

**Measurement of the Muon-Neutrino Charged-Current
Cross Section on Water with Zero Outgoing Pions**

by

Tianlu Yuan

B.S., Stanford University, 2010

M.S., University of Colorado Boulder, 2013

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2016

This thesis entitled:
Measurement of the Muon-Neutrino Charged-Current Cross Section on Water with Zero Outgoing
Pions
written by Tianlu Yuan
has been approved for the Department of Physics

Prof. Alysia Marino

Prof. Eric Zimmerman

Prof. John Cumalat

Prof. Kevin Stenson

Prof. Michael Shull

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Yuan, Tianlu (Ph.D. Physics)

Measurement of the Muon-Neutrino Charged-Current Cross Section on Water with Zero Outgoing Pions

Thesis directed by Prof. Alysia Marino

The Tokai to Kamioka (T2K) experiment is a 295-km long-baseline neutrino experiment aimed towards the measurement of neutrino oscillation parameters θ_{13} and θ_{23} . Precise measurement of these parameters requires accurate knowledge of neutrino cross sections. We present a flux-averaged double differential measurement of the charged-current cross section on water with zero pions in the final state using the T2K off-axis near detector, ND280. A selection of ν_μ charged-current events occurring in the Pi-Zero subdetector (PØD) of ND280 is performed with 5.8×10^{20} protons on target. The charged outgoing tracks are required to enter and be identified by the ND280 Tracker. Detector inefficiencies are corrected using an unfolding technique. By separating the dataset into time periods when the PØD water layers are filled with water and when they are empty, a subtraction method provides a distribution of ν_μ interactions on water only. Systematic uncertainties on the neutrino flux, interaction model, and detector simulation are propagated numerically towards the final result.

Dedication

To my mom, who was my first teacher

Acknowledgements

In hindsight, if I could tell my younger self what graduate school would be like, I would say that it's humbling. You spend many years accumulating knowledge, only to reach the end and realize how much there is left to do. While there's still much I do not know, one thing I know for certain is that this work would not be possible without the many people who helped me along the way. I owe them my thanks.

I've had the opportunity to meet a tremendous number of smart and passionate physicists on T2K. While the entire collaboration has influenced me in some way, I must single out a few who played a particularly strong role on my work. Raj, Walter, and Erez introduced me to PØD analyses and we met every Friday morning to discuss new progress and ideas. Jay Jo, Alex, and Damon were great partners in maintaining the PØD, talking about the PØD, and helping to make the PØD better. Xianguo showed me another perspective with transverse kinematics and was always eager to give excellent advice on life beyond graduate school. Kendall, one of the most enthusiastic and supportive members in the cross section group, helped guide my analysis from the very start. Anselmo and Martin designed the core of the analysis packages and never failed to answer my questions about the code or to implement my requests for new features. Enrico and I had endless discussions about error propagation, unfolding, and implementation details in the software. He always kept an eye on my commits!

I would also like to thank my officemates, who have patiently sat next to me for a long time. I had many interesting discussions on neutrino physics, general physics, and statistics with Scott and Andy. Since we all worked on different aspects of neutrino experiments, I learned much about the

parts of T2K that I wasn't personally involved with. Both of them provided valuable help as I was finishing up and needed to have some semblance of knowledge about T2K as a whole.

I must thank other members of the CU neutrino group, past and present. Eric was always happy to give sound advice on talks, physics, and careers. Rob, the first postdoc I worked with, set up much of the initial T2K software infrastructure at CU and taught me how to use ROOT and Linux, two of the most important tools of the trade. Before Rob, I didn't really know how to use a computer. Along that vein, Doug helped me debug many computing issues, catered to my many requests, and ensured that I never had to worry about running out of computing resources. He is a large reason I'm graduating sort of on time. Sorry for killing the cluster a few times in the beginning!

The person I worked most closely with on this cross section measurement was Jeremy, and this analysis would not have been possible without his competence. He effortlessly understood how the detector worked, how the software worked, and how to solve many of the difficult questions that arose in the process of making this measurement. His rigorous studies of many systematics help not only this analysis but future ones as well.

When I started out, Alysia often sat down to help debug code and answer my many questions about T2K and neutrino physics. Though we've stopped staring at code together, she still leaves her door open for all my questions. She taught me lots, but perhaps more importantly, she also let me figure things out on my own. This independence helped me develop useful computing skills and allowed me to contribute to T2K in ways that I found meaningful. That said, she also ensured I never lost sight of the end goal. To Alysia, thanks for giving me the freedom to explore new avenues of research, but always keeping me on track. Thanks for your patience in the early years, your encouragement in the last few, and your guidance throughout.

Finally, to my parents. Thanks for your everlasting support. All the parts of me that I value most come from you.

Contents

Chapter

1	The History and Physics of Neutrinos	1
1.1	Historical Prelude	1
1.2	The Standard Model	5
1.2.1	A Relativistic Theory of Quantum Mechanics	6
1.2.2	Gauge Invariance	7
1.2.3	Mass in the Standard Model	9
1.3	Neutrino Physics Beyond the Standard Model	11
1.3.1	Massive Neutrinos	12
1.3.2	Neutrino Oscillations	14
1.4	Neutrino Interactions	16
1.4.1	Point Particle Interactions	17
1.4.2	Neutrino-Nucleon Interactions	18
1.4.3	Nuclear Effects	22
1.5	Why Measure Cross Sections?	23
2	The T2K Experiment	26
2.1	Overview	26
2.2	Creating Neutrinos for T2K	29
2.2.1	The J-PARC Accelerator	29

2.2.2	Proton Beam Monitoring	29
2.2.3	The Secondary Beam Line	30
2.3	The Near Detector Facility	32
2.3.1	INGRID	33
2.3.2	The ND280 Off-Axis Detector	34
2.4	The ND280 Pi-zero Detector	37
2.4.1	Detector Description	37
2.4.2	Data Acquisition and Electronics	40
2.4.3	Detector Calibration	42
2.5	Super-Kamiokande	43
3	Monte Carlo Simulation and Reconstruction	47
3.1	Flux Simulation	48
3.2	Neutrino Interaction Modeling	50
3.2.1	Generators	50
3.2.2	GENIE Neutrino Monte Carlo	51
3.2.3	GENIE MC for T2K	52
3.3	ND280 Detector Simulation and Reconstruction	53
3.3.1	Digitization	55
3.3.2	Data Calibration	56
3.3.3	Event reconstruction	57
4	Analysis Frameworks	61
4.1	oaAnalysis Event Reduction	61
4.2	High Level Analysis at the Near Detector (highLAND)	62
4.2.1	Data Structures in highLAND2	64
4.2.2	Analysis in highLAND2	65
4.2.3	Validation of highLAND2	67

4.3	The Cross Section Tool Package (xsTool)	69
4.3.1	How the Cross Section Tool Works	70
4.3.2	Validation of xsTool	73
5	Data Selection	76
5.1	Software Versions	76
5.2	Event Selection	77
5.2.1	Signal Selection	78
5.2.2	Sidebands Selection	86
6	Systematic Errors and Corrections	91
6.1	Flux	92
6.1.1	Flux Correction	92
6.1.2	Flux Uncertainty	93
6.2	Interaction Modeling	94
6.2.1	Corrections to the Default NEUT MC	95
6.2.2	Corrections due to the EMC Effect	95
6.3	Detector	97
6.3.1	PØD Momentum Scale and Momentum Resolution	98
6.3.2	PØD OOFV	99
6.3.3	Mass of the PØD Fiducial Volume	102
6.3.4	PØD-TPC Track Matching	103
6.3.5	Tracker Systematics	105
6.4	Pile-up Correction	107
7	From Selection to Cross Section	109
7.1	Binning	110
7.2	Bayesian Unfolding	112

7.2.1	Number of Iterations	114
7.3	Background Treatment	117
7.3.1	Without Sidebands	118
7.3.2	With Sidebands	119
7.4	Propagation of Uncertainties	123
7.4.1	Propagation of Statistical Errors	125
7.4.2	Propagation of Detector Errors	126
7.4.3	Propagation of Theory Errors	127
7.4.4	Sidebands and Error Propagation	128
7.4.5	Subtraction and Error Propagation	131
7.4.6	Examples of Single Parameter Variations	133
7.5	Model Dependence and Bias	142
7.5.1	Dependence on Background Treatment	143
7.5.2	Tests of Model Dependence	144
7.6	Alternative Methods	148
8	Double Differential CC0π Cross Section	150
8.1	Results from T2K	150
8.2	Concluding Remarks	151
	Bibliography	156
	Appendix	
A	Ideal Water-in POT for Reduction of Statistical Errors	161
B	Purity Correction and Unfolding	163

C	Parameters for Flux and Theory Systematics	166
C.1	Flux Parameters	166
C.2	FSI Parameters	167
C.3	Cross Section Parameters	168
D	Detailed Cross Section Results	169

Tables

Table

1.1	Best-fit oscillation parameters	15
1.2	Errors on the SK ν_μ CC selection	25
2.1	Summary of T2K Runs 1–4	41
3.1	Target MC POT	53
3.2	GENIE interaction rates	54
5.1	Detailed POT rates for the data taking periods used by this analysis.	77
5.2	Overview of software packages	78
5.3	Composition of the CC0 π selection broken down by the true particle.	82
5.4	Composition of the CC0 π selection broken down by the true reaction channel	83
5.5	Composition of the CC0 π selection broken down by the true topology.	83
5.6	The χ^2/N_{DOF} values obtained from comparison of the sideband data and MC distributions	88
5.7	Overall sideband normalization scale factors	88
6.1	Percentage of OOFV events originating in different PØD regions	99
6.2	Results from fits to the central ECal region	103
6.3	Corrections and uncertainties on the PØD mass	103
6.4	TPC track reconstruction efficiencies	106

6.5	Pile up probabilities for each T2K Run	108
7.1	Contributions to the error on the total cross section due to statistics and systematics	113
7.2	Fractional uncertainties on the total cross section due to the flux	129
7.3	Summary of the behavior of the propagated flux error.	131
8.1	Results from a χ^2 test across all bins in the double differential measurement	151
8.2	Fractional uncertainties on the total cross section	154
D.1	The ν_μ CC0 π cross section on water in the first $\cos\theta_\mu$ slice.	169
D.2	The ν_μ CC0 π cross section on water in the second $\cos\theta_\mu$ slice.	169
D.3	The ν_μ CC0 π cross section on water in the third $\cos\theta_\mu$ slice.	170
D.4	The ν_μ CC0 π cross section on water in the fourth $\cos\theta_\mu$ slice.	170
D.5	The ν_μ CC0 π cross section on water in the fifth $\cos\theta_\mu$ slice.	170
D.6	The ν_μ CC0 π cross section on water in the sixth $\cos\theta_\mu$ slice.	171
D.7	The ν_μ CC0 π cross section on water in the seventh $\cos\theta_\mu$ slice.	171
D.8	The ν_μ CC0 π cross section on water in the eighth $\cos\theta_\mu$ slice.	171

Figures

Figure

1.1	Elementary particles and their interactions	9
1.2	Charged and neutral-current Feynman diagrams	16
(a)	Charged-current	16
(b)	Neutral-current	16
1.3	The Feynman diagram for $\nu_\mu + e^- \rightarrow \nu_e + \mu^-$	17
1.4	Charged-current cross section	19
1.5	Feynman diagram of CCQE interaction	20
2.1	Schematic of the T2K neutrino beam	26
2.2	The J-PARC facility and an outline of the T2K beam line.	27
(a)	Aerial view of the J-PARC complex.	27
(b)	The T2K neutrino beam line.	27
2.3	Accumulated protons on target at T2K	28
2.4	T2K secondary beam line	31
2.5	Schematic of the T2K target	32
(a)	T2K target schematic.	32
(b)	Target and first horn.	32
2.6	Unoscillated flux prediction at SK for various horn currents.	33
2.7	Near detector facility	34

2.8	INGRID and ND280	35
	(a) INGRID	35
	(b) ND280	35
2.9	ND280 TPC	37
	(a) Schematic of TPC detector	37
	(b) TPC dE/dx resolution.	37
2.10	Beam triggered event in ND280	38
2.11	Side view of the PØD	39
2.12	A scintillator bar used in PØDule construction	40
2.13	Three dimensional graphic of PØD	42
2.14	Digitized MPPC dark noise spectrum	43
2.15	Rendered diagram of Super-Kamiokande	44
2.16	Off-axis effect at T2K	45
2.17	SK event display	45
	(a) Electron-like	45
	(b) Muon-like	45
3.1	Unoscillated flux prediction at SK	49
3.2	Flux prediction at SK by parent particle	49
3.3	Workflow of ND280 Software Suite	55
3.4	Simulated neutrino interaction in the PØD	59
4.1	Workflow of the highLAND package	63
4.2	Test failures reported by <code>highland2</code> validation	68
	(a)	68
	(b)	68
4.3	Summary of automated tests of analysis framework	69
4.4	Comparison of simple quantities using <code>xsTool</code> and <code>highland2</code>	73

4.5	Comparison of the relative detector error	74
4.6	Example <code>xsTool</code> cross section extraction	75
5.1	Schematic of an interaction in the PØD	78
5.2	Data and MC distributions of the CC0 π selection	80
	(a) Water-in	80
	(b) Water-in	80
	(c) Water-out	80
	(d) Water-out	80
5.3	MC predicted backgrounds in signal selection	81
5.4	Two dimensional distribution of CC0 π selection	82
5.5	Two dimensional distribution of all CC0 π in PØD fiducial volume	83
5.6	Selection efficiency and purity	84
	(a) Water-in	84
	(b) Water-in	84
	(c) Water-out	84
	(d) Water-out	84
5.7	Comparison of the selection from data with a tuned and sideband constrained MC	85
	(a) Water-in	85
	(b) Water-out	85
5.8	CC1 π^+ and CCOther sideband distributions	87
	(a) Water-in CC1 π^+ Sideband	87
	(b) Water-in CCOther Sideband	87
	(c) Water-out CC1 π^+ Sideband	87
	(d) Water-out CCOther Sideband	87
5.9	A comparison of the water-in data and MC sideband distributions	89
	(a) Water-in CC1 π^+	89

(b)	Water-in CCOther	89
5.10	A comparison of the water-out data and MC sideband distributions	90
(a)	Water-out CC1 π^+	90
(b)	Water-out CCOther	90
6.1	Flux tuning weights	93
6.2	Fractional uncertainties on the flux prediction	94
6.3	Data on the EMC effect	96
(a)	σ_A/σ_D	96
(b)	σ_{Fe}/σ_C	96
6.4	Fits for the central ECal region	102
6.5	Event rates for in-beam cosmic rays and sand-muons	107
(a)	Cosmics	107
(b)	Sand-muons	107
7.1	The binning and global bin numbering scheme used for this measurement	111
7.2	Unstacked fractional uncertainties from statistics and systematics	112
7.3	Cartoon diagram of smearing between a true distribution and reconstructed distribution.	113
7.4	Example of multiple iterations and unsmearing	114
7.5	A test of the unfolding implementation	115
7.6	The fractional bias and error vs number of iterations	115
(a)	Water-in	115
(b)	Water-out	115
7.7	A study of the double differential result for one and four iterations	117
7.8	Test of background treatment without sidebands	120
7.9	Test of background treatment with pure sidebands	122
7.10	A study of the simultaneous unfolding method	123
7.11	Test of background treatment with pure sidebands and six iterations	124

7.12 Comparison of covariance built from throws and input covariance	127
(a) Build from Throws	127
(b) Input Covariance	127
7.13 Distribution of throws of the total background and purity vs the integrated flux, without sidebands	129
(a) B vs F	129
(b) p vs F	129
7.14 Distribution of throws of the total background and purity vs the integrated flux, with sidebands	130
(a) B' vs F	130
(b) p' vs F	130
7.15 Comparison of the true oxygen cross section to a truth-subtracted cross section . . .	132
7.16 Subtraction of two identical samples	134
7.17 Subtraction of two independent water-out samples	134
7.18 Individual throws of the NXSec_MaCCQE parameter	136
7.19 Ratios to the nominal of throws on NXSec_MaCCQE.	137
7.20 Individual throws of the NXSec_CA5RES parameter	138
7.21 Ratios to the nominal of throws on NXSec_CA5RES.	139
7.22 Individual throws of the NIWG_cccohe0 parameter	140
7.23 Ratios to the nominal of throws on NIWG_cccohe0.	141
7.24 Study using two different background treatment methods	144
7.25 Comparison of the true CC0 π watercross section for NEUT and GENIE	145
7.26 Comparison of selection efficiencies for NEUT and GENIE	146
7.27 Model dependence test using a biased MC	147
7.28 Similar to Fig. 7.27 but with the biased MC truth plotted	147
7.29 Comparison of the extracted cross section against the MC prior and toy-MC truth .	148
(a) NEUT MC and GENIE toy-MC	148

(b)	GENIE MC and NEUT toy-MC	148
7.30	Unfolded result with cross section and statistical errors	149
(a)	Result Comparison	149
(b)	Normalized Residuals	149
(c)	Distribution of Normalized Residuals	149
8.1	The double differential ν_μ CC0 π cross section	152
8.2	The fractional error on the double differential result	153
A.1	Plot of the water-in POT that minimizes the statistical uncertainty on the subtracted result	161
B.1	Comparison of background correction in unfolding to purity correction	165

Chapter 1

The History and Physics of Neutrinos

The first evidence for the existence of a new particle arose from an inconsistency. In basic two-body decay, the energy and momentum of the outgoing particles is uniquely defined. However, a continuous β -decay spectrum, discovered by James Chadwick in 1914, hinted at the existence of a third, neutral particle [1]. For several years after, the problem of missing energy retained its mystery as physicists acknowledged the possibility that either a new particle must exist, or conservation of energy does not hold absolutely. It wasn't until 1930 that Wolfgang Pauli proposed the existence of a weakly interacting fermion, calling it the **neutron**. Later coined **neutrino** by Enrico Fermi after Chadwick's discovery of the neutron, it took another couple of decades and the development of quantum electrodynamics (QED) before Frederick Reines and Clyde Cowan conclusively demonstrated its existence. The initial step was thus taken. Reines would go on to win the 1995 Nobel Prize, and the questions that his discovery portended continue to motivate us today.

1.1 Historical Prelude

In 1956 a team lead by Cowan and Reines at the Savannah River Plant designed an experiment to detect the inverse β -decay reaction

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

caused by fission antineutrinos escaping the reactor [2]. The detector consisted of two cadmium chloride doped water targets sandwiched between three scintillator detectors. 1400 liters of liquid

scintillator was completely enclosed by a lead shield and placed underground to reduce the cosmic ray background. A signal rate 20 times the background was measured. This confirmed the existence of the neutrino.

Around the same time, George Sudarshan and Robert Marshak formulated the V-A theory of weak interactions [3]. In an unfortunate but perhaps instructive twist, Marshak and Sudarshan mentioned this work to Gell-Mann in 1957 prior to publishing their conclusions, thereby allowing Feynman and Gell-Mann to usurp their place in history as the discoverers of V-A [4, 5]. Regardless of political intrigue, progress in the theory of weak interactions along with the observation of parity violation in β [6] and K^+ [7] decay, spurred tremendous interest towards further experiments probing the weak interaction and by 1958 Goldhaber, Grodzins and Sunyar had discovered that the ν_e was left-handed by measuring its polarization in e^- capture [8]. A few years later in 1962 the ν_μ would be discovered at Brookhaven National Lab [9] and became the first serious accelerator neutrino experiment [1]. In the BNL experiment, 15 GeV protons striking a beryllium target produced π^+ 's and π^- 's which decay in flight via the reaction

$$\pi^\pm \rightarrow \mu^\pm + \bar{\nu}^{(-)}$$

and its conjugate. These neutrinos were found to produce muons instead of electrons, evidence for a second flavor [9].

