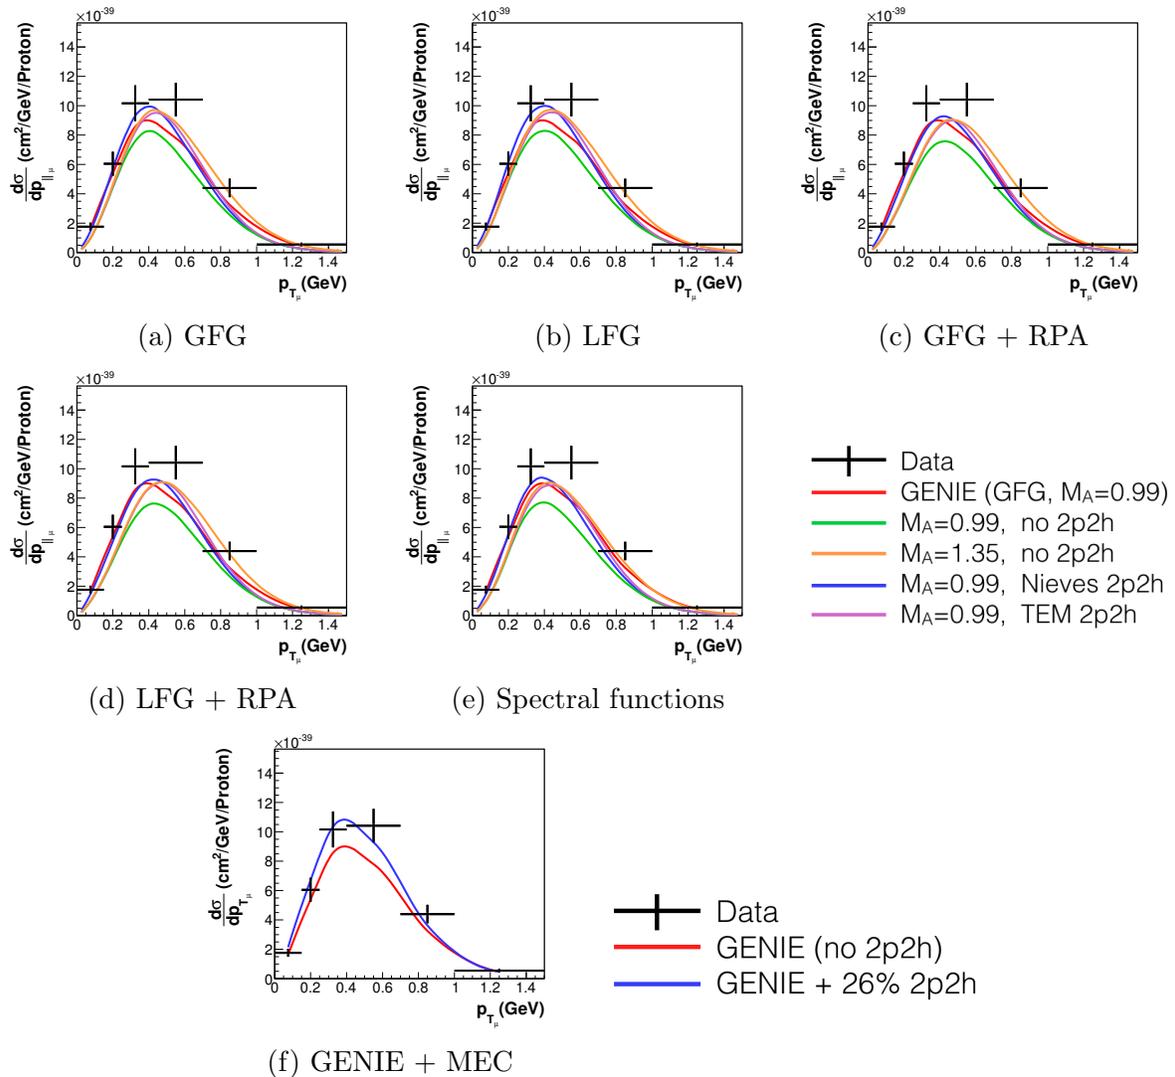


Figure 7.11. Comparison of quasi-elastic-like cross section $d\sigma/dp_T$ to GENIE for data and models. a) Global Fermi Gas without RPA; b) Local Fermi Gas without RPA; c) Global Fermi Gas with RPA; d) Local Fermi Gas with RPA; e) Spectral Functions; f) GENIE with and without MEC

Figure 7.13 shows a comparison of our measurements of the single-differential cross sections $d\sigma/dp_T$ and $d\sigma/dp_{||}$, for both quasi-elastic-like and true CCQE signal definitions, with the GENIE prediction without 2p2h effects, and with the corresponding NuWro

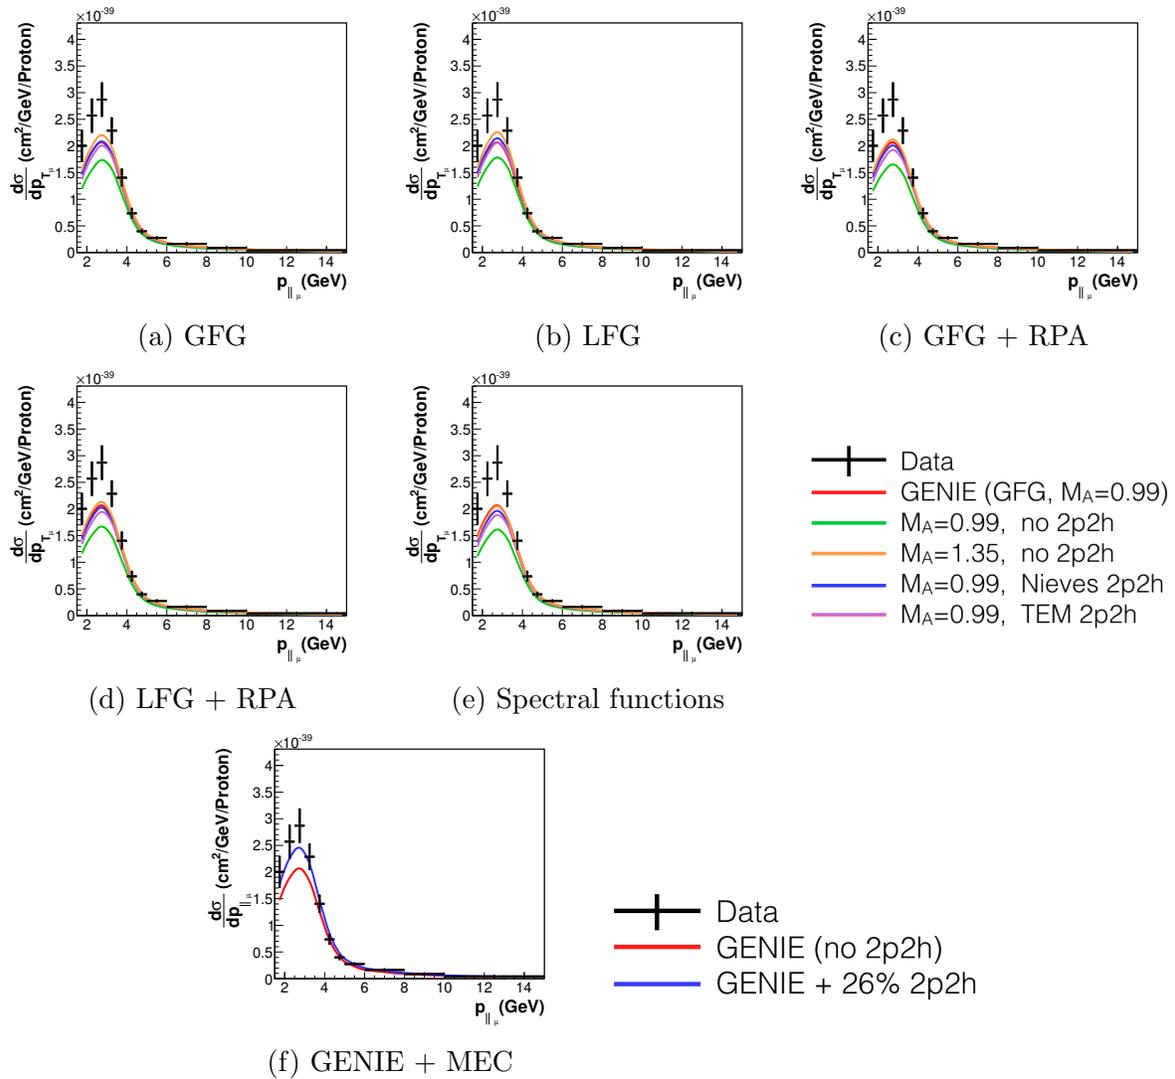


Figure 7.12. Comparison of quasi-elastic-like cross section $d\sigma/dp_{||}$ to GENIE for data and models. a) Global Fermi Gas without RPA; b) Local Fermi Gas without RPA; c) Global Fermi Gas with RPA; d) Local Fermi Gas with RPA; e) Spectral Functions f) GENIE with and without MEC

prediction. In each case, the difference between GENIE and NuWro is more pronounced when predicting a quasi-elastic-like cross section than when predicting a CCQE cross section. This suggests that the discrepancies between the models are mostly due to the

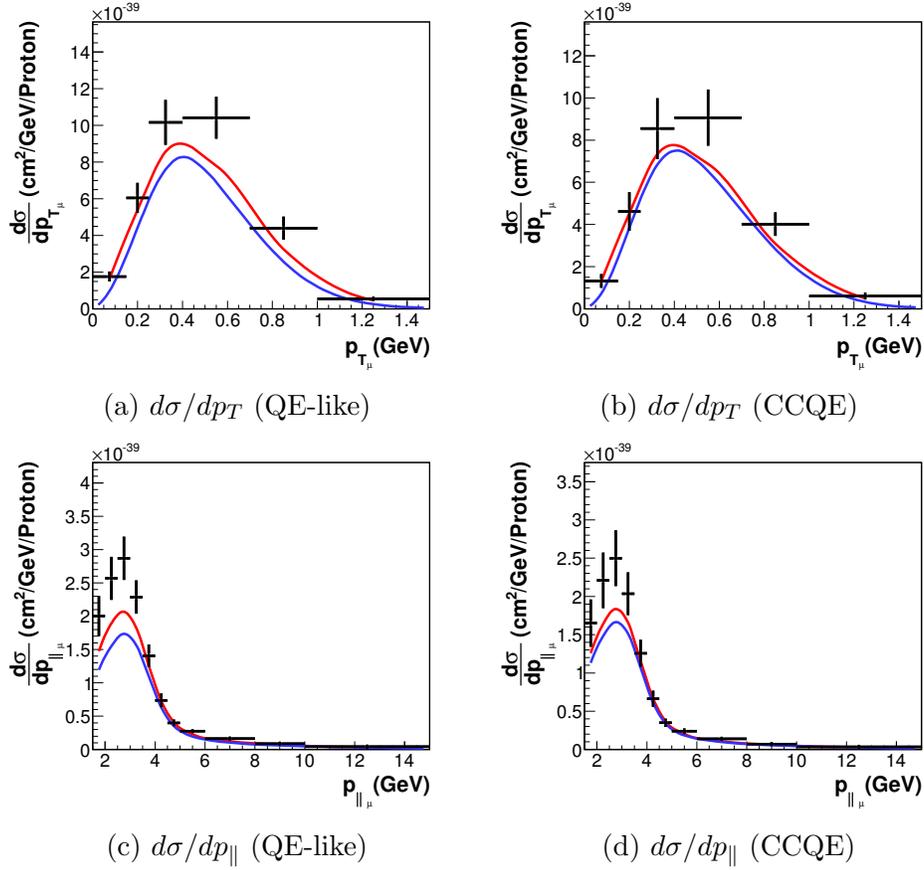


Figure 7.13. Comparison of $d\sigma/dp_T$ and $d\sigma/dp_{||}$ measurements with GENIE's (red) and NuWro's (blue) implementations of the Global Fermi Gas without RPA, for both the QE-like and true CCQE signal definitions

FSI modeling, or to the resonant cross section. This is as expected, as both generators ostensibly use the same parameters to model the CCQE cross section, whereas they use different methods to simulate final-state interactions.

A probably explanation for this is in the differing treatments of final-state interactions between NuWro and GENIE. GENIE's FSI simulates more protons in the final state than does NuWro's; in order to conserve energy, this necessarily means that GENIE's FSI protons will tend to have lower energies than NuWro's. This analysis defines the signal

to include events with any number of protons with kinetic energy below 120 MeV; it is therefore particularly sensitive to the energy distribution of final-state protons. As NuWro's protons generated by FSI tend to have higher energies, resonant production followed by FSI is more likely to generate background events (where one or more protons has kinetic energy above 120 MeV) in NuWro than in GENIE. [120] GENIE's FSI model produces quasi-elastic-like distributions that agree better with our data measurement than the distributions produced by NuWro.

Figure 7.14 shows a comparison of our CCQE measurement with GENIE's prediction both with and without the Nieves 2p2h effect. As with the QE-like measurement shown in figure 7.9, the data appear to agree better with the prediction that includes 2p2h effects.

7.4. Conclusions from model comparisons

To summarize, we can draw the following conclusions from comparing our measured quasi-elastic-like and true CCQE cross sections with models:

Our data prefer GENIE's global Fermi gas prediction with 2p2h over the prediction without 2p2h. As shown in 7.9, the quasi-elastic-like prediction with 2p2h is a better match for our quasi-elastic data than the prediction with no 2p2h. This is also borne out by the true CCQE predictions in 7.14.

Our data prefer GENIE's final-state interaction model to NuWro's. Our quasi-elastic-like signal definition, which allows final states containing the positive muon characteristic of a charged-current antineutrino interaction, all neutrons, and protons with kinetic energy below a 120 MeV threshold, is very sensitive to final-state interaction modeling. While GENIE and NuWro's CCQE models are very similar, they differ in their

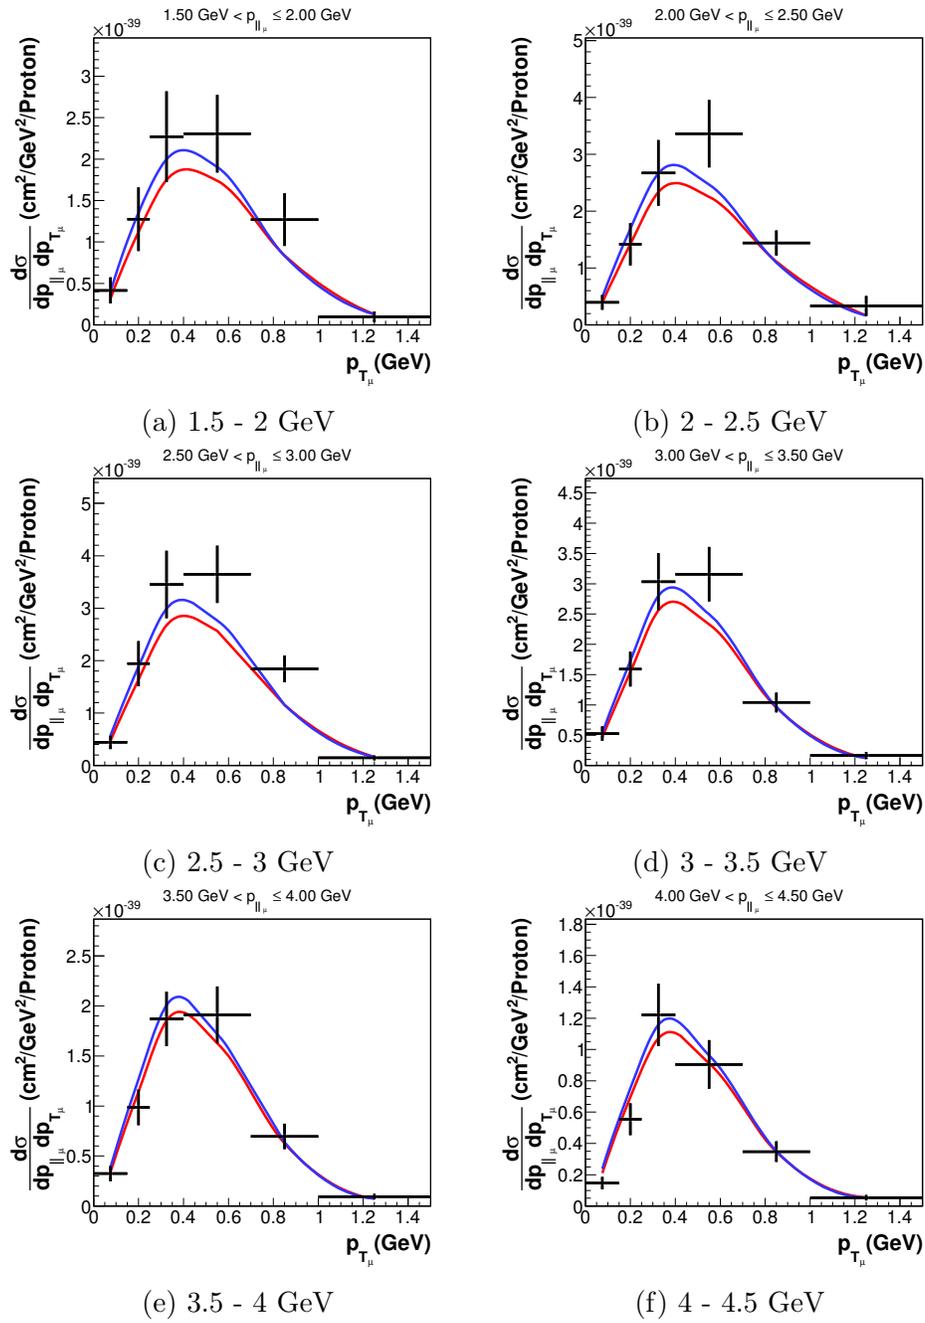


Figure 7.14. Comparison of measured true-CCQE cross section and GENIE with 26% 2p2h effect (blue), to GENIE with no 2p2h (red) (continued in next figure)

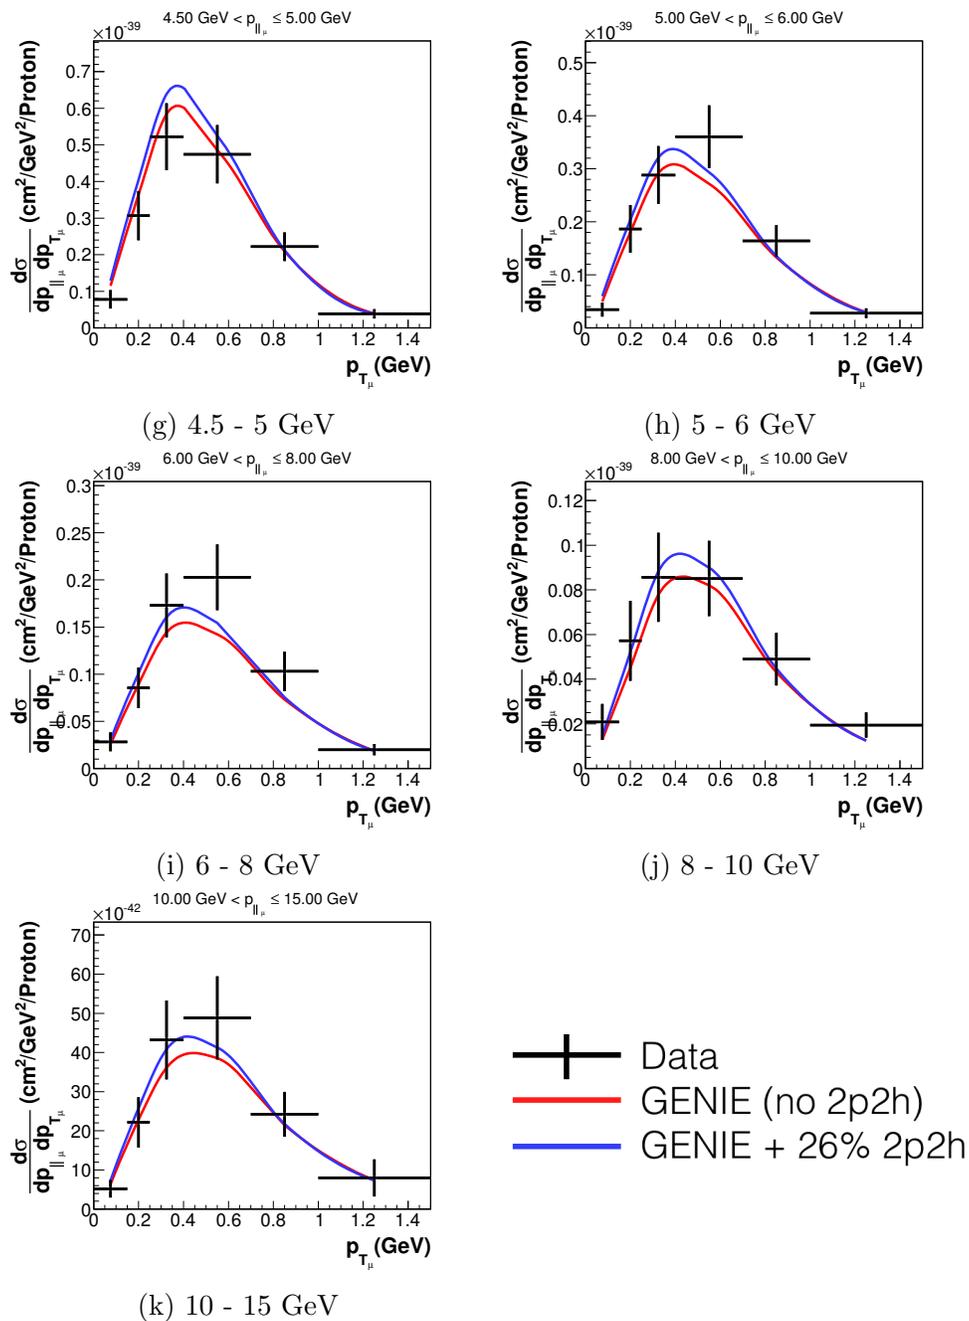


Figure 7.14. Comparison of measured true-CCQE cross section and GENIE with 26% 2p2h effect (blue), to GENIE with no 2p2h (red) (continued)

treatment of final states, causing more significant differences in their quasi-elastic-like predictions, as shown in 7.13. For the global Fermi gas model, our data agree better with GENIE’s model than NuWro’s. The enhancements generated by changing NuWro’s nuclear model to a local Fermi gas, and by adding 2p2h effects to NuWro’s prediction, as shown in 7.5, improve the agreement with data, but still under-predict for low p_{\parallel} .

More GENIE models will allow for better testing. Currently, 2p2h effects are simulated using the Nieves model. We do not yet have access to GENIE simulations for the TEM modeling of 2p2h, for the local Fermi gas or spectral function nuclear models, or for the RPA effect. A later version of GENIE will include these effects, allowing for further comparisons with the data. The reduced cross section for all NuWro distributions due to NuWro’s different FSI model makes it hard to compare the different 2p2h models in NuWro to the data.

A χ^2 comparison can provide more answers. To make a complete comparison, one can calculate a full two-dimensional χ^2 fit between the data and the different models, taking into account the correlations between bins. The data are in place and await a full suite of theoretical predictions for true-CCQE and QE-like signatures. These are currently being generated.

7.5. Comparison with previous and external results

7.5.1. Previous MINERvA CCQE antineutrino scattering results

In 2013, MINERvA published antineutrino scattering cross sections $d\sigma/dQ_{QE}^2$. [96] This measurement differed from the primary result of this analysis, in that it attempted to

measure the cross section for true quasi-elastic scattering, while this study looks at quasi-elastic-like scattering that includes a component generated by pion production followed by final-state interactions. The 2013 analysis was less complicated by the identification of final states with multiple nucleons, and its measurement was less dependent on final-state interaction modeling. However it was particularly dependent on GENIE's CCQE model for acceptance correction, a model that we know to be incomplete.

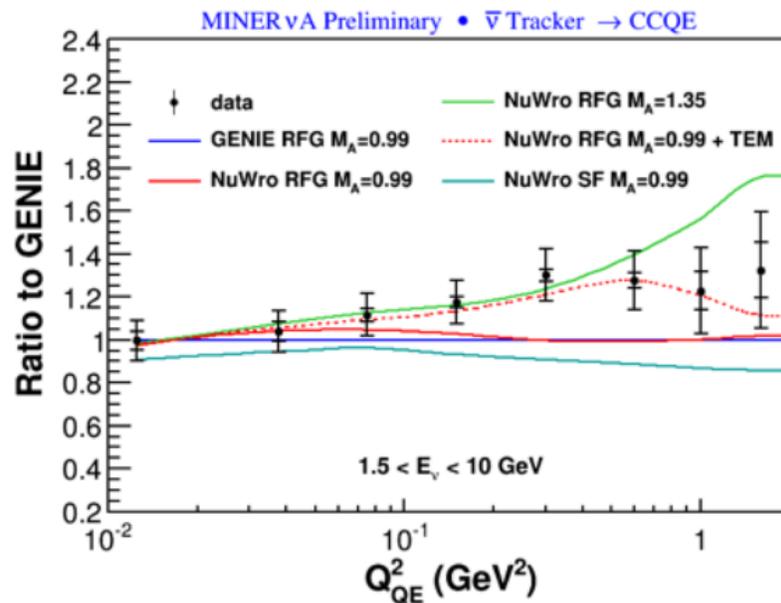
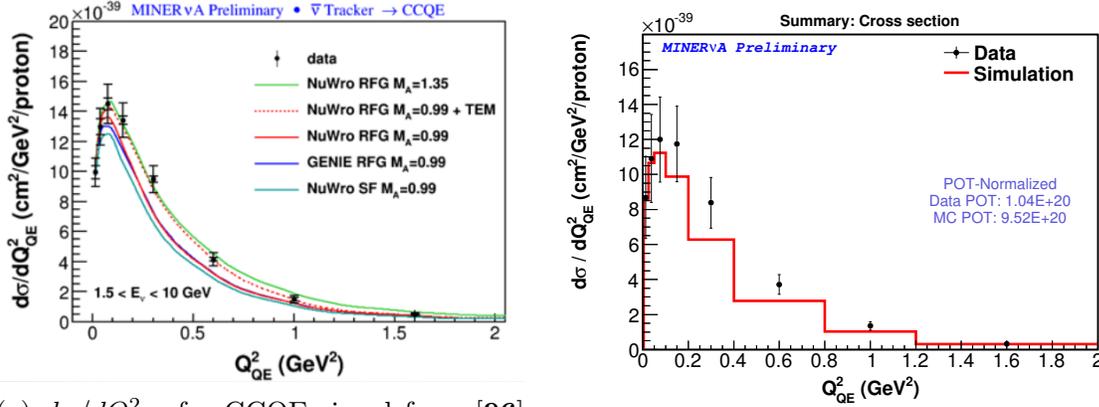


Figure 7.15. Data / model comparison of the 2013 $d\sigma/dQ_{QE}^2$ measurement, taken as a ratio to GENIE, from [96], but scaled to use the updated PPFX1 NuMI flux [36], by the method of [105]

Figure 7.15 shows the ratio of the cross section measurement from the 2013 paper to GENIE, along with several models from NuWro. Note that, for the 2013 analysis, there was much closer agreement between GENIE and NuWro than in the analysis discussed in this thesis. That supports the theory that the difference between the two generators comes in for the events that are quasi-elastic-like, but that are not quasi-elastic; in the



(a) $d\sigma/dQ_{QE}^2$ for CCQE signal from [96] but scaled to the updated PPFX1 NuMI flux [36], by the method of [105]

(b) $d\sigma/dQ_{QE}^2$ for CCQE signal, from this analysis

Figure 7.16. Comparison between 2013 CCQE $d\sigma/dQ_{QE}^2$ for a true CCQE signal, and the current measurement of $d\sigma/dQ_{QE}^2$ for a true CCQE signal. The two measurements use a different integrated flux.

2013 analysis those would have been considered background, and were not included in the model distributions shown here.

The measurement in [96] was taken in 2013; since then, there have been significant advances in the prediction of the NuMI flux, thanks to the PPFX flux simulation [36]. The cross section measurement from the 2013 analysis was re-weighted to the new flux prediction using the method described in [105], to give an estimate of the cross section using the PPFX flux. This method also recalculates the covariance matrix.

Figure 7.16 shows a comparison of the 2013 measurement, scaled to the improved PPFX flux measurement [36] using the method of [105] (figure 7.16a) and this measurement (figure 7.16b) of $d\sigma/dQ_{QE}^2$, projected from the double-differential cross section. Both cross sections are flux-averaged.

There are some differences between the analyses; the 2013 analysis looked at a flux range of $1.5 < E_\nu < 10 \text{ GeV}$, and introduced this neutrino energy range limit on its

signal definition and cuts; we place no limit on neutrino energy. Our full integrated flux is $2.36008 \times 10^{-8} \text{cm}^{-2}$ per POT; if we limit it to 1.5-10 GeV, it would be decreased to $2.02829 \times 10^{-8} \text{cm}^{-2}$, a 14% decrease. Our increased flux corresponds to a lower cross section (inversely proportional to the integrated flux) for the current 2016 measurement (figure 7.16b) compared to the updated 2013 measurement (figure 7.16a). There are also minor differences in the event selection and reconstruction: we have a maximum p_{\parallel} limit of 15 GeV, and have an updated tracking efficiency and mass model, and a correction for a muon angle bias.

7.5.2. MiniBooNE double-differential cross section

The neutrino oscillation experiment MiniBooNE has also measured an antineutrino double-differential scattering cross section [30]. There are several differences between MiniBooNE's measurement and ours. MiniBooNE has a lower beam energy [28] (with a mean energy of 665 MeV, while ours peaks around 3 GeV). They also use a very different detector technology; the MiniBooNE detector [77] is a spherical mineral oil Cherenkov detector, with no limit on the angular acceptance. The mineral oil target, while also a hydrocarbon, has a higher fraction of hydrogen than MINERvA's polystyrene (mineral oil has approximately 2 hydrogen atoms per carbon atom ($C_n H_{2n+2}$ for $n \approx 20$); polystyrene has approximately 1).

While MiniBooNE's detector is good at identifying muons, it has poor sensitivity to nucleons. Thus it identifies quasi-elastic antineutrino scattering events by looking for a muon and no pions in the final state. Any number of nucleons are allowed, as they are undetectable.

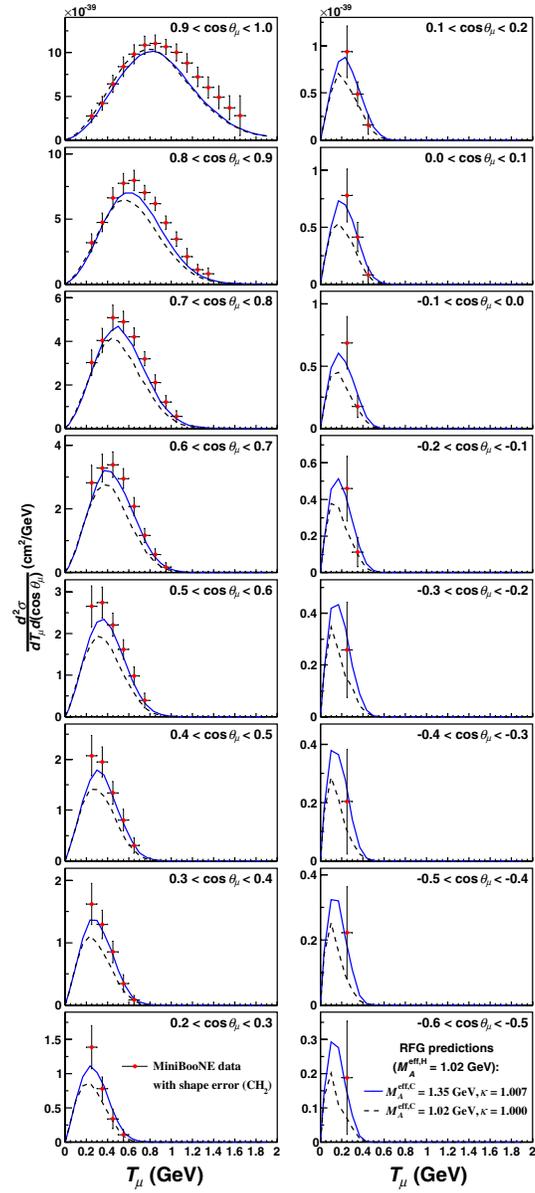


Figure 7.17. Double differential $\bar{\nu}_\mu$ scattering cross section from MiniBooNE[30]

As MiniBooNE benefits from angular acceptance in all directions, they choose to make their double-differential cross section measurement in the muon angle θ_μ with respect to the beam, and its kinetic energy T_μ . As shown in figure 7.17, the measured cross section

is a poor fit to the standard global Fermi gas model with $M_A = 0.99$ GeV. Instead, it provides a better fit to a version with $M_A = 1.35$ GeV, the best-fit value from their neutrino-mode quasi-elastic analysis [29].

7.5.3. NOMAD antineutrino quasi-elastic cross section

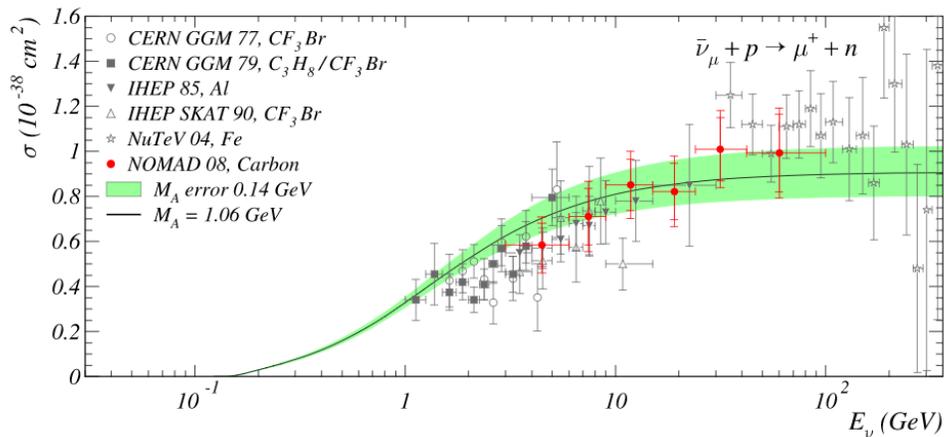


Figure 7.18. NOMAD’s antineutrino-carbon CCQE cross section measurements (red) compared with those from previous bubble chamber experiments (black). The data were consistent with a global relativistic Fermi gas model, with $M_A = 1.06 \pm 0.14$ GeV. Reprinted from [146]

NOMAD (Neutrino Oscillation Magnetic Detector), a $\nu_\mu \rightarrow \nu_\tau$ search which ran at CERN from 1995-8 [189], used a higher energy beam than either MINERvA or Mini-BooNE, with an average antineutrino energy of 17.6 GeV. The detector technology was also very different, consisting of 44 drift chamber modules, filled with an argon-ethane gas mixture, with the casing (average atomic mass 12.9) serving as a passive target, located in a 0.4 Tesla magnetic field. NOMAD measured an antineutrino CCQE cross section [146]. They did not take final-state interactions into account when choosing their signal definition, but instead made a true CCQE measurement; their event selection was on a

single positive muon (charge-identified thanks to the magnetic field) and no other tracks. A value of the axial mass, M_A , was extracted from the $d\sigma/dQ^2$ distribution shape, and from the total cross section.

NOMAD reported a total cross section $\langle\sigma_{QE}\rangle_{\bar{\nu}} = (0.81 \pm 0.05(\text{stat}) \pm 0.09(\text{sys}) \times 10^{-38} \text{ cm}^2$ on the 3-100 GeV energy range. They also extracted a best-fit value of the axial mass $M_A = 1.06 \pm 0.14 \text{ GeV}$.

7.5.4. Comparison of this analysis' results to NOMAD and MiniBooNE's

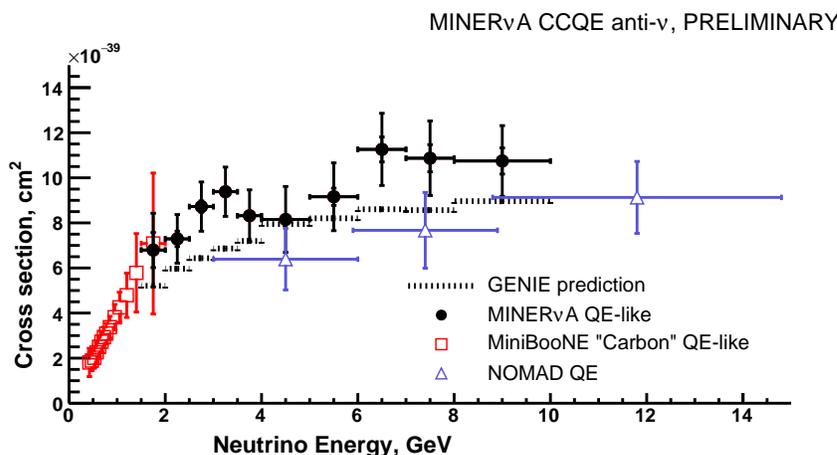


Figure 7.19. Preliminary plot showing the cross section $\sigma(E_\nu)$ from this analysis (black) compared with measurements from MiniBooNE [30] (red squares) and NOMAD [146] (blue triangles) and GENIE's CCQE-like prediction (black dashed line). Our measurements show statistical (inner tick) and total (full error bar) uncertainties.

Figure 7.19 gives an indication of how the MINERvA results compare to those of MiniBooNE, at lower energies, and NOMAD, at higher. (MiniBooNE's results, originally measured on mineral oil, have had the hydrogen component subtracted off to give a cross section on carbon.) The black points show the data from this analysis, unfolded to true

$E\nu$. The plot appears to show agreement with the MiniBooNE distribution, but to show tension with NOMAD's results. Note that, as with this analysis, MiniBooNE was looking at a quasi-elastic-like distribution, while NOMAD studied true CCQE.

CHAPTER 8

Conclusions

We have presented MINERvA's first double-differential cross section for antineutrino scattering, and its first quasi-elastic-like antineutrino scattering cross section. We studied scattering of muon antineutrinos on doped polystyrene scintillator, producing a double-differential flux-integrated cross section in the muon transverse and longitudinal momentum $d^2\sigma/dp_T dp_{\parallel}$. We introduced a novel quasi-elastic-like definition, where our signal was a positive muon in the final state, along with any number of neutrons and any number of protons with less than 120MeV of kinetic energy. We also produced an energy-dependent cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. By projecting these double-differential cross sections, we also generated the flux-integrated single-differential cross sections $d\sigma/dp_T$, $d\sigma/dp_{\parallel}$ and $d\sigma/dQ_{QE}^2$. Additionally, we calculated the total scattering cross section $\sigma(E_{\nu})$ as a function of neutrino energy. A comparison with different generator's models showed better agreement with the GENIE generator than with NuWro. The differences between these generators' predictions were shown to be due to the component of the quasi-elastic-like cross section that came from final-state interactions, rather than due to the quasi-elastic model itself.

An excess of data over GENIE suggests that additional nuclear 2-particle-2-hole effects such as meson exchange currents may be present. Adding a meson exchange current model to GENIE's prediction produced distributions that were a better match to our

measurement. An upcoming version of GENIE will allow us to model several different nuclear effects and extensions to the Fermi Gas model and test these against our data.

Our total cross section results were also compared with those from the MiniBooNE and NOMAD experiments, whose results had previously shown tension with each other. The MINERvA result appears to show greater agreement with MiniBooNE; however, it should be noticed that our signal definition is more similar to MiniBooNE's than it is to NOMAD's, which may affect the result.

A comparable double differential cross section is being calculated at MINERvA for quasi-elastic-like neutrino scattering; it will be interesting to see whether these results show similar patterns when compared to generators' models. It will also be interesting to compare with the models from the next GENIE version. If we are able to use the data to constrain the strength of correlation effects in the quasi-elastic like channel, it will be a great help to oscillation experiments who need to reduce their current levels of uncertainty on interaction models.

In future, we plan to study the difference between true quasi-elastic and quasi-elastic-like signal definitions, which gives us information about final-state interaction effects. MINERvA is also currently collecting data at a higher energy range with a neutrino beam, and plans to collect higher-energy antineutrino data; we also have approximately 9×10^{19} POT of antineutrino data taken with a partial "frozen" MINERvA detector [72], which could be incorporated into a future version of this analysis. An analysis on quasi-elastic-like scattering is underway for MINERVA's passive nuclear targets, which will provide cross sections on lead and iron as well as on carbon. Furthermore, the planned CAPTAIN-MINERvA extension [79], where the CAPTAIN liquid argon time-projection

chamber will be placed upstream of MINERvA, in place of its current helium target. This will be extremely valuable as the DUNE next-generation oscillation experiment will use a liquid-argon detector; CAPTAIN-MINERvA will provide neutrino-argon scattering cross sections in the DUNE energy range to help reduce their uncertainties, and to study how neutrino scattering on argon compares with that on carbon.

MINERvA has a rich program for measurements of quasi-elastic scattering for neutrinos and antineutrinos, with a variety of targets, and over a broad energy range that is very relevant to the current and future neutrino oscillation program. These measurements will be key to improving modeling of neutrino scattering on heavy nuclei, and to understanding how the nucleus affects scattering.

References

- [1] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. III. Results for the 2002~2007 data-taking period. *Phys. Rev. C*, 80:015807, Jul 2009.
- [2] K. Abe et al. Indication of electron neutrino appearance from an accelerator-produced off-axis muon neutrino beam. *Phys. Rev. Lett.*, 107:041801, Jul 2011.
- [3] K. Abe et al. Precise measurement of the neutrino mixing parameter θ_{23} from muon neutrino disappearance in an off-axis beam. *Phys. Rev. Lett.*, 112:181801, May 2014.
- [4] K. Abe et al. Measurement of the ν_μ charged-current quasielastic cross section on carbon with the nd280 detector at t2k. *Phys. Rev. D*, 92:112003, Dec 2015.
- [5] K. Abe et al. Measurement of the ν_μ charged-current quasielastic cross section on carbon with the ND280 detector at T2K. *Phys. Rev. D*, 92:112003, Dec 2015.
- [6] K. Abe et al. Measurement of the ν_μ charged current quasielastic cross section on carbon with the T2K on-axis neutrino beam. *Phys. Rev. D*, 91:112002, Jun 2015.
- [7] K. Abe and others. Measurement of double-differential muon neutrino charged-current interactions on C_8H_8 without pions in the final state using the T2K off-axis beam. arXiv 1602.03652 (physics.hep-ex), 2016.

- [8] Y. Abe et al. Reactor $\bar{\nu}_e$ disappearance in the Double Chooz experiment. *Phys. Rev. D*, 86:052008, Sep 2012.
- [9] Y. Abe et al. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector. *JHEP*, 10:086, 2014. [Erratum: JHEP02,074(2015)].
- [10] N. Abgrall et al. Measurements of cross sections and charged pion spectra in proton-carbon interactions at 31 GeV/c. *Phys. Rev. C*, 84:034604, Sep 2011.
- [11] N. Abgrall et al. Time Projection Chambers for the T2K Near Detectors. *Nucl. Instrum. Meth.*, A637:25–46, 2011.
- [12] E. Ables et al. P-875: A Long baseline neutrino oscillation experiment at Fermilab. Technical report, Fermilab, 1995.
- [13] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF. arXiv 1512.06148 (physics.ins-det), 2015.
- [14] P. Adamson et al. Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector. *Phys. Rev. D*, 81:072002, Apr 2010.
- [15] P. Adamson et al. Search for sterile neutrino mixing in the MINOS long-baseline experiment. *Phys. Rev. D*, 81:052004, Mar 2010.
- [16] P. Adamson et al. First direct observation of muon antineutrino disappearance. *Phys. Rev. Lett.*, 107:021801, Jul 2011.

- [17] P. Adamson et al. Improved Search for Muon-Neutrino to Electron-Neutrino Oscillations in MINOS. *Phys. Rev. Lett.*, 107:181802, Oct 2011.
- [18] P. Adamson et al. Measurement of the neutrino mass splitting and flavor mixing by MINOS. *Phys. Rev. Lett.*, 106:181801, May 2011.
- [19] P. Adamson et al. Electron Neutrino and Antineutrino Appearance in the Full MINOS Data Sample. *Phys. Rev. Lett.*, 110:171801, Apr 2013.
- [20] P. Adamson et al. Measurement of neutrino and antineutrino oscillations using beam and atmospheric data in MINOS. *Phys. Rev. Lett.*, 110:251801, Jun 2013.
- [21] P. Adamson et al. Combined analysis of ν_μ disappearance and $\nu_\mu \rightarrow \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos. *Phys. Rev. Lett.*, 112:191801, May 2014.
- [22] P. Adamson et al. Observation of seasonal variation of atmospheric multiple-muon events in the MINOS Near and Far Detectors. *Phys. Rev. D*, 91:112006, Jun 2015.
- [23] P. Adamson et al. First measurement of electron neutrino appearance in NOvA. arXiv:1601.05022 (physics.hep-ex), 2016.
- [24] P. Adamson et al. First measurement of muon-neutrino disappearance in NOvA. arXiv:1601.05037 (physics.hep-ex), 2016.
- [25] T. Adye. Unfolding algorithms and tests using RooUnfold. *arXiv:1105.1160 [physics.data-an]*, 2011.

- [26] N. Agafonova et al. Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment. *Phys. Rev. Lett.*, 115:121802, Sep 2015.
- [27] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [28] A. A. Aguilar-Arevalo et al. Neutrino flux prediction at MiniBooNE. *Phys. Rev. D*, 79:072002, Apr 2009.
- [29] A. A. Aguilar-Arevalo et al. First measurement of the muon neutrino charged current quasielastic double differential cross section. *Phys. Rev. D*, 81:092005, May 2010.
- [30] A. A. Aguilar-Arevalo et al. First measurement of the muon antineutrino double-differential charged-current quasielastic cross section. *Phys. Rev. D*, 88:032001, Aug 2013.
- [31] B. Aharmim et al. Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory. *Phys. Rev. C*, 88:025501, Aug 2013.
- [32] J. K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, 2012.
- [33] J. K. Ahn et al. Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, May 2012.
- [34] L. A. Ahrens et al. Measurement of Neutrino - Proton and Anti-neutrino - Proton Elastic Scattering. *Phys. Rev.*, D35:785, 1987.

- [35] L.A. Ahrens et al. A study of the axial-vector form factor and second-class currents in antineutrino quasielastic scattering. *Physics Letters B*, 202(2):284 – 288, 1988.
- [36] L. Aliaga. *Neutrino flux prediction for the NuMI beamline*. PhD thesis, College of William and Mary, 2016. In preparation.
- [37] L Aliaga et al. Design, Calibration, and Performance of the MINERvA Detector. *Nucl. Inst. and Meth.*, 130:28–46, 2014.
- [38] L. Aliaga, L. Fields, M. Kiveni, M. Kordosky, and A. Norrick. TN004: A brief documentation of the flux, 2015. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7634>.
- [39] C. Alt et al. Inclusive production of charged pions in p + C collisions at 158 GeV/c beam momentum. *The European Physical Journal C*, 49(4):897–917, 2007.
- [40] J. E. Amaro et al. Meson-exchange currents and quasielastic neutrino cross sections in the SuperScaling Approximation model. *Phys. Lett.*, B696:151–155, 2014.
- [41] J.E. Amaro, C. Maieron, J. Nieves, and M. Valverde. Equivalence between local Fermi gas and shell models in inclusive muon capture from nuclei. *The European Physical Journal A - Hadrons and Nuclei*, 24(3):343–353, 2005.
- [42] F. P. An et al. Improved measurement of electron antineutrino disappearance at Daya Bay. *Chinese Physics C*, 37(1):011001, 2013.

- [43] F. P. An et al. New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay. *Phys. Rev. Lett.*, 115(11):111802, 2015.
- [44] C. Anderson et al. The ArgoNeuT Detector in the NuMI Low-Energy beam line at Fermilab. *JINST*, 7:P10019, 2012.
- [45] R. Arnold et al. Probing New Physics Models of Neutrinoless Double Beta Decay with SuperNEMO. *Eur. Phys. J.*, C70:927–943, 2010.
- [46] J. Arrington, D.W. Higinbotham, G. Rosner, and M. Sargsian. Hard probes of short-range nucleon-nucleon correlations. *Prog.Part.Nucl.Phys.*, 67:898–938, 2012.
- [47] D. S. Ayres et al. The NOvA Technical Design Report. Technical report, Fermi National Accelerator Laboratory, 2007.
- [48] L. Bagby, H. Budd, M. Bustamante Rosell, and D. Ruggiero. Alcove 4 Muon Monitor ORC, 2014. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=8917>.
- [49] D. S. Baranov et al. An estimate for the formation length of hadrons in neutrino interactions. *PHE*, 4, 1984.
- [50] S. J. Barish et al. Study of neutrino interactions in hydrogen and deuterium: Description of the experiment and study of the reaction $\nu + d \rightarrow \mu^- + p + p_s$. *Phys. Rev. D*, 16:3103–3121, Dec 1977.

- [51] G. Barrand et al. GAUDI - The software architecture and framework for building LHCb data processing applications. In *Proceedings, 11th International Conference on Computing in High-Energy and Nuclear Physics (CHEP 2000)*, pages 92–95, 2000.
- [52] P. Barreau et al. Coulomb sum rule on ^{12}C . *Il Nuovo Cimento A*, 76:361–368, 1983.
- [53] D. S. Barton et al. Experimental study of the a dependence of inclusive hadron fragmentation. *Phys. Rev. D*, 27:2580–2599, Jun 1983.
- [54] G. Bellini et al. Precision Measurement of the ^7Be Solar Neutrino Interaction Rate in Borexino. *Phys. Rev. Lett.*, 107:141302, Sep 2011.
- [55] O. Benhar, A. Fabrocini, S. Fantoni, and I. Sick. Spectral function of finite nuclei and scattering of GeV electrons. *Nucl. Phys.*, A579:493–517, 1994.
- [56] M. Betancourt. MINOS Tracking Efficiencies for ME Energy, 2015. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11642>.
- [57] A Bodek, S Avvakumov, R Bradford, and H Budd. Extraction of the axial nucleon form factor from neutrino experiments on deuterium. *Journal of Physics: Conference Series*, 110(8):082004, 2008.
- [58] A. Bodek, S. Avvakumov, R. Bradford, and H. Budd. Vector and axial nucleon form factors: A duality constrained parameterization. *The European Physical Journal C*, 53(3):349–354, 2008.

- [59] A. Bodek and J. L. Ritchie. Fermi-motion effects in deep-inelastic lepton scattering from nuclear targets. *Phys. Rev. D*, 23:1070–1091, Mar 1981.
- [60] Arie Bodek, Inkyu Park, and Un-ki Yang. Improved low Q^2 model for neutrino and electron nucleon cross sections in few GeV region. *Nuclear Physics B - Proceedings Supplements*, 139:113 – 118, 2005. Proceedings of the Third International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region.
- [61] Bodek, A., Budd, H.S., and Christy, M.E. Neutrino quasielastic scattering on nuclear targets. *Eur. Phys. J. C*, 71(9):1726, 2011.
- [62] T.T. Böhlen, F. Cerutti, M.P.W. Chin, A. Fass, A. Ferrari, P.G. Ortega, A. Mairani, P.R. Sala, G. Smirnov, and V. Vlachoudis. The {FLUKA} code: Developments and challenges for high energy and medical applications. *Nuclear Data Sheets*, 120:211 – 214, 2014.
- [63] R. Bradford, A. Bodek, H. Budd, and J. Arrington. A New Parameterization of the Nucleon Elastic Form Factors. *Nuclear Physics B - Proceedings Supplements*, 159:127–132, September 2006.
- [64] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth.*, A389:81–86, 1997.
- [65] H. Budd, A. Bodek, and J. Arrington. Modeling Quasi-elastic Form Factors for Electron and Neutrino Scattering. *arxiv:hep-ex/0308005*, page 13, August 2003.

- [66] A. V. Butkevich and S. P. Mikheyev. Test of fermi gas model and plane-wave impulse approximation against electron-nucleus scattering data. *Phys. Rev. C*, 72:025501, Aug 2005.
- [67] Nicola Cabibbo. Unitary symmetry and leptonic decays. *Phys. Rev. Lett.*, 10:531–533, Jun 1963.
- [68] J. Carlson, J. Jourdan, R. Schiavilla, and I. Sick. Longitudinal and transverse quasielastic response functions of light nuclei. *Phys. Rev. C*, 65:024002, Jan 2002.
- [69] D. Casper. The nuance neutrino physics simulation, and the future. *Nuclear Physics B - Proceedings Supplements*, 112(1?3):161 – 170, 2002.
- [70] R. Cenni, T. W. Donnelly, and A. Molinari. Relativistic electromagnetic charge response: Finite versus infinite systems, exclusive versus inclusive processes. *Phys. Rev. C*, 56:276–291, Jul 1997.
- [71] Seonho Choi et al. Axial and pseudoscalar nucleon form factors from low energy pion electroproduction. *Phys. Rev. Lett.*, 71:3927–3930, Dec 1993.
- [72] Jesse John Chvojka. *Anti-Neutrino Charged Current Quasi-Elastic Scattering in MINERνA*. PhD thesis, Rochester U., 2012.
- [73] C. Ciofi degli Atti and S. Simula. Realistic model of the nucleon spectral function in few- and many-nucleon systems. *Phys. Rev. C*, 53:1689–1710, Apr 1996.