The late 60s and early 70s saw developments that illuminated our understanding of nature and results that cast further questions. The progress made by Steven Weinberg and Abdus Salam on Glashow's $SU(2) \times U(1)$ gauge model would lead to one of modern physics' crowning achievements in the GWS Standard Model of particle physics [1]. A more in-depth discussion of the Standard Model is warranted, and will be provided in Section 1.2. For now, we note that the theory predicted the existence of weak neutral-currents, which was experimentally verified by the discovery of neutral-current neutrino interactions in the Gargamelle experiment [1, 10]. In parallel, the Homestake experiment at the Homestake Gold Mine in Lead, South Dakota, USA began collecting data on solar neutrinos in 1970. The solar neutrino flux had been calculated by John Bahcall as early as

1962 [1, 11] using several Standard Solar Models (SSM). However, results from the initial Homestake measurements placed an upper limit on the solar ν_e flux that was about a factor of two less than model predictions [12]. Already, neutrinos gave hints of new physics beyond the Standard Model. The problem of the missing solar neutrino flux became known as the solar neutrino problem (SNP). It would last for over three decades before experimental confirmation of neutrino oscillations and the Mikheev-Smirnov-Wolfenstein (MSW) effect finally put it to rest.

As early as 1962 Maki, Nakagawa and Sakata had proposed a theory of two-neutrino mixing [13]. In 1967, motivated by K^0 oscillations, Bruno Pontecorvo suggested that two-neutrino oscillation may occur between $\nu_\mu \leftrightarrow \nu_e$ [14]. The theory for oscillation was formally developed during 1975 and 1976 by a number of physicists, including Pontecorvo [1]. One of the requirements for this phenomenon is the non-zero mass of the neutrino and a direct consequence is leptonic flavor violation. Both of these concepts are not explained by the SM and thus neutrinos became a probe of new physics. Additionally, as neutrinos play a crucial role in the evolution of the universe, experiments in cosmology have sought to probe neutrinos using astrophysical data. Astrophysical sources of neutrinos include supernovas and relic neutrinos from the Big Bang. One notable example of the former occurred in 1987 with SN 1987A. Neutrinos escaping the supernova were detected by the Kamiokande and IMB detectors. Over the course of 13 s, Kamiokande detected 11 electron events with energies between 7.5–36 MeV [15] while IMB detected 8 events over 6 s with energies between 20–40 MeV [16].

In the late 1980s, the Kamiokande and IMB experiments discovered the atmospheric neutrino anomaly [1]. Similar to the SNP, the atmospheric anomaly involved a deviation from the predicted ν_μ flux produced by cosmics interacting with Earth’s atmosphere. This, along with further experimental evidence for the SNP by Kamiokande, Super-Kamiokande, SAGE and other experiments, spurred progress towards finding the precise values of oscillation parameters within the framework of three-neutrino mixing.¹ By the turn of the century, the Super-Kamiokande experiment would resolve the

¹ Although the ν_τ was not observed until the year 2000 by the DONUT experiment at Fermilab [17], its existence was implied upon the discovery of the τ in 1975.

atmospheric neutrino anomaly in favor of neutrino oscillation by measuring an up-down asymmetry in the atmospheric neutrino flux [1, 18]. A few years later, the Sudbury Neutrino Observatory (SNO) provided measurements of neutral-current interactions caused by ^8B solar ν_e 's which showed ν_e 's were changing flavors. This served as direct evidence favoring neutrino oscillations and the MSW matter effect within the Sun [19]. With the oscillation problem solved, neutrino experiments moved in the direction of precision measurements and high-intensity experiments.

Over the past decade, numerous experiments have narrowed down the uncertainties on neutrino oscillation parameters. Different experimental methods, such as long-baseline accelerator experiments and short-baseline reactor experiments, have provided mostly complementary results on several parameters. Just recently, θ_{13} was observed to be non-zero by the Daya Bay experiment [20] and θ_{23} has been most precisely measured by the Tokai to Kamioka (T2K) experiment [21]. In 2013, Planck measurements on the cosmic microwave background (CMB) yielded additional constraints on the number and mass of neutrinos [22, 23]. The 2014 PDG reports upper limits for $m_{\nu_e} < 2\text{ eV}$ at 95 % confidence from tritium decay and other electron based experiments, $m_{\nu_\mu} < 0.19\text{ MeV}$ at 90 % confidence as measured by muon-based experiments and $m_{\nu_\tau} < 18.2\text{ MeV}$ at 95 % confidence as measured by tau-based experiments [24].

The tremendous progress in neutrino physics over the course of the last century leaves us with five major unanswered questions:

- The absolute mass of the neutrinos
- Mass hierarchy and sign of Δm_{31}^2
- δ_{CP} phase and CP violation in the lepton sector
- Majorana vs Dirac, or whether neutrinos are their own antiparticles
- Existence of additional “sterile” neutrinos

Answers to each of these questions seem within our grasp. Ongoing and planned experiments have made them the top priority. Solutions will provide us with a deeper understanding of matter-

antimatter asymmetry within the framework of leptogenesis, further confirmation of predictions from the SM in the form of Majorana leptons and perhaps hints of a solution to the Dark Matter question in cosmology [25]. In order to gain a better grasp on these implications, we provide an overview of the current state of neutrino theory and its place within the Standard Model.

1.2 The Standard Model

The Standard Model of particle physics is a gauge theory. In general, descriptions of physical phenomena are vastly simplified by assuming a symmetry. This powerful concept is well described by Noether’s elegant theorem, and has been exploited by generations of brilliant professors and naïve students alike in solving problems that would be otherwise incredibly more complicated. The formulation of the SM was no different and it hinges upon the principle of local gauge invariance (LGI). In quantum electrodynamics (QED) the requirement that the Lagrangian remain invariant under a local phase transformation imposes the mediating boson field: the photon. Following successful experimental confirmations of QED, expectations arose that local gauge invariance could also be utilized to describe the weak and strong forces. This ultimately resulted in an extremely successful theory that has predicted the existence of several particles, as well as provided the most accurate predictions of the magnetic moment of the electron and muon. The discovery of the SM-predicted Higgs boson in 2012 at the LHC by both the CMS and ATLAS experiments filled in one of the last missing pieces [26, 27]. Perhaps this can be said to be the greatest example of its success.

Nevertheless, we know that SM is incomplete. Besides failing to describe gravitation, it does not predict the masses of the fundamental particles and requires a seemingly arbitrary number of parameters to describe couplings. Furthermore, it does not describe cold dark matter and dark energy. Neutrinos were assumed to be massless due to the fact that only left-handed neutrinos have been observed. In this aspect, they differ from their massive flavor counterparts. For massive matter particles, the mass term couples left and right-handed states and we would expect to observe both helicities. However, it is possible to incorporate massive neutrinos as an extension to the SM.

Many of the unsolved questions in particle physics are driven by this incompleteness. Neutrinos are amongst the least understood fundamental particles and hence provide a tool for probing beyond the SM.

1.2.1 A Relativistic Theory of Quantum Mechanics

In order for quantum mechanics to satisfy special relativity, a relativistic form for the Schrödinger equation was suggested by Oskar Klein, Walter Gordon and others, motivated by the relativistic energy-momentum equation $E^2 - p^2 = m^2$,

$$(\square + m^2)\Psi = 0. \quad (1.1)$$

If Ψ is interpreted analogous to a non-relativistic wavefunction, we find that the probability density of solutions to Eq. (1.1) is not positive definite [28]. In order to remedy this, a linear equation like Schrödinger's equation is needed. This is the Dirac equation, which achieves linearity at the cost of added dimensions in the form of a matrix equation,

$$(i\not{A} - m)\Psi = 0 \quad (1.2)$$

where we have used $\not{A} \equiv \gamma^\mu A_\mu$ and γ^μ are the Dirac matrices. Solutions of the Dirac equation automatically satisfy the Klein-Gordon equation, and thus the Dirac equation satisfies the relativistic energy equation. Furthermore, the probability density as defined by

$$\rho(x) = \Psi^\dagger(x)\Psi(x) \quad (1.3)$$

is positive definite and the probability current is conserved [1].

Solutions to the Dirac equation represent spin- $\frac{1}{2}$ point particles such as the electron and the neutrino. It is adequate for single-particle descriptions. However, additional formalism is required for a many-particle approach. This is made possible by the canonical quantization (second quantization) of the wavefunctions into field operators by imposing the fermionic anticommutation relationship [1]

$$\left\{ \Psi_\alpha(t, \mathbf{x}), \Psi_\beta^\dagger(t, \mathbf{y}) \right\} = \delta^3(\mathbf{x} - \mathbf{y})\delta_{\alpha\beta} \quad (1.4)$$

$$\left\{ \Psi_\alpha^\dagger(t, \mathbf{x}), \Psi_\beta^\dagger(t, \mathbf{y}) \right\} = 0. \quad (1.5)$$

In doing so, we can express the fields in terms of creation and annihilation operators in \mathbf{k} -space and the quantized fields themselves become operators. This, along with a perturbative approach to interactions, formulates the basis of QFT and allows the calculation of cross sections in a rigorous, quantized framework.

1.2.2 Gauge Invariance

As noted earlier, symmetries are transformations that leave the physics unchanged. A simple example is the addition of a constant term to the electric potential in the classical formulation of electromagnetism. In quantum field theory (QFT), the symmetry group fixes the properties of **gauge bosons** (aside from couplings), which mediate the interactions of that particular field theory [1]. If we begin with the Dirac Lagrangian,

$$\mathcal{L} = \bar{\Psi}(i\not{\partial} - m)\Psi, \quad (1.6)$$

where $\bar{\Psi} = \Psi^\dagger \gamma^0$, it should be obvious that \mathcal{L} is invariant under the transformation

$$\Psi(x) \rightarrow e^{i\alpha}\Psi(x). \quad (1.7)$$

This is a **global** phase transformation and it implies the existence of a conserved current,

$$\partial_\mu j^\mu = 0 \quad (1.8)$$

with $j^\mu(x) = \bar{\Psi}(x)\gamma^\mu\Psi(x)$ [1]. If we further impose the condition that \mathcal{L} remains invariant under a **local** phase transformation,

$$\Psi(x) \rightarrow e^{i\alpha(x)}\Psi(x). \quad (1.9)$$

a modification to the Lagrangian needs to be made. The important difference between Eq. (1.7) and Eq. (1.9) is the phase parameter α , which is dependent on x in case of a local gauge transformation but not in the global case. We see that Eq. (1.6) is no longer invariant due to the derivative on Ψ , which becomes

$$\partial_\mu \Psi \rightarrow e^{i\alpha(x)}\partial_\mu \Psi + ie^{i\alpha(x)}\Psi\partial_\mu \alpha. \quad (1.10)$$

In order to preserve invariance, a correction term must be added to form a modified derivative

$$D_\mu = \partial_\mu - ieA_\mu \quad (1.11)$$

with the requirement that

$$A_\mu \rightarrow A_\mu + \frac{1}{e}\partial_\mu\alpha. \quad (1.12)$$

This is the heart of local gauge invariance. If we substitute Eq. (1.11) into Eq. (1.6) we have,

$$\mathcal{L} = \bar{\Psi}(i\not{D} - m)\Psi \quad (1.13)$$

$$= \bar{\Psi}(i\not{\partial} - m)\Psi + e\bar{\Psi}\gamma^\mu\Psi A_\mu. \quad (1.14)$$

Thus, the assumption of local gauge invariance forces the creation of a gauge field, A_μ which couples to the fermion field Ψ . The addition of a kinetic term $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ satisfies Eq. (1.12) and leads us to the QED Lagrangian [28],

$$\mathcal{L}_{\text{QED}} = \bar{\Psi}(i\not{\partial} - m)\Psi + e\bar{\Psi}\gamma^\mu\Psi A_\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}. \quad (1.15)$$

QED is a U(1) gauge theory, which means it has one generator, corresponding to the photon, that generates the group of local transformations. We can extend this formalism to describe electroweak interactions, which is necessary for neutrino interactions. Again local gauge invariance is invoked, but this time the transformation is taken from the $\text{SU}(2)_L \times \text{U}(1)_Y$ electroweak group, where Y is hypercharge and L indicates that only left-handed fields are affected by the SU(2) transformations. The generalized, modified derivative is [1]

$$D_\mu = \partial_\mu + ig\mathbf{A}_\mu \cdot \mathbf{I} + ig'B_\mu \frac{Y}{2}, \quad (1.16)$$

where \mathbf{I} is a “vector” corresponding to the three generators of $\text{SU}(2)_L$ and $\frac{Y}{2}$ corresponds to the generator for $\text{U}(1)_Y$. These generators each correspond to a massless gauge boson field which mediates the interaction of the fermion fields, but they are not the W^+ , W^- , Z^0 , and γ fields that actually mediate interactions of the physical fermionic fields measured by experiments. The physical gauge fields for the weak interaction are massive, but an additional mass term (e.g. $m^2 W^\mu W_\mu$) would

break the local gauge invariance [28]. Luckily, another mechanism called the *Higgs mechanism* allows them to acquire mass via spontaneous symmetry breaking without sacrificing local gauge invariance. The full electroweak Lagrangian can then be derived using the modified derivative in Eq. (1.16) and adding the gauge boson self-interaction terms, the Higgs self-interaction, and mass terms arising via the Higgs-fermion couplings [1]. In a similar manner, quantum chromodynamics (QCD) interactions are fixed by the $SU(3)_C$ gauge group and the tensor product of $SU(3)_C \times SU(2)_L \times U(1)$ is the transformation group upon which rests the SM.

1.2.3 Mass in the Standard Model




Figure 1.1: Elementary particles and their interactions. The lines indicate couplings that can occur within the Standard Model.

The particles that are known to exist within the Standard Model can be grouped by their spin into fermions (half-integer) and bosons (integer). Fermions are grouped by their interaction type into quarks, which undergo electroweak and strong interactions, and leptons, which are only affected by the electroweak force. The gauge bosons mediate interactions and the Higgs boson mediates the mass coupling as shown in Fig. 1.1. Aside from the gluon and the photon, all elementary particles are massive.

As mentioned earlier, gauge bosons are fields that arise due to local gauge invariance of the Lagrangian. The symmetry group fixes the number and properties of the gauge bosons, leaving three unknown, independent parameters (the coupling constants) that must be determined from experiment. The $SU(3)_C$ group for example has eight generators corresponding to the eight different types of gluons. Properties of the Higgs and the fermions are unconstrained by the theory [1]. These indeterminate parameters are evidence that the SM is an effective field theory, valid only at low energies relative to the Planck scale.

Recall that the three generators of $SU(2)_L$ in Eq. (1.16) correspond to **massless** gauge bosons. That is, the electroweak Lagrangian as formulated based on LGI alone contains no mass terms for the vector bosons. In order to explain the masses of the W^\pm and Z^0 bosons we have to invoke the Higgs mechanism and spontaneous symmetry breaking by the vacuum [1]. The Higgs doublet is defined as

$$\Phi(x) = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.17)$$

where ϕ^+ is a charged, scalar field and ϕ^0 is a neutral, scalar field. The Higgs Lagrangian is

$$\mathcal{L}_{Higgs} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2. \quad (1.18)$$

where μ and λ correspond to the quadratic and quartic self-coupling parameters. It can be shown that \mathcal{L}_{Higgs} preserves LGI if the Higgs doublet transforms trivially under gauge transformations [1]. Now, the potential term of \mathcal{L}_{Higgs} is given as

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2, \quad (1.19)$$

and its minimum determines the vacuum expectation value (VEV) of the Higgs field. Crucially, we assume $\mu^2 < 0$. The minimum of V is found to be

$$\Phi^\dagger \Phi = \frac{v^2}{2}, \quad (1.20)$$

where $v \equiv \sqrt{-\frac{\mu^2}{\lambda}}$. With the additional constraint that the vacuum cannot be charged, the VEV $\langle \Psi \rangle$

is

$$\langle \Psi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \quad (1.21)$$

The predictive power of the Lagrangians in particle physics depends on perturbation theory, which in turn requires a choice of minimum about which to expand the potential. This choice is fixed by the Higgs' VEV, and it is precisely this that hides the symmetry of the Lagrangian [28]. The Higgs doublet can be written as

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (1.22)$$

where $H(x)$ represents the physical Higgs field [1]. Substituting this back into \mathcal{L}_{Higgs} gives the required mass terms for the W^\pm and Z^0 gauge bosons.

Now we have a massive scalar particle whose coupling to the gauge bosons gives the required mass terms in the SM Lagrangian while preserving LGI. The fermion masses arise through Yukawa coupling to the Higgs field [1]. Although neutrinos are massless in the SM, observation of neutrino oscillation leads to the conclusion that neutrinos are in fact massive particles.

1.3 Neutrino Physics Beyond the Standard Model

Neutrino oscillation is one of the most intriguing phenomena associated with neutrinos. It arises from the fact that neutrinos have slightly different masses and are measured in a superposition of their mass eigenstates. The theory of neutrino oscillations was formulated in the 1970s but experimental confirmation remained elusive until the turn of the millennium. Since then, experimental measurements of oscillation parameters have become ever more precise, providing stringent limits on mass differences. However, the absolute mass remains one of the most important unsolved questions in neutrino physics. This section will give a brief exposition of massive neutrinos before moving on to a description of neutrino oscillations.

1.3.1 Massive Neutrinos

As shown in the previous section, coupling to the Higgs field determines mass. For neutrinos, however, the precise form of the coupling will depend on whether neutrinos are Dirac or Majorana particles. In the case of Dirac masses the neutrino mixing is similar to that of quark mixing, but for quarks the mass and flavor fields coincide. Quark mixing is determined by the Cabbibo-Kobayashi-Maskawa (CKM) matrix, which contains four independent parameters in the case of three-generation mixing [1]. The analogous mixing matrix for neutrinos is the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix, which is parameterized with three mixing angles and one phase,

$$\mathbf{U} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.23)$$

The PMNS matrix gives the relationship between neutrino flavor and mass states,

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.24)$$

$$\boldsymbol{\nu}_L = \mathbf{U}\mathbf{n}_L, \quad (1.25)$$

where L denotes left-handedness. Neutrino flavor is thus **defined** in order to match the flavor of the charged lepton with which it couples to in weak charged-current interactions,

$$j_{W,L}^\rho = 2\bar{\nu}_L\gamma^\rho\mathbf{l}_L. \quad (1.26)$$

This is similar to how up-generation quarks couple to a superposition of down-generation quarks. The only difference is that, since quark mass differences are so much larger, we do not define a quark’s flavor by its coupling. Instead, a quark’s mass determines its flavor and quark flavors “mix” via the weak interaction.

In the case of Dirac masses, the Higgs fields couples left and right handed neutrinos in their

mass eigenstates,

$$\mathcal{L}_{H,\nu} = -\left(\frac{v+H}{\sqrt{2}}\right)\bar{\mathbf{n}}_L \mathbf{Y} \mathbf{n}_R + h.c. \quad (1.27)$$

Here, \mathbf{Y} is the diagonalized Yukawa matrix with elements that determine the neutrino mass [1],

$$m_k = \frac{Y_k v}{\sqrt{2}}. \quad (1.28)$$

Again, note that the Y_k 's are not fixed by theory but must be measured. Measuring the neutrino mass is the obvious method to fix these parameters.