- [74] Bruce T. Cleveland et al. Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector. *The Astrophysical Journal*, 496(1):505, 1998.
- [75] Frank T. Cole, Edwin L. Goldwasser, and Robert Rathbun Wilson. National Accelerator Laboratory design report January 1968. Technical report, Fermilab, 1968.
- [76] MINERvA Collaboration. Minerva neutrino detector response measured with test beam data. *arXiv:1501.06431 [physics.ins-det]*, April 2015.
- [77] MiniBooNE Collaboration. The MiniBooNE Detector. *Nucl.Instrum.Meth.*, A599:28–46, 2009.
- [78] T2K Collaboration. The T2K Neutrino Flux Prediction. *Phys. Rev. D*, 87:28–46, 2013.
- [79] The MINERvA Collaboration. CAPTAIN-MINERvA Proposal, 2015. FERMILAB-PROPOSAL-1061 available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11031>.
- [80] E. D. Commins and P.H. Bucksbaum. *Weak Interactions of Leptons and Quarks*. Cambridge University Press, 1983.
- [81] G. D’Agostini. Improved iterative Bayesian unfolding. *arXiv:1010.0632 [physics.data-an]*, 2010.

- [82] D B Day, J S McCarthy, T W Donnelly, and I Sick. Scaling in inclusive electron-nucleus scattering. *Annual Review of Nuclear and Particle Science*, 40(1):357–410, 1990.
- [83] H. De Vries, C. W. De Jager, and C. De Vries. Nuclear charge and magnetization density distribution parameters from elastic electron scattering. *Atom. Data Nucl. Data Tabl.*, 36:495–536, 1987.
- [84] J. Devan. TN051: Test beam proton calorimetry note. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=9986>.
- [85] J. Devan. Low Nu Muon Angle Approved Plots, 2013. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=9471>.
- [86] Nicholas Edward Devenish. *Muon Antineutrino Disappearance in the MINOS Experiment*. PhD thesis, Sussex U., 2015.
- [87] National Diagnostics. Mechanism of liquid scintillation counting. <https://www.nationaldiagnostics.com/liquid-scintillation/article/mechanism-liquid-scintillation-counting?osCsid=0f352a9e0b698613ea3618c7e086941a>.
- [88] M. Dorman and MINOS Collaboration. Preliminary results for ccqe scattering with the minos near detector. *AIP Conference Proceedings*, 1189(1):133–138, 2009.
- [89] S. Dytman. INTRANUKE-hA. In *AIP Conf Proc*, volume 896, pages 178–184, 2007.

- [90] B. Eberly et al. Charged pion production in ν_μ interactions on hydrocarbon at $\langle E_\nu \rangle = 4.0$ GeV. *Phys. Rev. D*, 92:092008, Nov 2015.
- [91] F. J. Ernst, R. G. Sachs, and K. C. Wali. Electromagnetic form factors of the nucleon. *Phys. Rev.*, 119:1105–1114, Aug 1960.
- [92] X. Espinal and F. Sanchez. Measurement of the axial vector mass in neutrino-carbon interactions at K2K. *AIP Conf. Proc.*, 967:117–122, 2007.
- [93] C. Andreopoulos et al. The GENIE neutrino Monte Carlo generator. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 614(1):87 – 104, 2010.
- [94] D.G. Michael et al. The magnetized steel and scintillator calorimeters of the MINOS experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 596(2):190 – 228, 2008.
- [95] K.S. Egiyan et al. Measurement of two- and three-nucleon short-range correlation probabilities in nuclei. *Phys. Rev. Lett.*, 96:082501, Mar 2006.
- [96] L. Fields et al. Measurement of Muon Antineutrino Quasielastic Scattering on a Hydrocarbon Target at $E_\nu \sim 3.5$ GeV. *Physical Review Letters*, 111(2):022501, July 2013.
- [97] N. Fomin et al. New measurements of high-momentum nucleons and short-range structures in nuclei. *Phys. Rev. Lett.*, 108:092502, Feb 2012.

- [98] P. Adamson et al. The NuMI Neutrino Beam. *arXiv:1507.06690 [physics.acc-ph]*, July 2015.
- [99] R. Shneor et al. Investigation of proton-proton short-range correlations via the $^{12}\text{C}(e, e'pp)$ reaction. *Phys. Rev. Lett.*, 99:072501, Aug 2007.
- [100] R. Subedi et al. Probing cold dense nuclear matter. *Science*, 320(5882):1476–1478, 2008.
- [101] Alfredo Ferrari, Paola R. Sala, Alberto Fasso, and Johannes Ranft. FLUKA: A multi-particle transport code (Program version 2005). *CERN-2005-010, SLAC-R-773, INFN-TC-05-11*, 2005.
- [102] R. P. Feynman and M. Gell-Mann. Theory of the fermi interaction. *Phys. Rev.*, 109:193–198, Jan 1958.
- [103] L. Fields. TN010 An Analysis of $\bar{\nu}$ Charged Current Quasi-Elastic Interactions. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=8716>.
- [104] L. Fields. Neutrino Electron Scattering Plot Approval, 2015. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11668>.
- [105] L. Fields, K. McFarland, and C. Patrick. Approved Results PRL CCQE Flux Update, 2013. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11594>.

- [106] John M. Finn, Robert W. Lourie, and Bruce H. Cottman. Scaling violation in the separated response functions of ^{12}C . *Phys. Rev. C*, 29:2230–2238, Jun 1984.
- [107] G.A. Fiorentini et al. Measurement of Muon Neutrino Quasielastic Scattering on a Hydrocarbon Target at $E_\nu \sim 3.5$ GeV. *Physical Review Letters*, 111(2):022502, July 2013.
- [108] G.A. Fiorentini and D. Schmitz. QE Analysis Update. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7644>.
- [109] D. V. Forero, M. Tórtola, and J. W. F. Valle. Neutrino oscillations refitted. *Phys. Rev. D*, 90:093006, Nov 2014.
- [110] J. A. Formaggio and G. P. Zeller. From ev to eev: Neutrino cross sections across energy scales. *Rev. Mod. Phys.*, 84:1307–1341, Sep 2012.
- [111] Patrick Fuhrmann. dCache, the Overview. Technical report, Deutsches Elektronen Synchrotron.
- [112] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998.
- [113] H. Gallagher. Reweightable GENIE Uncertainties, 2012. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7451>.
- [114] H. Gallagher, G. Garvey, and G. P. Zeller. Neutrino-nucleus interactions. *Ann. Rev. Nucl. Part. Sci.*, 61:355–378, 2011.

- [115] S. Galster, H. Klein, J. Moritz, K.H. Schmidt, D. Wegener, and J. Bleckwenn. Elastic electron-deuteron scattering and the electric neutron form factor at four-momentum transfers $5 \text{ fm}^2 < q^2 < 14 \text{ fm}^2$. *Nuclear Physics B*, 32(1):221–237, January 1971.
- [116] A. Gando et al. Reactor On-Off Antineutrino Measurement with KamLAND. *Phys. Rev.*, D88(3):033001, 2013.
- [117] S.S. Gershtein and Y.B. Zeldovich. Meson corrections in the theory of beta decay. *Sov. Phys. JETP*, 2:576, 1956.
- [118] S. L. Glashow, J. Iliopoulos, and L. Maiani. Weak interactions with lepton-hadron symmetry. *Phys. Rev. D*, 2:1285–1292, Oct 1970.
- [119] T. Golan. *Modeling nuclear effects in NuWro Monte Carlo neutrino event generator*. PhD thesis, University of Wroclaw, 2014.
- [120] T. Golan. Private communication, 2016.
- [121] T. Golan, J. T. Sobczyk, and J. Zmuda. NuWro: the Wroclaw Monte Carlo Generator of Neutrino Interactions. *Nucl. Phys. Proc. Suppl.*, 229-232:499, 2012.
- [122] M. L. Goldberger and S. B. Treiman. Decay of the Pi Meson. *Phys. Rev.*, 5:1178–1184, 1958.
- [123] Maury Goodman. The Deep Underground Neutrino Experiment,. *Advances in High Energy Physics*, 2015, 2015. Article ID 256351.

- [124] Kurt Gottfried. The determination of the nuclear pair correlation function and momentum distribution. *Annals of Physics*, 21(1):29 – 46, 1963.
- [125] R. Gran. TN045: Pion calorimetry from MINERvA testbeam data. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=9474>.
- [126] R. Gran, J. Nieves, F. Sanchez, and M. J. Vicente Vacas. Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV. *Phys. Rev. D*, 88:113007, Dec 2013.
- [127] R. Gran and others. Measurement of the quasielastic axial vector mass in neutrino interactions on oxygen. *Phys. Rev. D*, 74:052002, Sep 2006.
- [128] Joe M. Grange. *First Measurement of the Muon Anti-Neutrino Charged Current Quasielastic Double-Differential Cross-Section*. PhD thesis, Florida U., 2013.
- [129] David Griffiths. *Introduction to elementary particles*. Wiley-VCH Verlag GmbH, 2009.
- [130] A Heikkinen, N Stepanov, and J.P. Wellisch. Bertini intra-nuclear cascade implementation in Geant4. arXiv nucl-th/0306008, 2003.
- [131] J. Hosaka et al. Solar neutrino measurements in Super-Kamiokande-I. *Phys. Rev. D*, 73:112001, Jun 2006.
- [132] E. L. Hubbard. Booster Synchrotron. Technical report, Fermilab, 1973.
- [133] R.A. Illingworth, M.W. Mengel, and K Lato. SAM User Guide. Available at <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5483>.

- [134] Cezary Juszczak, Jan T. Sobczyk, and Jakub Żmuda. Extraction of the axial mass parameter from miniboone neutrino quasielastic double differential cross-section data. *Phys. Rev. C*, 82:045502, Oct 2010.
- [135] B. Kayser, F. Perrier, and F. Gibrat-Debu. *The physics of massive neutrinos*. World Scientific lecture notes in physics. World Scientific, 1989.
- [136] J. J. Kelly. Simple parametrization of nucleon form factors. *Physical Review C*, 70(6):068202, December 2004.
- [137] Hamamatsu Photonics K.K. *Photomultiplier Tubes - Basics and Applications*. Hamamatsu Photonics K.K., Electron Tube Division, 2007.
- [138] J. Kleykamp. MEU/LY/PE as a function of time plots, 2016. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11844>.
- [139] Makoto Kobayashi and Toshihide Maskawa. CP Violation in the Renormalizable Theory of Weak Interaction. *Prog. Theor. Phys.*, 49:652–657, 1973.
- [140] K.S. Kuzmin, V.V. Lyubushkin, and V.A. Naumov. Quasielastic axial-vector mass from experiments on neutrino-nucleus scattering. *Eur. Phys. J. C*, 54:517–538, 2008.
- [141] Olga Lalakulich, Emmanuel A. Paschos, and Giorgi Piranishvili. Resonance production by neutrinos: The second resonance region. *Phys. Rev. D*, 74:014009, Jul 2006.

- [142] Andrey V. Lebedev. *Ratio of pion kaon production in proton carbon interactions*. PhD thesis, Harvard U., 2007.
- [143] T. Leitner, O. Buss, L. Alvarez-Ruso, and U. Mosel. Electron- and neutrino-nucleus scattering from the quasielastic to the resonance region. *Phys. Rev. C*, 79:034601, Mar 2009.
- [144] L.C. Liu. Pauli blocking and final-state interaction in electron-nucleus quasielastic scattering. *Phys. Rev. C*, 79:014605, Jan 2009.
- [145] C.H. Llewellyn Smith. Neutrino reactions at accelerator energies. *Physics Reports*, 3(5):261–379, 1972.
- [146] V Lyubushkin et al. A Study of quasi-elastic muon neutrino and antineutrino scattering in the NOMAD experiment. *Eur. Phys. J.*, C63:355–381, 2009.
- [147] C. Mahaux and R. Sartor. Theoretical approaches to the momentum distribution of a normal fermi liquid. *Physics Reports*, 211(2):53 – 111, 1992.
- [148] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962.
- [149] Vahe Mamyán. *Measurements of F_2 and $R = \sigma_L/\sigma_T$ on Nuclear Targets in the Nucleon Resonance Region*. PhD thesis, University of Virginia, 2010.
- [150] Bastian Märkisch. Experimental Status of ν_{ud} from Neutron Decay. *arXiv:1107.3422 [nucl-ex]*, Jul 2011.

- [151] M. Martini, M. Ericson, and G. Chanfray. Energy reconstruction effects in neutrino oscillation experiments and implications for the analysis. *Phys. Rev.*, D87(1):013009, 2013.
- [152] M. Martini, M. Ericson, G. Chanfray, and J. Marteau. Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei. *Phys. Rev. C*, 80:065501, Dec 2009.
- [153] Joachim Alexander Maruhn, Paul-Gerhard Reinhard, and Eric Suraud. *Simple Models of Many-Fermion Systems*. Springer-Verlag Berlin Heidelberg, 2009.
- [154] K. McFarland. Neutron Interaction Reweighting. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7518>.
- [155] Susanne Mertens. Status of the KATRIN Experiment and Prospects to Search for keV-mass Sterile Neutrinos in Tritium β -decay. *Physics Procedia*, 61:267 – 273, 2015. 13th International Conference on Topics in Astroparticle and Underground Physics, {TAUP} 2013.
- [156] D.G. Michael et al. The magnetized steel and scintillator calorimeters of the MINOS experiment . *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 596(2):190 – 228, 2008.
- [157] A. Mislevic. Last LE Coherent Pre-Approval Talk, 2015. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=11567>.

- [158] J Morfín, J Nieves, and J Sobczyk. Recent developments in neutrino/antineutrino-nucleus interactions. *Advances in High Energy Physics*, 2012:934597, 2012.
- [159] Sir N. F. Mott and Sir H Massey. *The theory of atomic collisions*. Oxford : Clarendon Press, 2nd ed edition, 1933.
- [160] John W. Negele. Structure of Finite Nuclei in the Local-Density Approximation. *Phys. Rev.*, C1:1260–1321, 1970.
- [161] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasielastic charged-current neutrino-nucleus reactions. *Phys. Rev. C*, 70:055503, Nov 2004.
- [162] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. Inclusive charged-current neutrino-nucleus reactions. *Phys. Rev. C*, 83:045501, Apr 2011.
- [163] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014. 2015 update.
- [164] J.W Van Orden and T.W Donnelly. Mesonic processes in deep-inelastic electron scattering from nuclei. *Annals of Physics*, 131(2):451 – 493, 1981.
- [165] J. Park. Neutrino-Electron Scattering in MINERvA for Constraining the NuMI Neutrino Flux, 2014. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=9407>.
- [166] J. Park et al. Measurement of Neutrino Flux from Neutrino-Electron Elastic Scattering. arXiv:1512.07699 (physics.ins-det), 2015.

- [167] J. Park et al. Measurement of Neutrino Flux from Neutrino-Electron Elastic Scattering. arXiv:1512.07699 (physics.ins-det), 2016.
- [168] J. Park and K. McFarland. Nu-E At Last, 2015. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=10907>.
- [169] R. B. Patterson. The NO ν A Experiment: Status and Outlook. *Nucl. Phys. Proc. Suppl.*, 151:235–236, 2013.
- [170] Zarko Pavlovic. *Observation of Disappearance of Muon Neutrinos in the NuMI Beam*. PhD thesis, Texas U., 2008.
- [171] G. N. Perdue et al. The MINER ν A Data Acquisition System and Infrastructure. *Nucl. Instrum. Meth.*, A694:179–192, 2012.
- [172] B. Pontecorvo. Mesonium and antimesonium. *Sov. Phys. JETP*, 6:429, 1957.
- [173] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958. [*Zh. Eksp. Teor. Fiz.*34,247(1957)].
- [174] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov. Phys. JETP*, 26:984–988, 1968. [*Zh. Eksp. Teor. Fiz.*53,1717(1967)].
- [175] Hyper-Kamiokande Proto-Collaboration et al. Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande. *Progress of Theoretical and Experimental Physics*, 2015(5), 2015.

- [176] J. Rademacker. An exact formula to describe the amplification process in a photomultiplier tube. *Nucl. Instrum. Meth.*, A484:432–443, 2002.
- [177] R. Ransome. MINERvA Physical Description. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=6016>.
- [178] Dieter Rein and Lalit M. Sehgal. Neutrino Excitation of Baryon Resonances and Single Pion Production. *Annals Phys.*, 133:79–153, 1981.
- [179] P. Rodrigues, C. Wilkinson, and K. McFarland. Constraining the GENIE model of neutrino-induced single pion production using reanalyzed bubble chamber data. arXiv 1601.01888 (physics.hep-ex), 2016.
- [180] S.M. Ross. Peirce’s criterion for the elimination of suspect experimental data. *Journal of Engineering Technology*, 20(2):38–41, 2003.
- [181] R. G. Sachs. High-energy behavior of nucleon electromagnetic form factors. *Phys. Rev.*, 126:2256–2260, Jun 1962.
- [182] R. Schiavilla, R. B. Wiringa, Steven C. Pieper, and J. Carlson. Tensor forces and the ground-state structure of nuclei. *Phys. Rev. Lett.*, 98:132501, Mar 2007.
- [183] D. Schmitz. TN016: Normalization Corrections for Titan Production. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=8646>.
- [184] R.A. Smith and E.J. Moniz. Neutrino reactions on nuclear targets. *Nuclear Physics B*, 43:605 – 622, 1972.

- [185] M. Storms. Summary of MINERvA/MINOS Positioning, 2012. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7832>.
- [186] T. Sumikama et al. Test of the conserved vector current hypothesis by a β -ray angular distribution measurement in the mass-8 system. *Phys. Rev. C*, 83:065501, Jun 2011.
- [187] N. Tagg et al. Arachne - A web-based event viewer for MINERvA. *Nucl. Instrum. Meth.*, 676:44–49, 2012.
- [188] Brian George Tice. *Measurement of Nuclear Dependence in Inclusive Charged Current Neutrino Scattering*. PhD thesis, Rutgers U., Piscataway, 2014-01.
- [189] F. Vannucci. The NOMAD Experiment at CERN. *Adv. High Energy Phys.*, 2014:129694, 2014.
- [190] J.P. Velasquez, S. Dytman, and A. Gago. Geant4 Hadronic Uncertainties. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=8249>.
- [191] T. Walton, M. Betancourt, et al. Measurement of muon plus proton final states in ν_μ interactions on hydrocarbon at $\langle E_\nu \rangle = 4.2$ GeV. *Phys. Rev. D*, 91:071301, Apr 2015.
- [192] R. R. Whitney, I. Sick, J. R. Ficenec, R. D. Kephart, and W. P. Trower. Quasielastic electron scattering. *Phys. Rev. C*, 9:2230–2235, Jun 1974.

- [193] Denys H. Wilkinson. Analysis of Neutron Beta Decay. *Nucl. Phys.*, A377:474–504, 1982.
- [194] D.H. Wilkinson. Limits to second-class nucleonic currents. *Nucl. Inst. and Meth.*, 455:656–659, 2000.
- [195] J. Wolcott. Cross-talk Measurement and Systematics. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7528>.
- [196] J. Wolcott. TN015: Systematic uncertainties from Feynman scaling applied to NA49, 2013. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=8628>.
- [197] T. Yang. AGKY — transitions between KNO-based and JETSET. In *AIP Conf Proc*, volume 967, pages 269–275, 2007.
- [198] A. Zee. Parametrizing the neutrino mixing matrix. *Phys. Rev. D*, 68:093002, Nov 2003.
- [199] B. Ziemer. Reconstruction Systematics, 2012. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7635>.
- [200] B. Ziemer. Vertex Resolution Study, 2012. Available at <http://minerva-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=7581>.

APPENDIX A

Preliminary tables of cross section measurements

Note: these measurements are **preliminary** and may change slightly before publication.

A.1. Quasi-elastic-like cross sections

The tables in this section list our cross section measurements for the quasi-elastic-like signal definition explained in section 4.1.3. In each case, we show a table of values corresponding to the cross section, followed by a table of statistical uncertainties, then one of systematic uncertainties. Units are explained in the captions to the tables.

A.1.1. Muon parallel and transverse momentum

The cross section vs. muon parallel and transverse momentum, $d^2\sigma/dp_T dp_{\parallel}$ is shown in table A.1. The statistical uncertainty on these measurements is in table A.2, and the systematic uncertainty in table A.3.

Table A.4 shows the differential cross section $d\sigma/dp_T$, generated by projecting the two-dimensional measurement onto the p_T axis. Table A.5 shows the differential cross section $d\sigma/dp_{\parallel}$, generated by projecting the two-dimensional measurement onto the p_{\parallel} axis.

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	59.93	179.79	286.27	280.38	143.00	6.57
2-2.5	57.69	197.56	327.00	397.44	161.32	23.43
2.5-3	62.74	258.48	410.63	420.40	192.86	12.17
3-3.5	63.16	198.63	352.40	350.22	112.64	15.38
3.5-4	42.10	126.39	216.03	206.86	74.15	9.93
4-4.5	17.71	65.48	137.29	99.12	37.53	5.78
4.5-5	8.82	39.11	61.77	53.60	23.69	4.21
5-6	4.59	23.90	33.81	41.31	17.68	2.98
6-8	3.69	9.79	20.22	24.06	11.66	1.92
8-10	2.46	7.62	10.71	9.81	6.03	2.53
10-15	0.68	2.69	5.10	5.84	2.90	0.94

Table A.1. Measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	7.41	17.06	19.88	18.57	26.34	4.23
2-2.5	6.72	15.46	17.52	16.63	15.13	9.34
2.5-3	6.61	17.26	18.63	14.86	12.57	2.81
3-3.5	6.49	14.39	16.38	12.38	7.56	3.13
3.5-4	5.00	10.76	11.98	8.88	5.65	1.88
4-4.5	2.86	6.96	9.42	5.80	3.53	1.11
4.5-5	1.96	5.47	5.73	3.90	2.61	0.84
5-6	1.19	3.47	3.46	2.86	1.88	0.61
6-8	0.85	1.58	2.05	1.62	1.15	0.35
8-10	0.75	1.45	1.46	1.01	0.79	0.48
10-15	0.23	0.48	0.67	0.50	0.37	0.20

Table A.2. Statistical uncertainty on the measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

A.1.2. True neutrino energy and Q_{QE}^2

The energy-dependent cross section vs Q_{QE}^2 , $d\sigma(E_{\nu})/dQ_{QE}^2$ is shown in table A.6. The statistical uncertainty on these measurements is in table A.7, and the systematic uncertainty in table A.8.

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	13.33	34.89	46.89	41.59	28.45	1.41
2-2.5	9.72	31.11	46.88	48.65	19.88	6.54
2.5-3	9.56	37.19	53.32	45.18	22.79	2.93
3-3.5	8.54	22.52	38.12	37.73	15.43	4.28
3.5-4	5.34	13.94	22.07	25.17	11.82	2.23
4-4.5	2.82	7.45	16.44	14.49	5.92	1.40
4.5-5	1.55	4.40	6.91	6.36	3.08	1.15
5-6	0.77	2.82	3.69	4.63	2.43	0.74
6-8	0.49	1.36	2.42	2.76	1.82	0.45
8-10	0.37	1.14	1.44	1.22	1.10	0.54
10-15	0.11	0.30	0.96	0.72	0.49	0.38

Table A.3. Systematic uncertainty on the measured double differential quasi-elastic-like cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
0-0.15	176.36	7.97	25.87
0.15-0.25	604.88	18.53	79.71
0.25-0.4	1016.84	21.08	120.36
0.4-0.7	1042.29	17.60	112.82
0.7-1	440.14	17.66	60.21
1-1.5	55.29	5.90	12.29

Table A.4. Differential quasi-elastic-like cross section $d\sigma/dp_T$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_T bins are in GeV.

The flux-integrated single-differential cross section $d\sigma/dQ_{QE}^2$ is shown in table A.9. This is generated by dividing the (background-subtracted, unfolded, efficiency-corrected) event distribution dN/dQ_{QE}^2 by the total flux Φ .

Table A.10 shows the energy-dependent cross section $\sigma(E_{\nu})$. This is generated by dividing the number of events in each energy bin $N_i(E_{\nu})$ by the flux in that bin $\Phi_i(E_{\nu})$.

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	200.21	10.53	28.14
2-2.5	256.80	8.81	30.93
2.5-3	286.92	6.92	31.65
3-3.5	228.74	5.52	24.37
3.5-4	140.62	3.98	16.57
4-4.5	73.68	2.67	10.03
4.5-5	39.79	1.81	4.78
5-6	27.34	1.25	3.16
6-8	16.24	0.72	1.95
8-10	8.75	0.54	1.14
10-15	4.23	0.24	0.55

Table A.5. Differential quasi-elastic-like cross section $d\sigma/dp_{\parallel}$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_{\parallel} bins are in GeV.

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	3498.11	4098.42	3587.67	3204.90	1882.18	585.97	145.49	0.00
2-2.5	2582.14	3462.87	3259.15	2973.10	2278.96	778.20	199.09	0.00
2.5-3	2984.61	3852.23	3824.71	3566.23	2379.89	965.92	325.91	41.33
3-3.5	2932.52	3534.99	3672.29	3460.72	2462.35	1175.62	378.99	88.44
3.5-4	2602.07	2879.70	3491.03	2982.39	2131.48	1003.12	398.14	83.95
4-5	1000.55	1298.97	1610.57	1568.23	1037.29	478.51	182.42	60.64
5-6	1045.86	1961.67	1584.68	1483.00	1345.00	578.34	150.77	65.62
6-7	1221.27	1853.47	2224.07	1794.67	1705.64	690.83	231.40	60.65
7-8	1395.01	1351.66	1738.52	2058.22	1395.95	740.43	221.33	76.95
8-10	805.77	1072.27	1042.58	928.46	556.75	317.68	156.95	55.24

Table A.6. Measured energy-dependent quasi-elastic-like cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_{ν} (GeV).

A.1.3. Neutrino energy reconstructed in the CCQE hypothesis, and Q_{QE}^2

The reconstructed energy-dependent cross section vs Q_{QE}^2 , $d\sigma(E_{\nu}^{QE})/dQ_{QE}^2$ is shown in table A.11. The statistical uncertainty on these measurements is in table A.12, and the systematic uncertainty in table A.13. The total reconstructed energy-dependent cross

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	486.84	561.95	409.15	327.46	257.58	200.56	288.22	0.00
2-2.5	294.74	363.61	259.26	194.36	147.61	91.23	105.01	0.00
2.5-3	282.07	332.15	240.39	176.40	113.79	63.10	58.36	21.22
3-3.5	271.52	304.04	225.63	164.64	104.35	57.21	44.03	20.80
3.5-4	287.41	291.86	240.93	165.81	102.59	51.99	38.23	15.15
4-5	133.81	149.31	124.67	91.96	54.80	27.06	19.16	9.36
5-6	220.72	330.36	210.25	147.70	105.91	49.98	23.10	12.00
6-7	316.60	399.39	334.89	205.62	159.75	70.68	39.14	14.10
7-8	407.92	280.89	310.54	271.13	150.05	84.61	40.94	17.49
8-10	228.50	222.31	173.16	115.61	64.60	35.81	25.65	12.00

Table A.7. Statistical uncertainty on the measured quasi-elastic-like energy-dependent cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	969.37	1067.58	772.20	706.64	393.98	121.76	93.22	0.00
2-2.5	480.75	629.37	513.87	453.95	325.30	100.90	59.43	0.00
2.5-3	502.88	563.41	539.69	480.41	294.04	118.48	62.68	12.74
3-3.5	421.24	450.58	455.24	424.31	290.21	137.96	65.31	16.08
3.5-4	387.47	391.82	462.70	389.82	294.47	144.74	73.82	24.03
4-5	205.44	223.68	272.64	267.14	183.06	87.74	43.35	13.76
5-6	190.97	333.09	249.07	232.75	209.65	101.73	31.26	16.16
6-7	200.61	334.32	351.08	241.64	226.70	101.85	46.32	12.95
7-8	213.68	219.82	233.72	300.70	189.95	122.90	39.99	18.41
8-10	110.16	207.16	162.95	122.12	74.79	46.66	28.48	14.02

Table A.8. Systematic uncertainty on the measured energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV).

section $\sigma(E_\nu^{QE})$, generated by dividing the event count $N_i(E_\nu^{QE})$ in each bin by that energy bin's flux $\Phi(E_\nu)$, is shown in table A.14. Note the difference between true E_ν (a property of the incoming neutrino) and E_ν^{QE} , in which we reconstruct an approximation of E_ν from the muon kinematics, using the quasi-elastic hypothesis, as explained in appendix B. The

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
0-0.025	1146.41	49.58	169.06
0.025-0.05	1454.29	57.91	195.82
0.05-0.1	1493.60	43.12	184.31
0.1-0.2	1382.50	32.43	163.35
0.2-0.4	973.90	23.16	108.07
0.4-0.8	416.82	16.88	44.58
0.8-1.2	127.83	10.50	18.75
1.2-2	25.86	2.65	4.80

Table A.9. Flux-integrated quasi-elastic-like differential cross section $d\sigma/dQ_{QE}^2$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. The Q_{QE}^2 bins are in GeV^2 .

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	679.40	77.81	143.73
2-2.5	729.05	34.10	102.43
2.5-3	872.28	25.31	106.36
3-3.5	938.22	22.63	106.99
3.5-4	831.90	20.97	112.72
4-5	815.18	22.37	144.48
5-6	915.86	38.12	145.51
6-7	1126.08	55.30	150.32
7-8	1086.87	60.33	154.01
8-10	1074.65	58.02	145.77

Table A.10. Energy-dependent quasi-elastic-like cross section $\sigma(E_\nu)$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν bins are in GeV .

ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$ is shown in table A.15, with the statistical and systematic uncertainties in tables A.16 and A.17 respectively.

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	3492.24	4245.94	3778.16	3415.43	2120.79	745.85	0.00	0.00
2-2.5	2687.21	3677.30	3368.53	3143.17	2388.83	840.00	187.09	0.00
2.5-3	3022.38	3872.41	3837.07	3551.01	2359.51	974.87	340.26	4.53
3-3.5	2760.38	3284.92	3486.76	3323.43	2381.88	1136.16	360.80	94.08
3.5-4	2347.31	2611.69	3282.67	2754.17	1984.54	934.42	385.77	80.17
4-5	913.13	1192.03	1463.71	1469.71	963.34	450.88	175.41	61.05
5-6	926.76	1875.77	1525.16	1368.65	1298.58	556.74	147.22	64.98
6-7	1198.66	1735.20	2118.60	1796.22	1667.76	663.96	219.75	59.56
7-8	1387.64	1365.39	1713.50	1972.86	1321.36	720.99	219.32	81.27
8-10	803.37	1027.70	1002.36	913.42	543.70	309.62	155.43	53.63

Table A.11. Measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	452.28	560.26	406.34	339.12	281.21	264.08	0.00	0.00
2-2.5	306.86	375.30	264.43	200.82	152.57	99.84	117.39	0.00
2.5-3	284.08	339.79	241.93	177.93	113.69	65.78	67.58	2.30
3-3.5	262.59	291.80	219.33	160.30	103.09	57.59	44.81	27.02
3.5-4	268.02	275.28	232.13	156.10	99.14	51.50	41.44	15.01
4-5	125.30	139.70	116.47	87.02	51.93	25.90	18.80	9.69
5-6	204.88	313.96	205.46	139.07	104.31	48.22	22.64	12.38
6-7	320.39	380.10	321.52	203.89	157.98	68.74	38.32	13.89
7-8	417.28	289.00	306.50	260.04	142.86	82.08	41.03	19.18
8-10	230.53	211.15	168.50	115.10	63.61	35.39	25.81	11.95

Table A.12. Statistical uncertainty on the measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	817.40	934.45	717.97	673.81	406.74	148.44	0.00	0.00
2-2.5	473.43	624.12	507.66	466.78	310.00	95.08	72.63	0.00
2.5-3	490.82	534.71	521.70	469.85	265.52	104.28	56.86	1.84
3-3.5	362.31	380.19	387.49	361.02	249.33	121.18	47.52	22.53
3.5-4	310.08	320.70	391.08	316.73	252.87	125.67	67.14	19.47
4-5	168.85	187.73	229.05	232.45	159.69	78.92	39.34	13.00
5-6	164.87	284.42	225.41	202.09	189.39	96.75	29.66	15.43
6-7	184.90	351.92	285.34	236.37	212.26	95.06	42.32	12.25
7-8	211.19	217.36	239.55	296.93	172.24	104.61	40.49	17.83
8-10	107.96	189.32	159.14	119.49	71.66	45.13	27.78	14.37

Table A.13. Systematic uncertainty on the measured reconstructed energy-dependent quasi-elastic-like cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	723.20	63.66	138.00
2-2.5	765.23	36.93	102.24
2.5-3	860.45	25.14	98.29
3-3.5	904.12	23.55	92.57
3.5-4	776.32	20.66	96.83
4-5	764.81	21.53	128.00
5-6	876.47	37.07	133.26
6-7	1093.58	54.23	139.80
7-8	1057.20	58.92	142.17
8-10	1049.80	57.44	141.65

Table A.14. Reconstructed energy-dependent quasi-elastic-like cross section $\sigma(E_\nu^{QE})$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν^{QE} bins are in GeV .

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	1.00	1.04	1.05	1.07	1.13	1.27	0.00	0.00
2-2.5	1.04	1.06	1.03	1.06	1.05	1.08	0.94	0.00
2.5-3	1.01	1.01	1.00	1.00	0.99	1.01	1.04	0.11
3-3.5	0.94	0.93	0.95	0.96	0.97	0.97	0.95	1.06
3.5-4	0.90	0.91	0.94	0.92	0.93	0.93	0.97	0.95
4-5	0.91	0.92	0.91	0.94	0.93	0.94	0.96	1.01
5-6	0.89	0.96	0.96	0.92	0.97	0.96	0.98	0.99
6-7	0.98	0.94	0.95	1.00	0.98	0.96	0.95	0.98
7-8	0.99	1.01	0.99	0.96	0.95	0.97	0.99	1.06
8-10	1.00	0.96	0.96	0.98	0.98	0.97	0.99	0.97

Table A.15. Ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV²), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	0.19	0.20	0.17	0.15	0.21	0.63	0.00	0.00
2-2.5	0.17	0.16	0.12	0.10	0.10	0.18	0.77	0.00
2.5-3	0.13	0.12	0.09	0.07	0.07	0.09	0.28	0.08
3-3.5	0.12	0.11	0.08	0.07	0.06	0.07	0.16	0.39
3.5-4	0.14	0.13	0.09	0.07	0.06	0.07	0.14	0.25
4-5	0.17	0.15	0.10	0.08	0.07	0.08	0.14	0.22
5-6	0.27	0.23	0.18	0.13	0.11	0.12	0.21	0.26
6-7	0.37	0.29	0.20	0.16	0.13	0.14	0.23	0.32
7-8	0.42	0.30	0.25	0.18	0.14	0.16	0.26	0.35
8-10	0.40	0.28	0.23	0.17	0.16	0.16	0.23	0.30

Table A.16. Statistical uncertainty on the ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV²), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	0.08	0.05	0.04	0.05	0.05	0.24	0.00	0.00
2-2.5	0.05	0.03	0.01	0.01	0.03	0.04	0.19	0.00
2.5-3	0.02	0.03	0.01	0.01	0.02	0.03	0.07	0.05
3-3.5	0.03	0.02	0.02	0.02	0.02	0.03	0.07	0.13
3.5-4	0.03	0.02	0.03	0.03	0.02	0.02	0.03	0.07
4-5	0.03	0.04	0.02	0.03	0.03	0.03	0.03	0.04
5-6	0.02	0.03	0.02	0.02	0.02	0.02	0.02	0.04
6-7	0.02	0.04	0.04	0.01	0.01	0.02	0.02	0.03
7-8	0.02	0.04	0.04	0.03	0.02	0.05	0.02	0.04
8-10	0.02	0.02	0.02	0.01	0.01	0.02	0.01	0.02

Table A.17. Systematic uncertainty on the ratio of $d\sigma(E_\nu^{QE})/dQ_{QE}^2$ to $d\sigma(E_\nu)/dQ_{QE}^2$. Columns represent bins of Q_{QE}^2 (GeV²), rows are bins of E_ν^{QE} (GeV).

A.2. True CCQE cross sections

The tables in this section list our cross section measurements for the true quasi-elastic signal definition explained in section 4.1.2. In each case, we show a table of values corresponding to the cross section, followed by a table of statistical uncertainties, then one of systematic uncertainties. Units are explained in the captions to the tables.

A.2.1. Muon parallel and transverse momentum

The cross section vs. muon parallel and transverse momentum, $d^2\sigma/dp_T dp_{\parallel}$ is shown in table A.18. The statistical uncertainty on these measurements is in table A.19, and the systematic uncertainty in table A.20.

Table A.21 shows the differential cross section $d\sigma/dp_T$, generated by projecting the two-dimensional measurement onto the p_T axis. Table A.22 shows the differential cross section $d\sigma/dp_{\parallel}$, generated by projecting the two-dimensional measurement onto the p_{\parallel} axis.

A.2.2. True neutrino energy and Q_{QE}^2

The energy-dependent true CCQE cross section vs Q_{QE}^2 , $d\sigma(E_{\nu})/dQ_{QE}^2$ is shown in table A.23. The statistical uncertainty on these measurements is in table A.24, and the systematic uncertainty in table A.25. The flux-integrated single-differential cross section $d\sigma/dQ_{QE}^2$ is shown in table A.26. Table A.27 shows the energy-dependent cross section $\sigma(E_{\nu})$.

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	41.63	127.56	226.94	230.63	126.94	9.72
2-2.5	39.75	141.78	267.19	336.24	144.06	33.34
2.5-3	44.30	194.17	345.37	364.72	184.25	14.58
3-3.5	52.63	159.07	303.57	315.70	104.01	16.47
3.5-4	32.45	98.67	187.01	191.08	69.51	9.37
4-4.5	14.69	55.46	122.14	90.43	34.78	5.36
4.5-5	7.80	30.67	52.24	47.44	22.19	3.85
5-6	3.45	18.67	28.82	36.06	16.40	2.79
6-8	2.83	8.55	17.30	20.28	10.30	1.99
8-10	2.09	5.72	8.57	8.51	4.90	1.94
10-15	0.52	2.22	4.32	4.88	2.42	0.80

Table A.18. Measured double differential true CCQE cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	5.80	13.47	16.45	15.91	24.48	6.15
2-2.5	5.26	12.52	14.96	14.76	14.47	13.78
2.5-3	5.39	14.44	16.39	13.50	12.47	3.69
3-3.5	5.99	12.60	14.76	11.53	7.31	3.55
3.5-4	4.41	9.40	10.88	8.38	5.50	1.96
4-4.5	2.65	6.49	8.76	5.39	3.40	1.14
4.5-5	1.96	4.73	5.12	3.57	2.58	0.88
5-6	1.01	3.04	3.10	2.59	1.85	0.61
6-8	0.74	1.50	1.83	1.43	1.11	0.40
8-10	0.71	1.25	1.24	0.90	0.71	0.46
10-15	0.19	0.44	0.61	0.44	0.33	0.20

Table A.19. Statistical uncertainty on the measured double differential true CCQE cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

A.2.3. Neutrino energy reconstructed in the CCQE hypothesis, and Q_{QE}^2

The reconstructed energy-dependent true CCQE cross section vs Q_{QE}^2 , $d\sigma(E_{\nu}^{QE})/dQ_{QE}^2$ is shown in table A.28. The statistical uncertainty on these measurements is in table

	0-0.15	0.15-0.25	0.25-0.4	0.4-0.7	0.7-1	1-1.5
1.5-2	14.63	35.90	52.26	44.13	20.17	1.85
2-2.5	11.84	34.87	55.96	57.36	16.66	10.91
2.5-3	11.77	40.52	62.15	52.83	22.33	2.91
3-3.5	10.30	25.79	44.57	43.46	14.42	4.62
3.5-4	6.07	15.11	24.91	27.08	11.47	1.87
4-4.5	2.82	7.76	17.92	14.55	5.62	1.30
4.5-5	1.58	4.76	7.58	7.11	2.96	0.93
5-6	0.85	3.33	4.49	5.31	2.35	0.70
6-8	0.67	1.53	2.85	3.18	1.75	0.45
8-10	0.38	1.28	1.56	1.43	0.94	0.33
10-15	0.12	0.47	0.80	0.97	0.47	0.43

Table A.20. Systematic uncertainty on the measured double differential true CCQE cross section $d^2\sigma/dp_T dp_{\parallel}$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of p_T (GeV), rows are bins of p_{\parallel} (GeV).

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
0-0.15	176.36	7.97	25.87
0.15-0.25	604.88	18.53	79.71
0.25-0.4	1016.84	21.08	120.36
0.4-0.7	1042.29	17.60	112.82
0.7-1	440.14	17.66	60.21
1-1.5	55.29	5.90	12.29

Table A.21. Differential true CCQE cross section $d\sigma/dp_T$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_T bins are in GeV.

A.29, and the systematic uncertainty in table A.30. By projecting, we can get the total reconstructed energy-dependent cross section $\sigma(E_{\nu}^{QE})$, shown in table A.31.

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	200.21	10.53	28.14
2-2.5	256.80	8.81	30.93
2.5-3	286.92	6.92	31.65
3-3.5	228.74	5.52	24.37
3.5-4	140.62	3.98	16.57
4-4.5	73.68	2.67	10.03
4.5-5	39.79	1.81	4.78
5-6	27.34	1.25	3.16
6-8	16.24	0.72	1.95
8-10	8.75	0.54	1.14
10-15	4.23	0.24	0.55

Table A.22. Differential true CCQE cross section $d\sigma/dp_{\parallel}$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}/\text{proton}$. The p_{\parallel} bins are in GeV.

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	2486.00	2942.41	2674.31	2731.36	1718.83	612.40	165.39	0.00
2-2.5	1807.76	2535.17	2733.30	2534.40	1962.21	738.15	217.15	0.00
2.5-3	2203.80	2888.97	3062.89	2999.53	2011.61	883.46	308.43	57.12
3-3.5	2256.33	2597.41	2950.67	2874.98	2096.97	1023.97	383.27	108.57
3.5-4	1897.96	1916.82	2597.81	2485.24	1814.04	861.18	384.21	97.56
4-5	756.11	1032.50	1274.83	1289.62	868.23	415.31	168.29	65.30
5-6	744.45	1561.85	1195.15	1196.51	1128.38	487.99	142.53	61.32
6-7	871.60	1401.54	1826.69	1574.08	1413.66	593.37	211.66	58.61
7-8	998.60	974.68	1303.28	1635.18	1063.63	599.33	208.39	78.97
8-10	680.52	728.05	725.94	744.01	457.66	247.87	138.89	57.54

Table A.23. Measured energy-dependent true CCQE cross section $d\sigma(E_{\nu})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_{ν} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	363.43	441.39	324.65	276.18	231.86	207.28	277.29	0.00
2-2.5	239.06	298.73	226.97	167.96	128.87	87.01	113.52	0.00
2.5-3	236.63	281.35	207.44	153.89	99.35	59.58	57.09	30.28
3-3.5	239.14	255.30	197.35	142.72	92.58	52.45	47.83	26.36
3.5-4	244.51	229.83	201.19	144.14	90.73	46.95	39.87	17.99
4-5	115.85	132.73	107.54	78.90	47.47	24.41	19.56	10.59
5-6	178.96	294.14	175.72	125.08	91.33	44.18	23.96	12.50
6-7	263.10	346.17	298.30	184.70	137.07	63.85	39.19	15.55
7-8	362.44	234.64	252.42	229.96	120.41	72.72	41.64	18.56
8-10	207.42	167.26	135.62	96.00	53.86	29.50	24.75	13.57

Table A.24. Statistical uncertainty on the measured energy-dependent true CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	896.62	904.51	722.84	725.27	409.80	145.19	75.53	0.00
2-2.5	510.83	647.31	602.94	514.30	387.97	127.12	87.34	0.00
2.5-3	612.46	659.93	618.49	554.25	349.37	134.57	56.24	11.95
3-3.5	554.01	509.19	553.89	502.78	334.06	146.19	52.70	17.74
3.5-4	454.88	415.72	514.47	437.53	316.23	135.14	63.30	20.14
4-5	213.87	241.62	281.18	269.43	180.21	81.32	36.91	13.38
5-6	216.03	394.63	262.96	247.76	226.04	92.52	25.90	13.27
6-7	272.38	360.72	374.09	282.58	252.42	111.83	40.05	10.71
7-8	249.77	214.34	237.41	327.38	186.64	106.65	39.66	15.17
8-10	132.94	196.31	153.46	129.14	78.38	44.73	20.78	13.05

Table A.25. Systematic uncertainty on the measured energy-dependent true CCQE cross section $d\sigma(E_\nu)/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν (GeV).

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
0-0.025	1146.41	49.58	169.06
0.025-0.05	1454.29	57.91	195.82
0.05-0.1	1493.60	43.12	184.31
0.1-0.2	1382.50	32.43	163.35
0.2-0.4	973.90	23.16	108.07
0.4-0.8	416.82	16.88	44.58
0.8-1.2	127.83	10.50	18.75
1.2-2	25.86	2.65	4.80

Table A.26. Flux-integrated true CCQE differential cross section $d\sigma/dQ_{QE}^2$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. The Q_{QE}^2 bins are in GeV^2 .

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	679.40	77.81	143.73
2-2.5	729.05	34.10	102.43
2.5-3	872.28	25.31	106.36
3-3.5	938.22	22.63	106.99
3.5-4	831.90	20.97	112.72
4-5	815.18	22.37	144.48
5-6	915.86	38.12	145.51
6-7	1126.08	55.30	150.32
7-8	1086.87	60.33	154.01
8-10	1074.65	58.02	145.77

Table A.27. Energy-dependent true CCQE cross section $\sigma(E_\nu)$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν bins are in GeV .

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	2476.72	2921.28	2678.35	2723.60	1723.41	614.08	0.00	0.00
2-2.5	1797.27	2608.45	2753.07	2539.04	1951.23	737.86	229.93	0.00
2.5-3	2226.43	2873.04	3069.11	2978.50	2009.43	867.50	295.76	10.16
3-3.5	2241.41	2564.85	2914.48	2889.10	2099.49	1027.20	370.35	117.69
3.5-4	1880.15	1887.07	2597.00	2477.97	1808.05	858.48	395.86	98.58
4-5	754.60	1030.20	1266.95	1296.74	873.71	421.12	174.51	67.38
5-6	718.59	1562.92	1184.83	1178.02	1133.04	491.66	145.16	64.88
6-7	871.72	1336.74	1833.63	1575.58	1408.35	596.83	208.94	57.70
7-8	1023.25	1013.68	1298.89	1645.34	1081.07	598.71	210.89	85.47
8-10	681.88	721.18	712.22	742.16	456.44	247.16	137.21	56.53

Table A.28. Measured reconstructed energy-dependent cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	359.71	436.72	320.03	276.70	235.66	226.29	0.00	0.00
2-2.5	238.78	302.93	227.87	167.59	128.55	91.17	151.55	0.00
2.5-3	237.04	281.48	208.04	154.45	100.30	60.97	62.61	4.64
3-3.5	238.08	253.00	196.53	143.38	93.60	54.04	50.03	36.30
3.5-4	241.36	227.26	201.61	144.00	91.80	48.65	44.60	19.59
4-5	116.10	131.89	107.07	79.13	47.80	24.73	20.39	11.21
5-6	178.13	291.91	175.49	123.49	92.55	44.27	24.11	13.62
6-7	267.15	335.48	297.82	181.57	137.05	64.23	39.01	14.87
7-8	370.82	236.33	254.36	231.79	121.99	73.22	42.11	21.21
8-10	209.78	166.87	133.24	96.47	53.96	29.54	24.92	13.55

Table A.29. Statistical uncertainty on the measured reconstructed energy-dependent true CCQE cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

	0-0.025	0.025-0.05	0.05-0.1	0.1-0.2	0.2-0.4	0.4-0.8	0.8-1.2	1.2-2
1.5-2	872.39	913.43	710.10	664.11	386.76	188.42	0.00	0.00
2-2.5	541.68	684.11	635.23	534.00	358.72	111.03	90.23	0.00
2.5-3	614.28	611.99	601.92	551.67	324.05	115.89	43.04	4.04
3-3.5	447.89	434.08	459.83	423.23	296.67	133.96	48.29	21.00
3.5-4	349.48	311.68	412.80	351.99	272.96	128.32	64.33	18.44
4-5	159.34	190.90	224.91	230.18	159.85	77.91	37.80	13.47
5-6	187.75	330.52	223.39	213.00	201.89	91.32	26.27	14.75
6-7	239.25	323.30	312.73	267.67	229.46	111.44	40.73	11.43
7-8	230.44	200.32	238.41	309.73	178.71	101.69	37.58	15.52
8-10	128.52	184.73	146.57	124.40	78.44	43.53	20.22	11.69

Table A.30. Systematic uncertainty on the measured reconstructed energy-dependent true CCQE cross section $d\sigma(E_\nu^{QE})/dQ_{QE}^2$. Units are $10^{-41}\text{cm}^2/\text{GeV}^2/\text{proton}$. Columns represent bins of Q_{QE}^2 (GeV^2), rows are bins of E_ν^{QE} (GeV).

Bin	Cross section	Statistical uncertainty	Systematic uncertainty
1.5-2	723.20	63.66	138.00
2-2.5	765.23	36.93	102.24
2.5-3	860.45	25.14	98.29
3-3.5	904.12	23.55	92.57
3.5-4	776.32	20.66	96.83
4-5	764.81	21.53	128.00
5-6	876.47	37.07	133.26
6-7	1093.58	54.23	139.80
7-8	1057.20	58.92	142.17
8-10	1049.80	57.44	141.65

Table A.31. Reconstructed energy-dependent true CCQE cross section $\sigma(E_\nu^{QE})$, along with statistical and systematic uncertainties. Units are $10^{-41}\text{cm}^2/\text{proton}$. The E_ν^{QE} bins are in GeV.

APPENDIX B

Quasi-elastic kinematics

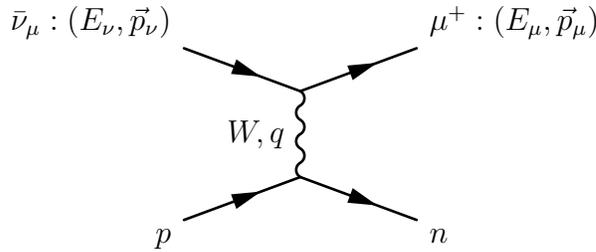
B.1. Q_{QE}^2 calculation in the quasi-elastic hypothesis

Figure B.1. Quasi-elastic scattering kinematics

Calculation of squared 4-momentum transfer from the leptonic to the hadronic system. Assuming quasi-elastic scattering stationary initial-state proton, this can be calculated purely from the muon kinematics. In

this calculation, E_ν and E_μ are the neutrino and muon energy respectively. Muon momentum is represented by p_μ , and θ_{mu} represents the angle between the outgoing muon and the incoming neutrino. Four-momentum transfer q is the difference in 4-momentum or the final state muon and the initial-state neutrino. As neutrino mass is negligible (less than 1eV), we take $m_\nu = 0$, meaning $E_\nu = |\vec{p}_\nu|$. Muon mass is represented by m_μ .

$$\begin{aligned}
Q^2 &= -q^2 \\
&= -(p_\mu^\mu - p_\nu^\mu)^2 \\
&= (\vec{p}_\mu - \vec{p}_\nu)^2 - (E_\mu - E_\nu)^2 \\
&= p_\mu^2 + p_\nu^2 - 2\vec{p}_\mu \cdot \vec{p}_\nu - (E_\mu^2 + E_\nu^2 - 2E_\nu E_\mu) \\
&= (E_\mu^2 - m_\mu^2) + E_\nu^2 - 2E_\nu p_\mu \cos \theta_\mu - E_\mu^2 - E_\nu^2 + 2E_\nu E_\mu \\
&= 2E_\nu(E_\mu - p_\mu \cos \theta_\mu) - m_\mu^2
\end{aligned}$$

B.2. E_ν^{QE} in the quasi-elastic hypothesis

The calculation below is for the quasi-elastic hypothesis with an antineutrino incident upon a proton at rest within a nucleus, with a binding energy E_b . The interaction produces a positively-charged muon and a recoil neutron. Under the quasi-elastic assumption, no energy is lost to the rest of the nucleus - its only effect is to provide the binding energy that lowers the initial state energy of the stationary proton. The same terminology is used as in the previous calculation; additionally, E_n , p_n and θ_n refer to the energy, momentum and angle of the neutron with respect to the incoming neutrino.

$$p_\nu = E_\nu = p_\mu \cos \theta_\mu + p_n \cos \theta_n$$

$$0 = p_\mu \sin \theta_\mu + p_n \sin \theta_n$$

Square and add the p_n terms to eliminate θ_n :

$$\begin{aligned} p_n^2 &= p_\mu^2 \sin^2 \theta_\mu + p_\mu^2 \cos^2 \theta_\mu + E_\nu^2 - 2E_\nu p_\mu \cos \theta_\mu \\ &= p_\mu^2 + E_\nu^2 - 2E_\nu p_\mu \cos \theta_\mu \end{aligned}$$

Now conserve energy. Initially we have our incident neutrino and a proton bound in the nucleus with binding energy E_b ; afterwards we have a muon and a free neutron.

$$E_\nu + (m_p - E_b) = E_n + E_\mu$$

$$E_n = E_\nu + (m_p - E_b) - E_\mu$$

Square both sides and substitute in for $E_n^2 = m_n^2 + p_n^2$ with p_n^2 as given above:

$$m_n^2 + (p_\mu^2 + E_\nu^2 - 2E_\nu p_\mu \cos \theta_\mu) = (E_\nu + (m_p - E_b) - E_\mu)^2$$

$$\begin{aligned} m_n^2 + p_\mu^2 + E_\nu^2 - 2E_\nu p_\mu \cos \theta_\mu &= E_\nu^2 + (m_p - E_b)^2 + E_\mu^2 - 2E_\mu E_\nu \\ &\quad + 2E_\nu(m_p - E_b) - 2E_\mu(m_p - E_b) \end{aligned}$$

$$E_\nu(2p_\mu \cos \theta_\mu - 2E_\mu + 2(m_p - E_b)) = m_n^2 - E_\mu^2 + E_\mu^2 - m_\mu^2$$

$$+ 2(m_p - E_b)E_\mu - (m_p - E_b)^2$$

$$E_\nu^{QE} = \frac{m_n^2 - (m_p - E_b)^2 - m_\mu^2 + 2(m_p - E_b)E_\mu}{2(m_p - E_b - E_\mu + p_\mu \cos \theta_\mu)}$$

APPENDIX C

Ntuple branches**C.1. Quantities for cross section measurement**

This section explains the combination of ntuple branches used to calculate the quantities used in the cross section measurement.