Neutrino masses can also be Majorana, which obviates the simultaneous existence of both left and right-handed neutrinos. By assuming that the fields Ψ_R and Ψ_L are not independent, Ettore Majorana discovered that it was possible to express the field of a massive particle in terms of a 2-component spinor [1],

$$\Psi = \Psi_L + \Psi_R \quad (1.29)$$

$$= \Psi_L + \mathcal{C}\bar{\psi}_L^T \quad (1.30)$$

$$= \Psi_L + \Psi_L^C. \quad (1.31)$$

This implies the particle and antiparticle fields are degenerate,

$$\Psi = \Psi^C, \quad (1.32)$$

thus only neutral fermions (the mass term vanishes if fields commute) can be Majorana [1]. The Majorana mass term contributes to the Lagrangian as,

$$\mathcal{L}_{mass}^M = -\frac{1}{2}\bar{\mathbf{n}}^C \mathbf{M} \mathbf{n} + h.c., \quad (1.33)$$

where \mathbf{M} is the diagonalized mass matrix with elements corresponding to the masses of ν_k .

Currently, we do not know if neutrinos are their own antiparticles and we also do not know if right-handed neutrinos exist. If they are not their own antiparticles then only the Dirac mass \mathcal{L}_{mass}^D exists. If they are their own antiparticles but right-handed neutrinos do not exist then only the Majorana mass \mathcal{L}_{mass}^M exists. However, if they are both their own antiparticles and right-handed neutrinos exist, then it is possible to have both Dirac and Majorana contributions.

1.3.2 Neutrino Oscillations

The evidence for massive neutrinos came in the form of neutrino oscillations. Not only does oscillation explain SNP and the atmospheric neutrino anomaly, it also implies that neutrinos come in a variety of masses. In this section we will show why this must be the case. Several simplifying assumptions are made but it can be shown in the full field-theoretic derivation that the results are identical.

Equation (1.24) describes neutrino flavor eigenstates as a superposition of their mass eigenstates. In bra-ket notation we write this as,

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle. \quad (1.34)$$

This can be inverted to write,

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle. \quad (1.35)$$

From the Schrödinger equation we know that the time-evolution of $|\nu_\alpha\rangle$ is given by,

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle, \quad (1.36)$$

where $E_k = \sqrt{p^2 + m_k^2}$. Substituting Eq. (1.35) into Eq. (1.36),

$$|\nu_\alpha(t)\rangle = \sum_\beta \left(\sum_k U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle. \quad (1.37)$$

This makes it clear that the time-evolved state $\nu_\alpha(t)$ can be written as a superposition of all flavor states. From this, the transition probability can be easily derived to be [1],

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (1.38)$$

Since neutrinos are ultrarelativistic, $E_k \approx E + \frac{m_k^2}{2E}$ and $t \approx L$, or the distance the neutrino travels, so

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{kj}^2 L}{2E}}. \quad (1.39)$$

Flavor change is essentially a quantum mechanical effect dependent on the parameter $\frac{L}{E}$ and the existence of multiple neutrino mass eigenstates. The basic idea of an accelerator-based neutrino

oscillation experiment is to create neutrinos of one flavor and then see if it has changed at a remote detector some distance L away from the neutrino source.

Neutrino oscillation experiments are categorized by their average L as short-baseline (SBL), long-baseline (LBL), very-long-baseline (VLB), and solar experiments. Terrestrial experiments typically use reactors or accelerators as sources. In addition to the mixing parameters, oscillation experiments are sensitive to Δm^2 and thus can constrain neutrino masses. As we move into the era of precision neutrino experiments, several important questions are just beginning to be probed. One example is CP violation which has been observed in the quark sector but not in leptons. It can be said to occur if a physical process exhibits differences under the reversal of both parity and charge conjugation. The most natural method to encode this asymmetry is within the δ_{CP} phase in Eq. (1.23), which switches sign under charge conjugation. Another example is the mass hierarchy question, which asks whether ν_1 or ν_3 is the lightest eigenstate. For the normal hierarchy, $m_1 < m_2 < m_3$, whereas the inverted has $m_3 < m_1 < m_2$ (solar neutrino experiments have shown that $m_1 < m_2$). The value of δ_{CP} as well as the mass hierarchy can both be probed by oscillation experiments and a compilation of all the oscillation parameters with their best-fit values is given in Table 1.1. In order to produce accurate results, a detailed understanding of neutrino production and interaction is crucial.

Parameter	Value	
	Normal	Inverted
$\Delta m_{21}^2 [10^{-5} \text{ eV}^2]$	$7.54^{+0.26}_{-0.22}$	same
$ \Delta m^2 [10^{-3} \text{ eV}^2]$	2.43 ± 0.06	2.38 ± 0.06
$\sin^2 \theta_{12}$	0.308 ± 0.017	same
$\sin^2 \theta_{23}$	$0.437^{+0.033}_{-0.023}$	$0.455^{+0.039}_{-0.031}$
$\sin^2 \theta_{13}$	$0.0234^{+0.0020}_{-0.0019}$	$0.0240^{+0.0019}_{-0.0022}$
$\delta_{CP}/\pi \ (2\sigma)$	$1.39^{+0.38}_{-0.27}$	$1.31^{+0.29}_{-0.33}$

Table 1.1: Best-fit values with 1σ ranges of the various neutrino oscillation parameters for both normal and inverted hierarchies. Here $\Delta m^2 \equiv m_3^2 - (m_2^2 + m_1^2)/2$ so Δm^2 greater than (less than) 0 corresponds to the normal (inverted) hierarchy [24].




Figure 1.2: Charged and neutral-current Feynman diagrams of the weak interaction.

1.4 Neutrino Interactions

The primary goal of this thesis is to give a measurement of a specific neutrino-nuclear interaction: charged-current quasielastic (CCQE) on oxygen. Since neutrinos only interact via the weak force, the propagator can be either the W^\pm or the Z^0 boson. The former interaction is categorized as **charged-current** and the latter as **neutral-current** as shown in Fig. 1.2. Both are well explained by the SM, and purely leptonic interactions such as neutrino-electron scattering are calculable exactly. Complications arise, however, when calculating the cross section for neutrino-nucleon or neutrino-nuclear interactions. For the former, form factors are used to describe the hadronic current, for the latter, nuclear effects play a role in nucleon selection and intranuclear final state interactions (FSI) must also be taken into account.

In long-baseline neutrino experiments such as T2K, the primary goal is to provide a measurement of neutrino oscillation parameters. However, in order to do this precisely, an in depth understanding of several complex aspects of neutrino generation, interaction and oscillation must be well understood. These include neutrino flux simulations and cross section modeling. With a near-detector such as T2K's ND280, both flux and cross section uncertainties can be internally constrained. In addition, the ND280 also provides direct measurements of neutrino cross sections in specific interaction channels. This section reviews the theory behind neutrino interactions on point particles, free nucleons, and atomic nuclei.




Figure 1.3: The Feynman diagram for $\nu_\mu + e^- \rightarrow \nu_e + \mu^-$.

1.4.1 Point Particle Interactions

It is easiest to begin with a description of neutrino-lepton scattering. In order to ignore interference terms between charged and neutral-currents we focus on $\nu_\mu + e^- \rightarrow \nu_e + \mu^-$ scattering. This interaction proceeds only via the charged current channel as shown in Fig. 1.3.

Using the Feynman rules, we can construct the invariant amplitude as [28]

$$\mathcal{M} = \frac{G}{\sqrt{2}} (\bar{u}(k') \gamma^\mu (1 - \gamma^5) u(p)) (\bar{u}(p') \gamma_\mu (1 - \gamma^5) u(k)). \quad (1.40)$$

In the high-energy limit, the masses of the charged leptons can be neglected. Squaring \mathcal{M} and summing over spins and averaging we get

$$\overline{|\mathcal{M}|^2} = 64G^2(k \cdot p)(k' \cdot p') \quad (1.41)$$

$$= 16G^2 s^2. \quad (1.42)$$

This allows us to calculate the cross section, and in the center of mass frame we have,

$$\frac{d\sigma}{d\Omega} = \frac{1}{64\pi^2 s} \overline{|\mathcal{M}|^2} = \frac{G^2 s}{4\pi^2}, \quad (1.43)$$

where $s = (k + p)^2 \approx 2k \cdot p$.

Neutrino-quark scattering proceeds identically with the cross section for $\nu_\mu + d \rightarrow \mu^- + u$ given by Eq. (1.43). Due to helicity conservation and the left-handed nature of the weak interaction, the analogous antineutrino interaction's ($\bar{\nu}_\mu + u \rightarrow \mu^+ + d$) phase space is modified by a factor of $\frac{1}{4}(1 - \cos \theta)^2$, where θ is defined as the angle between incoming antineutrino and outgoing μ^+ in the CM frame. With this definition, $\theta = 0$ corresponds to a backwards going μ^+ and the cross section

vanishes. The total neutrino/antineutrino-quark cross section is,

$$\sigma_\nu = \frac{G^2 s}{\pi} \quad (1.44)$$

$$\sigma_{\bar{\nu}} = \frac{G^2 s}{3\pi}. \quad (1.45)$$

1.4.2 Neutrino-Nucleon Interactions

Composite particles such as protons and neutrons make up the vast majority of ordinary matter. Neglecting nuclear effects and final state interactions (FSI) for the moment, neutrino-nucleon scattering is the basis on which experimental measurements connect with theory. Different processes dominate at various energy scales. As this thesis uses T2K data, this section will discuss the different interaction channels dominant in the range of 0.1–30 GeV. At these intermediate energy scales, complexities arise that make direct calculation difficult. All theoretical predictions are based on models with inherent assumptions that depend on values only accessible via experiment.

The three main categories are [29]:

- Elastic and quasielastic (QE)
- Resonance scattering (RES)
- Deep inelastic scattering (DIS)

in increasing order of energy scale. Figure 1.4 shows the contributions of each interaction mode to the total neutrino cross section.

1.4.2.1 Quasielastic Scattering

A neutrino can scatter elastically off a nucleon via a neutral-current interaction. The analogous charged-current process is referred to as **quasielastic** scattering (CCQE) and is shown in Fig. 1.5.

$$\begin{aligned} \nu_l + n &\rightarrow p + l^- \\ \bar{\nu}_l + p &\rightarrow n + l^+ \end{aligned} \quad (1.46)$$




Figure 1.4: The neutrino charged-current cross section per nucleon as a function of neutrino energy. The primary interaction modes are shown with data points compiled from a variety of experiments. Predictions (dashed and solid lines) come from the NUANCE generator [30]. The CCQE contribution is normalized over both neutrons and protons and thus decreased by a factor of two. A description of the different interaction channels is provided in Section 1.4. Modified from [29] with the approximate T2K peak energy of 0.6 GeV highlighted.

The CCQE amplitude depends on several nucleon form factors that are needed to describe the hadronic current. In the Llewellyn-Smith formalism, the CCQE cross section is

$$\frac{d\sigma}{dQ^2} = \frac{G^2 M^2}{8\pi E_\nu^2} \left[A \pm \frac{s-u}{M^2} B + \frac{(s-u)^2}{M^4} C \right], \quad (1.47)$$

where $(-)+$ refers to (anti)neutrino scattering and Q^2 is the squared four-momentum transfer [29].

The invariants s and u are Mandelstam variables and correspond to the square of the center of mass energy and the square of the four-momentum transfer, respectively. A , B , and C are composed of combinations of the Dirac (F_1), Pauli (F_2), axial (G_A), and pseudoscalar (G_p) nucleon form factors [29]. F_1 and F_2 are vector contributions and are well understood from EM form factor relations and electron scattering experiments. G_p is suppressed by a factor of m_ν^2/M^2 . This leaves us with only the axial-vector form factor G_A , which can only be probed through weak interactions [1, 29]. Measurements of the neutrino CCQE cross section can thus act to constrain G_A .




Figure 1.5: Feynman diagram of the ν_μ CCQE interaction, which is the primary interaction channel at T2K.

The most accepted expression for G_A is the dipole form,

$$G_A(Q^2) = \frac{g_A}{\left(1 + \frac{Q^2}{m_A^2}\right)^2}. \quad (1.48)$$

$g_A \approx 1.3$ is known from measurements of the neutron lifetime [1] so we are left with only a single unknown parameter m_A , the axial mass. Thus, measurements of the CCQE neutrino cross section essentially help constrain m_A . Until recently, the best guess for m_A was (1.026 ± 0.021) GeV [31], but tension between results from ongoing experiments have raised more questions about the QE channel [29].

1.4.2.2 Resonant Pion Production

At higher energies, neutrinos can inelastically scatter off of nucleons. One such case is resonant scattering whereby the neutrino excites a nucleon to a baryon resonance before decaying to a single or multiple pions and its ground state.

$$\nu_l + N \rightarrow N^* + l^- \rightarrow N' + l^- + \pi's, \quad (1.49)$$

where N and N' can be either n or p and N^* is the resonant state. The interaction shown in Eq. (1.49) is CC, but resonant pion production occurs in NC interactions as well. Resonant scattering has historically been calculated using the Rein-Sehgal model [32]. By treating the nucleon as a bound state of three quarks, they show that excited states can be produced by the weak current via neutrino interactions. These states can then decay into a pion-nucleon system as shown in Eq. (1.49).

1.4.2.3 Deep Inelastic Scattering

DIS occurs when the energy of the incoming neutrino is high enough to resolve the individual constituent quarks. In a sense, the vector boson interacts directly with a quark in the nucleon. The CC process can be written as

$$\begin{aligned}\nu_l + N &\rightarrow X + l^- \\ \bar{\nu}_l + N &\rightarrow X + l^+, \end{aligned} \tag{1.50}$$

where X symbolizes a hadronic shower resulting from the breaking apart of the nucleon.

The formalism for calculating the DIS amplitude is based on the quark-parton model which makes the following assumptions [1]:

- Nucleons are composed of quarks.
- Interactions between quarks can be neglected in the DIS regime.
- In the Breit frame, constituent quarks have momenta in the same direction as the nucleon.
- In the Breit frame, quark masses are negligible.

The Breit frame is a uniquely defined reference frame in which $q^0 = 0$ and $\mathbf{q} = -2x\mathbf{p}_N$, where $x = \frac{Q^2}{2p_N \cdot q}$ is the Bjorken Scaling factor [1]. The regime in which DIS dominates corresponds to $p_N^2 \gg m_N^2$. Under these assumptions it is possible to write the momentum of an individual quark, p_i , as a fraction of the nucleon's momentum,

$$\mathbf{p}_i = x\mathbf{p}_N,$$

and since m_q is negligible, the relationship holds for the four-momenta as well $p_i = xp_N$ [1]. Now, for each value of x the cross section must be dependent on the probability of hitting a quark with fractional nucleon momentum x . This probability is called the parton distribution function (PDF) and can be obtained through a combination of experiment and QCD. The PDFs then contribute to the form of the structure functions that determine the DIS cross section [1]. Full Monte Carlo simulations of DIS interactions typically incorporate modifications from the Bodek-Yang model,

which accounts for target mass corrections through the use of a new scaling factor and also corrects parton distributions at low Q^2 [33].

1.4.3 Nuclear Effects

Neutrino detectors are typically made of heavy nuclear targets, not single nucleons. Nuclear effects and FSI introduce additional complexities for cross section measurements. This proves especially painful when it comes to RES interactions as pions are likely to rescatter, be absorbed, or charge exchange within the nucleus [29]. The exiting pion, if it exists, may have different kinematics and identity from the original. All this forces a reconsideration of what truly defines an experimental measurement and what is simply another test of a complex theoretical model.

Several nuclear models exist for describing neutrino-nuclear interactions. Perhaps the oldest and most commonly used is the relativistic Fermi-gas (RFG) approximation by Smith and Moniz [34]. More recent models include Spectral Functions (SF) and Random Phase Approximation (RPA) [35, 36]. These models prescribe how a specific nucleon within the nucleus interacts with an incoming neutrino. Furthermore, recent experimental discrepancies have led to an emphasis on understanding the effects of multiple-nucleon scattering. A consequence of neutrino interactions on multiple-nucleon is the 2-particle 2-hole (2p-2h) effect, whereby two nucleons are emitted from the interaction. This process is mediated primarily by the meson exchange current (MEC) [37]. In the case of MEC, a pair of nucleons can exchange a virtual meson, typically a pion, and this virtual meson couples to the weak vector boson mediating the interaction. In a sense, the neutrino interacts with a pair of nucleons as opposed to a single nucleon. MEC has been shown to produce results in agreement with QE data from MiniBooNE and is believed to be responsible for the dip-region observed in electron scattering [37]. Finally, intranuclear hadron transport must be modeled before a particle can exit the nucleus. Recently, much work has been done on neutrino Monte Carlo generators to incorporate new nuclear models and fine-tune the FSI [33], but many uncertainties remain.

For these reasons, experimental measurements of neutrino cross sections are often quoted in channels of their final state topology, not the true single-nucleon reaction. Nuclear effects make

it virtually impossible to perfectly distinguish QE from RES or RES from DIS. Experiments can only measure what the detector detects, and for neutrino interactions this reduces to the particles that exit the nucleus. These topologies are given terms such as $CC0\pi$, $CC1\pi^+$, or $CCOther$, which correlate with $CCQE$, $CCRES$, and $CCDIS$ but are not identical. The $CC0\pi$ topology corresponds to events with a single charged lepton from the neutrino interaction, zero mesons, and any number of nucleons, the $CC1\pi^+$ topology is similar but requires a single π^+ in the final state, and $CCOther$ corresponds to any other charged-current interaction topology.

1.5 Why Measure Cross Sections?

In general, neutrino-nucleon interactions provide a useful mechanism for studying the weak interaction and the axial vector component of the weak force Lagrangian. Additionally, neutrino-nucleus interactions can shed insight on intranuclear effects and hadron transport within the nucleus. These interactions will differ across target nuclei and thus it is important to measure neutrino cross sections for a variety of different targets.

A better knowledge of neutrino cross sections can also help reduce uncertainties for neutrino oscillation analyses. Neutrino oscillation experiments typically measure an event rate, which is largely dependent on the product of the flux and the cross section. A better knowledge of the cross section allows for a better constraint on the neutrino flux, which is the key ingredient for all oscillation analyses. With neutrino beams coming of age in the last decade or so, cross sections in the intermediate energy range have come under sharper focus. Much of the previous measurements in Fig. 1.4 come from bubble chamber or spark chamber data gathered in the 1970s and 1980s [29]. However, due to the discovery of neutrino oscillations and the importance of the intermediate energy regime for accelerator-based neutrino experiments, there has been a renewed effort towards measuring cross sections in these energy ranges. Experiments such as ArgoNeuT [38], MINERvA [39], MiniBooNE [40], and SciBooNE [41] at Fermilab, and K2K [42] and T2K [43] in Japan, have all started to study this regime [29].

Figure 1.4 shows neutrino-nucleon interaction channels across a wide band of neutrino energies.

At the T2K peak energy of approximately 0.6 GeV, the QE channel dominates. This means the T2K near detector is capable of providing a high-statistics sample of CCQE interactions, which allows us to perform a double-differential cross section measurement across many bins in outgoing muon kinematics.

The reason for a measurement in muon kinematics and not the neutrino energy is that the muon kinematics are the observables. A measurement in terms of the neutrino energy would introduce additional model dependence. Even in the simple case of a pure CCQE interaction, the outgoing proton often escapes detection, and a reconstructed neutrino energy based purely on the muon momentum is an approximation. In the case of interactions with FSIs, a reconstruction-based neutrino energy will be subject entirely to the model. A measurement in the observable lepton kinematics allows more flexibility for model builders to test their models, given the T2K flux.