Transverse muon momentum

Reconstructed:

$$\text{CCQEAntiNuTool_p_mu} \sin(\text{CCQEAntiNuTool_theta_mu})/1000$$

True:

$$\sqrt{(\text{mc_primFSLepton}[2] \sin \theta_b + \text{mc_primFSLepton}[1] \cos \theta_b)^2 + \text{mc_primFSLepton}[0]^2}/1000$$

where θ_b is the beam angle, 0.05887 radians. Values are in GeV.

Longitudinal muon momentum

Reconstructed:

$$\text{CCQEAntiNuTool_p_mu} \cos \text{CCQEAntiNuTool_theta_mu}/1000$$

True:

$$(\text{mc_primFSLepton}[2] \cos \theta_b - \text{mc_primFSLepton}[1] \sin \theta_b)/1000$$

where θ_b is the beam angle, 0.05887 radians. Values are in GeV.

Note: for data only, a bias in the beam angle is corrected for by shifting the angle in the $y - z$ plane by -1 milliradians and in the $x - z$ plane by -3 milliradians. These are then combined to give a corrected value for CCQEAntiNuTool_theta.mu.

Neutrino energy (quasi-elastic hypothesis)

Reconstructed:

CCQEAntiNuTool_E/1000

True:

truth_E_nu_tmk/1000

Note that the true value here is the value you would get if you reconstructed the neutrino energy using the quasi-elastic hypothesis, but from the true muon kinematics. This is not necessarily the same as the actual true neutrino energy, as the assumptions of the quasi-elastic hypothesis do not always hold. Value is in GeV.

Q_{QE}^2 (quasi-elastic hypothesis)

Reconstructed:

CCQEAntiNuTool_Q2/1000000.

Reconstructed:

truth_Q2_tmk/1000000.

Note that the true value here is the value you would get if you reconstructed Q_{QE}^2 using the quasi-elastic hypothesis, but from the true muon kinematics. This is not necessarily

the same as the actual true Q^2 , as the assumptions of the quasi-elastic hypothesis do not always hold. Value is in GeV^2 .

C.2. Reconstructed variables used in central value analysis

C.2.1. Variables used for cuts

phys_n_dead_discr_pair_upstream_prim_track_proj: used for the dead-time cut - must be less than 2.

n_extra_outgoing_tracks: Used for track cut - must be 0.

CCQEAntiNuTool_charge_mu: Muon charge

CCQEAntiNuTool_vtx: Array corresponding to vertex (x,y,z) position in the detector in mm

CCQEAntiNuTool_is_plausible: True if event could be reconstructed as a charged-current event

CCQEAntiNuTool_vtx_fiducial: True if event is reconstructed in the fiducial volume

CCQEAntiNuTool_blob_dispersed_energy_nohcal: Recoil energy (not in isolated energy blobs) - used in recoil cut

CCQEAntiNuTool_blob_isolated_energy_nohcal: Recoil energy (in isolated blobs) - used in recoil cut. Recoil is the sum of the dispersed and isolated energy.

C.2.2. Variables used for differential cross section

CCQEAntiNuTool_E: Reconstructed E_ν^{QE}

CCQEAntiNuTool_Q2: Reconstructed Q_{QE}^2

CCQEAntiNuTool_E_mu: Reconstructed muon energy

CCQEAntiNuTool_p_mu: Muon momentum

CCQEAntiNuTool_theta_mu: Muon angle

truth_genie_wgt_Rvn1pi: Used in simulation only to calculate a correct event weight

C.3. Variables used in systematic uncertainty variations (Monte Carlo only)

C.3.1. GENIE model uncertainties

truth_genie_wgt_MaCCQEshape: Modifies the shape of the CCQE signal cross section model for a shift in axial mass M_A

truth_genie_wgt_VecFFCCQEshape: Modifies the shape of the CCQE signal cross section model for a shift in vector form factor

truth_genie_wgt_NormCCQE: Modifies the normalization of the CCQE signal cross section model

truth_genie_wgt_CCQEPauliSupViaKF: Used in the Pauli suppression uncertainty for CCQE signal model

truth_genie_wgt_EtaNCEL: Used in systematic to vary η for neutral current elastic scattering background

truth_genie_wgt_MaNCEL: Modifies neutral current elastic scattering background for a shift in axial mass M_A

truth_genie_wgt_MaRES: Modifies resonant background distribution if resonant axial form factor changes

truth_genie_wgt_MvRES: Modifies resonant background distribution if resonant vector form factor changes

truth_genie_wgt_NormDISCC: Modifies deep-inelastic scattering cross section normalization

truth_genie_wgt_Rvn1pi: Modifies cross section for non-resonant inelastic background processes

truth_genie_wgt_Rvn2pi: Modifies cross section for non-resonant inelastic background processes

truth_genie_wgt_Rvp1pi: Modifies cross section for non-resonant inelastic background processes

truth_genie_wgt_Rvp2pi: Modifies cross section for non-resonant inelastic background processes

C.3.2. Recoil uncertainties

truth_neutron_kevinwgtvardown, truth_neutron_kevinwgtvarup: Used in the neutron path length systematic uncertainty

CCQEAntiNuTool_blob_dispersed_energy_nohcal_otherfrac: Used in other-particle (EM) response uncertainty

CCQEAntiNuTool_blob_dispersed_energy_nohcal_pionfrac: Used in pion response uncertainty

CCQEAntiNuTool_blob_dispersed_energy_nohcal_protonfrac: Used in proton response uncertainty

CCQEAntiNuTool_blob_dispersed_energy_nohcal_xtalkfrac: Used in cross-talk uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_otherfrac: Used in other-particle

(EM) response uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_constrained_pionfrac: Used in pion response uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_unconstrained_pionfrac: Used in pion response uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_low_protonfrac: Used in proton response uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_med_protonfrac: Used in proton response uncertainty

CCQEAntiNuTool_blob_isolated_energy_nohcal_high_protonfrac: Used in proton response uncertainty

C.3.3. Hadronization (final-state interaction) uncertainties

truth_genie_wgt_AGKYxF1pi: Used for the deep inelastic scattering hadronization model uncertainty

truth_genie_wgt_FrAbs_N: Used in nucleon absorption systematic (FSI)

truth_genie_wgt_FrAbs_pi: Used in pion absorption systematic (FSI)

truth_genie_wgt_FrCEX_N: Used in nucleon charge exchange systematic (FSI)

truth_genie_wgt_FrCEX_pi: Used in pion charge exchange systematic (FSI)

truth_genie_wgt_FrElas_N: Used in nucleon elastic scattering systematic (FSI)

truth_genie_wgt_FrElas_pi: Used in pion elastic scattering systematic (FSI)

truth_genie_wgt_FrInel_N: Used in nucleon inelastic scattering systematic (FSI)

truth_genie_wgt_FrInel_pi: Used in pion inelastic scattering systematic (FSI)

truth_genie_wgt_FrPiProd_N: Used in nucleon pion production scattering systematic (FSI)

truth_genie_wgt_FrPiProd_pi: Used in pion pion production scattering systematic (FSI)

truth_genie_wgt_MFP_N: Used in systematic for nucleon mean free path in the nucleus (FSI)

truth_genie_wgt_MFP_pi: Used in systematic for pion mean free path in the nucleus (FSI)

truth_genie_wgt_RDecBR1gamma: Modifies branching ratio for background resonant event decaying to photons

truth_genie_wgt_Theta_Delta2Npi: Modifies delta decay angle for resonant backgrounds.

C.3.4. Muon reconstruction uncertainties

CCQEAntiNuTool_p_mu_minos: Muon momentum in MINOS. Used in the tracking efficiency correction, where the MnvNormalizer tool computes a momentum-dependent efficiency.

CCQEAntiNuTool_sys_E_minerva_energy_scale_(down/up): Shifted neutrino energy used to calculate the MINERvA energy scale systematic uncertainty

CCQEAntiNuTool_sys_E_mu_minerva_energy_scale_(down/up): Shifted muon energy used to calculate the MINERvA energy scale systematic uncertainty

CCQEAntiNuTool_sys_E_theta_bias_(down/up): Shifted neutrino energy used to calculate the muon angle bias uncertainty

CCQEAntiNuTool_sys_E_theta_smear: Shifted neutrino energy used to calculate the muon angle smearing uncertainty

CCQEAntiNuTool_sys_Q2_minerva_energy_scale_(down/up): Shifted Q_{QE}^2 used to calculate the MINERvA energy scale systematic uncertainty

CCQEAntiNuTool_sys_Q2_theta_bias_(down/up): Shifted Q_{QE}^2 used to calculate the muon angle bias uncertainty

CCQEAntiNuTool_sys_Q2_theta_smear: Shifted Q_{QE}^2 used to calculate the muon angle smearing uncertainty

CCQEAntiNuTool_sys_biased_theta_mu_(down/up): Shifted muon angle for the muon angle bias uncertainty

CCQEAntiNuTool_sys_p_mu_minerva_energy_scale_(down/up): Shifted muon momentum used to calculate the MINERvA energy scale systematic uncertainty

CCQEAntiNuTool_sys_smeared_theta_mu: Muon angle shifted by a random value from a 1 sigma distribution, used to calculate the theta smear uncertainty

CCQEAntiNuTool_sys_smeared_vertex_(x,y,z): Event vertex position shifted by a random value from a 1 sigma distribution, used to calculate vertex smear uncertainty

C.4. True GENIE information (Monte Carlo only)

truth_is_fiducial: Used for the fiducial volume cut for the true distributions used in the efficiency/acceptance correction.

truth_qelike: Used for the signal cut for the QE-like signal definition, to identify events with a muon, nucleons and no other particles in the final state.

truth_E_nu_tmkn: Neutrino energy, constructed in the CCQE hypothesis from true

muon kinematics. In MeV.

truth_Q2_tmK: Q_{QE}^2 , constructed in the CCQE hypothesis from true muon kinematics. In MeV^2 .

truth_muon_theta_tmK: :True muon angle - equivalent to the long complicated formula we are using

truth_dist_to_plane_center: Distance in mm from the centre of the plane, in the z direction

mc_intType: Interaction type: 1=elastic (signal for the true CCQE definition); 2=resonant, 3=DIS

mc_current: 1 for charged-current events, 0 for neutral current. Must be 1 for this analysis' signal.

mc_charm: Charm events are not part of our signal - we only want events where this is false

mc_incoming: PDF code of incoming particle type. 14 is a muon-neutrino; -14 is a muon-antineutrino. Used to select appropriate signal events

mc_incomingE: Generated incoming neutrino energy - actual, not calculated in CCQE hypothesis

mc_Q2: Generated interaction Q^2 (not Q_{QE}^2 calculated in CCQE hypothesis)

mc_vtx: Vector corresponding to vertex (x,y,z) position

mc_primFSLepton: Vector corresponding to outgoing muon 4-momentum ($p_{x\mu}$, $p_{y\mu}$, $p_{z\mu}$, E_μ)

C.5. Standard MINERvA information

ev_run: MINERvA run number in which this event took place (data)

ev_subrun: MINERvA subrun number (within the run) in which this event took place (data)

ev_gate: MINERvA gate number (within the subrun) in which this event took place (data)

mc_run: Generated MINERvA run number in which this event was simulated (simulation)

mc_subrun: Generated MINERvA subrun number in which this event was simulated (simulation)

mc_nthEvtInFile: Position in simulation file (simulation)

APPENDIX D

Validating the unfolding method

As we do not know the true distribution of events in our data (that is what we are trying to measure!), we must use Monte Carlo to evaluate the ability of our unfolding method to reproduce a generated distribution from a reconstructed one. To validate our unfolding, we used our reconstructed signal Monte Carlo to generate a migration matrix, and used this to unfold various Monte Carlo subsamples.

By definition, if there is no bug in the unfolding procedure, unfolding the reconstructed distribution used to generate the migration matrix should exactly reproduce the true distribution used to generate the matrix, with a single iteration of Bayesian unfolding. It was confirmed that this was the case, giving a discrepancy of less than one part in 10^{-10} in each bin.

D.1. Optimizing number of iterations

Ten subsamples were generated by dividing up the Monte Carlo into ten based on the gate number in the file

$$(D.1) \quad \text{sample number} = (\text{event number}/2)\%10$$

Each of these subsamples has roughly the same statistics as our data. The results of unfolding the ten subsamples were recorded for up to nine iterations of Bayesian unfolding. For each subsample, in each bin, the pull between the unfolded and true event counts

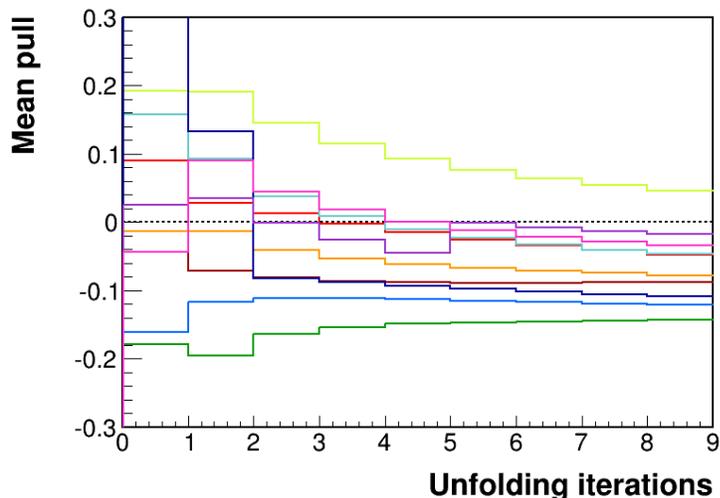


Figure D.1. Average pull of the 66 bins in 10 unfolded Monte Carlo sub-samples, vs number of Bayesian unfolding iterations

was calculated (that is unfolded value minus true value, divided by the uncertainty on the unfolded value). As statistical uncertainty increases with number of unfolding iterations, the goal is to select the smallest number of iterations for which the pulls converge to a small value. To study this, the mean pull of the 66 non-overflow and -underflow bins was plotted vs the number of iterations, for each of the 10 samples. The results of this are shown in figure D.1.

Based on this plot, we conclude that the pulls have stabilized by four iterations. The mean pull in each bin after four iterations is shown in figure D.2.

D.2. Ensemble test

D.2.1. Test of statistical uncertainties

For each bin, the number of events reconstructed in each of the ten samples was histogrammed and fitted to a Gaussian. The width of the Gaussian should be equal (within

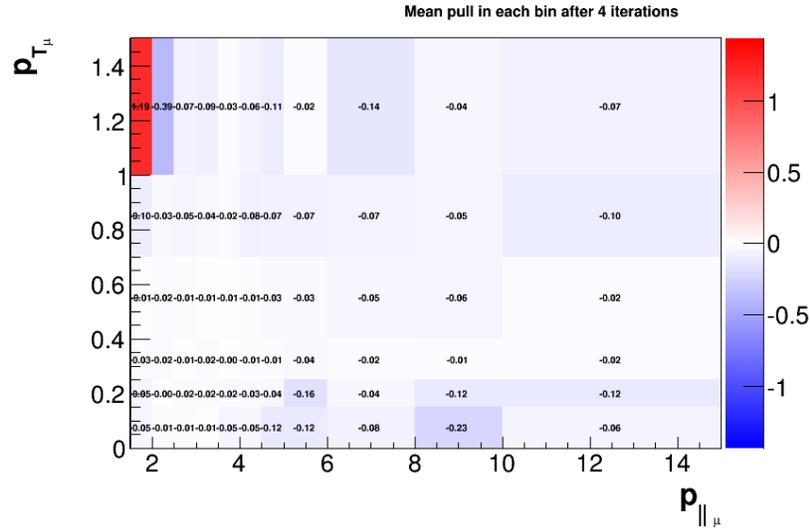


Figure D.2. Pull averaged over the 10 unfolded Monte Carlo sub-samples, after four iterations. Note that the bins with the largest pulls correspond to those with very low acceptance (for example, at high p_T and low p_{\parallel} , where the high-angle trajectory makes MINOS matching close to impossible)

errors) to the average statistical uncertainty on the number of events in the bin for the individual samples. To confirm this, the pull of each was calculated: the difference between average statistical uncertainty and the fit width, in units of the combined uncertainty on the two measurements (added in quadrature). The results of this test for reconstructed and true distributions are shown in figure D.3

D.2.2. Test of unfolding

Each of these samples was unfolded using four iterations of Bayesian unfolding, based on the migration matrix generated from the complete sample. To check the statistical uncertainties reported by the unfolding procedure, the same procedure was performed on the results of the unfolding, as shown in figure D.4.

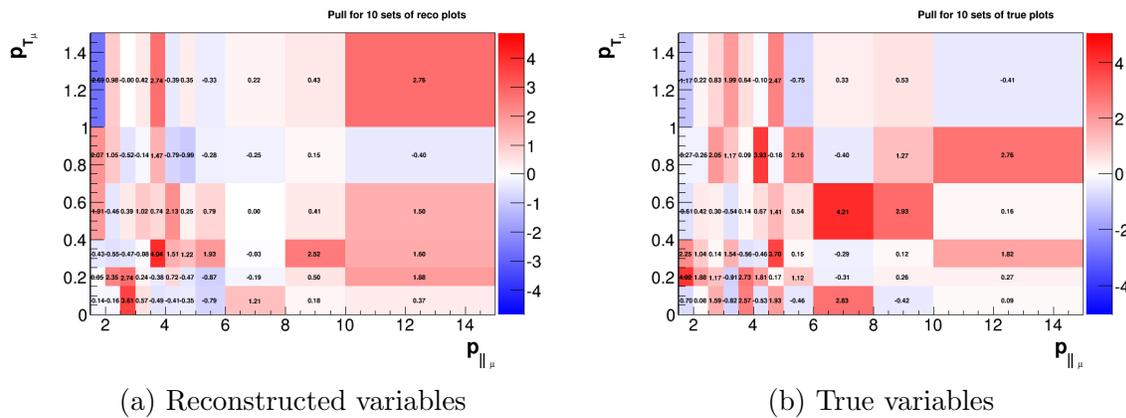


Figure D.3. Pull between the average statistical uncertainty of 10 Monte Carlo sub-samples, and the width of a Gaussian fit to a histogram of the number of events in each sample

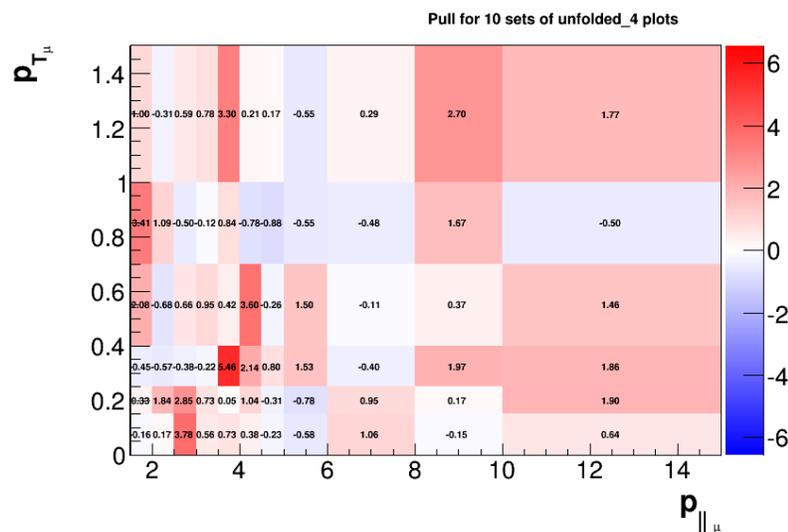


Figure D.4. Pull between the average statistical uncertainty of 10 unfolded Monte Carlo sub-samples, as reported by the unfolding procedure, and the width of a Gaussian fit to a histogram of the number of events in each sample

To test whether the unfolding procedure reproduced the true values, the pull between the average unfolded and true values was calculated for each bin. In each bin, the ten values corresponding to the ten datasets were histogrammed and fitted to a Gaussian to

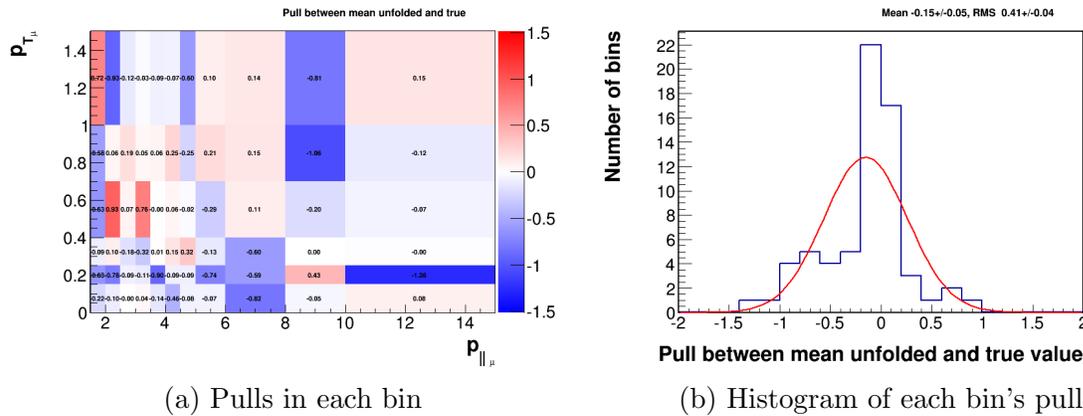


Figure D.5. Pull between mean unfolded and true values in each bin.

obtain a mean and RMS for that bin. This procedure was performed on both the true and unfolded histograms. The pull was calculated for each bin by measuring the difference between the mean true value and the mean unfolded value, divided by the combined uncertainty on these two means (added in quadrature). Figure D.5 shows these pulls. The pull values in each bin are histogrammed (figure D.5). We expect a mean value of 0 and a width of 1. The plot shows that there is an average negative pull of 14%, however this is not significant enough to warrant assigning a correction or systematic uncertainty to the procedure.

To test the uncertainties reported by the unfolding procedure, we compare the RMS of the values of the 10 unfolded subsamples with the uncertainty reported on the unfolded values. Figure D.6 shows the ratio of the RMS test to the error given by the unfolding procedure.

Within errors, we expect this ratio to be 1; the histogram confirms this to be the case, validating the uncertainties returned by the unfolding.

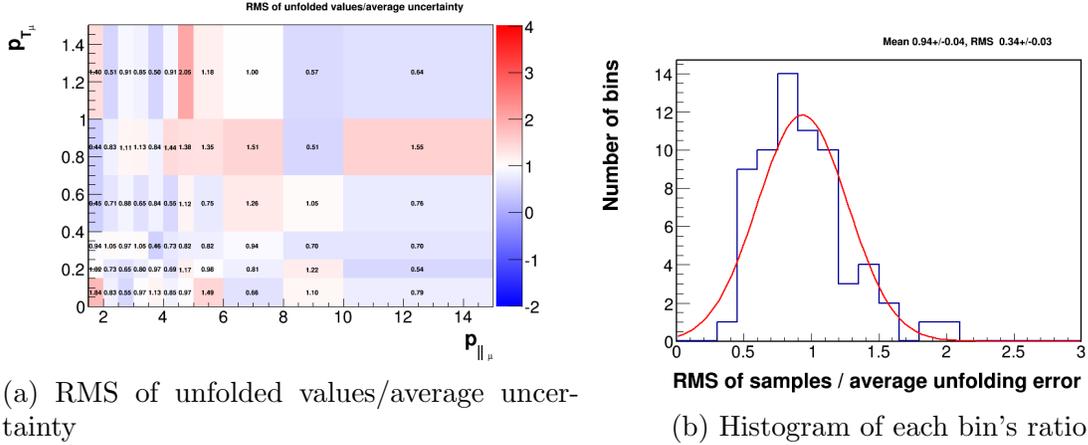
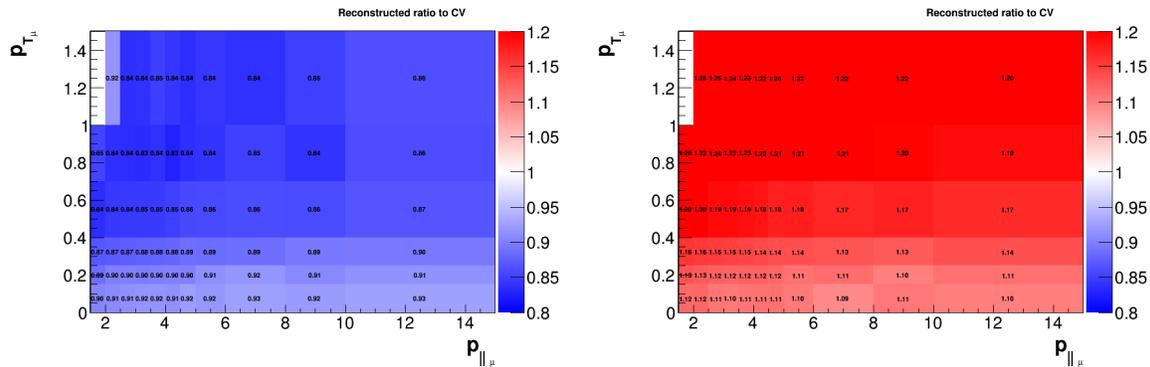


Figure D.6. Ratio of the RMS of the 10 unfolded values to the average uncertainty reported by the unfolding procedure

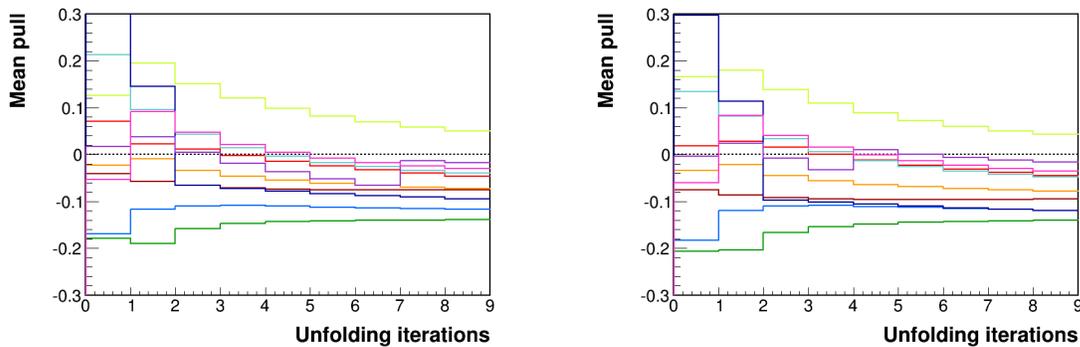
D.3. Warped Monte Carlo unfolding test

Each of the ten subsamples used in the ensemble test corresponded to approximately one tenth of the original Monte Carlo sample, and was made up of events from that sample; thus the reconstructed and true event distributions in each sample were close to being proportional to the distribution used to generate the unsmearing matrix. Our data distribution, while similar, does not follow the exact shape of the Monte Carlo. To check that the unfolding procedure is able to unfold a distribution that does not exactly mimic that of the original simulation, two sets of ten “warped” samples were generated. In each case, weights were applied to each event, corresponding to changes in the quasi-elastic model. For the first sample, the events were weighted to simulate a distribution in which the value of the axial mass M_A was reduced by 1σ , while also reducing the normalization for all CCQE events by 1σ . For the second sample, both M_A and the CCQE event normalization were increased by 1σ . More information about the effect of these changes can be seen in the chapter describing systematic uncertainties. Figure D.7 shows the



(a) Subsample 8, warped with M_A decreased (b) Subsample 3, warped with M_A increased

Figure D.7. Ratio of the warped to the central value for two of the warped subsamples



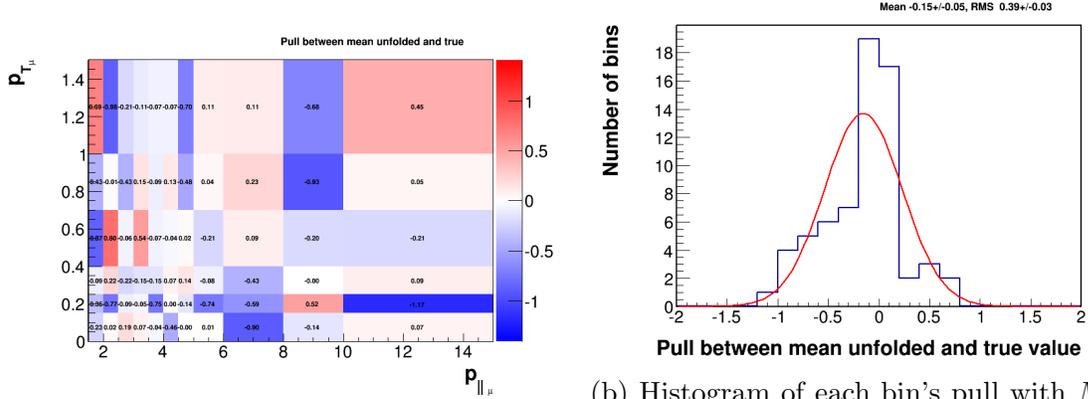
(a) Warped with M_A decreased

(b) Warped with M_A increased

Figure D.8. Average pull between unfolded and true values of the warped samples, vs number of iterations

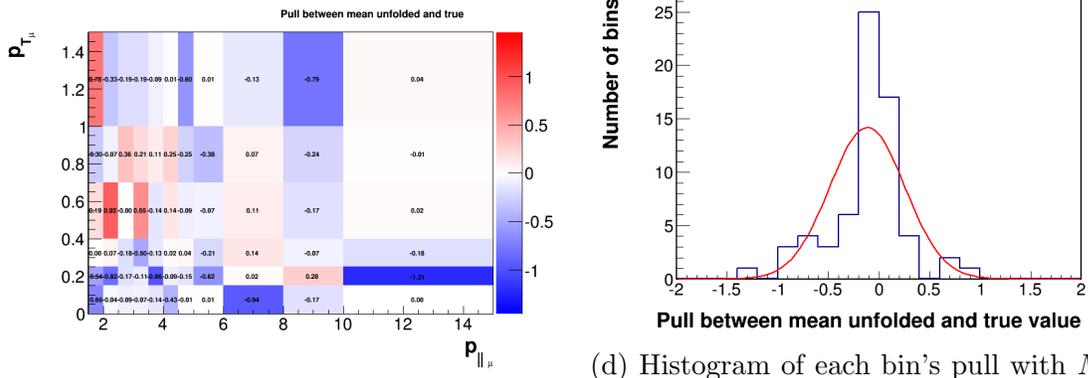
effect of warping on two of our samples, by taking a ratio of the warped to the central value for the subsample. As can be seen, the effect is around 15-20%, and varies over the phase space.

The same procedure as was used in the ensemble test was also performed on these warped samples. As with the unwarped samples, the unfolding procedure stabilizes by four iterations.



(a) Pulls in each bin with M_A reduced

(b) Histogram of each bin's pull with M_A reduced



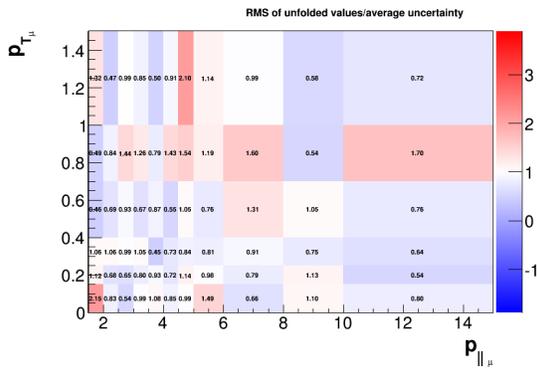
(c) Pulls in each bin with M_A increased

(d) Histogram of each bin's pull with M_A increased

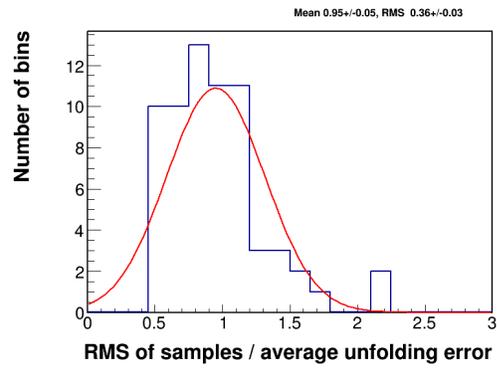
Figure D.9. Pull between mean unfolded and true values of the warped samples in each bin.

The pull between the mean unfolded and true values of the ten samples is shown in figure D.9

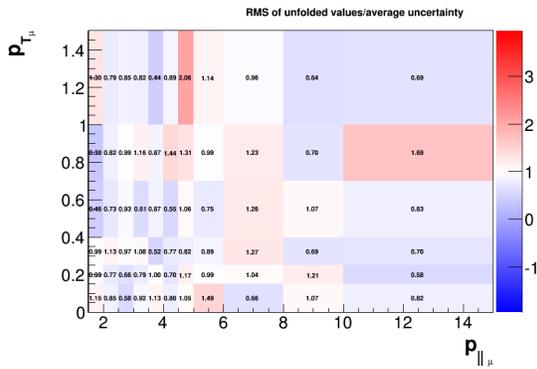
The mean of the 66 bins' average pulls is -15% for both samples; almost exactly the same as for the unwarped samples. We also check the error reported by the unfolding procedure, as show in D.10. Again, we confirm that these distributions are centered at 1, verifying the unfolding uncertainties.



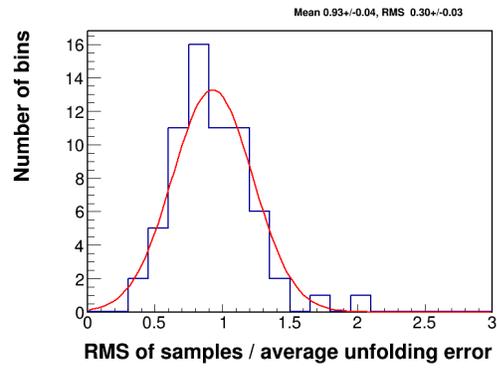
(a) Ratio each bin with M_A reduced



(b) Ratios in each bin with M_A reduced



(c) Ratio in each bin with M_A increased



(d) Ratios in each bin with M_A increased

Figure D.10. Ratio of the RMS of the 10 unfolded values in each bin to the average uncertainty reported by the unfolding procedure

APPENDIX E

Tables of systematic uncertainties

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Flux
< 1.5	< 0	0.0%
< 1.5	0 - 0.15	8.3%
< 1.5	0.15 - 0.25	7.9%
< 1.5	0.25 - 0.4	8.1%
< 1.5	0.4 - 0.7	7.7%
< 1.5	0.7 - 1	7.5%
< 1.5	1 - 1.5	0.0%
< 1.5	> 1.5	0.0%
1.5 - 2	< 0	0.0%
1.5 - 2	0 - 0.15	7.8%
1.5 - 2	0.15 - 0.25	7.8%
1.5 - 2	0.25 - 0.4	7.8%
1.5 - 2	0.4 - 0.7	7.7%
1.5 - 2	0.7 - 1	7.8%
1.5 - 2	1 - 1.5	8.1%
1.5 - 2	> 1.5	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Flux
2 - 2.5	< 0	0.0%
2 - 2.5	0 - 0.15	7.8%
2 - 2.5	0.15 - 0.25	7.8%
2 - 2.5	0.25 - 0.4	7.8%
2 - 2.5	0.4 - 0.7	7.7%
2 - 2.5	0.7 - 1	7.7%
2 - 2.5	1 - 1.5	7.6%
2 - 2.5	> 1.5	0.0%
2.5 - 3	< 0	0.0%
2.5 - 3	0 - 0.15	7.8%
2.5 - 3	0.15 - 0.25	7.7%
2.5 - 3	0.25 - 0.4	7.8%
2.5 - 3	0.4 - 0.7	7.7%
2.5 - 3	0.7 - 1	7.7%
2.5 - 3	1 - 1.5	7.7%
2.5 - 3	> 1.5	0.0%
3 - 3.5	< 0	0.0%
3 - 3.5	0 - 0.15	7.8%
3 - 3.5	0.15 - 0.25	7.7%
3 - 3.5	0.25 - 0.4	7.8%
3 - 3.5	0.4 - 0.7	7.7%
3 - 3.5	0.7 - 1	7.7%
3 - 3.5	1 - 1.5	7.6%
3 - 3.5	> 1.5	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Flux
3.5 - 4	< 0	0.0%
3.5 - 4	0 - 0.15	7.8%
3.5 - 4	0.15 - 0.25	7.7%
3.5 - 4	0.25 - 0.4	7.8%
3.5 - 4	0.4 - 0.7	7.7%
3.5 - 4	0.7 - 1	7.7%
3.5 - 4	1 - 1.5	7.9%
3.5 - 4	> 1.5	0.0%
4 - 4.5	< 0	0.0%
4 - 4.5	0 - 0.15	7.8%
4 - 4.5	0.15 - 0.25	7.7%
4 - 4.5	0.25 - 0.4	7.8%
4 - 4.5	0.4 - 0.7	7.8%
4 - 4.5	0.7 - 1	7.9%
4 - 4.5	1 - 1.5	8.3%
4 - 4.5	> 1.5	12.0%
4.5 - 5	< 0	0.0%
4.5 - 5	0 - 0.15	7.7%
4.5 - 5	0.15 - 0.25	7.8%
4.5 - 5	0.25 - 0.4	7.8%
4.5 - 5	0.4 - 0.7	7.8%
4.5 - 5	0.7 - 1	7.9%
4.5 - 5	1 - 1.5	8.4%
4.5 - 5	> 1.5	8.5%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Flux
5 - 6	< 0	0.0%
5 - 6	0 - 0.15	7.8%
5 - 6	0.15 - 0.25	7.7%
5 - 6	0.25 - 0.4	7.8%
5 - 6	0.4 - 0.7	7.8%
5 - 6	0.7 - 1	7.9%
5 - 6	1 - 1.5	8.0%
5 - 6	> 1.5	9.8%
6 - 8	< 0	0.0%
6 - 8	0 - 0.15	7.9%
6 - 8	0.15 - 0.25	7.9%
6 - 8	0.25 - 0.4	7.9%
6 - 8	0.4 - 0.7	7.8%
6 - 8	0.7 - 1	7.9%
6 - 8	1 - 1.5	8.1%
6 - 8	> 1.5	8.3%
8 - 10	< 0	0.0%
8 - 10	0 - 0.15	8.4%
8 - 10	0.15 - 0.25	7.9%
8 - 10	0.25 - 0.4	7.4%
8 - 10	0.4 - 0.7	8.0%
8 - 10	0.7 - 1	8.0%
8 - 10	1 - 1.5	7.8%
8 - 10	> 1.5	7.7%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Flux
10 - 15	< 0	0.0%
10 - 15	0 - 0.15	9.0%
10 - 15	0.15 - 0.25	6.9%
10 - 15	0.25 - 0.4	3.5%
10 - 15	0.4 - 0.7	8.0%
10 - 15	0.7 - 1	8.2%
10 - 15	1 - 1.5	7.7%
10 - 15	> 1.5	1.5%
> 15	< 0	0.0%
> 15	0 - 0.15	161.8%
> 15	0.15 - 0.25	250.6%
> 15	0.25 - 0.4	154.6%
> 15	0.4 - 0.7	1651.9%
> 15	0.7 - 1	608.8%
> 15	1 - 1.5	269.0%
> 15	> 1.5	0.0%

Table E.1. Absolute fractional flux uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	E_{μ} scale	θ bias	θ smear	Vtx smear	Tracking eff.	Mass model
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	54.8%	1.6%	0.6%	0.0%	1.2%	1.4%
< 1.5	0.15 - 0.25	57.8%	0.4%	1.0%	1.0%	1.2%	1.4%
< 1.5	0.25 - 0.4	38.3%	0.6%	1.4%	0.7%	1.2%	1.4%
< 1.5	0.4 - 0.7	31.2%	1.5%	0.3%	0.3%	1.2%	1.4%
< 1.5	0.7 - 1	12.5%	2.7%	2.4%	0.1%	1.2%	1.4%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	11.0%	0.0%	0.1%	0.3%	1.2%	1.4%
1.5 - 2	0.15 - 0.25	11.3%	0.6%	0.1%	0.1%	1.2%	1.4%
1.5 - 2	0.25 - 0.4	9.7%	0.3%	0.4%	0.1%	1.2%	1.4%
1.5 - 2	0.4 - 0.7	7.5%	0.2%	0.3%	0.1%	1.2%	1.4%
1.5 - 2	0.7 - 1	4.6%	0.2%	0.6%	0.5%	1.2%	1.4%
1.5 - 2	1 - 1.5	7.4%	1.0%	5.0%	0.7%	1.0%	1.4%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	5.5%	0.4%	0.1%	0.0%	1.3%	1.4%
2 - 2.5	0.15 - 0.25	6.9%	0.6%	0.6%	0.2%	1.2%	1.4%
2 - 2.5	0.25 - 0.4	6.4%	0.0%	0.2%	0.0%	1.2%	1.4%
2 - 2.5	0.4 - 0.7	3.5%	0.2%	0.3%	0.1%	1.2%	1.4%
2 - 2.5	0.7 - 1	4.3%	0.8%	0.1%	0.1%	1.2%	1.4%
2 - 2.5	1 - 1.5	4.6%	3.6%	1.9%	0.2%	1.2%	1.4%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	E_{μ} scale	θ bias	θ smear	Vtx smear	Tracking eff.	Mass model
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	4.8%	0.6%	0.0%	0.2%	1.2%	1.4%
2.5 - 3	0.15 - 0.25	6.7%	0.2%	0.1%	0.0%	1.2%	1.4%
2.5 - 3	0.25 - 0.4	4.7%	0.0%	0.3%	0.1%	1.2%	1.4%
2.5 - 3	0.4 - 0.7	1.3%	0.0%	0.1%	0.1%	1.2%	1.4%
2.5 - 3	0.7 - 1	4.9%	0.0%	0.2%	0.0%	1.2%	1.4%
2.5 - 3	1 - 1.5	13.3%	0.6%	1.6%	0.1%	1.2%	1.4%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	5.1%	0.2%	0.8%	0.1%	1.2%	1.4%
3 - 3.5	0.15 - 0.25	1.1%	0.5%	0.0%	0.0%	1.2%	1.4%
3 - 3.5	0.25 - 0.4	0.9%	0.1%	0.3%	0.1%	1.2%	1.4%
3 - 3.5	0.4 - 0.7	3.7%	0.1%	0.1%	0.1%	1.2%	1.4%
3 - 3.5	0.7 - 1	8.0%	0.4%	0.2%	0.1%	1.2%	1.4%
3 - 3.5	1 - 1.5	11.0%	0.2%	2.0%	0.3%	1.2%	1.4%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	4.2%	0.1%	1.2%	0.2%	1.2%	1.4%
3.5 - 4	0.15 - 0.25	0.9%	0.4%	0.7%	0.1%	1.2%	1.4%
3.5 - 4	0.25 - 0.4	1.2%	0.5%	0.2%	0.0%	1.2%	1.4%
3.5 - 4	0.4 - 0.7	7.6%	0.2%	0.1%	0.1%	1.1%	1.4%
3.5 - 4	0.7 - 1	12.2%	0.0%	0.2%	0.1%	1.1%	1.4%
3.5 - 4	1 - 1.5	15.8%	0.3%	0.6%	1.1%	1.0%	1.4%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	E_{μ} scale	θ bias	θ smear	Vtx smear	Tracking eff.	Mass model
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	9.4%	1.3%	0.8%	0.7%	0.8%	1.4%
4 - 4.5	0.15 - 0.25	4.9%	0.8%	0.0%	0.0%	0.8%	1.4%
4 - 4.5	0.25 - 0.4	6.0%	0.2%	0.3%	0.1%	0.8%	1.4%
4 - 4.5	0.4 - 0.7	11.0%	0.3%	0.5%	0.2%	0.8%	1.4%
4 - 4.5	0.7 - 1	11.6%	0.2%	0.0%	0.2%	0.7%	1.4%
4 - 4.5	1 - 1.5	19.3%	0.8%	0.7%	0.5%	0.6%	1.4%
4 - 4.5	> 1.5	2.7%	5.5%	11.0%	4.8%	0.5%	1.4%
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	5.1%	1.8%	1.0%	0.5%	0.4%	1.4%
4.5 - 5	0.15 - 0.25	2.1%	0.3%	0.0%	0.3%	0.4%	1.4%
4.5 - 5	0.25 - 0.4	4.3%	0.6%	0.0%	0.0%	0.4%	1.4%
4.5 - 5	0.4 - 0.7	5.9%	0.2%	0.6%	0.1%	0.4%	1.4%
4.5 - 5	0.7 - 1	6.8%	0.1%	0.4%	0.4%	0.4%	1.4%
4.5 - 5	1 - 1.5	17.3%	1.2%	0.9%	0.1%	0.4%	1.4%
4.5 - 5	> 1.5	21.5%	3.0%	2.5%	0.2%	0.4%	1.4%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	4.0%	1.4%	0.2%	0.1%	0.4%	1.4%
5 - 6	0.15 - 0.25	2.0%	0.5%	0.4%	0.1%	0.4%	1.4%
5 - 6	0.25 - 0.4	2.8%	0.7%	0.1%	0.1%	0.4%	1.4%
5 - 6	0.4 - 0.7	2.9%	0.6%	0.1%	0.0%	0.4%	1.4%
5 - 6	0.7 - 1	8.1%	0.8%	0.5%	0.1%	0.4%	1.4%
5 - 6	1 - 1.5	15.9%	0.5%	3.2%	0.1%	0.4%	1.4%
5 - 6	> 1.5	0.7%	1.5%	1.0%	2.6%	0.4%	1.4%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	E_{μ} scale	θ bias	θ smear	Vtx smear	Tracking eff.	Mass model
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	2.7%	0.3%	2.9%	0.3%	0.3%	1.4%
6 - 8	0.15 - 0.25	1.6%	0.3%	1.6%	0.2%	0.4%	1.4%
6 - 8	0.25 - 0.4	1.6%	0.2%	0.8%	0.1%	0.4%	1.4%
6 - 8	0.4 - 0.7	4.3%	0.0%	0.4%	0.0%	0.4%	1.4%
6 - 8	0.7 - 1	8.8%	0.4%	0.7%	0.2%	0.4%	1.4%
6 - 8	1 - 1.5	12.0%	1.9%	1.9%	0.2%	0.4%	1.4%
6 - 8	> 1.5	36.1%	9.7%	21.1%	1.3%	0.5%	1.4%
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.6%	3.1%	5.9%	1.2%	0.3%	1.4%
8 - 10	0.15 - 0.25	3.5%	0.1%	3.5%	0.4%	0.3%	1.4%
8 - 10	0.25 - 0.4	0.6%	0.6%	0.2%	0.3%	0.4%	1.4%
8 - 10	0.4 - 0.7	1.5%	0.0%	1.0%	0.1%	0.4%	1.4%
8 - 10	0.7 - 1	7.0%	1.4%	2.2%	0.5%	0.4%	1.4%
8 - 10	1 - 1.5	10.6%	0.7%	2.0%	0.1%	0.4%	1.4%
8 - 10	> 1.5	27.2%	4.6%	9.6%	0.3%	0.4%	1.4%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	1.9%	0.4%	2.6%	0.0%	0.4%	1.4%
10 - 15	0.15 - 0.25	0.2%	0.4%	2.9%	0.2%	0.3%	1.4%
10 - 15	0.25 - 0.4	0.5%	0.3%	0.5%	0.4%	0.4%	1.4%
10 - 15	0.4 - 0.7	3.2%	0.5%	0.6%	0.0%	0.4%	1.4%
10 - 15	0.7 - 1	9.1%	2.1%	2.3%	0.3%	0.4%	1.4%
10 - 15	1 - 1.5	1.1%	14.4%	31.2%	0.6%	0.4%	1.4%
10 - 15	> 1.5	130.5%	78.4%	159.9%	0.7%	0.5%	1.4%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	NormDISCC	Rvn1pi	Rvn2pi	Rvp1pi	Rvp2pi	GENIE variations
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.0%	0.0%	0.1%	0.2%	0.2%	7.5%
1.5 - 2	0.15 - 0.25	0.0%	0.0%	0.2%	0.0%	0.2%	1.0%
1.5 - 2	0.25 - 0.4	0.0%	0.1%	0.1%	0.0%	0.2%	2.0%
1.5 - 2	0.4 - 0.7	0.0%	0.1%	0.4%	0.1%	0.3%	3.1%
1.5 - 2	0.7 - 1	0.0%	0.3%	0.2%	0.5%	0.9%	13.6%
1.5 - 2	1 - 1.5	0.0%	0.4%	0.3%	0.0%	0.8%	8.4%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.0%	0.0%	0.1%	0.1%	0.1%	0.5%
2 - 2.5	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.2%	1.7%
2 - 2.5	0.25 - 0.4	0.0%	0.0%	0.1%	0.0%	0.0%	0.8%
2 - 2.5	0.4 - 0.7	0.0%	0.1%	0.3%	0.0%	0.3%	1.4%
2 - 2.5	0.7 - 1	0.0%	0.2%	0.6%	0.1%	0.4%	3.1%
2 - 2.5	1 - 1.5	0.0%	0.2%	0.5%	0.3%	0.8%	22.6%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.0%	0.0%	0.0%	0.1%	0.1%	2.4%
2.5 - 3	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.1%	1.2%
2.5 - 3	0.25 - 0.4	0.0%	0.0%	0.1%	0.0%	0.1%	1.3%
2.5 - 3	0.4 - 0.7	0.0%	0.1%	0.3%	0.0%	0.3%	0.1%
2.5 - 3	0.7 - 1	0.0%	0.2%	0.5%	0.1%	0.5%	2.5%
2.5 - 3	1 - 1.5	0.0%	0.3%	0.8%	0.2%	0.3%	17.0%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	NormDISCC	Rvn1pi	Rvn2pi	Rvp1pi	Rvp2pi	GENIE variations
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.0%	0.0%	0.0%	0.1%	0.1%	0.8%
3 - 3.5	0.15 - 0.25	0.0%	0.0%	0.1%	0.0%	0.1%	2.4%
3 - 3.5	0.25 - 0.4	0.0%	0.0%	0.1%	0.1%	0.1%	0.5%
3 - 3.5	0.4 - 0.7	0.0%	0.1%	0.2%	0.0%	0.2%	1.0%
3 - 3.5	0.7 - 1	0.0%	0.2%	0.4%	0.0%	0.3%	4.5%
3 - 3.5	1 - 1.5	0.0%	0.3%	0.8%	0.4%	0.8%	22.7%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.0%	0.0%	0.0%	0.1%	0.1%	2.8%
3.5 - 4	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.0%	1.2%
3.5 - 4	0.25 - 0.4	0.0%	0.0%	0.1%	0.0%	0.1%	1.1%
3.5 - 4	0.4 - 0.7	0.0%	0.1%	0.2%	0.0%	0.2%	0.0%
3.5 - 4	0.7 - 1	0.0%	0.2%	0.3%	0.1%	0.2%	0.2%
3.5 - 4	1 - 1.5	0.0%	0.4%	0.7%	0.7%	0.8%	6.7%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.0%	0.0%	0.0%	0.0%	0.1%	6.1%
4 - 4.5	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.1%	2.2%
4 - 4.5	0.25 - 0.4	0.0%	0.0%	0.1%	0.1%	0.1%	2.4%
4 - 4.5	0.4 - 0.7	0.0%	0.1%	0.2%	0.0%	0.2%	0.3%
4 - 4.5	0.7 - 1	0.0%	0.2%	0.3%	0.1%	0.6%	1.5%
4 - 4.5	1 - 1.5	0.0%	0.5%	0.8%	0.5%	1.3%	4.3%
4 - 4.5	> 1.5	0.0%	1.3%	0.3%	1.5%	3.3%	37.6%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	NormDISCC	Rvn1pi	Rvn2pi	Rvp1pi	Rvp2pi	GENIE variations
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.0%	0.0%	0.0%	0.1%	0.1%	11.2%
4.5 - 5	0.15 - 0.25	0.0%	0.0%	0.1%	0.2%	0.2%	1.3%
4.5 - 5	0.25 - 0.4	0.0%	0.0%	0.1%	0.1%	0.2%	0.3%
4.5 - 5	0.4 - 0.7	0.0%	0.1%	0.3%	0.1%	0.2%	0.8%
4.5 - 5	0.7 - 1	0.0%	0.2%	0.6%	0.1%	0.7%	0.1%
4.5 - 5	1 - 1.5	0.0%	0.4%	0.4%	0.4%	0.8%	14.4%
4.5 - 5	> 1.5	0.0%	0.4%	0.1%	1.9%	0.2%	53.2%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.0%	0.0%	0.0%	0.1%	0.1%	10.1%
5 - 6	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.1%	0.2%
5 - 6	0.25 - 0.4	0.0%	0.0%	0.1%	0.1%	0.2%	1.8%
5 - 6	0.4 - 0.7	0.0%	0.1%	0.3%	0.2%	0.2%	0.1%
5 - 6	0.7 - 1	0.0%	0.2%	0.5%	0.0%	0.6%	2.1%
5 - 6	1 - 1.5	0.0%	0.4%	0.7%	0.1%	1.1%	9.8%
5 - 6	> 1.5	0.0%	0.6%	1.8%	0.0%	1.0%	68.2%
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	0.0%	0.0%	0.0%	0.4%	0.1%	0.3%
6 - 8	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.1%	5.2%
6 - 8	0.25 - 0.4	0.0%	0.0%	0.1%	0.2%	0.2%	4.1%
6 - 8	0.4 - 0.7	0.0%	0.1%	0.3%	0.0%	0.4%	0.1%
6 - 8	0.7 - 1	0.0%	0.2%	0.5%	0.3%	0.4%	4.9%
6 - 8	1 - 1.5	0.0%	0.5%	1.2%	0.2%	1.5%	8.4%
6 - 8	> 1.5	0.0%	1.1%	2.6%	1.7%	3.6%	4.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	NormDISCC	Rvn1pi	Rvn2pi	Rvp1pi	Rvp2pi	GENIE variations
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.0%	0.0%	0.0%	0.0%	0.1%	5.7%
8 - 10	0.15 - 0.25	0.0%	0.0%	0.1%	0.1%	0.2%	4.7%
8 - 10	0.25 - 0.4	0.0%	0.0%	0.1%	0.0%	0.2%	1.3%
8 - 10	0.4 - 0.7	0.0%	0.1%	0.2%	0.2%	0.5%	4.1%
8 - 10	0.7 - 1	0.0%	0.3%	0.7%	0.2%	1.0%	8.8%
8 - 10	1 - 1.5	0.0%	0.5%	1.2%	0.3%	1.7%	0.7%
8 - 10	> 1.5	0.0%	0.9%	2.1%	0.7%	3.8%	40.6%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.0%	0.0%	0.0%	0.2%	0.0%	7.8%
10 - 15	0.15 - 0.25	0.0%	0.0%	0.1%	0.2%	0.2%	1.6%
10 - 15	0.25 - 0.4	0.0%	0.0%	0.2%	0.4%	0.3%	14.4%
10 - 15	0.4 - 0.7	0.0%	0.1%	0.2%	0.2%	0.3%	2.5%
10 - 15	0.7 - 1	0.0%	0.3%	0.7%	0.1%	0.2%	0.7%
10 - 15	1 - 1.5	0.0%	0.6%	1.0%	1.2%	1.6%	6.4%
10 - 15	> 1.5	0.0%	2.2%	5.5%	4.1%	10.2%	177.3%
> 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	0.0%	0.0%	0.0%	0.5%	0.0%	5.6%
> 15	0.15 - 0.25	0.0%	0.0%	0.1%	0.3%	0.0%	56.4%
> 15	0.25 - 0.4	0.0%	0.0%	0.2%	0.8%	0.3%	31.1%
> 15	0.4 - 0.7	0.0%	0.2%	0.2%	0.1%	0.4%	53.0%
> 15	0.7 - 1	0.0%	0.1%	0.4%	2.1%	0.3%	14.8%
> 15	1 - 1.5	0.0%	0.4%	1.3%	1.0%	1.1%	14.6%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	NormDISCC	Rvn1pi	Rvn2pi	Rvp1pi	Rvp2pi	GENIE variations
--------------------------	------------------	-----------	--------	--------	--------	--------	------------------