The current knowledge of CCQE neutrino interactions suffers from a lack of agreement on the axial mass parameter m_A , which governs Eq. (1.48). The world average neutrino-deuteron scattering and pion scattering data has $m_A \approx 1.0 \text{ GeV}/c^2$ [44]. However, both the K2K and MiniBooNE experiments have measured a larger than expected axial mass of around 1.2–1.4 GeV/c^2 . The MINERvA experiment, however, has more recently measured $m_A = 1.00 \text{ GeV}/c^2$, in agreement with previous measurements [44]. As these experiments all involve nuclear targets, attempts to resolve their discrepancies have revolved around better modeling of the nucleus and additional nuclear effects, such as RPA and multinucleon interactions [36, 44]. Recently, a combined fit of the MiniBooNE and MINERvA data returned a best fit of $m_A = (1.15 \pm 0.03) \text{ GeV}/c^2$ using a RFG model with relativistic RPA and multinucleon effects [44]. However, significant tension between the datasets means that a more conclusive result requires additional data. This involves cross section measurements on different nuclear targets, which can help further understanding of nuclear effects.

The closest measurable topology to CCQE is $\text{CC}0\pi$. The $\text{CC}0\pi$ cross section has been measured on carbon but not oxygen, and a detector with water targets provides an ideal venue towards measurements on oxygen. Water-based detectors are used in a variety of neutrino experiments. This includes the T2K far detector, Super-Kamiokande, which is a water Cherenkov detector. As

mentioned previously, neutrino cross sections may exhibit dependence on the nuclear targets, and thus there are several parameters that govern specific nuclei in most neutrino interaction models. These, such as the binding energy E_b and the Fermi momentum p_F , are subject to uncertainties that can be constrained by the near detector in oscillation analyses, provided that the near detector dataset corresponds to the same target as the far detector. While the T2K near detector, ND280, does contain water targets to match SK, events on those targets have yet to be incorporated into the oscillation analysis [45]. Therefore, the largest remaining source of uncertainty in the event selection at SK is due to those independent, unconstrained cross section parameters, as shown in Table 1.2. A cross section measurement in the most dominant interaction channel for many neutrino oscillation experiments on a target used by many neutrino detectors is therefore of crucial importance.

Source	Relative Error [%]
Flux and ND constrained cross sections	2.7
Independent cross sections	5.0
SK detector systematic	4.0
Nuclear effects	3.0

Table 1.2: The relative error on the SK ν_μ CC selection used for the oscillation analysis after ND280 constraints [45]. The largest single source of uncertainty are the cross section parameters which are not shared between ND280 and SK, which includes the parameters governing the oxygen nucleus.

This thesis reports a flux-averaged, double-differential measurement of the muon-neutrino $\text{CC}0\pi$ cross section on water. The analysis relies on an unfolding technique to correct for detector misreconstructions and inefficiencies. In order to extract measurements on water, the T2K Pi-zero detector (PØD, Section 2.4) was run in two different water-fill configurations. This allows for a subtraction of the measured rate during the water-in and water-out periods, which ultimately provides a measurement on water only. In order to remain unbiased by the data, the entire sequence of steps outlined in Chapter 7 was first performed and tested entirely on Monte-Carlo simulations. Only after the measurement procedure was approved by an internal review committee did we unblind the analysis.

Chapter 2

The T2K Experiment

2.1 Overview

In the overall scheme of neutrino experiments the T2K experiment falls in the LBL accelerator category. As shown in Fig. 2.1, neutrinos produced at the Japan Proton Accelerator Research Complex (J-PARC) in Tokai-mura, Ibaraki, Japan are directed towards first a near-detector (ND) site 280 m from the production target before traveling across Japan to SK 295 km away. Unlike previous LBL experiments, however, T2K's primary detectors are placed 2.5° off-axis from the beam line. This allows for detection of narrow-band muon-neutrinos with a peak energy of approximately 0.6 GeV, which maximizes the oscillation probability at SK as shown in Fig. 2.16 [46]. Designed with the primary goal of measuring the mixing parameter θ_{13} , T2K has also produced precision measurements of $\sin^2(2\theta_{23})$ and Δm_{23}^2 [21].




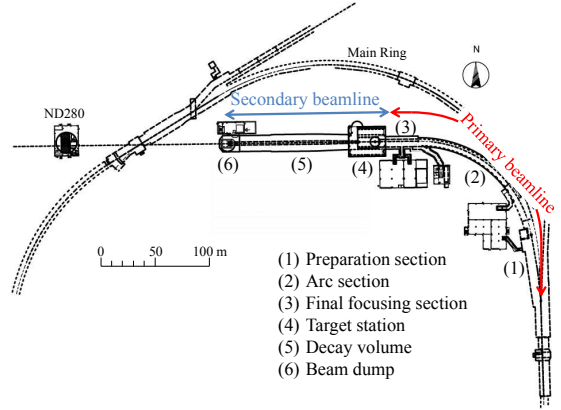
Figure 2.1: Schematic of the T2K neutrino beam with J-PARC, ND280, and SK labeled. The beam travels across Japan from east to west.

At J-PARC, protons are injected into the main ring (MR) synchrotron and accelerated to 30 GeV. Each MR spill is subdivided into eight bunches and fast extracted for the neutrino beam

line [46]. The MR beam has a designed power of ~ 750 kW with a beam intensity of $\sim 3 \times 10^{14}$ p/spill. Both the J-PARC facility and the T2K neutrino beam line are shown in Fig. 2.2. After extraction, the protons are directed towards the target station and impinge on a graphite target producing secondary pions. A set of three focusing horns act together to select either π^- 's or π^+ 's, depending on whether we want neutrinos or antineutrinos. The pions then decay primarily into muons and muon-neutrinos in the ~ 96 m long decay volume.



(a) Aerial view of the J-PARC complex.



(b) The T2K neutrino beam line.

Figure 2.2: The J-PARC facility and an outline of the T2K beam line.

T2K's near detector site is situated at J-PARC 280 m from the target. It houses both an on-axis detector (INGRID) and an 2.5° off-axis detector (ND280). INGRID provides measurements of the on-axis neutrino direction and confirms beam integrity. The ND280 consists of several subdetectors placed within a magnetic field. It allows for measurements of both the ν_μ and ν_e flux at the same off-axis angle as T2K's far-detector but prior to oscillation. This helps characterize the expected ν_e background contamination in ν_e appearance analyses and the overall ν_μ signal in disappearance analyses [46]. Additionally, both INGRID and ND280 provide neutrino cross section measurements on various detector targets. Our current understanding of neutrino interactions, especially on heavy nuclei, leaves open a lot of questions. Thus, it is desirable to utilize all possible datasets to help

reduce uncertainties.

Finally, neutrinos produced at J-PARC travel across Japan to T2K's far detector SK. Super-Kamiokande is located 1000 m underground in the Mozumi mine, below Mt. Ikenoyama [46]. It is a water Cherenkov detector, consisting of 13 000 photomultiplier tubes (PMTs) that line the inside of a 39 m diameter by 42 m tall cylinder [46]. With a water capacity of 50 000 t, SK can detect and identify the outgoing charged lepton from a CC interaction, making it ideal for a LBL oscillation experiment.

The first T2K physics run began in January 2010. The total accumulated protons on target (POT) and beam power up to the time of this writing are shown in Fig. 2.3. In 2014, after completing its objective of observing ν_e appearance, T2K flipped the polarity of its focusing horns and switched to running in antineutrino mode for the first time.




Figure 2.3: The accumulated POT and beam intensity at T2K over all run periods as of the writing of this thesis. The data sample used for this analysis will be most of T2K Runs 2–4. Two disruptions, over which T2K had no control, occurred that affected the data taking periods.

2.2 Creating Neutrinos for T2K

2.2.1 The J-PARC Accelerator

As mentioned earlier, accelerator neutrinos are produced via pion-decay, and in T2K pions are produced via energetic proton collisions in a graphite target. The relatively newly constructed J-PARC facility in Tokai consists of three accelerators: the linear accelerator (LINAC), the rapid-cycling synchrotron (RCS), and the main ring (MR) synchrotron. Initially, a H^- beam is accelerated by the LINAC to ~ 400 MeV [46]. The H^- beam is converted to H^+ by charge-stripping foils and then injected into the RCS, which further accelerates the beam to ~ 3 GeV. A fraction of these bunches are supplied to the MR, which provides the protons needed for both T2K, via fast extraction, and the hadron facility, via slow extraction [46]. For T2K, the MR accelerates protons to ~ 30 GeV in eight bunches, all of which are extracted in one turn into the T2K beam line using five kicker magnets at a spill cycle of ~ 0.5 Hz. A typical extraction contains $\sim 10^{14}$ POT.

2.2.2 Proton Beam Monitoring

The T2K beam line consists of two sections, the primary beam line and the secondary beam line, both of which are shown in Fig. 2.2. The primary beam line turns protons towards Kamioka, dipping 3.64° downwards with respect to the horizontal. The secondary beam line contains the target and focusing horns for generating and directing secondary pions before their decay into neutrinos [46]. The distribution of the proton beam can have an important effect on the generated neutrino flux, and a sophisticated system was designed and put in place in the primary beam line to measure and verify beam integrity prior to impinging the target. The beam monitoring system consists of five current transformers (CTs), 21 electrostatic monitors (ESMs), 19 segmented secondary emission monitors (SSEMs) and 50 beam loss monitors (BLMs) for monitoring the beam intensity, position, profile, and loss respectively [46]. Each CT is a 50 turn toroidal coil that measures the current generated by the beam as it passes. It has an absolute intensity sensitivity of 2 % and a timing precision of 10 ns [46]. The ESMs monitor the beam position with four electrodes surrounding the proton

beam, thus allowing it to measure the beam-induced current at four corners and non-destructively monitor the top-bottom and left-right asymmetry [46]. ESMs measure beam position to a precision of less than $450\text{ }\mu\text{m}$. The SSEMs monitor beam profile using a HV foil sandwiched between two thin titanium foils that are stripped horizontally and vertically. As the proton beam passes through the strips it produces secondary electrons which then drift along the electric field inducing currents along the strips [46]. SSEMs cause beam loss and are only inserted during beam tuning. The beam width uncertainty is less than $200\text{ }\mu\text{m}$. Finally, beam loss is monitored along the primary beam line with BLMs, wire proportional counters. Its beam loss sensitivity is below 16 mW [46].

2.2.3 The Secondary Beam Line

The secondary beam line shown in Fig. 2.4 follows the primary and contains the target station, decay volume, and beam dump. Secondary pions produced from the proton beam impinging the target are focused by three magnetic horns and decay in flight within the decay volume into muons and muon-neutrinos. The beam dump stops most of the muons below 5 GeV and those that pass through are monitored to further characterize the neutrino flux. The neutrinos pass through the beam dump and are used for physics measurements.

Figure 2.5 shows a schematic of the T2K target system. Within the target station, the target, horns, and optical transition radiation (OTR) monitor are located within a helium vessel. The OTR is a last line beam profile monitoring device consisting of a thin titanium foil placed at 45° relative to the incident proton beam. Transition radiation is produced as the proton beam enters and exits and is used to measure the beam profile just before target incidence [46].

The T2K target consists of a graphite core surrounded by a graphite tube all sealed within a titanium case placed within the inner conductor of the first magnetic horn [46]. The core is a cylinder with length 91.4 cm and diameter 2.6 cm . Helium gas cycles between the core and tube, and tube and case to cool the entire assembly. At the designed beam power of 750 kW the center of the target is expected to reach 700°C [46].

To get a neutrino or antineutrino beam T2K selects the charge of the pions using three focusing

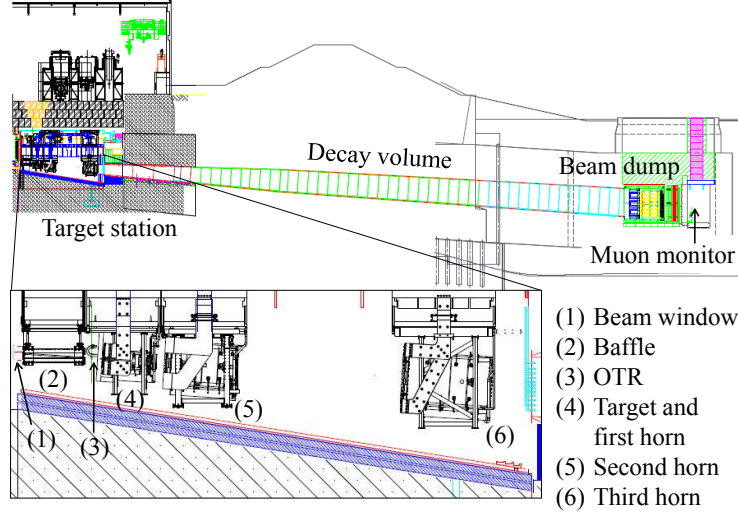


Figure 2.4: An overview of T2K’s secondary beam line with objects of importance labeled. The beam enters on the left.

horns. In neutrino mode the horns focus π^+ ’s and defocus π^- ’s, and vice versa in antineutrino mode. Each horn relies on two coaxial conductors, an inner and an outer, made of aluminum alloy to generate a toroidal magnetic field that varies as the inverse of the distance from the horn axis [47]. The uncertainty on the magnetic field strength is $\sim 2\%$. Each horn was designed to be capable of operating independently, but currently they are not all powered separately. So far the current on all three have been set identically for most physics runs. With a 250 kA current, the peak neutrino flux (~ 0.6 GeV) at SK is ~ 17 times higher than with 0 kA [47]. See Fig. 2.6 for a comparison between the predicted unoscillated flux at SK for different horn currents. It should be noted that T2K began taking antineutrino data in 2013 and this involved reversing the horn currents to -250 kA. However, the results in this thesis will focus on data collected in the neutrino-mode runs.

The decay volume and beam dump comprise the last parts of the secondary beam line. Pions are focused into the decay volume, which is a 96 m long channel with an upstream cross section of 1.4 m wide by 1.7 m high and a downstream cross section of 3.0 m wide by 5.0 m high. It is filled with helium at 1 atm to suppress pion absorption [46]. The 75 t graphite beam dump at the end of the decay volume stops decay muons below 5 GeV. Higher energy muons escape the beam dump



(a) T2K target schematic.



(b) Target and first horn.

Figure 2.5: Schematic of the T2K target and picture of the target outside the first horn.

into the downstream muon pit. Since muons are typically produced in conjunction with neutrinos, their direction is highly correlated with the neutrino beam. For this reason, a muon monitor is placed after the beam dump, 118 m from the target. Consisting of two types of detectors, ionization chambers and silicon PIN photodiodes, over a $150\text{ cm} \times 150\text{ cm}$ area, the muon monitor measures the direction of the neutrino beam to a precision of 0.25 mrad [47]. Any remaining muons stop in the $\sim 180\text{ m}$ of sand between the dump and the near detector.

2.3 The Near Detector Facility

Even with all the sophisticated beam monitoring devices mentioned in Section 2.2, our knowledge of the absolute neutrino flux still carries a large uncertainty. Although primary protons in the beam line are measured in detail, uncertainties in the hadron production normalizations and kinematics contribute to an overall absolute flux uncertainty of $\sim 10\%$. Detailed beam Monte Carlo (MC) simulations provide our best estimates of the neutrino flux exiting the T2K beam line but additional data inputs, both external and internal, are required to constrain the flux uncertainty. Thin target data from the dedicated hadron production experiment NA61/SHINE at CERN serves as an external input [48]. Measurements made at ND provide on-axis and off-axis data to constrain the




Figure 2.6: Unoscillated flux prediction at SK for various horn currents.

unoscillated flux. These can also be made into independent cross section measurements. However, because the detected events result from a product of flux and cross section, care must be taken to ensure that the same data is not doubly applied.

The ND facility contains two detectors, INGRID and ND280, for on-axis and off-axis measurements respectively. The detectors are housed in a pit 37 m deep and three stories high, buried in the sandy earth of J-PARC. The right-handed ND280 coordinate system is oriented such that z lies horizontally along the nominal neutrino direction and x and y are horizontal and vertical respectively. A drawing of the complex is shown in Fig. 2.7.

2.3.1 INGRID

The on-axis detector, INGRID (Interactive Neutrino Grid), monitors beam intensity and direction by measuring interactions on iron. Shown in Fig. 2.8, it uses 16 identical detectors arranged horizontally and vertically into a cross around the beam center [46, 47]. Distributions of neutrino events across the modules provide a beam center resolution of ~ 0.4 mrad at ND.

Each INGRID module comprises a sandwich structure of nine iron planes interspersed between 11 scintillator planes. The total iron target mass is 7.1 t per module and the tracking scintillator planes are composed of 24 horizontal and 24 vertical bars [46]. Each bar is read out using a




Figure 2.7: The ND facility with both INGRID and ND280 depicted. Note their positions relative to the incoming neutrino beam. INGRID is placed along the beam axis, ND280 is 2.5° off-axis.

wavelength-shifting (WLS) fiber attached to a multi-pixel photon counter (MPPC). Both of these technologies are used extensively in the ND280 as well.

Finally, between the two central modules lies an additional Proton Module. It uses differently sized scintillator bars to provide improved tracking abilities and increased QE identification efficiencies [46]. Recently, INGRID has been used to produce a CCQE cross section measurement on carbon [49], in addition to its beam monitoring tasks.

2.3.2 The ND280 Off-Axis Detector

As ND280 sits along the direction towards the far detector, its spectrum ideally resembles the SK spectrum. Neutrinos that pass through on their way to SK sometimes interact within its array of detectors, allowing measurements that constrain the unoscillated flux, background ν_e contamination, and independent cross sections on various targets. The subcomponents of ND280 are all shown in Fig. 2.8. The primary tracking components fit in a $6.5\text{ m} \times 2.6\text{ m} \times 2.5\text{ m}$ (length \times width \times height) container called the “basket” which consist of the π^0 detector (PØD) and the tracker, itself consisting

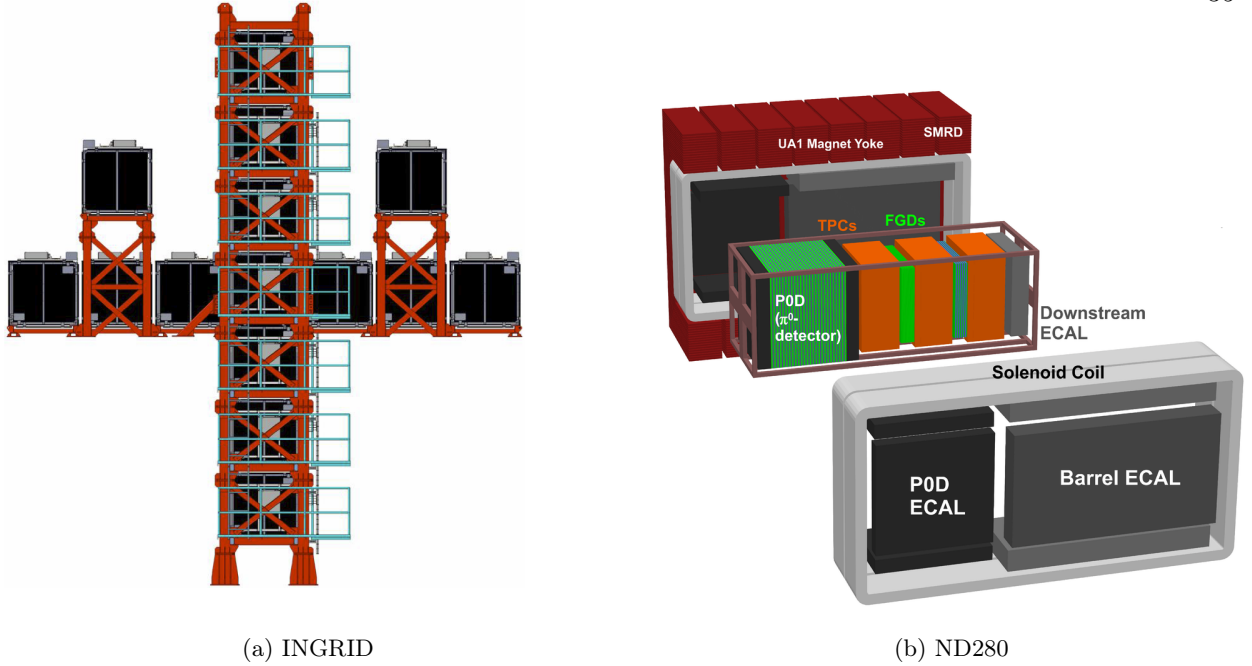


Figure 2.8: The INGRID detector comprises 16 identical neutrino detectors arranged vertically and horizontally around the beam center. The center of the INGRID cross consists of two overlapping modules and the two off-axis modules checks the axial symmetry of the neutrino beam [46]. ND280 is composed of several subdetectors all surrounded by UA1 magnet. In this figure, the beam points out of the page for INGRID and left-right for ND280.

of three time projection chambers (TPCs) interspersed with two fine grained detectors (FGDs). Surrounding both the tracker and PØD are electromagnetic calorimeters (ECals) designed to provide coverage for side-exiting tracks and showers. This entire assembly is enclosed within the refurbished UA1/NOMAD magnet with scintillators that act as a side muon range detector (SMRD).

The UA1 magnet produces a 0.2 T magnetic field which allows for particle momentum resolution within the basket. It is driven by a nominal current of 2900 A and is composed of two mirror-symmetric halves which can be opened for inner detector access during maintenance periods. In between the steel plates lining the outside of the magnet are 440 scintillator panels that make up the SMRD. Its main purpose is to provide momentum measurements of high-angle muons that escape through the sides of the ND280. Additionally, the SMRD can act as a cosmics trigger and a veto of external beam interactions [46].

Two ECals lie immediately within the magnet structure and surround the basket. Constructed

out of scintillator panels sandwiched with lead sheets for absorption, the ECals detect photons and charged particles and measure energy deposition and direction. Readout is done using WLS-fibers attached to MPPCs, similar to INGRID and the PØD. The PØD-ECal surrounds the four sides of the PØD along the z-axis, the Barrel-ECal the tracker, and a downstream module (Ds-ECal) covers downstream exit of the tracker [46]. Together with the inner detectors, they are designed to provide complete phase space coverage for interactions within the basket.

The FGDs are two $2300\text{ mm} \times 2400\text{ mm} \times 365\text{ mm}$ (width \times height \times length along beam direction) scintillator based detectors sandwiched between three TPCs. The active region consists of scintillating bars oriented perpendicular to the beam in alternating x-y layers. These XY modules act as a target for neutrino interactions and their alternating pattern provides tracking capabilities for outgoing particles. The more upstream FGD, FGD-1, is composed of 15 XY modules. The downstream FGD-2 contains seven XY modules alternating with six modules filled with water [46]. Measurements from both detectors can be used in conjunction to extract separate cross sections on carbon and water.

At T2K's energy, the most common reaction is CCQE. Measurement of this interaction gives information on the flux and helps to predict the spectrum at SK [50]. The tracker was designed to distinguish CCQE reactions from other charged-current reactions and thus predict unoscillated event rates at ND and SK [50, 51]. This is possible by selecting events that contain only an outgoing muon and proton. T2K uses three argon-based TPCs numbered TPC1, TPC2, and TPC3 from upstream to downstream. In conjunction with the FGD's electron tagging and vertex resolution capabilities, the 3D tracking capabilities of the TPCs provide three crucial ingredients necessary towards this end. First, the TPC provides excellent positional resolution which allows for precise track reconstruction. Second, the magnetic field in which the TPC is placed allows for momentum reconstruction of charged particles. Third, ionization and energy loss over the course of the track gives particle identification (PID) information [46]. Each TPC consists of an inner box filled with 1 atm argon gas surrounded by an outer box filled with insulating CO_2 . As charged particles traverse through, they knock off ionization electrons that then drift along the electric field towards the

readout plane. There, the electrons are multiplied and sampled with micromegas detectors with $7.0\text{ mm} \times 9.8\text{ mm}$ (vertical \times horizontal) anode pad segmentation [46]. These components are shown in Fig. 2.9. They give a spatial resolution of 0.7 mm and momentum scale resolution of less than 2% [51].

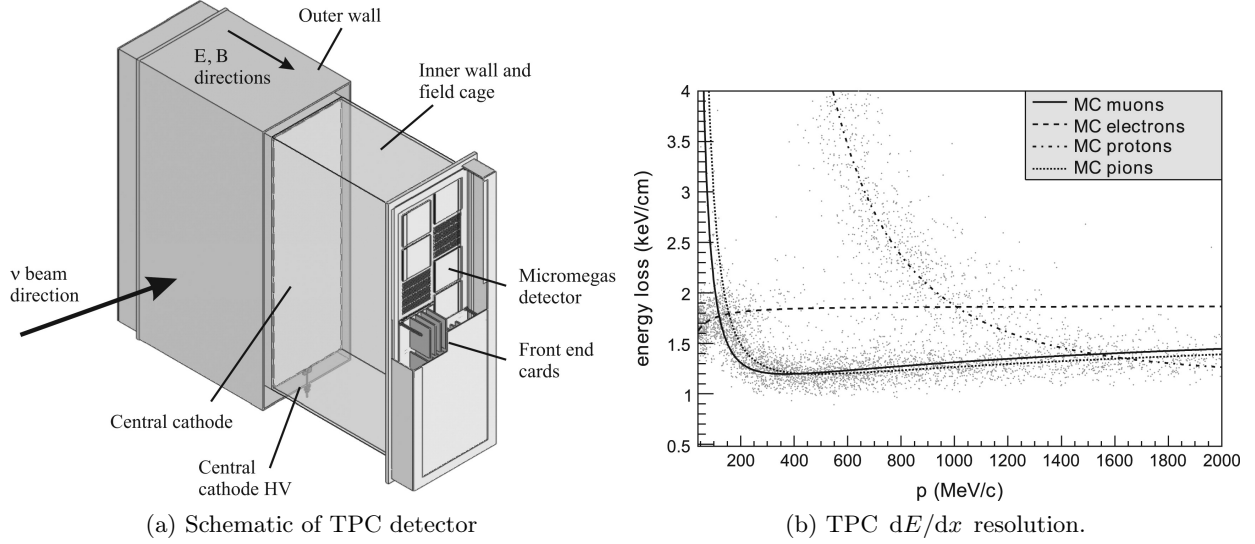


Figure 2.9: The TPC provides excellent tracking capabilities and particle identification.

The results of this thesis are based on the combined reconstruction from the PØD and tracker. The former provides the necessary target mass to extract a high statistics measurement, the latter precise kinematic reconstruction to probe a large region of phase space. This is illustrated by the beautiful event shown in Fig. 2.10. As the interactions occur inside the PØD, these results are said to be a PØD-based measurement. Therefore, a detailed description of the PØD is provided in Section 2.4.

2.4 The ND280 Pi-zero Detector

2.4.1 Detector Description

The PØD lies at the most upstream end of the basket, immediately upstream of TPC1. In the ν_e appearance analysis, a dominant source of background is π^0 from $\nu_\mu \text{ NC}\pi^0$ ($\nu_\mu + N \rightarrow \nu_\mu + N + \pi^0 + X$) interactions. The PØD has been designed precisely to measure this interaction [52], but it is capable




Figure 2.10: A beam triggered event in ND280. This event occurred during run 3 on January 3, 2012. All detectors in the basket region registered hits for this event. The PØD is in brown, the three TPCs are in blue, and the two FGDs are maroon and sandwiched between the TPCs.

of measuring other reaction channels as well. Since the far detector is water Cherenkov, the PØD was outfitted with passive water layers to allow it to measure reaction rates on water for comparison. These water layers lie interspersed between active scintillation modules as shown in Fig. 2.11. At the upstream and downstream ends of the PØD are the upstream and central ECals comprising of alternating layers of scintillator and lead. The ECals act as an effective veto of external tracks originating from upstream of the PØD, and provide effective shower containment [52].

The scintillation modules are composed of two perpendicular (XY) arrays of triangular scintillator bars, 134 horizontal and 126 vertical, and provide readout through WLS-fibers and MPPCs [52]. Each scintillator XY-module is called a PØDule. It is composed of triangular bars made of polystyrene and coated with a reflective layer of TiO_2 to trap the blue scintillation light. A picture of a single scintillation bar is shown in Fig. 2.12. Each bar is an isosceles triangle 17 mm in height with a 33 mm base [52]. Each bar also contains a ~ 2.6 mm hole in which a WLS-fiber is inserted and capped at one end by the readout MPPC. The bars are glued together, in an up-down pattern, to create a single PØDulelayer. The PØDules themselves are then organized into four regions called Super-PØDules as shown in Fig. 2.11. The upstream ECal (USECal) and central ECal (CECal) Super-PØDules sandwich the upstream water target (USWT) and central water




Figure 2.11: Side cut-away view of the PØD and its four Super-PØDules [52]. The blue water PØDules can be filled and drained. The beam travels left-right.

target (CWT) Super-PØDules. Both ECal Super-PØDules contain seven PØDules alternating with seven stainless steel-clad 4.5 mm thick lead sheets [52]. The USWT (CWT) contains 13 PØDule interspersed with 13 (12) water layers (28 mm thick) and 13 (12) brass sheets (1.28 mm thick) [52]. In total, the PØD contains 40 PØDules each with 260 triangular scintillator bars, giving a total of 10 400 channels across the entire PØD. By looking for intersecting X and Y-bars that light up in unison, the PØD can provide hit positions for track reconstruction.

The two water Super-PØDules serve as the main target region of the PØD, which has overall dimensions of $2103 \text{ mm} \times 2239 \text{ mm} \times 2400 \text{ mm}$ along the x , y , and z axes [52]. The water containers are thin bladders made of high-density polyethylene (HDPE), which provide enough strength and support but do not add excess target material. Both the top and bottom of the water bags are sealed with HDPE frames with ports for monitoring apparatus. As the bags themselves do not maintain shape the pressure on both sides of the scintillator PØDule must cancel. Therefore, pressure sensors and depth sensors are crucial from both engineering and analysis perspectives.



Figure 2.12: A scintillator bar used in PØDule construction. The hole for the WLS-fiber is visible in the center of the triangle.

The PØD was designed to provide datasets for water filled and empty periods. In such a way, a subtraction can be performed to extract a water-only measurement, which is the technique used in this thesis. The water fill and drain system sits outside the ND280 detector. In order to provide an evenly distributed amount of water in each layer for pressure balance, the water level in each bag must be carefully monitored such that the water level remains approximately the same in all bags at all times. This thesis relies on T2K Runs 2–4 data, during which 2.3×10^{20} PoT (3.5×10^{20} POT) were delivered during PØD water-in (water-out) operations.

2.4.2 Data Acquisition and Electronics

The PØD, along with the SMRD, ECals, and INGRID, uses a Trip-t based front-end electronics readout [46]. Each board (TFB) contains four 32-channel Trip-t ASIC chips serving 64 MPPC channels [52]. The signal from the MPPC is multiplexed through two dual-channel 10-bit ADCs that digitize the data [46]. A high gain channel is reserved for signals up to 50 p.e. (photoelectrons) and a low gain channel for signals up to 500 p.e. [52]. The TFBs are secured on the top and side of the PØD as shown in Fig. 2.13.

The Trip-t chips divide each beam spill into twenty-three 580 ns cycles which are synchronized with the eight bunches contained in each spill. Each bunch falls into a separate integration cycle [52] with data being stored in a capacitor array [46]. The last 100 ns in the window is a reset period during which no hits can be recorded. This reset period is necessary for the analog Trip-t data to

T2K Run	Dates	POT Delivered	POT (good DQ)
1	March 19, 2010–June 26, 2010	3.11×10^{19}	1.65×10^{19}
2	November 18, 2010–March 11, 2011	1.12×10^{20}	7.89×10^{19}
3	February 27, 2012–June 9, 2012	1.59×10^{20}	1.57×10^{20}
4	October 19, 2012–May 8, 2013	3.60×10^{20}	3.45×10^{20}

Table 2.1: Summary of dates and POT corresponding to T2K runs 1–4. The PØD was filled for run 1 and part of runs 2 and 4.

be digitized by an ADC and moved off to the back-end. Beam bunches are ensured to fall within the active time window, but out-of-bunch events such as Michel electrons and cosmic rays can sometimes fall into the deadtime window. A field programmable gate array (FPGA) controlling the TFB provides a time stamp resolution of 2.5 ns [52].

There are 29 TFBs each for the USECal and CECal and 58 for each of the water targets for a total of 174 TFBs on the PØD [52]. These interface with the back-end electronics via 6 readout merger modules (RMMs) serving 29 TFBs each. The RMMs act as the bridge between the data coming from the TFBs and the ND280 data acquisition (DAQ) system by synchronizing clock and trigger signals in one direction and data the other [52].

In addition to the RMMs, the back-end also consists of a shared cosmic trigger module (CTM), a slave clock module (SCM), and a master clock module (MCM) common to the entire ND280 detector. The detector can therefore trigger on cosmics as well as beam events. The CTM triggers on coincidences between the USECal and trigger primitives from the SMRD and the DsECal, which occur when cosmic ray muons traverse the sides of ND280 [52]. Appropriately, the cosmic ray tracks recorded from this type of trigger are called “Trip-t cosmics”, in contrast to cosmic rays triggered by the FGD which are called “FGD cosmics”. The MCM handles the spill timing and trigger signal transmission, and provides clock synchronization to the RMMs and TFBs via the SCM [52].

The ND280 DAQ relies on the MIDAS framework operating on computing nodes running the Scientific Linux OS. Each front-end processor node (FPN) serves two RMMs and control data coming from the Trip-T detectors [46, 52]. The PØD uses three FPNs. All data collected by the




Figure 2.13: Three dimensional graphic of the PØD with surrounding electronics.

FPNs gets transferred via a Gigabit Ethernet (GbE) network to the backend DAQ run control node. In addition the FPNs interface with the online monitoring (OM) system, which provides real time calibration histograms using the ROOT framework, and global slow control (GSC) system, which relies on a MySQL database to collect monitoring data such as power supply voltages and TFB temperatures [46, 52].

2.4.3 Detector Calibration

Detector calibration is crucial to ensuring accurate physics data and is performed at the end of each run. Charge calibration in the PØD occurs in three steps to convert the raw ADC into photoelectron units: pedestal subtraction, electronics non-linearity correction, and MPPC gain correction [52]. Due to MPPC dark noise, the pedestal acts as the baseline value which must be subtracted to give the actual 1 p.e. to ADC conversion [52]. For each MPPC channel, the ADC distribution across all 23 integration cycles is plotted and fitted to a Gaussian as shown in Fig. 2.14 [52]. The difference between the 1 p.e. peak and pedestal peak is the MPPC gain and gives the ADC to p.e. conversion factor. Furthermore, the ADC response must be corrected for electronic non-linearities before conversion into p.e. units. Each channel's response is measured while

the MPPCs are powered on and non-linearities are parameterized as a bi-cubic polynomial¹. The parameterization is then used offline to correct the raw ADC values [52].




Figure 2.14: Digitized MPPC dark noise spectrum with the pedestal and 1 p.e. peak each fitted to a Gaussian. The difference between the means gives the ADC to 1 p.e. conversion.

In order to calibrate the overall PØD layer response to through-going particles, the PØD relies on minimum ionizing particles (MIPs), which have an approximately constant energy loss for different energies [52]. Cosmic ray muons and muons from beam neutrino interactions in the rock or sand outside the ND280 (sand-muons) are ideal candidates to calibrate the PØD’s MIP light yield. Each channel in the PØD can be calibrated to the MIP mean energy deposition per unit length, thus ensuring uniform detector response.

Finally, the PØD was constructed with a light injection system (LIS) for all 10 400 channels. The LIS exposes each MPPC to a varying scale of light intensities and allows controlled monitoring of photosensor response. For calibration purposes, the LIS cycles through ten amplitudes to ensure correct MPPC performance [52]. A similar procedure is used for other Trip-T detectors.

2.5 Super-Kamiokande

The results presented in this thesis do not use T2K’s far detector. Nevertheless, a brief overview of Super-Kamiokande can motivate much of the work done on T2K and at ND in particular.

The largest water Cherenkov detector in the world with 50 000 t of water (Fig. 2.15), SK sits

¹ A bi-cubic polynomial is a function of the form $\sum_{i=0}^3 \sum_{j=0}^3 a_{ij} x^i y^j$.




Figure 2.15: Rendered diagram of the Super-Kamiokande detector with PMTs clearly visible. For a sense of scale, note the size of humans in the walkway above.

2.5° off-axis from the neutrino beam line, allowing for a peak energy that maximizes the oscillation probability at its 295 km baseline as shown in Fig. 2.16. Constructed underground in the Mozumi mine of the Kamioka Mining and Smelting Company in Gifu, Japan, the detector lies below 1000 m of rock or 2700 meter water equivalent (m.w.e) [53]. This provides shielding from cosmic ray muons with energy less than 1.3 TeV. Commissioned in 1996, SK underwent repair after an accident in 2001. Since 2006, the PMT surface area coverage has been restored to 40 % [53].

SK detects neutrino interactions via Cherenkov light emitted by the produced charged particles. Equipped with both inward and outward facing PMTs, SK can identify entering particles and thus distinguish backgrounds from fully contained (FC), neutrino induced events [53]. Furthermore, the shape, size, and orientation of the Cherenkov pattern on the inner detector categorizes the track as single-ring electron-like, single-ring muon-like, or multi-ring [53]. Two example FC events from neutrino interactions within the inner volume are shown in Fig. 2.17. Since electrons shower and muons do not, they produce distinctive patterns of Cherenkov light.




Figure 2.16: The power of an off-axis neutrino beam is illustrated here. At an off-axis angle of 2.5° the neutrino flux peaks at the disappearance maximum. The narrow-band flux minimizes backgrounds from unoscillated ν_μ 's.