Table E.3. Absolute fractional primary interaction uncertainties (inelastic) vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	EtaNCEL	MaNCEL	MaRES	MvRES	NormNCRES
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	0.0%	0.0%	1.3%	0.0%	0.0%
< 1.5	0.15 - 0.25	0.0%	0.0%	2.4%	0.1%	0.0%
< 1.5	0.25 - 0.4	0.0%	0.0%	2.7%	0.2%	0.0%
< 1.5	0.4 - 0.7	0.0%	0.0%	4.0%	0.4%	0.0%
< 1.5	0.7 - 1	0.0%	0.0%	4.3%	0.4%	0.0%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.0%	0.1%	0.4%	0.0%	0.0%
1.5 - 2	0.15 - 0.25	0.0%	0.0%	0.9%	0.0%	0.0%
1.5 - 2	0.25 - 0.4	0.0%	0.0%	1.5%	0.1%	0.0%
1.5 - 2	0.4 - 0.7	0.0%	0.0%	2.4%	0.4%	0.0%
1.5 - 2	0.7 - 1	0.0%	0.0%	3.6%	0.4%	0.0%
1.5 - 2	1 - 1.5	0.0%	0.0%	8.1%	1.0%	0.0%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	EtaNCEL	MaNCEL	MaRES	MvRES	NormNCRES
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.0%	0.0%	0.4%	0.0%	0.0%
2 - 2.5	0.15 - 0.25	0.0%	0.0%	0.7%	0.0%	0.0%
2 - 2.5	0.25 - 0.4	0.0%	0.0%	1.3%	0.1%	0.0%
2 - 2.5	0.4 - 0.7	0.0%	0.0%	1.8%	0.3%	0.0%
2 - 2.5	0.7 - 1	0.0%	0.0%	1.6%	0.1%	0.0%
2 - 2.5	1 - 1.5	0.0%	0.0%	5.6%	0.4%	0.0%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.0%	0.0%	0.3%	0.0%	0.0%
2.5 - 3	0.15 - 0.25	0.0%	0.0%	0.6%	0.0%	0.0%
2.5 - 3	0.25 - 0.4	0.0%	0.0%	1.0%	0.1%	0.0%
2.5 - 3	0.4 - 0.7	0.0%	0.0%	1.4%	0.2%	0.0%
2.5 - 3	0.7 - 1	0.0%	0.0%	0.9%	0.0%	0.0%
2.5 - 3	1 - 1.5	0.0%	0.0%	0.0%	0.3%	0.0%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.0%	0.0%	0.2%	0.0%	0.0%
3 - 3.5	0.15 - 0.25	0.0%	0.0%	0.5%	0.0%	0.0%
3 - 3.5	0.25 - 0.4	0.0%	0.0%	0.7%	0.1%	0.0%
3 - 3.5	0.4 - 0.7	0.0%	0.0%	1.1%	0.2%	0.0%
3 - 3.5	0.7 - 1	0.0%	0.0%	1.1%	0.1%	0.0%
3 - 3.5	1 - 1.5	0.0%	0.0%	1.6%	0.0%	0.0%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	EtaNCEL	MaNCEL	MaRES	MvRES	NormNCRES
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.0%	0.0%	0.1%	0.0%	0.0%
3.5 - 4	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
3.5 - 4	0.25 - 0.4	0.0%	0.0%	0.6%	0.0%	0.0%
3.5 - 4	0.4 - 0.7	0.0%	0.0%	0.9%	0.1%	0.0%
3.5 - 4	0.7 - 1	0.0%	0.0%	0.8%	0.0%	0.0%
3.5 - 4	1 - 1.5	0.0%	0.0%	4.2%	0.4%	0.0%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.0%	0.0%	0.1%	0.0%	0.0%
4 - 4.5	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
4 - 4.5	0.25 - 0.4	0.0%	0.0%	0.5%	0.0%	0.0%
4 - 4.5	0.4 - 0.7	0.0%	0.0%	0.9%	0.2%	0.0%
4 - 4.5	0.7 - 1	0.0%	0.0%	0.7%	0.1%	0.0%
4 - 4.5	1 - 1.5	0.0%	0.0%	2.8%	0.1%	0.0%
4 - 4.5	> 1.5	0.0%	0.0%	28.1%	4.2%	0.0%
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.0%	0.0%	0.1%	0.0%	0.0%
4.5 - 5	0.15 - 0.25	0.0%	0.0%	0.5%	0.0%	0.0%
4.5 - 5	0.25 - 0.4	0.0%	0.0%	0.7%	0.1%	0.0%
4.5 - 5	0.4 - 0.7	0.0%	0.0%	1.1%	0.2%	0.0%
4.5 - 5	0.7 - 1	0.0%	0.0%	0.7%	0.1%	0.0%
4.5 - 5	1 - 1.5	0.0%	0.0%	4.4%	0.5%	0.0%
4.5 - 5	> 1.5	0.0%	0.0%	8.4%	0.7%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	EtaNCEL	MaNCEL	MaRES	MvRES	NormNCRES
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
5 - 6	0.25 - 0.4	0.0%	0.0%	0.6%	0.1%	0.0%
5 - 6	0.4 - 0.7	0.0%	0.0%	1.0%	0.2%	0.0%
5 - 6	0.7 - 1	0.0%	0.0%	0.6%	0.0%	0.0%
5 - 6	1 - 1.5	0.0%	0.0%	4.2%	0.6%	0.0%
5 - 6	> 1.5	0.0%	0.0%	10.8%	1.6%	0.0%
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
6 - 8	0.25 - 0.4	0.0%	0.0%	0.9%	0.1%	0.0%
6 - 8	0.4 - 0.7	0.0%	0.0%	1.3%	0.2%	0.0%
6 - 8	0.7 - 1	0.0%	0.0%	1.9%	0.2%	0.0%
6 - 8	1 - 1.5	0.0%	0.0%	5.6%	0.7%	0.0%
6 - 8	> 1.5	0.0%	0.0%	25.3%	3.5%	0.0%
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.0%	0.0%	0.1%	0.0%	0.0%
8 - 10	0.15 - 0.25	0.0%	0.0%	0.6%	0.0%	0.0%
8 - 10	0.25 - 0.4	0.0%	0.0%	1.1%	0.1%	0.0%
8 - 10	0.4 - 0.7	0.0%	0.0%	1.3%	0.2%	0.0%
8 - 10	0.7 - 1	0.0%	0.0%	2.8%	0.3%	0.0%
8 - 10	1 - 1.5	0.0%	0.0%	5.9%	0.7%	0.0%
8 - 10	> 1.5	0.0%	0.0%	14.0%	1.5%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	EtaNCEL	MaNCEL	MaRES	MvRES	NormNCRES
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.0%	0.0%	0.1%	0.0%	0.0%
10 - 15	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
10 - 15	0.25 - 0.4	0.0%	0.0%	0.7%	0.0%	0.0%
10 - 15	0.4 - 0.7	0.0%	0.0%	1.5%	0.3%	0.0%
10 - 15	0.7 - 1	0.0%	0.0%	3.5%	0.5%	0.0%
10 - 15	1 - 1.5	0.0%	0.0%	7.7%	0.9%	0.0%
10 - 15	> 1.5	0.0%	0.0%	48.4%	6.7%	0.0%
> 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	0.0%	0.0%	0.5%	0.0%	0.0%
> 15	0.15 - 0.25	0.0%	0.0%	0.1%	0.0%	0.0%
> 15	0.25 - 0.4	0.0%	0.0%	0.7%	0.1%	0.0%
> 15	0.4 - 0.7	0.0%	0.0%	1.0%	0.2%	0.0%
> 15	0.7 - 1	0.0%	0.0%	2.0%	0.4%	0.0%
> 15	1 - 1.5	0.0%	0.0%	13.1%	1.6%	0.0%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

Table E.4. Absolute fractional primary interaction shape uncertainties (resonant and neutral current) vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Vector FF	M_A shape	Pauli Supp	Normalization	MEC
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	1.1%	2.4%	4.7%	5.0%	8.3%
< 1.5	0.15 - 0.25	0.3%	1.5%	3.9%	4.0%	11.4%
< 1.5	0.25 - 0.4	0.1%	0.9%	2.8%	2.9%	7.1%
< 1.5	0.4 - 0.7	0.9%	0.8%	0.5%	2.8%	0.9%
< 1.5	0.7 - 1	3.5%	5.5%	1.1%	6.9%	4.3%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.5%	1.2%	2.4%	2.8%	12.0%
1.5 - 2	0.15 - 0.25	0.2%	0.9%	2.5%	2.5%	10.3%
1.5 - 2	0.25 - 0.4	0.1%	0.6%	1.7%	2.2%	7.0%
1.5 - 2	0.4 - 0.7	0.5%	0.3%	0.4%	2.1%	5.2%
1.5 - 2	0.7 - 1	1.9%	0.4%	0.6%	4.1%	0.6%
1.5 - 2	1 - 1.5	1.3%	1.5%	1.1%	6.6%	2.4%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.6%	1.1%	1.9%	2.6%	9.7%
2 - 2.5	0.15 - 0.25	0.2%	0.6%	1.8%	1.7%	9.0%
2 - 2.5	0.25 - 0.4	0.0%	0.4%	1.3%	1.8%	7.5%
2 - 2.5	0.4 - 0.7	0.3%	0.2%	0.3%	1.6%	5.4%
2 - 2.5	0.7 - 1	1.5%	0.2%	0.5%	2.4%	0.1%
2 - 2.5	1 - 1.5	2.1%	1.6%	0.7%	5.1%	0.7%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Vector FF	M_A shape	Pauli Supp	Normalization	MEC
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.4%	0.7%	1.5%	1.8%	9.0%
2.5 - 3	0.15 - 0.25	0.2%	0.5%	1.4%	1.6%	7.6%
2.5 - 3	0.25 - 0.4	0.0%	0.3%	1.1%	1.5%	6.8%
2.5 - 3	0.4 - 0.7	0.2%	0.1%	0.2%	1.3%	4.5%
2.5 - 3	0.7 - 1	1.2%	0.3%	0.4%	2.1%	1.1%
2.5 - 3	1 - 1.5	2.4%	0.6%	0.7%	2.8%	0.9%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.3%	0.5%	1.0%	1.4%	7.0%
3 - 3.5	0.15 - 0.25	0.2%	0.4%	1.2%	1.3%	5.6%
3 - 3.5	0.25 - 0.4	0.0%	0.2%	0.8%	1.0%	5.6%
3 - 3.5	0.4 - 0.7	0.2%	0.0%	0.1%	1.0%	4.2%
3 - 3.5	0.7 - 1	1.0%	0.4%	0.4%	2.0%	1.2%
3 - 3.5	1 - 1.5	2.2%	0.2%	0.6%	3.6%	0.7%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.3%	0.5%	1.1%	1.2%	6.1%
3.5 - 4	0.15 - 0.25	0.1%	0.3%	0.8%	0.9%	6.2%
3.5 - 4	0.25 - 0.4	0.0%	0.2%	0.6%	0.9%	5.0%
3.5 - 4	0.4 - 0.7	0.1%	0.0%	0.1%	0.7%	3.5%
3.5 - 4	0.7 - 1	0.9%	0.5%	0.4%	1.7%	1.8%
3.5 - 4	1 - 1.5	2.4%	0.4%	1.0%	5.3%	0.7%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Vector FF	M_A shape	Pauli Supp	Normalization	MEC
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.2%	0.3%	0.6%	0.9%	5.8%
4 - 4.5	0.15 - 0.25	0.1%	0.2%	0.8%	0.8%	4.6%
4 - 4.5	0.25 - 0.4	0.0%	0.1%	0.5%	0.7%	5.3%
4 - 4.5	0.4 - 0.7	0.1%	0.1%	0.1%	0.7%	3.3%
4 - 4.5	0.7 - 1	1.0%	0.9%	0.7%	1.9%	1.3%
4 - 4.5	1 - 1.5	3.1%	1.5%	1.5%	4.9%	0.8%
4 - 4.5	> 1.5	4.9%	2.6%	3.7%	18.7%	10.5%
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.3%	0.4%	1.0%	1.2%	6.5%
4.5 - 5	0.15 - 0.25	0.1%	0.4%	1.3%	1.3%	5.5%
4.5 - 5	0.25 - 0.4	0.0%	0.2%	0.7%	1.0%	5.1%
4.5 - 5	0.4 - 0.7	0.1%	0.0%	0.1%	0.9%	3.9%
4.5 - 5	0.7 - 1	1.0%	0.9%	0.8%	2.1%	1.5%
4.5 - 5	1 - 1.5	2.7%	1.5%	1.4%	5.2%	0.7%
4.5 - 5	> 1.5	0.4%	1.6%	1.3%	4.7%	0.6%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.3%	0.6%	1.1%	1.4%	7.3%
5 - 6	0.15 - 0.25	0.1%	0.3%	1.1%	1.2%	6.9%
5 - 6	0.25 - 0.4	0.0%	0.2%	0.8%	1.0%	5.5%
5 - 6	0.4 - 0.7	0.1%	0.0%	0.1%	0.9%	4.8%
5 - 6	0.7 - 1	1.1%	0.9%	0.7%	2.0%	1.8%
5 - 6	1 - 1.5	2.5%	1.3%	1.4%	5.8%	1.2%
5 - 6	> 1.5	3.0%	2.6%	2.7%	9.7%	0.8%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Vector FF	M_A shape	Pauli Supp	Normalization	MEC
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	0.4%	0.5%	1.3%	1.4%	6.7%
6 - 8	0.15 - 0.25	0.2%	0.4%	1.3%	1.3%	7.9%
6 - 8	0.25 - 0.4	0.0%	0.2%	0.8%	1.3%	5.7%
6 - 8	0.4 - 0.7	0.1%	0.1%	0.0%	1.0%	4.3%
6 - 8	0.7 - 1	1.0%	0.6%	0.7%	2.6%	2.0%
6 - 8	1 - 1.5	2.7%	1.3%	1.4%	7.0%	0.0%
6 - 8	> 1.5	3.3%	1.8%	3.1%	19.3%	4.3%
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.4%	0.7%	1.7%	1.7%	4.0%
8 - 10	0.15 - 0.25	0.2%	0.6%	2.4%	1.8%	7.7%
8 - 10	0.25 - 0.4	0.0%	0.2%	0.9%	1.5%	9.0%
8 - 10	0.4 - 0.7	0.1%	0.1%	0.0%	0.9%	5.9%
8 - 10	0.7 - 1	1.1%	0.6%	0.7%	3.3%	2.6%
8 - 10	1 - 1.5	2.6%	1.1%	1.2%	6.9%	0.6%
8 - 10	> 1.5	2.8%	2.1%	2.6%	13.3%	5.2%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.5%	0.9%	2.1%	1.9%	6.7%
10 - 15	0.15 - 0.25	0.1%	0.4%	1.0%	1.5%	5.7%
10 - 15	0.25 - 0.4	0.0%	0.2%	1.0%	1.4%	8.9%
10 - 15	0.4 - 0.7	0.1%	0.1%	0.0%	1.1%	5.2%
10 - 15	0.7 - 1	1.3%	0.7%	0.9%	3.7%	2.6%
10 - 15	1 - 1.5	3.1%	1.2%	1.3%	8.5%	0.2%
10 - 15	> 1.5	8.0%	2.5%	5.1%	34.4%	2.4%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	Vector FF	M_A shape	Pauli Supp	Normalization	MEC
> 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	1.1%	2.7%	5.8%	5.1%	9.4%
> 15	0.15 - 0.25	0.5%	0.4%	5.1%	0.7%	7.5%
> 15	0.25 - 0.4	0.2%	0.3%	2.0%	1.5%	6.4%
> 15	0.4 - 0.7	0.2%	0.7%	0.2%	0.9%	17.1%
> 15	0.7 - 1	0.7%	0.1%	0.5%	3.1%	0.3%
> 15	1 - 1.5	2.6%	1.3%	1.5%	8.0%	2.4%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

Table E.5. Absolute fractional signal model uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	AGKYxF1pi	MFP_N	MFP_pi	RDecBR1gamma	Theta_Delta2Npi
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	0.1%	0.6%	2.4%	0.0%	0.0%
< 1.5	0.15 - 0.25	0.4%	1.0%	2.0%	0.0%	0.0%
< 1.5	0.25 - 0.4	0.6%	0.1%	1.7%	0.0%	0.0%
< 1.5	0.4 - 0.7	0.5%	0.9%	1.1%	0.0%	0.1%
< 1.5	0.7 - 1	1.2%	3.7%	2.8%	0.0%	0.0%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	AGKYxF1pi	MFP_N	MFP_pi	RDecBR1gamma	Theta_Delta2Npi
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.1%	0.3%	1.5%	0.0%	0.1%
1.5 - 2	0.15 - 0.25	0.1%	0.0%	0.7%	0.0%	0.0%
1.5 - 2	0.25 - 0.4	0.3%	0.5%	1.2%	0.0%	0.0%
1.5 - 2	0.4 - 0.7	0.3%	1.2%	1.0%	0.0%	0.0%
1.5 - 2	0.7 - 1	0.2%	1.0%	0.9%	0.0%	0.0%
1.5 - 2	1 - 1.5	0.4%	0.2%	2.5%	0.0%	0.1%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.1%	0.1%	1.2%	0.0%	0.0%
2 - 2.5	0.15 - 0.25	0.0%	0.3%	0.9%	0.0%	0.0%
2 - 2.5	0.25 - 0.4	0.1%	0.3%	0.9%	0.0%	0.0%
2 - 2.5	0.4 - 0.7	0.3%	0.4%	0.8%	0.0%	0.0%
2 - 2.5	0.7 - 1	0.2%	1.2%	0.9%	0.0%	0.0%
2 - 2.5	1 - 1.5	0.9%	4.4%	1.5%	0.0%	0.3%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.0%	0.5%	1.0%	0.0%	0.0%
2.5 - 3	0.15 - 0.25	0.0%	0.0%	0.7%	0.0%	0.0%
2.5 - 3	0.25 - 0.4	0.1%	0.3%	0.7%	0.0%	0.0%
2.5 - 3	0.4 - 0.7	0.2%	0.5%	0.8%	0.0%	0.0%
2.5 - 3	0.7 - 1	0.2%	1.1%	0.7%	0.0%	0.0%
2.5 - 3	1 - 1.5	0.2%	1.5%	0.7%	0.0%	0.0%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	AGKYxF1pi	MFP_N	MFP_pi	RDecBR1gamma	Theta_Delta2Npi
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.0%	0.3%	0.8%	0.0%	0.0%
3 - 3.5	0.15 - 0.25	0.1%	0.1%	0.4%	0.0%	0.0%
3 - 3.5	0.25 - 0.4	0.0%	0.2%	0.6%	0.0%	0.0%
3 - 3.5	0.4 - 0.7	0.1%	0.4%	0.6%	0.0%	0.0%
3 - 3.5	0.7 - 1	0.2%	1.4%	0.7%	0.0%	0.0%
3 - 3.5	1 - 1.5	0.2%	2.8%	0.3%	0.0%	0.0%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.0%	0.4%	0.6%	0.0%	0.0%
3.5 - 4	0.15 - 0.25	0.0%	0.0%	0.4%	0.0%	0.0%
3.5 - 4	0.25 - 0.4	0.1%	0.2%	0.5%	0.0%	0.0%
3.5 - 4	0.4 - 0.7	0.1%	0.4%	0.4%	0.0%	0.0%
3.5 - 4	0.7 - 1	0.2%	0.9%	0.7%	0.0%	0.0%
3.5 - 4	1 - 1.5	1.0%	2.6%	0.6%	0.0%	0.1%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.1%	0.0%	0.2%	0.0%	0.0%
4 - 4.5	0.15 - 0.25	0.0%	0.2%	0.2%	0.0%	0.0%
4 - 4.5	0.25 - 0.4	0.0%	0.3%	0.7%	0.0%	0.0%
4 - 4.5	0.4 - 0.7	0.2%	0.6%	0.6%	0.0%	0.0%
4 - 4.5	0.7 - 1	0.4%	1.3%	0.7%	0.0%	0.0%
4 - 4.5	1 - 1.5	0.1%	1.8%	0.2%	0.0%	0.1%
4 - 4.5	> 1.5	0.7%	7.8%	2.0%	0.1%	0.1%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	AGKYxF1pi	MFP_N	MFP_pi	RDecBR1gamma	Theta_Delta2Npi
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.1%	0.4%	0.5%	0.0%	0.0%
4.5 - 5	0.15 - 0.25	0.1%	0.2%	0.4%	0.0%	0.0%
4.5 - 5	0.25 - 0.4	0.1%	0.1%	0.5%	0.0%	0.0%
4.5 - 5	0.4 - 0.7	0.2%	0.3%	0.7%	0.0%	0.0%
4.5 - 5	0.7 - 1	0.3%	1.0%	0.9%	0.0%	0.0%
4.5 - 5	1 - 1.5	0.8%	0.1%	0.8%	0.0%	0.1%
4.5 - 5	> 1.5	0.0%	13.4%	0.1%	0.1%	0.3%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.1%	0.0%	1.2%	0.0%	0.0%
5 - 6	0.15 - 0.25	0.1%	0.1%	0.6%	0.1%	0.0%
5 - 6	0.25 - 0.4	0.0%	0.1%	0.6%	0.0%	0.0%
5 - 6	0.4 - 0.7	0.2%	0.6%	0.7%	0.0%	0.0%
5 - 6	0.7 - 1	0.3%	1.2%	0.7%	0.0%	0.1%
5 - 6	1 - 1.5	1.1%	1.1%	1.1%	0.0%	0.1%
5 - 6	> 1.5	6.6%	1.3%	0.8%	0.5%	0.6%
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	0.2%	0.1%	0.4%	0.0%	0.1%
6 - 8	0.15 - 0.25	0.1%	0.1%	0.5%	0.0%	0.0%
6 - 8	0.25 - 0.4	0.1%	0.3%	0.6%	0.0%	0.1%
6 - 8	0.4 - 0.7	0.3%	0.4%	0.6%	0.0%	0.0%
6 - 8	0.7 - 1	0.3%	0.2%	1.3%	0.0%	0.0%
6 - 8	1 - 1.5	1.1%	2.2%	1.1%	0.0%	0.2%
6 - 8	> 1.5	1.9%	3.2%	2.8%	0.0%	0.1%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	AGKYxF1pi	MFP_N	MFP_pi	RDecBR1gamma	Theta_Delta2Npi
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.0%	0.7%	0.3%	0.0%	0.1%
8 - 10	0.15 - 0.25	0.1%	0.2%	1.0%	0.0%	0.0%
8 - 10	0.25 - 0.4	0.1%	0.2%	0.7%	0.0%	0.0%
8 - 10	0.4 - 0.7	0.2%	0.5%	1.0%	0.0%	0.0%
8 - 10	0.7 - 1	0.3%	0.5%	1.6%	0.0%	0.1%
8 - 10	1 - 1.5	1.4%	2.4%	2.0%	0.0%	0.1%
8 - 10	> 1.5	0.8%	0.7%	0.3%	0.2%	0.2%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.1%	0.1%	0.1%	0.0%	0.2%
10 - 15	0.15 - 0.25	0.0%	0.6%	0.9%	0.0%	0.1%
10 - 15	0.25 - 0.4	0.2%	0.7%	0.8%	0.0%	1.0%
10 - 15	0.4 - 0.7	0.4%	0.7%	0.9%	0.0%	0.1%
10 - 15	0.7 - 1	0.9%	1.4%	1.2%	0.0%	0.0%
10 - 15	1 - 1.5	1.4%	1.0%	1.0%	0.2%	0.0%
10 - 15	> 1.5	5.3%	4.5%	0.3%	0.0%	0.0%
> 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	0.2%	1.2%	3.4%	0.0%	0.2%
> 15	0.15 - 0.25	0.1%	1.3%	0.7%	0.0%	0.2%
> 15	0.25 - 0.4	0.3%	0.9%	0.8%	0.0%	0.7%
> 15	0.4 - 0.7	0.3%	3.5%	0.9%	0.0%	0.2%
> 15	0.7 - 1	0.9%	2.3%	0.5%	0.0%	0.1%
> 15	1 - 1.5	1.2%	2.8%	1.1%	0.1%	0.0%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%