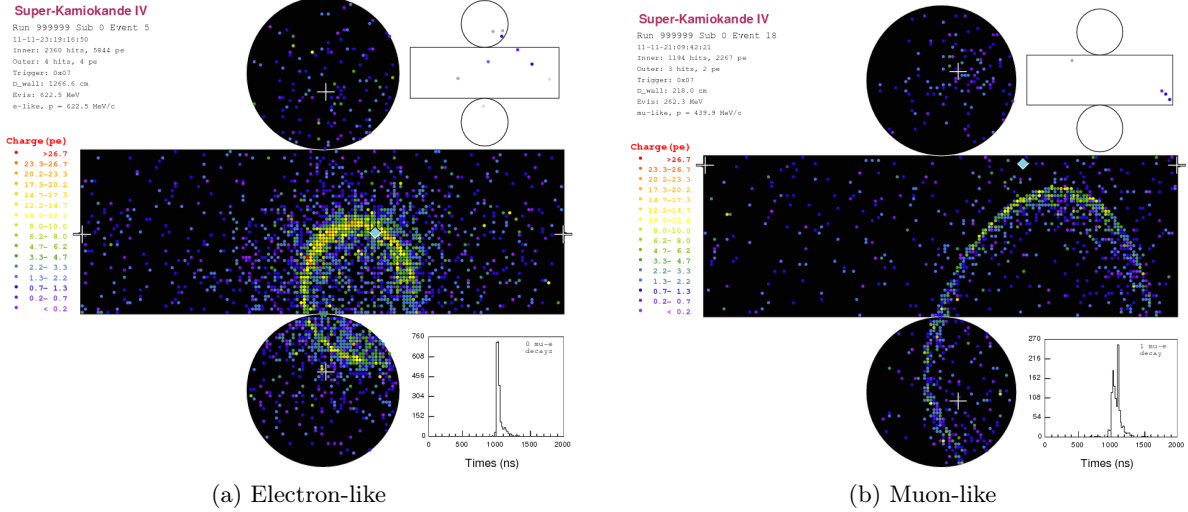


Figure 2.17: The Cherenkov radiation patterns produced along the wall of the inner-facing detector allow for particle identification. Here, the right plot shows a ν_e and the left a ν_μ FC event. Since electrons shower, e-like patterns appear fuzzier than their μ -like counterparts.

Water acts as the primary target for neutrino interactions at SK. While several measurements

of cross sections on carbon have been performed in the past by K2K, MiniBooNE, SciBooNE, and MINER ν A [54, 55, 56, 57], neutrino cross sections on water are less well known. T2K’s near-detector contains two detectors with water as a passive target: the FGD2 and the PØD. One major purpose of both detectors is to precisely measure various neutrino interaction channels on water, thus providing useful signal and background predictions at SK. Additionally, direct water cross section measurements act as tests of interaction models and nuclear effects as discussed in Section 1.4. Modern HEP experiments rely on accurate MC simulations to encapsulate our understanding of nature in its entirety, of which much depends on empirical parameters with large uncertainties. Before moving on to results, we discuss the computational aspects of making a cross section measurement in the next chapter.

Chapter 3

Monte Carlo Simulation and Reconstruction

To measure something that cannot be directly detected is no simple task. A neutrino cross section measurement requires knowledge of several factors that must be inferred indirectly, such as the flux and the total interaction rate. Due to this incomplete knowledge, assumptions must be made on the basis of existing theories. Since experimental data is typically interpreted in terms of counts or event rates dependent on some underlying probability distribution, Monte Carlo (MC) simulations are an ideal tool with which to test these assumptions. MC attempts to encapsulate all the available knowledge on neutrino production and interaction. In this way, any differences between MC and data can be said to contain new physics, or at the very least hint at inadequacies in the current models.

The T2K MC can be roughly divided into three stages: neutrino flux generation, neutrino interaction modeling, and detector simulation. Each plays an important role for both cross section measurements and oscillation analyses. In this chapter we provide an overview of each part in the MC chain, followed by a discussion of the analysis framework used for data reduction, selection, and cross section extraction.

Throughout the remainder of this thesis, a **typewriter** font will be used when referencing a computing package, command, or tool. However, normal font will be used when package names are meant to be interpreted as acronyms or in a non-computing context.

3.1 Flux Simulation

Accurate and precise generation of the neutrino flux is the first step in producing a detailed simulation of most neutrino experiments. In long-baseline experiments like T2K, flux comparisons are the basis upon which oscillation parameters are extracted. For absolute cross section measurements, both the event rate and the overall normalization are dependent on the neutrino flux. The flux simulation includes modeling of hadronic interactions within the target based on proton beam inputs, followed by decays of the secondary hadrons which produces the neutrino beam.

The secondary beam line, discussed in Section 2.2.3, encompasses the target station, decay volume, and beam dump. Beginning with the target, a package for multi-particle transport called FLUKA [58] is used to model interactions of 30 GeV primary protons on the graphite target [47]. After this stage, secondary hadrons, such as pions and kaons, are produced and passed to the beam line simulation.

The FLUKA output is transferred to a GEANT3 [59] based MC simulation called JNUBEAM. JNUBEAM simulates secondary beam line geometries and is ultimately the MC that transports the neutrino flux to INGRID, ND280, and SK. JNUBEAM uses GCALOR to simulate additional hadronic interactions outside of the target. The geometries and relative positions of these various components are based on mechanical drawings of the beam line and survey results [47]. Particles are allowed to propagate in the helium-filled decay volume until they interact or decay. Neutrinos from these decays comprise the flux shown in Fig. 3.1, produced in neutrino-mode running with horn currents set to 250 kA.

To save computing time, neutrinos produced in JNUBEAM are always projected towards SK or the near detector plane. At the ND280, there are two planes for which the flux is provided. One encompasses the entire ND280 including the magnet and is referred to as the “magnet” production. The other is restricted to the ND280 basket and is referred to as the “basket” production. Both geometries are described in Section 2.3.2. It is then possible to calculate interaction probabilities under the assumed direction. These are saved along with the neutrino kinematics such that the




Figure 3.1: T2K unoscillated flux prediction at SK with bands indicating the systematic uncertainty prior to ND fitting. The primary contribution is ν_μ , but there is an $\sim 3\%$ wrong-sign contamination and $\sim 1.2\%$ intrinsic beam ν_e background.

proper flux and energy spectrum can be obtained by reweighting [47]. The neutrino's parent particle tends to be a pion at low energies and a kaon at high energies, as shown in Fig. 3.2. Note how the ν_μ peak around 0.6 GeV comes almost entirely from pions. This is due to the off-axis effect.

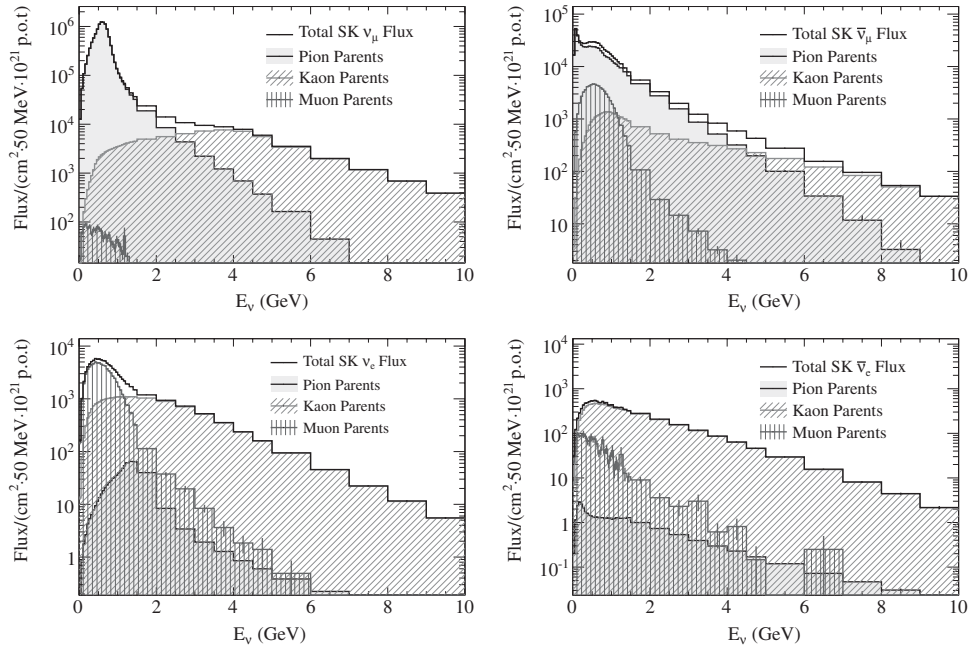


Figure 3.2: Flux predictions at SK broken down by neutrino flavor and parent particle [47].

3.2 Neutrino Interaction Modeling

If it were possible to exactly measure the kinematics of all neutrino interactions within a detector, then only the flux would be needed for a cross section measurement. Practically, this is impossible due to FSI and imperfect detector response. Accurate interaction models are therefore necessary to properly evaluate post-reconstruction selection efficiencies and purities. Correct efficiencies are crucial in predicting the full neutrino interaction spectra. Unfortunately, neutrino-nucleus cross section models are often poorly constrained due to limited data and nuclear model complexities. This section will describe how we attempt to model neutrino interactions with T2K MC.

3.2.1 Generators

Given the flux, event rates from neutrino interactions can be calculated from the cross section models discussed in Section 1.4. Several software packages for generating neutrino interactions exist including NEUT [60], GENIE [35], NuWRO [61], and GiBUU [62]. These are referred to as “generators”. T2K relies on NEUT as a primary generator (since it is the MC used by SK) and GENIE as a secondary generator for cross-checks. Differences between NEUT and GENIE in both the models used and their implementations allow for useful tests with one sample as fake-data and the other as MC, as is done for various tests used in this thesis.

Both NEUT and GENIE have evolved over the course of many years. They serve as independent frameworks that can be used, not just for T2K, but for general simulation of neutrino interactions. The generator takes the JNUBEAM flux and a detector geometry as input and simulates interactions within the specified detector. Additional FSIs are also handled by the generator. Finally, an output of the outgoing particles and kinematics is given to the GEANT-based detector simulation for passage through matter [63].

Official production MC for T2K Productions 5 and 6 using the GENIE [35] generator took place on the University of Colorado, Boulder HEP computing cluster. To give a sense of scale, the

Production 6 took approximately 51 CPU years of processing time and over 110 TB of disk space, of which 40 TB were transferred to long-term storage at TRIUMF.

Although NEUT is the primary generator, its FORTRAN implementation is less well documented than GENIE's. Furthermore, GENIE MC was the focus for production on the CU HEP computing cluster. Therefore, to give a sense of how a neutrino interaction generator works, an overview of the important components of the GENIE neutrino Monte Carlo is given in Section 3.2.2.

3.2.2 GENIE Neutrino Monte Carlo

The cross section calculations in GENIE are built from a set of splines containing the cross section for a particular combination of the interaction channel, neutrino flavor, and target atom. Free nucleon cross sections are calculated via numerical integration, as are some of the simpler nuclear cross sections (e.g. CCQE). More complex nuclear cross sections are approximated based on the free nucleon cross sections [63]. Since GENIE is under active development, it is difficult to describe all its current models in detail. Detailed descriptions can be found in [35] so only a brief mention is given here. Quasielastic interactions rely on the Llewellyn Smith model [64]. Elastic neutral current scattering is based off the Ahrens model [65]. Baryon resonance interactions rely on the model given by Rein and Sehgal [32] and proceeds in two separate steps: the first for resonance excitation, the second for its decay. Non-resonance inelastic scattering, including DIS, is calculated according to suggestions by Bodek and Yang [66]. Additional interaction channels in GENIE include coherent neutrino-nucleus scattering (where the neutrino scatters off the nucleus as a whole), charm production, inverse muon decay (IMD), and MEC (described in Section 1.4.3). In older implementations, nuclear models relied on the relativistic Fermi gas (RFG) model given by Smith and Moniz [34] but current revisions use a spectral function model that more accurately describes nucleon distributions within a nucleus. Brief phenomenological overviews of the standard neutrino cross section models are given in Section 1.4.

Beyond cross section models, generators need to take into account hadron production and transport. The hadronization/fragmentation model determines the final state particles and kinematics

from neutrino-nucleon interactions. This is important even for CCQE-like analyses since background contributions from resonance and inelastic channels are typically non-negligible. GENIE relies on the PYTHIA/JETSET model for hadronic showers at high invariant mass. For low invariant mass showers, a new hadronization model called the “AKGY” model is used [35]. The hadronization models are tuned to experimental data in the few-GeV energy range and exhibit good agreement for charged and neutral pions. Intranuclear transport of the produced hadrons are propagated via a cascade model within a GENIE subpackage called INTRANUKE. Since FSIs are practically irreducible sources of systematic uncertainty, a trustworthy MC is crucial for neutrino cross section measurements. For example, GENIE sees that approximately 20 % of pion production events are absorbed [35]. Exact quantum mechanical models for hadron transport would be ideal but are technically challenging. Therefore, semi-classical, intranuclear cascade (INC) models are used [35]. FSI models are a necessary part of neutrino cross section measurements, and typically are validated using hadron-beam data [35].

3.2.3 GENIE MC for T2K

As mentioned earlier, the GENIE ND280 MC was produced locally on the HEP computing cluster. The version of GENIE used for the production 6B ND280 MC is 2.8.0. The production chain is driven by sets of Python and shell scripts which act as wrappers around the heavy-duty, compiled Monte Carlo code. GENIE itself provides drivers for the JPARC neutrino flux used at SK, ND280 and INGRID. Here, we focus on the specifics of the GENIE neutrino vector generation at ND280.

JNUBEAM flux files provided by the beam group are used to first generate a set of precalculated interaction probabilities for neutrino vectors in flux. This procedure is done for both NEUT and GENIE. These probabilities, or “flux probs”, are regenerated for each combination of the flux and detector geometries, and rely on the cross section splines provided by GENIE (splines can also be regenerated on an as-needed basis). Neutrino vectors from the flux file are cycled through and GENIE decides whether to actually use the vector based on a random throw [63]. Flux probs

generation is important as it only needs to be done once for each flux/geometry and afterwards the probabilities can be used to generate a large number of independent fluxes.

Run	Geometry	Target MC POT
1	2010-02-water	a 9.0×10^{20}
2	2010-11-water	b 1.2×10^{21}
2	2010-11-air	b 9.2×10^{20}
3	2010-11-air	c 3.1×10^{21}
4	2010-11-water	c 3.5×10^{21}
4	2010-11-air	c 3.5×10^{21}

Table 3.1: Target MC POT for combinations of T2K geometries and beam specifications, along with the corresponding real data taking run period. The “water” and “air” tags in the geometry refers to the PØD water status, and the beam specifications correspond to intensities as measured in data.

Once flux probs are available the actual GENIE vector production begins in earnest. For each run, the target MC POT is typically an order of magnitude higher than the actual data POT. Due to the large amount of data involved, the final output is spread across thousands of files. Table 3.1 gives detailed numbers regarding the GENIE 6B MC generation. The differences between combinations of geometries and beam specifications are most obviously seen in Table 3.2, which shows how increasing the beam intensity also increases the POT per spill and hence the number of interactions per spill.

The default GENIE output is stored in a custom data table called GHEP [35]. For each event, the neutrino that caused the interaction, the generated particles, as well as additional information regarding specific details of the interaction are stored. The GHEP format differs from NEUT’s and a conversion module called `genieConvert` is run over it to standardize the generator outputs into a “rooTracker” format before passing them to the ND280 detector simulation. Detailed kinematics of each particle along the simulation chain is saved by GENIE, but only the final state particles are passed to GEANT.

3.3 ND280 Detector Simulation and Reconstruction

A final component to the T2K MC chain takes the generator outputs and processes them into a format suitable for comparison with data. Generally, this procedure involves running a detector

Geometry	Beam Spec	Evts/POT	Interactions/Spill
2010-02-water	a	1.00e-13	3.72
2010-11-water	b	1.06e-13	8.51
2010-11-air	b	1.06e-13	8.49
2010-11-air	c	1.06e-13	10.0
2010-11-water	c	1.06e-13	10.0

Table 3.2: GENIE interaction rates for combinations of ND280 geometries and beam specifications. In general, the interaction rate increases with the beam intensity.

simulation on the final state particles from generators and then reconstructing the (simulated) detector read-outs. It follows then that INGRID, ND280, and SK have separate detector simulations, and the focus here will be on the software used by ND280. The version of the ND280 software used for this analysis is `nd280 v11r31`, which corresponds to the production 6B MC sample.

In terms of a cross section measurement, detector MC provides a best estimate for the actual detector response. While data selections are based on reconstructed particle kinematics and topologies, and are typically impure and inefficient, the physics of interest should not be dependent on the detector used to measure it. Accurate detector simulation allows us to correct for this artifact of a real measurement device.

The ND280 software is a suite of packages that handles everything from detector calibration to electronics simulation to event reconstruction. Primarily written in C++, the code is dependent on several external packages, including ROOT, GEANT, CLHEP, and MYSQL. Package dependencies are handled by the Configuration Management Tool (CMT) and source code revision control is handled by the Concurrent Versioning System (CVS). In addition to simulation and reconstruction packages, analysis frameworks and computing tools exist to help process and organize the data. This section will focus on the detector simulation and reconstruction, while analysis specific tools will be discussed in Chapter 4.

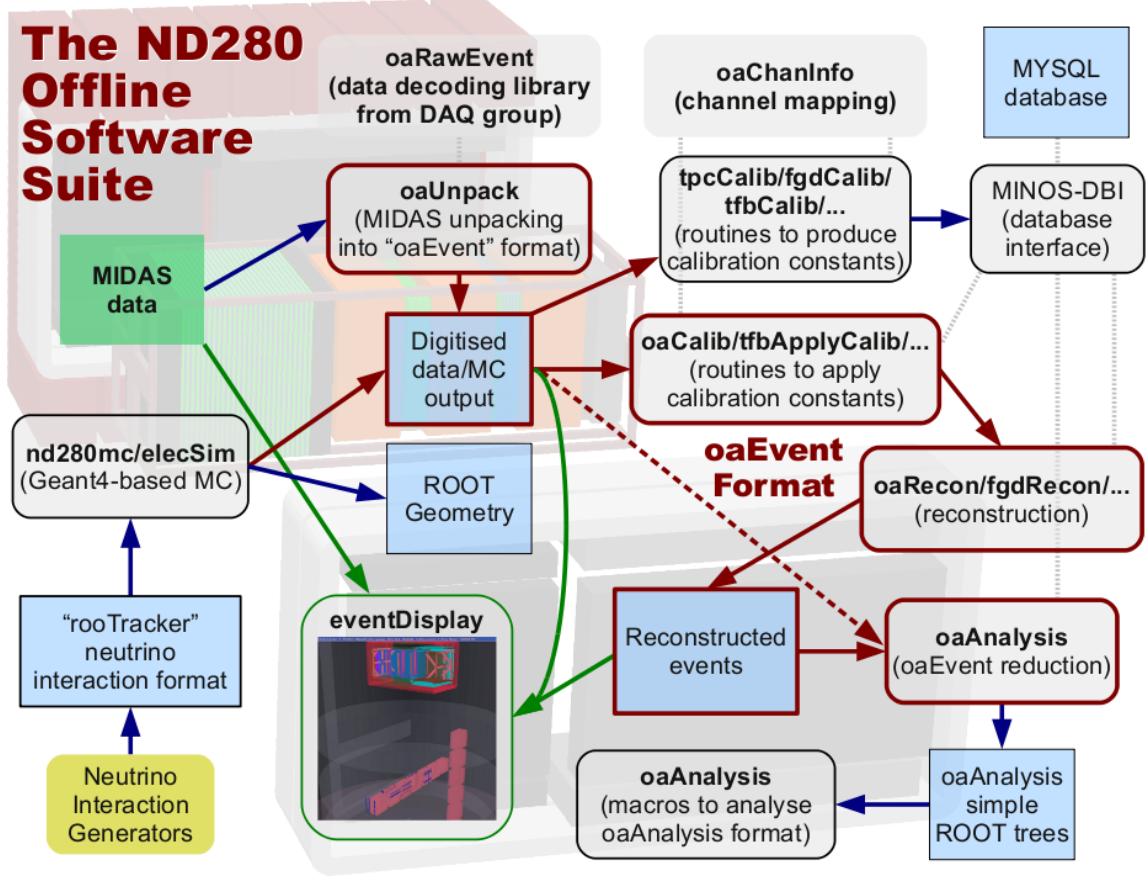


Figure 3.3: The ND280 Software Suite with process flow shown via arrows [67]. Not all package dependencies are shown.

3.3.1 Digitization

Once the rooTracker generator outputs are produced, they are passed to a GEANT4 based detector simulation of the ND280. GEANT4 is a CERN-managed package specifically for simulating the passage of particles through matter [68]. It allows for the construction of complex geometries using many types of materials. As discussed in Section 3.2.1, T2K relies on two different neutrino generators and thus nd280mc was developed to be generator agnostic. Final state particles passed on from the generators are propagated through the detector by the nd280mc. The spill structure discussed in Section 2.4.2 is also introduced at this stage. As the particle passes through detector active regions, information such as the true energy deposited, position, and time is stored [69]. This

presents a sequence of truth trajectories in the ND280 which are then passed through an electronics simulation package `elecSim` to convert the truth information into raw detector signals.

For a scintillator-based detector such as the PØD, this digitization takes into account the size of the scintillator bars, attenuation length, and chip type (which differs for the FGD and other scintillation detectors) [69]. The energy deposited according to the GEANT4 truth information is converted into a photon count as transmitted by the WLS fibers to the MPPCs.

All digitized data and MC up to the analysis reduction stage is stored in a complicated data format for event tracking called `oaEvent`. The `oaEvent` package inherits from ROOT’s `TObject` class, and all data is saved in “.root” files. An “event” has a rather broad definition within the package, and can range from individual tracks to entire spills [69]. The actual beam data is initially stored in the MIDAS format, which differs from the output of `elecSim`. An internal digitization package, `oaUnpack`, is used to convert the MIDAS data into “digits”, which serves as a store of all the MIDAS information, just in a different format [70]. At this point the digitized data output and the MC should be identical as shown in Fig. 3.3.

3.3.2 Data Calibration

As digits are simply representations of signals on the electronics channels, they must be calibrated into representations of the energy deposition, or “hits”. In the ND280 software, the `THit` class is used to store the information related to the hit, such as charge, position, and time. Conversion of digits into hits occurs in the `oaCalib` package, which performs calibration for the three types of ND280 subdetectors:

- TPC with MicroMEGAS readouts
- FGD with MPPC readouts based on the AFTER chip
- TFB with MPPC readouts based on the TRIPt chip

Each type of subdetector relies on somewhat different readout technology and therefore requires different sets of calibration constants. Since minimum ionizing particles (MIP) deposit a relatively

small and even amount of energy as they traverse the detector, cosmic ray muons are typically used for calibration [69]. The `oaCalib` package is a master package that contains subroutines such as `tfbApplyCalib` to perform the subdetector calibration.

3.3.3 Event reconstruction

Once the `THits` are created event reconstruction proceeds first via the subdetector reconstruction then via a global ND280 reconstruction. Reconstructed objects are stored in the `TReconObject` object class, which is a generic data structure for groups of `THits`. In a broad sense, the reconstruction takes a set of hits and produces a set of higher-level objects classified as follows, in increasing order of complexity [67, 69]:

Cluster A group of hits with position, energy, and time.

Shower A representation of energy deposition with additional direction and cone size information.

Track A curvilinear sequence of energy deposits with well defined start and end positions, as well as dE/dx information. Constructed from a list of Nodes.

Node A point along a Track with position, direction, and curvature.

PID Cluster, shower, or track with additional PID information

Vertex A list of reconstructed objects with common origin

A full description of all subdetector reconstruction algorithms is given in [71]. For the purposes of this analysis, an overview of the PØD and global reconstructions is given in this section.

Since the PØD is capable of acting as an individual detector, the PØD reconstruction is typically used as a standalone package. As mentioned in Section 2.4, the PØD was designed for the purpose of measuring π^0 s from CC and NC interactions. Keeping this in mind, the PØD reconstruction provides two separate routines. The first attempts to reconstruct MIP-like tracks using a Hough transform [72]. These are classified as “track-like”. Following this, hits associated with

these tracks are removed and a shower reconstruction routine proceeds via a clustering algorithm. These two routines are called the “track recon” and “shower recon”.

The inputs to the PØD reconstruction are all PØD hits in the spill. These hits are assumed to be independent across the 23 PØD cycles discussed in Section 2.4.2, of which eight (six for Run 1) align with the beam bunching described in Section 2.2.1. Thus, the first step involves separating hits into cycles. Both the track and shower recon are then run separately on hits in each cycle [71]. Separate from this, a muon-decay tagging routine runs over incoming hits across all cycles. The muon-decay search is performed on all final vertices and looks for clusters of hits more than 200 ns after the vertex [71].

Following cycle separation, noise hits are removed by imposing a minimum charge/neighboring distance cut. Only bunches with greater than five hits passing this criteria are kept for further reconstruction [71]. Following this, the track reconstruction proceeds first via a 2D-tracking algorithm, wherein the xz and yz views are processed separately with a Hough transform. The fully reconstructed 2D tracks are then processed with a 3D-tracking algorithm that attempts to match 2D track combinations from the two views. The 3D tracks are fitted by either a Kalman filter which starts at the downstream end of the track, or a Parametric fitter for high-angle or short tracks that fail the Kalman method [71]. Vertexing is performed on the full set of 2D and 3D tracks, and then a PID algorithm is run for the tracks. This analysis is only concerned with tracks that exit the PØD, all of which are classified by the PID as a `kLightTrack`, corresponding to either a muon or a charged pion [71].

The PØD shower reconstruction is not used by this analysis. In brief, the shower reconstruction takes the result of the track reconstruction with PIDs of `kEM` or `kOther`, and attempts to run a clustering algorithm on these PIDs. The shower vertex is determined via a grid search of hits that can be best clustered into angular bins [71].

Global reconstruction is overseen by the `oaRecon` package and performs track matching across subdetectors using RECPACK [73, 74]. RECPACK is a C++ toolkit that provides a set of common pattern recognition and fitting algorithms for general purpose usage. Designed to be modular,

RECPACK allows for the reconstruction of a dynamical system from experimental measurements, for example by fitting trajectories using a Kalman filter [73]. In the ND280 global reconstruction, subdetector inputs are provided as either track-like or shower-like objects. Showers are left unaltered while tracks are compared starting from the Tracker and iteratively merged pair-wise into a single object [71]. After each merger, RECPACK is used to refit the global trajectory with a Kalman filter and track kinematics are updated accordingly. Figure 3.4 shows an example of a global reconstructed track with the PØD track highlighted in green. The primary global refit is shown in yellow, although it seems to deviate slightly from the TPC3 and DSECAL subdetector reconstruction.

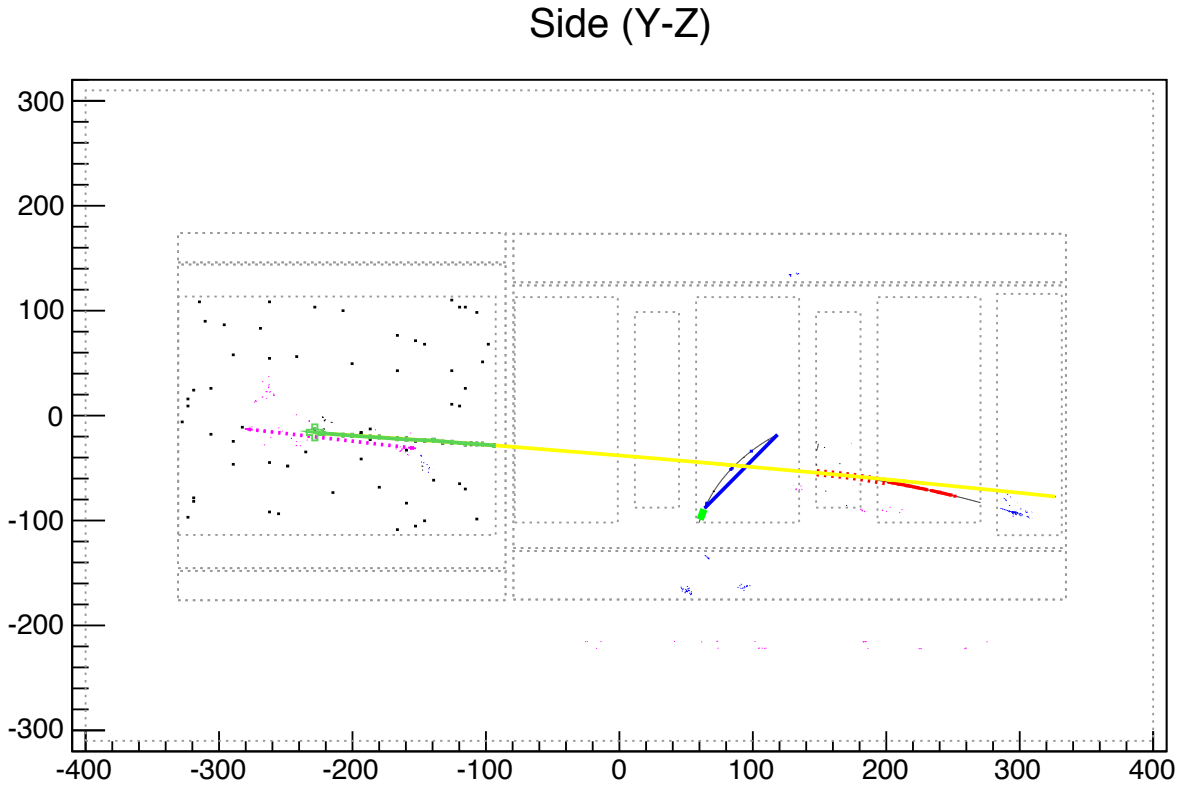


Figure 3.4: A side-view of reconstructed tracks from a neutrino interaction in the PØD with a single muon track originating in the PØD and entering the tracker. The beam direction is from left to right. The PØD reconstructed track is colored green and matched with the Tracker component to produce a global reconstructed track.

The measurement reported in this thesis is the first PØD analysis to rely on the ND280 global reconstruction as opposed to the PØD subdetector reconstruction. We chose to use the

global reconstruction because the measurement is based on a PØD+TPC selection. This involved implementing a new analysis package, including selections and detector systematics, within the global ND280 analysis framework discussed in Chapter 4.

Chapter 4

Analysis Frameworks

The ND280 Software Group provides several packages for simplifying and processing the data. Since the reconstruction output is stored in the `oaEvent` format, all the hit and low-level information is stored alongside reconstructed and truth objects. This makes the total dataset, especially the MC, which tends to be an order of magnitude larger than real data, cumbersome to analyze. Couple this to the complexities involved in processing `oaEvent` formatted data and the raw reconstructed output becomes impractical for users. This chapter details several tools to aid in analyzing ND280 data for cross section measurements. Broadly speaking there are three components. The first is the data reduction framework `oaAnalysis`. Selection and detector systematics are handled by the “High Level at the Near Detector” framework, or “highLAND”. Finally, a generic framework called `xsTool` is used for numerical error propagation and cross section extraction.

4.1 `oaAnalysis` Event Reduction

In order to summarize and reduce the data, the `oaAnalysis` package was developed to be part of the official computing chain. As shown in Fig. 3.3, it is the last component in the official ND280 Software Suite. Its outputs are referred to as `oaAnalysis` files, which are simplified ROOT files that contain most of the pertinent reconstructed, truth, and header data. Unlike the `oaEvent` files, `oaAnalysis` files are written in a “pure” ROOT format, which allows them to be more easily read and parsed. Typically `oaAnalysis` filesizes are reduced by a factor of four in comparison to `oaRecon`. Additionally, there is no need to load the `oaEvent` library for data processing. Thus, the

`oaAnalysis` package reduces the data into a flexible, user-friendly format at the expense of some loss of low-level information.

The `oaAnalysis` file contains several directories of which the most important are,

HeaderDir Contains beam summary and data quality information per spill as well as other metadata for the file.

TruthDir Contains truth trajectories and vertices as well as `rooTracker` information copied directly from the generators. Only applicable to MC.

ReconDir Contains results from global and subdetector reconstruction.

Each directory consists of several ROOT TTrees, which are basically optimized structures for efficient data storage and fast retrieval, similar to a large, complicated spreadsheet. The rows of a TTree are typically independent, individual events. These in turn are often built from classes designed to encapsulate the concept of trajectories or vertices.

As an example, the `ReconDir`'s `GlobalRecon` tree contains events stored in a `TGlobalPID` class. This class describes a global reconstructed PID object, and contains a set of attributes such as the charge and momentum of the track. To facilitate with writing analysis macros for `oaAnalysis`, a `nd280AnalysisTools` package is provided, which contains all the header files corresponding to each class in `oaAnalysis`. These TTrees can be parsed interactively on the ROOT CINT (ROOT's C/C++ interpreter) command line, and its events can fairly easily be drawn into distributions.

4.2 High Level Analysis at the Near Detector (`highLAND`)

Analyzers often find that selecting events from `oaAnalysis` files involves several generic steps. Selections typically involve a sequence of cuts, which in itself is simple enough, but can quickly become complicated by flux corrections, data quality cuts, beam bunch timing, track sorting per bunch, POT counting, etc. Many of these tasks require information external to `oaAnalysis`, and thus can be prone to error if used improperly. Add in a dose of systematic propagation due to

uncertainties in detector components and a once simple selection can quickly turn into a complex package with plenty of room for bugs. As many of the bookkeeping tasks and systematics are quite common between analyses, a generic framework called the “High Level at the Near Detector” framework, or “highLAND”, was developed to drive both oscillation and cross section analyses for several detectors at ND280. Now in its second iteration, the original framework has been separated into the `highland2` and the `psyche` (Propagation of Systematics and Characterization of Events) packages as shown in Fig. 4.1. `psyche` is a lightweight package used by oscillation analyses.

Analysis flow

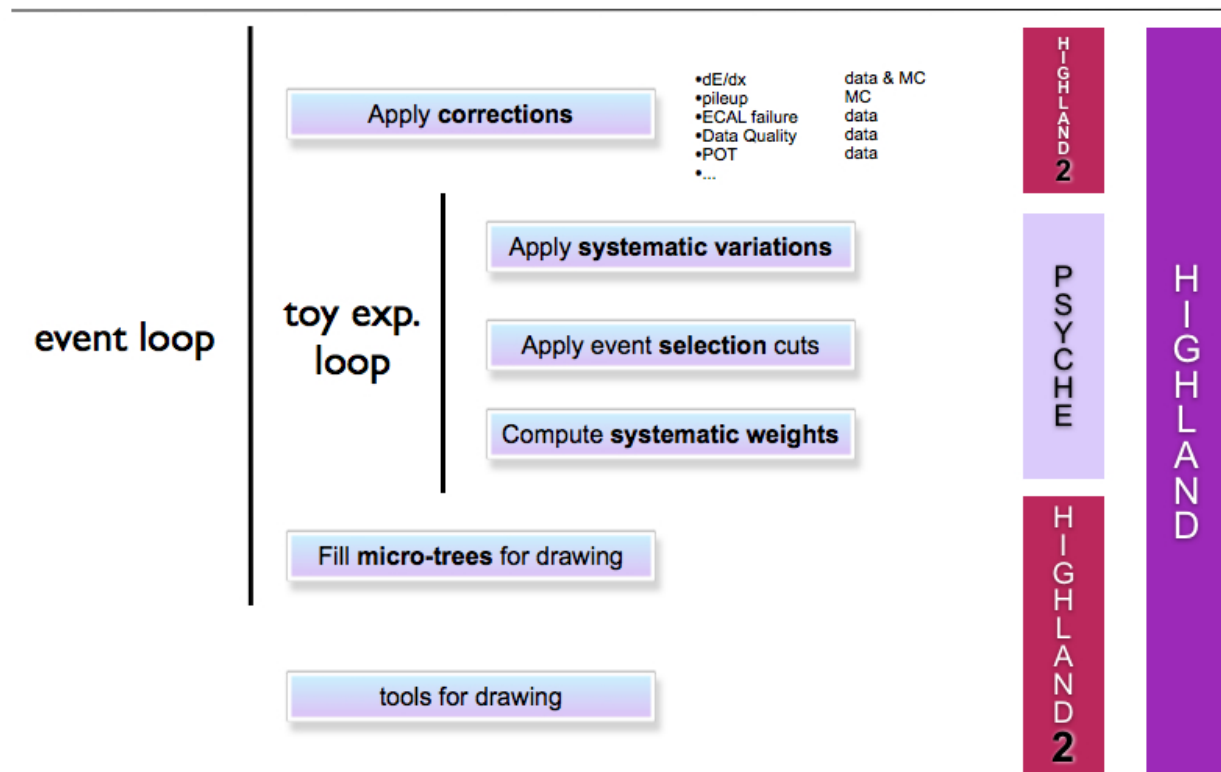


Figure 4.1: Analysis flow of the High Level Analysis and Development package [75]. A standard analysis proceeds from top to bottom, with the loops over events and toy-experiments as shown. `psyche` acts as a fast and more stable interface for the fitter used by oscillation analyses. This thesis relied primarily on `highland2` independently.

4.2.1 Data Structures in highLAND2

To facilitate the transition from `oaAnalysis` to `highland2`, a set of basic data classes are provided to organize the `oaAnalysis` events into spills, bunches, tracks, and particles. A converter library, `oaAnalysisTreeConverter`, takes info from the global reconstruction tree in the `oaAnalysis` file and saves it into the `highland2` data classes. More recently, this converter has been extended to work with subdetector reconstructions such as the PØD reconstruction.

The basic data classes in `highland2` contain information such as the reconstructed momentum, charge, and position, and in the case of MC the true reaction code, neutrino energy, and more. In `oaAnalysis`, the links between various reconstructed objects can be a bit unclear. For instance, the truth track associated with the reconstructed track is contained in a separate TTree and truth matching is left to the analyzer. Similar issues exist for subdetector reconstruction matching as well. The data structures in `highland2` attempt to restore the links between various `oaAnalysis` objects such that the data class for a reconstructed track contains a pointer to its true track.

In addition, `highland2` accounts for the beam bunching structure. Since events in `oaAnalysis` TTrees correspond to an entire spill, reconstructed objects are usually not bunch-separated in `oaAnalysis`. Instead the beam bunch structure handled within `highland2` and the tree converter based on timing information. Thus, the `highland2` data classes provide an extra layer of data reduction and organization on top of the `oaAnalysis` dataset.

In addition, `highland2` also provides executables designed to skim `oaAnalysis` files into “flattree” or “minitree” formats. These formats are even simpler than the `oaAnalysis` format as they essentially contain primitive data (such as floats, doubles, and integers). Not all `oaAnalysis` information is flattened and stored, so these formats save a lot of space and can be processed much more quickly. For the purposes of our measurement, however, these formats were not used primarily because they were designed for events reconstructed in the Tracker. Processing `oaAnalysis` directly proved to give more reliable results for the PØD selection.