Table E.6. Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrAbs_N	FrAbs_pi	FrCEX_N	FrCEX_pi	FrElas_N	FrElas_pi
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	0.9%	10.5%	2.0%	0.1%	2.4%	0.9%
< 1.5	0.15 - 0.25	0.9%	7.2%	0.1%	0.6%	1.2%	0.7%
< 1.5	0.25 - 0.4	0.8%	5.4%	1.9%	0.2%	1.8%	0.6%
< 1.5	0.4 - 0.7	0.2%	5.2%	0.6%	0.2%	1.1%	0.5%
< 1.5	0.7 - 1	4.5%	5.2%	3.5%	0.5%	2.2%	2.2%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.2%	6.4%	0.6%	0.4%	0.7%	0.6%
1.5 - 2	0.15 - 0.25	0.1%	5.5%	0.1%	0.2%	0.1%	0.7%
1.5 - 2	0.25 - 0.4	0.1%	4.7%	0.2%	0.3%	0.3%	0.5%
1.5 - 2	0.4 - 0.7	0.2%	4.8%	0.4%	0.3%	0.9%	0.5%
1.5 - 2	0.7 - 1	0.4%	5.2%	1.6%	0.2%	1.0%	0.6%
1.5 - 2	1 - 1.5	1.6%	5.9%	2.3%	1.5%	3.1%	0.1%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.5%	6.4%	0.0%	0.5%	1.2%	0.8%
2 - 2.5	0.15 - 0.25	0.0%	4.5%	0.1%	0.2%	0.7%	0.5%
2 - 2.5	0.25 - 0.4	0.2%	4.2%	0.0%	0.3%	0.9%	0.5%
2 - 2.5	0.4 - 0.7	0.0%	4.0%	0.5%	0.2%	0.7%	0.5%
2 - 2.5	0.7 - 1	0.2%	3.7%	0.4%	0.0%	1.3%	0.4%
2 - 2.5	1 - 1.5	1.6%	5.4%	1.3%	1.7%	0.8%	0.1%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrAbs_N	FrAbs_pi	FrCEX_N	FrCEX_pi	FrElas_N	FrElas_pi
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.5%	4.8%	0.3%	0.3%	0.2%	0.5%
2.5 - 3	0.15 - 0.25	0.1%	4.0%	0.0%	0.2%	0.2%	0.5%
2.5 - 3	0.25 - 0.4	0.2%	3.8%	0.1%	0.3%	0.7%	0.4%
2.5 - 3	0.4 - 0.7	0.1%	3.3%	0.2%	0.2%	0.9%	0.4%
2.5 - 3	0.7 - 1	0.1%	3.3%	0.5%	0.2%	1.8%	0.4%
2.5 - 3	1 - 1.5	0.8%	2.6%	0.8%	0.0%	0.8%	0.1%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.2%	4.1%	0.0%	0.2%	0.7%	0.4%
3 - 3.5	0.15 - 0.25	0.0%	3.4%	0.2%	0.3%	0.4%	0.4%
3 - 3.5	0.25 - 0.4	0.1%	2.9%	0.1%	0.2%	0.4%	0.3%
3 - 3.5	0.4 - 0.7	0.2%	2.7%	0.3%	0.2%	0.9%	0.3%
3 - 3.5	0.7 - 1	0.6%	3.0%	1.0%	0.0%	1.2%	0.4%
3 - 3.5	1 - 1.5	0.3%	3.3%	0.9%	0.1%	1.4%	0.1%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.2%	3.7%	0.6%	0.2%	0.7%	0.4%
3.5 - 4	0.15 - 0.25	0.1%	2.7%	0.3%	0.2%	0.3%	0.3%
3.5 - 4	0.25 - 0.4	0.0%	2.4%	0.1%	0.2%	0.4%	0.3%
3.5 - 4	0.4 - 0.7	0.2%	2.3%	0.2%	0.1%	0.7%	0.2%
3.5 - 4	0.7 - 1	0.3%	3.0%	0.7%	0.1%	1.4%	0.3%
3.5 - 4	1 - 1.5	0.5%	4.6%	1.2%	0.3%	2.7%	0.4%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrAbs_N	FrAbs_pi	FrCEX_N	FrCEX_pi	FrElas_N	FrElas_pi
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.2%	3.4%	0.1%	0.1%	0.2%	0.4%
4 - 4.5	0.15 - 0.25	0.3%	2.5%	0.0%	0.2%	1.1%	0.2%
4 - 4.5	0.25 - 0.4	0.1%	2.0%	0.0%	0.1%	0.5%	0.2%
4 - 4.5	0.4 - 0.7	0.2%	2.5%	0.2%	0.1%	0.8%	0.3%
4 - 4.5	0.7 - 1	0.1%	3.5%	1.3%	0.2%	2.0%	0.4%
4 - 4.5	1 - 1.5	0.1%	4.0%	1.0%	0.0%	0.6%	0.3%
4 - 4.5	> 1.5	3.8%	13.6%	5.4%	0.4%	2.7%	5.0%
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.0%	4.3%	0.6%	0.2%	0.0%	0.2%
4.5 - 5	0.15 - 0.25	0.3%	3.2%	0.3%	0.2%	1.4%	0.4%
4.5 - 5	0.25 - 0.4	0.0%	2.7%	0.3%	0.1%	0.1%	0.3%
4.5 - 5	0.4 - 0.7	0.2%	3.1%	0.6%	0.1%	0.7%	0.3%
4.5 - 5	0.7 - 1	0.3%	3.9%	1.3%	0.1%	1.6%	0.4%
4.5 - 5	1 - 1.5	0.0%	5.7%	1.1%	0.1%	1.3%	0.4%
4.5 - 5	> 1.5	8.2%	8.1%	8.0%	0.1%	8.4%	0.3%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.2%	4.1%	0.7%	0.5%	0.5%	0.2%
5 - 6	0.15 - 0.25	0.3%	2.9%	0.5%	0.2%	1.6%	0.3%
5 - 6	0.25 - 0.4	0.1%	2.5%	0.3%	0.2%	0.2%	0.4%
5 - 6	0.4 - 0.7	0.0%	3.4%	0.3%	0.1%	0.5%	0.3%
5 - 6	0.7 - 1	0.1%	3.8%	0.7%	0.1%	1.2%	0.4%
5 - 6	1 - 1.5	0.6%	6.2%	1.1%	0.4%	2.3%	0.7%
5 - 6	> 1.5	1.2%	4.0%	5.4%	1.4%	4.3%	0.2%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrAbs_N	FrAbs_pi	FrCEX_N	FrCEX_pi	FrElas_N	FrElas_pi
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	0.3%	4.4%	0.0%	0.3%	2.8%	0.6%
6 - 8	0.15 - 0.25	0.4%	3.7%	0.2%	0.2%	0.0%	0.3%
6 - 8	0.25 - 0.4	0.5%	3.5%	0.6%	0.5%	0.2%	0.4%
6 - 8	0.4 - 0.7	0.3%	3.5%	0.4%	0.2%	1.2%	0.5%
6 - 8	0.7 - 1	0.1%	4.6%	0.7%	0.3%	1.4%	0.5%
6 - 8	1 - 1.5	0.7%	6.6%	0.1%	0.6%	1.0%	0.5%
6 - 8	> 1.5	1.3%	11.9%	0.9%	0.3%	1.1%	1.2%
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	0.5%	4.5%	1.0%	0.1%	0.3%	0.7%
8 - 10	0.15 - 0.25	0.8%	4.1%	0.0%	0.4%	2.8%	0.5%
8 - 10	0.25 - 0.4	0.6%	3.9%	0.2%	0.1%	1.6%	0.4%
8 - 10	0.4 - 0.7	0.0%	3.6%	0.7%	0.1%	0.6%	0.4%
8 - 10	0.7 - 1	0.2%	6.0%	0.9%	0.2%	1.7%	0.5%
8 - 10	1 - 1.5	0.6%	7.1%	1.1%	0.0%	1.5%	0.5%
8 - 10	> 1.5	4.1%	6.2%	4.2%	0.7%	2.5%	0.4%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.5%	4.8%	0.8%	0.5%	0.2%	0.5%
10 - 15	0.15 - 0.25	0.3%	3.5%	0.2%	0.2%	0.2%	0.4%
10 - 15	0.25 - 0.4	0.0%	3.9%	1.9%	0.0%	1.3%	0.3%
10 - 15	0.4 - 0.7	0.3%	3.7%	0.2%	0.1%	0.9%	0.4%
10 - 15	0.7 - 1	0.2%	5.5%	0.8%	0.3%	2.1%	0.6%
10 - 15	1 - 1.5	0.6%	7.3%	1.0%	0.1%	1.4%	0.6%
10 - 15	> 1.5	1.0%	8.6%	1.8%	0.2%	0.2%	0.9%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrAbs_N	FrAbs_pi	FrCEX_N	FrCEX_pi	FrElas_N	FrElas_pi
> 15	< 0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	3.9%	10.6%	12.6%	0.4%	18.3%	0.8%
> 15	0.15 - 0.25	3.9%	2.4%	0.7%	0.2%	8.1%	0.3%
> 15	0.25 - 0.4	2.5%	4.1%	2.6%	0.1%	2.4%	0.4%
> 15	0.4 - 0.7	1.5%	3.5%	0.6%	0.0%	7.8%	0.3%
> 15	0.7 - 1	1.8%	5.6%	3.2%	0.3%	4.5%	0.5%
> 15	1 - 1.5	0.6%	5.8%	1.6%	0.1%	4.2%	0.5%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table E.7. Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrInel_N	FrInel_pi	FrPiProd_N	FrPiProd_pi
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	0.8%	11.1%	0.1%	0.2%
< 1.5	0.15 - 0.25	1.6%	6.7%	0.1%	0.1%
< 1.5	0.25 - 0.4	0.1%	5.3%	0.1%	0.1%
< 1.5	0.4 - 0.7	0.7%	5.2%	0.3%	0.1%
< 1.5	0.7 - 1	2.2%	3.1%	0.7%	0.1%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrInel_N	FrInel_pi	FrPiProd_N	FrPiProd_pi
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	0.1%	6.3%	0.0%	0.1%
1.5 - 2	0.15 - 0.25	0.1%	5.3%	0.0%	0.1%
1.5 - 2	0.25 - 0.4	0.7%	4.4%	0.0%	0.1%
1.5 - 2	0.4 - 0.7	0.0%	4.4%	0.1%	0.1%
1.5 - 2	0.7 - 1	1.5%	4.7%	0.9%	0.3%
1.5 - 2	1 - 1.5	0.9%	3.9%	2.9%	0.4%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	0.5%	5.6%	0.0%	0.1%
2 - 2.5	0.15 - 0.25	0.6%	4.4%	0.0%	0.0%
2 - 2.5	0.25 - 0.4	0.6%	3.8%	0.0%	0.1%
2 - 2.5	0.4 - 0.7	0.1%	3.7%	0.1%	0.1%
2 - 2.5	0.7 - 1	0.7%	3.8%	0.3%	0.1%
2 - 2.5	1 - 1.5	1.2%	3.7%	0.3%	0.2%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	0.2%	4.7%	0.0%	0.1%
2.5 - 3	0.15 - 0.25	0.2%	3.9%	0.0%	0.0%
2.5 - 3	0.25 - 0.4	0.2%	3.5%	0.0%	0.1%
2.5 - 3	0.4 - 0.7	0.3%	3.1%	0.0%	0.1%
2.5 - 3	0.7 - 1	0.4%	3.1%	0.4%	0.1%
2.5 - 3	1 - 1.5	0.6%	2.6%	1.2%	0.1%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrInel_N	FrInel_pi	FrPiProd_N	FrPiProd_pi
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	0.2%	4.2%	0.0%	0.0%
3 - 3.5	0.15 - 0.25	0.1%	3.2%	0.0%	0.0%
3 - 3.5	0.25 - 0.4	0.0%	2.7%	0.0%	0.0%
3 - 3.5	0.4 - 0.7	0.2%	2.5%	0.0%	0.0%
3 - 3.5	0.7 - 1	0.1%	2.9%	0.3%	0.1%
3 - 3.5	1 - 1.5	0.8%	3.7%	1.3%	0.1%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	0.2%	3.6%	0.0%	0.0%
3.5 - 4	0.15 - 0.25	0.3%	2.5%	0.0%	0.0%
3.5 - 4	0.25 - 0.4	0.2%	2.1%	0.0%	0.0%
3.5 - 4	0.4 - 0.7	0.2%	2.2%	0.0%	0.0%
3.5 - 4	0.7 - 1	0.4%	2.9%	0.3%	0.1%
3.5 - 4	1 - 1.5	1.0%	4.3%	1.5%	0.1%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	0.0%	3.6%	0.0%	0.0%
4 - 4.5	0.15 - 0.25	0.6%	2.5%	0.0%	0.0%
4 - 4.5	0.25 - 0.4	0.3%	1.9%	0.0%	0.0%
4 - 4.5	0.4 - 0.7	0.3%	2.4%	0.0%	0.1%
4 - 4.5	0.7 - 1	0.0%	3.2%	0.3%	0.1%
4 - 4.5	1 - 1.5	2.0%	4.2%	1.1%	0.1%
4 - 4.5	> 1.5	1.6%	8.2%	2.4%	0.5%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrInel_N	FrInel_pi	FrPiProd_N	FrPiProd_pi
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	0.7%	4.8%	0.0%	0.0%
4.5 - 5	0.15 - 0.25	0.3%	3.1%	0.0%	0.0%
4.5 - 5	0.25 - 0.4	0.3%	2.8%	0.0%	0.0%
4.5 - 5	0.4 - 0.7	0.4%	3.2%	0.0%	0.0%
4.5 - 5	0.7 - 1	0.2%	3.8%	0.3%	0.1%
4.5 - 5	1 - 1.5	1.4%	5.7%	1.0%	0.2%
4.5 - 5	> 1.5	46.1%	8.9%	9.5%	0.4%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	0.0%	4.2%	0.0%	0.0%
5 - 6	0.15 - 0.25	0.4%	2.8%	0.0%	0.0%
5 - 6	0.25 - 0.4	0.3%	2.1%	0.0%	0.0%
5 - 6	0.4 - 0.7	0.1%	3.5%	0.0%	0.1%
5 - 6	0.7 - 1	0.1%	3.8%	0.4%	0.1%
5 - 6	1 - 1.5	0.2%	5.2%	1.4%	0.2%
5 - 6	> 1.5	3.1%	2.0%	2.0%	0.6%
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	1.5%	3.8%	0.0%	0.0%
6 - 8	0.15 - 0.25	0.3%	3.7%	0.0%	0.0%
6 - 8	0.25 - 0.4	0.2%	2.8%	0.0%	0.0%
6 - 8	0.4 - 0.7	0.2%	3.2%	0.0%	0.1%
6 - 8	0.7 - 1	0.1%	4.2%	0.3%	0.1%
6 - 8	1 - 1.5	0.4%	5.8%	1.2%	0.3%
6 - 8	> 1.5	1.6%	10.7%	2.6%	0.7%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	FrInel_N	FrInel_pi	FrPiProd_N	FrPiProd_pi
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	1.0%	4.3%	0.0%	0.1%
8 - 10	0.15 - 0.25	1.3%	3.6%	0.0%	0.0%
8 - 10	0.25 - 0.4	0.9%	3.8%	0.0%	0.1%
8 - 10	0.4 - 0.7	0.3%	3.5%	0.0%	0.1%
8 - 10	0.7 - 1	0.3%	5.9%	0.4%	0.2%
8 - 10	1 - 1.5	0.0%	7.4%	0.8%	0.4%
8 - 10	> 1.5	4.5%	6.0%	1.2%	0.7%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	0.6%	4.0%	0.0%	0.0%
10 - 15	0.15 - 0.25	0.2%	3.3%	0.0%	0.0%
10 - 15	0.25 - 0.4	1.1%	4.7%	0.0%	0.1%
10 - 15	0.4 - 0.7	0.6%	3.7%	0.0%	0.1%
10 - 15	0.7 - 1	0.5%	4.7%	0.2%	0.2%
10 - 15	1 - 1.5	0.4%	7.5%	1.0%	0.2%
10 - 15	> 1.5	1.2%	8.5%	0.1%	0.6%
> 15	< 0	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	5.2%	9.7%	0.4%	0.6%
> 15	0.15 - 0.25	2.7%	2.3%	0.0%	0.0%
> 15	0.25 - 0.4	2.8%	3.9%	0.0%	0.1%
> 15	0.4 - 0.7	5.9%	3.8%	0.0%	0.1%
> 15	0.7 - 1	1.8%	5.1%	0.4%	0.2%
> 15	1 - 1.5	5.9%	5.9%	0.3%	0.3%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%

Table E.8. Absolute fractional hadron interaction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	n path length	p response	π response	Other response
< 1.5	< 0	0.0%	0.0%	0.0%	0.0%
< 1.5	0 - 0.15	3.8%	0.0%	1.7%	0.3%
< 1.5	0.15 - 0.25	3.1%	0.3%	0.7%	0.2%
< 1.5	0.25 - 0.4	1.6%	0.0%	0.8%	0.1%
< 1.5	0.4 - 0.7	2.9%	0.3%	0.5%	0.4%
< 1.5	0.7 - 1	8.3%	1.2%	3.0%	0.4%
< 1.5	1 - 1.5	0.0%	0.0%	0.0%	0.0%
< 1.5	> 1.5	0.0%	0.0%	0.0%	0.0%
1.5 - 2	< 0	0.0%	0.0%	0.0%	0.0%
1.5 - 2	0 - 0.15	2.8%	0.0%	0.7%	0.2%
1.5 - 2	0.15 - 0.25	1.8%	0.0%	0.6%	0.2%
1.5 - 2	0.25 - 0.4	1.4%	0.1%	0.5%	0.2%
1.5 - 2	0.4 - 0.7	2.4%	0.2%	0.5%	0.2%
1.5 - 2	0.7 - 1	5.3%	0.5%	1.0%	0.4%
1.5 - 2	1 - 1.5	5.5%	0.4%	0.1%	0.4%
1.5 - 2	> 1.5	0.0%	0.0%	0.0%	0.0%
2 - 2.5	< 0	0.0%	0.0%	0.0%	0.0%
2 - 2.5	0 - 0.15	2.4%	0.1%	1.1%	0.2%
2 - 2.5	0.15 - 0.25	1.4%	0.0%	0.6%	0.1%
2 - 2.5	0.25 - 0.4	1.4%	0.0%	0.5%	0.2%
2 - 2.5	0.4 - 0.7	2.2%	0.2%	0.5%	0.2%
2 - 2.5	0.7 - 1	3.8%	0.4%	0.6%	0.2%
2 - 2.5	1 - 1.5	5.2%	0.3%	0.0%	0.0%
2 - 2.5	> 1.5	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	n path length	p response	π response	Other response
2.5 - 3	< 0	0.0%	0.0%	0.0%	0.0%
2.5 - 3	0 - 0.15	1.9%	0.0%	1.0%	0.2%
2.5 - 3	0.15 - 0.25	1.4%	0.1%	0.8%	0.3%
2.5 - 3	0.25 - 0.4	1.2%	0.0%	0.4%	0.2%
2.5 - 3	0.4 - 0.7	2.0%	0.1%	0.5%	0.2%
2.5 - 3	0.7 - 1	3.1%	0.3%	0.7%	0.2%
2.5 - 3	1 - 1.5	2.8%	0.1%	0.4%	0.1%
2.5 - 3	> 1.5	0.0%	0.0%	0.0%	0.0%
3 - 3.5	< 0	0.0%	0.0%	0.0%	0.0%
3 - 3.5	0 - 0.15	1.7%	0.0%	0.8%	0.2%
3 - 3.5	0.15 - 0.25	1.1%	0.0%	0.5%	0.2%
3 - 3.5	0.25 - 0.4	1.0%	0.0%	0.4%	0.2%
3 - 3.5	0.4 - 0.7	1.6%	0.1%	0.4%	0.2%
3 - 3.5	0.7 - 1	3.1%	0.3%	0.4%	0.2%
3 - 3.5	1 - 1.5	3.4%	0.3%	0.4%	0.3%
3 - 3.5	> 1.5	0.0%	0.0%	0.0%	0.0%
3.5 - 4	< 0	0.0%	0.0%	0.0%	0.0%
3.5 - 4	0 - 0.15	1.3%	0.0%	0.5%	0.4%
3.5 - 4	0.15 - 0.25	1.0%	0.0%	0.6%	0.2%
3.5 - 4	0.25 - 0.4	0.9%	0.0%	0.4%	0.2%
3.5 - 4	0.4 - 0.7	1.5%	0.1%	0.4%	0.2%
3.5 - 4	0.7 - 1	3.2%	0.2%	0.4%	0.2%
3.5 - 4	1 - 1.5	5.2%	0.3%	1.0%	0.3%
3.5 - 4	> 1.5	0.0%	0.0%	0.0%	0.0%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	n path length	p response	π response	Other response
4 - 4.5	< 0	0.0%	0.0%	0.0%	0.0%
4 - 4.5	0 - 0.15	1.1%	0.0%	0.7%	0.4%
4 - 4.5	0.15 - 0.25	0.7%	0.0%	0.5%	0.3%
4 - 4.5	0.25 - 0.4	0.7%	0.0%	0.2%	0.2%
4 - 4.5	0.4 - 0.7	1.4%	0.1%	0.4%	0.2%
4 - 4.5	0.7 - 1	2.7%	0.2%	0.6%	0.2%
4 - 4.5	1 - 1.5	5.7%	0.0%	0.6%	0.4%
4 - 4.5	> 1.5	31.3%	0.7%	1.5%	0.5%
4.5 - 5	< 0	0.0%	0.0%	0.0%	0.0%
4.5 - 5	0 - 0.15	1.7%	0.0%	0.4%	0.2%
4.5 - 5	0.15 - 0.25	0.8%	0.0%	0.7%	0.2%
4.5 - 5	0.25 - 0.4	0.9%	0.0%	0.3%	0.2%
4.5 - 5	0.4 - 0.7	1.7%	0.1%	0.4%	0.2%
4.5 - 5	0.7 - 1	3.2%	0.1%	0.6%	0.3%
4.5 - 5	1 - 1.5	6.3%	0.3%	0.6%	0.4%
4.5 - 5	> 1.5	3.1%	0.5%	0.7%	0.4%
5 - 6	< 0	0.0%	0.0%	0.0%	0.0%
5 - 6	0 - 0.15	1.3%	0.1%	0.7%	0.5%
5 - 6	0.15 - 0.25	1.2%	0.0%	0.6%	0.1%
5 - 6	0.25 - 0.4	0.9%	0.1%	0.3%	0.2%
5 - 6	0.4 - 0.7	1.9%	0.1%	0.5%	0.2%
5 - 6	0.7 - 1	3.2%	0.2%	0.7%	0.4%
5 - 6	1 - 1.5	6.9%	0.3%	1.0%	0.4%
5 - 6	> 1.5	18.4%	0.2%	0.5%	0.4%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	n path length	p response	π response	Other response
6 - 8	< 0	0.0%	0.0%	0.0%	0.0%
6 - 8	0 - 0.15	1.5%	0.0%	1.3%	0.2%
6 - 8	0.15 - 0.25	1.2%	0.0%	0.7%	0.4%
6 - 8	0.25 - 0.4	1.1%	0.1%	0.3%	0.1%
6 - 8	0.4 - 0.7	1.9%	0.2%	0.5%	0.2%
6 - 8	0.7 - 1	4.1%	0.3%	0.7%	0.2%
6 - 8	1 - 1.5	8.6%	0.2%	1.5%	0.5%
6 - 8	> 1.5	14.5%	0.1%	0.7%	0.4%
8 - 10	< 0	0.0%	0.0%	0.0%	0.0%
8 - 10	0 - 0.15	2.2%	0.3%	0.8%	0.4%
8 - 10	0.15 - 0.25	1.3%	0.0%	0.7%	0.2%
8 - 10	0.25 - 0.4	1.4%	0.0%	0.6%	0.3%
8 - 10	0.4 - 0.7	1.8%	0.1%	0.6%	0.4%
8 - 10	0.7 - 1	4.9%	0.1%	0.9%	0.2%
8 - 10	1 - 1.5	7.0%	0.2%	1.1%	0.7%
8 - 10	> 1.5	8.7%	0.1%	1.0%	0.9%
10 - 15	< 0	0.0%	0.0%	0.0%	0.0%
10 - 15	0 - 0.15	1.4%	0.0%	0.2%	0.1%
10 - 15	0.15 - 0.25	1.0%	0.0%	0.5%	0.1%
10 - 15	0.25 - 0.4	1.5%	0.0%	0.8%	0.2%
10 - 15	0.4 - 0.7	2.3%	0.1%	0.4%	0.1%
10 - 15	0.7 - 1	4.5%	0.4%	1.1%	0.3%
10 - 15	1 - 1.5	8.2%	0.3%	1.5%	0.4%
10 - 15	> 1.5	16.8%	0.5%	1.6%	1.7%

$p_{\parallel\mu}$ (GeV)	$p_{T\mu}$ (GeV)	n path length	p response	π response	Other response
> 15	< 0	0.0%	0.0%	0.0%	0.0%
> 15	0 - 0.15	0.6%	0.0%	0.8%	0.1%
> 15	0.15 - 0.25	2.0%	0.0%	0.4%	0.0%
> 15	0.25 - 0.4	0.4%	0.0%	0.6%	1.3%
> 15	0.4 - 0.7	3.6%	0.0%	0.5%	0.1%
> 15	0.7 - 1	5.7%	0.2%	0.8%	0.2%
> 15	1 - 1.5	10.2%	0.4%	1.4%	0.5%
> 15	> 1.5	0.0%	0.0%	0.0%	0.0%

Table E.9. Absolute fractional recoil reconstruction uncertainties vs. $p_{\parallel\mu}$ and $p_{T\mu}$

Vita

Cheryl Patrick received her bachelor's degree in physics from the University of Oxford in 1997. She then worked as a software developer until 2010, when she commenced her doctoral studies at Northwestern University. She received a master's degree from Northwestern, also in physics, in 2012.