4.2.2 Analysis in highLAND2

The heart of `highland2` resides in an event loop as shown in Fig. 4.1. Functions used by the loop are contained in the `AnalysisLoop` class, which processes all events from the input (`oaAnalysis` or flattened) files and performs a specific selection defined by users for their analysis. The loop proceeds as follows.

```

for all spills in input do
    {Each spill corresponds to an entry in the input file}
    for all bunches per spill do
        {Bunches are filled and organized by the highLAND2 converter}
        for all enabled configurations do
            {Configurations can correspond to the nominal reconstruction or an altered version of the
            reconstruction due to detector systematic variations}
            for all throws per systematic do
                {Throws correspond to toy experiments for systematic configurations}
                process user selection
            end for
        end for
    end for
end for

```

All user-defined analyses rely on the `AnalysisLoop` through a set of rigid interfaces, or methods, that define, amongst other things, the selection, systematics, and filling of the output “microtree”. These interfaces are encapsulated in a class distinct from `AnalysisLoop` called `baseAnalysis`, which serves as a base class for all physics analyses in `highland2`. The base class tries to handle the bookkeeping tasks common to all analyses. It includes corrections from the latest flux tuning, as well as some common variables such as the run, subrun, event, and bunch numbers. Additionally, truth information involving the neutrino parent is stored as well for MC. Users can then derive

their own analyses from the base class, adding selections and variables as desired, provided that the necessary interface methods are properly overwritten.

Corrections and detector systematics are handled in `highland2` as well. Several generic corrections in `baseAnalysis` affect the final selection and true MC distributions. Typically, these are stored in the microtrees as event-by-event weights that depend on the specific interaction topology and kinematics. For example, flux correction weights, as discussed in Section 6.1.1, depend on the neutrino energy and flavor from MC truth.

Uncertainties due to detector reconstruction and simulation are handled by a configuration manager, which saves enabled systematics as toy-experiments, or throws, in separate TTrees, or configurations, from the nominal configuration. Unlike corrections which can apply to both the truth and selected events, detector systematics ideally should only affect the selection. Thus, only the selection is thrown for enabled systematics. In `highland2` there are two types of detector systematics, variation systematics which alter a continuous variable such as the momentum, and weight systematics which can change the overall normalization and reconstruction efficiencies. The user can preset the number of toy-experiments to generate for each systematic configuration. The nominal configuration consists of a single throw. Then, for each throw, an underlying parameter that either alters reconstructed variables or reweights the event is set, and the selection is redone as shown in the event loop algorithm.¹ The different throws can then be used to calculate a bin-by-bin uncertainty.

The `highland2` outputs are saved in a simple, standardized ROOT format and consists of TTrees corresponding to each enabled systematic configuration in addition to the default selection and truth trees. These are simple enough to be parsed directly, but to aid plotting comparisons across different datasets and calculating uncertainties, a package called `DrawingTools` is provided as part of `highland2`. This aids in creating nicely formatted plots of selections and systematics and can easily be expanded into complex analyses.

¹ Technically, redoing the selection is only performed for variation systematics that alter some reconstructed variable, as then the selected events may also change. Weight systematics do not alter the actual event selection.

4.2.3 Validation of highLAND2

The original idea of highLAND began around 2013 and the concept of highLAND2 originated in 2014. Unlike the other upstream components of the ND280 Software Suite, which have undergone many more revisions and are relatively stable, the `highland2` package remains more of a work-in-progress. Suffice it to say that not only do the user analysis packages undergo rapid development, but the core `highland2` packages often have major refactorings and updates from many contributors. This type of development requires some sort of checks to ensure consistency as code is updated, and an automated validation tool was developed for that purpose.

The idea of the validation package originated out of a desire to compare outputs from different revisions of `highland2`. Processing a constant set of input files should always produce the same outputs. Consistency across different versions of the software is of key concern to analyzers and `highland2` developers. Any differences should be understood and corrected as necessary.

To ensure consistency, a series of tests are run using a Python testing framework. Processing is first performed across both beam data and MC input files for several `highland2` analyses, many with specific detector systematics enabled. Then the `highland2` output is put through a series of tests. These perform simple comparison checks of the number of entries, the POT count, the number of good spills, as well as distributions of a variety of variables. If there is any discrepancy, the test is logged and marked as failed. Tests of distributions across kinematic variables are run by histogramming and then performing a bin-by-bin comparison of the two raw distributions. Figure 4.2 shows two test failures as found by the validation tool. The discrepancy is slight, only a couple events have migrated, but the cause was found to be rather serious, having to do with how tracks in a bunch used to be sorted by momentum (correct for selection) but no longer were in the updated version (incorrect for selection).

Over time, the validation became tedious to manually perform for every code update. Luckily there exist tools to aid with this type of automation, and the Jenkins Continuous Integration client² was used to automate the building, processing, and validating process. Jenkins works by polling the

² The Jenkins homepage is <https://jenkins-ci.org>




Figure 4.2: Test failures reported by the `highland2` validation tool. This shows a comparison of reconstructed muon momentum (`selmu_mom`) and polar angle (`selmu_theta`) distributions for two versions of `highland2`, selected using the PØD analysis package. The selection from a previously stable freeze is shown in black, and the HEAD is in red.

`highland2` repository for updates, and if there is an update it triggers a new build based on the latest version of the code, or what is referred to as the HEAD. Following a successful build, the newly compiled programs are run over a predefined set of input files, and then the a set of validation checks is performed on the output. The client can be deployed online and accessed publicly through any web browser.³

In addition to cross checks between different revisions of `highland2`, validation is also performed on the flattree and minitree outputs described in Section 4.2.1. Since these are skimmed from `oaAnalysis` files and then used as input for physics analyses, the eventual result must match the result processed using `oaAnalysis` directly. Thus, a similar suite of tests is performed to validate the flattree and minitree outputs.

³ Results of the `highland2` build and validation are live and hosted at <https://advance.colorado.edu/jenkins>




Figure 4.3: Summary of test results for the previous 120 builds as displayed on the Jenkins server. This allows for a continuous monitoring of the software. Blue indicates the number of passing tests, red the number of failing tests. The gaps are periods where tests failed to run due to processing issues. The time period shown here spanned a major reorganization of `highland2`, which explains the gaps and the test failures.

4.3 The Cross Section Tool Package (`xsTool`)

While `highland2` provides a mechanism for selecting events and throwing toy experiments from detector systematics, it alone does not provide a means for proper cross section measurements. The final piece is filled by the `xsTool` package, which was designed to be a general software framework for making cross section measurements. The `xsTool` has been implemented with `highland2` outputs in mind, but is capable of being extended to handle other types of inputs as well. Essentially, `xsTool` handles both the numerical propagation of systematic and statistical uncertainties, and the cross section extraction itself, which can be performed using a variety of “engines”. So, in summary, the conversion of `oaAnalysis` files into a cross section measurement proceeds roughly as follows.

- 1: `oaAnalysis` files are given as inputs to `highland2`
- 2: `highland2` filters events and selects those of interest to save in microtrees
- 3: `highland2` microtrees are given as inputs to `xsTool`
- 4: `xsTool` performs numerical error propagation and cross section extraction

This section attempts to give a general overview of the capabilities of **xsTool** with a software-oriented perspective. Details specific to the analysis presented in this thesis are deferred to a later chapter dedicated specifically to this cross section measurement.

4.3.1 How the Cross Section Tool Works

The **xsTool** package provides a set of modular and extensible classes, allowing users to define their own input reading class, additional types of engines for unfolding or fitting, and additional systematic uncertainties. Inputs are handled by **xsInputdata** derived classes and cross section extraction is handled by **xsEngine** derived classes which use those inputs. These are the core classes of the package. Users must first define a **xsInputdata** object with the binning, selection, and truth definitions included as part of the setup, and then pass that object to the engine of their choice. Again, a set of common interfaces are required for any **xsInputdata** object to interact with the engines. These are typically processed versions of the actual events, summarized as histograms, which are then used for the cross section extraction.

The core cross section framework revolves around a couple of additional packages. **RooUnfold** is an external package with several unfolding functionalities [76]. Now, unfolding and specifically Bayesian unfolding is a concept that will be discussed at length later. For now, the goal of an unfolding procedure can be thought of as correcting misreconstructions in a probabilistic manner using the MC truth information. The **RooUnfold** package provides several unfolding methods which can then be wrapped by **xsTool**. The **RooUnfold** unfolding methods currently available in **xsTool** are,

- Iterative Bayesian unfolding
- Singular value decomposition (SVD)
- Matrix inversion
- Bin-by-bin corrections

- ROOT's TUnfold interface

Each of these are discussed briefly in [76]. The measurement reported in this thesis uses the iterative Bayesian unfolding algorithm described in [77, 78].

Another dependency is a reweighting package developed internally by T2K collaborators called `T2KReWeight` [79]. Since the `xsTool` needs to be able to calculate systematic uncertainties from other sources in addition to the detector systematics provided by `highland2`, the `T2KReWeight` framework provides a means for propagating uncertainties on the base flux and interaction parameters onto the resulting distribution via event-by-event reweighting. As discussed in Chapter 3, this is computationally much faster than rerunning the full simulation chain for each perturbed value of the parameters.

Perhaps the most useful aspect of `xsTool` is its error propagation functionality. In addition to reading in detector systematic throws from `highland2`, `xsTool` handles reweighting and propagating uncertainties due to the flux, cross section model, and FSIs as discussed in Chapter 3. These are handled in a similar manner as the `highland2` detector systematics, using a set number of throws to emulate toy-experiments with tweaked parameter settings. Each throw is performed with `T2KReWeight` and applied to the MC. Unlike the detector systematics, flux and interaction uncertainties affect the truth distribution as well as the reconstructed distribution. Weights are generated and stored for each throw for each event using several useful utility programs in `xsTool`. This is necessary as `T2KReWeight` was designed to provide very flexible interfaces where each parameter can be set independently. In reality though, there are correlations within the error groups (for instance correlations between flux bins as shown in Fig. 7.12). The `xsTool` package handles all these correlations such that all parameters are varied with those constraints in place before being passed on to `T2KReWeight` for the actual event reweighting. A full list of these parameters is given in Appendix C.

Once event weights have been generated and stored, they should have a one-to-one correspondence with selected and truth events from `highland2`. The weights are then applied in the

numerical error propagation routine for each throw. This means slightly altered MC histograms are built in `xsInputdata`, and are then passed to the cross section extraction engine. Since the final cross section measurement depends on these constructed MC histograms for unfolding and efficiency corrections, the final results will also differ across throws. Thus, an error envelope can be built using the distribution of results, and a sample covariance matrix given by

$$\text{cov}(X_i, X_j) = \frac{1}{T-1} \sum_{t=1}^T (x_{ti} - \bar{x}_i)(x_{tj} - \bar{x}_j), \quad (4.1)$$

where X_i and X_j symbolize random variables corresponding to the result in bins i and j , T is the total number of throws, x_{ti} and x_{tj} are the measured values in those bins for throw t , and the overline denotes the sample mean across all throws, can be calculated for each error source. The final covariance matrices from the sources are summed to obtain the total uncertainty. The bin-by-bin uncertainty then is taken to be the square root of the diagonal of this covariance matrix. The interpretation that these correspond to 1σ confidence intervals is correct only under the Gaussian approximation. Nevertheless, the covariances between bins are important and only really useful under the multivariate Gaussian assumption, so the intervals are reported under this assumption as well.

Several methods for the final cross section measurement are implemented in `xsTool`. Unfolding engines mentioned earlier in this section provide a direct measurement of the cross section. In addition, an in-house fit engine exists that uses MINUIT to fit to the cross section model parameters directly. As this is an extremely model dependent method for extracting the cross section, it is not used for this analysis. For the direct cross section measurement, efficiency correction is handled by `RooUnfold` at the same time as the unfolding, while target and flux normalization are all handled by `xsTool`. These normalization factors are also varied by the mass, and flux systematic throws that affect them, respectively.




Figure 4.4: Comparison of reconstructed, data, sideband, and signal histograms using `xsTool` (lines) vs using `highland2` directly (dots) [80]. The sample used is a test sample from the ν_μ CCInclusive selection, and the same selection is used for all the different categories. The `xsTool` distributions match those from `highland2` exactly.

4.3.2 Validation of `xsTool`

Like `highland2`, the `xsTool` package also undergoes rapid development. Being a complicated tool, unforeseen issues can arise. Often, the attempt to fix an initial problem will lead to other problems. Because of this, a suite of tests have been written to validate the framework. The validation has also been automated using the Jenkins CI client.⁴ The whole suite of tests is run whenever updates to the package are committed to the repository.

Several tests are performed to ensure consistency with `highland2` and ensure that the inputs are read in properly [80]. Figure 4.4 shows that the histograms generated using `xsTool` matches those from `highland2` exactly. Furthermore, as shown in Fig. 4.5, the relative detector error for a test sample as calculated using `xsTool` matches that from the `DrawingTools` of `highland2`.

Internal consistency tests are also performed against reference distributions from previously

⁴ Results of the `xsTool` validation are live and hosted at <https://lheppc88.unibe.ch:8080>

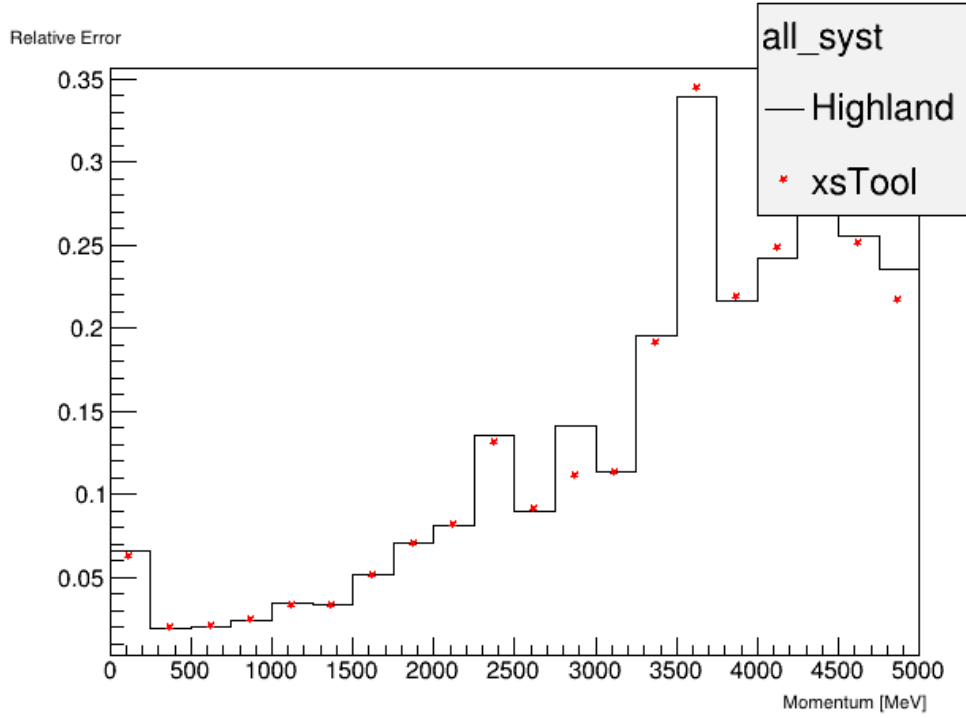


Figure 4.5: Comparison of the relative detector error for a test sample using `xsTool` vs using `highland2` directly. The two methods agree nicely.

frozen versions of the cross section tool. These, like the validation tests from `highland2`, ensure that any updates that result in changes to the resulting cross section measurement are caught and reported. Figure 4.6 shows the extracted cross section measurement from the previously released CCInclusive measurement on carbon [81]. The `xsTool` results can be compared to the official T2K results, and an internal consistency check is performed against a previously generated reference histogram to ensure that any updates to `xsTool` do not change the cross section measurement.

In addition to the consistency checks mentioned in this section, several low-level sanity checks are performed. These test specific but important features of the `xsTool` such as throw generation, handling of multiple runs, and covariance matrix generation. Many of these rely on using contrived inputs to test for an expected output. All tests can be found on the continuous integration website or the `xsTool` documentation page [80].

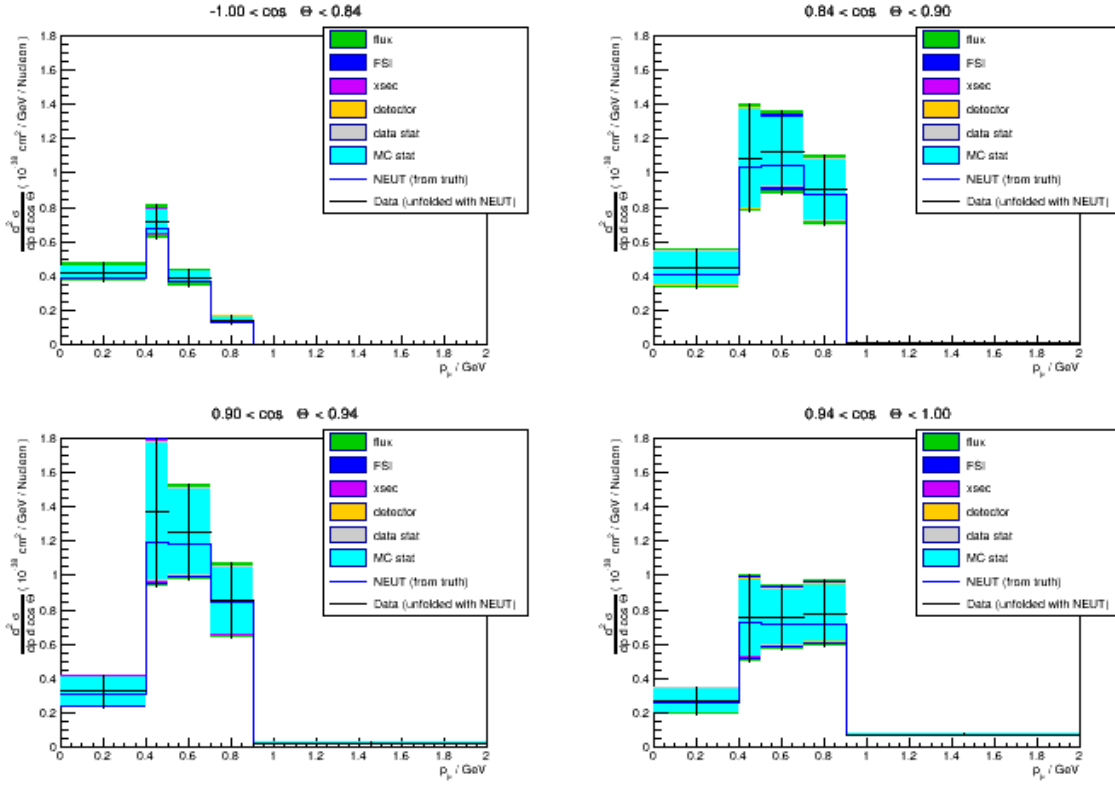


Figure 4.6: The full **xsTool** cross section extraction from the ν_μ CCInclusive analysis [81] using a subset of the total statistics. This is compared to a previously calculated reference result to ensure internal consistency as the **xsTool** code is updated.

Chapter 5

Data Selection

Figure 2.3 shows the total POT collected by T2K as of April 3, 2016. The data sample used for this analysis includes most of T2K Runs 2–4. The POT for these runs is listed in Table 5.1. These are the primary runs during which data was collected in the neutrino (forward horn current) mode and thus corresponds to neutrino data as opposed to antineutrino data. Run 1 is a low statistics (1.6494×10^{19} POT) run that is not used by this analysis or the tracker analyses. It contains global subdetector reconstruction issues and less accurate calibration constants. For this analysis specifically, the Run 1 global reconstruction does not store the momentum at the entrance position of a subdetector. This makes it impossible to evaluate certain PØD systematics dependent on the particle’s Tracker entrance momentum. Additionally, only a subset of Run 3, Run 3c, is used. The start of Run 3 corresponds to a test period during the resumption of data-taking after the great East Japan earthquake of 2011. For a flux-averaged measurement, it is better that the flux shape remains constant across the data sample periods.

The measurement is performed via a subtraction of the PØD water-in and water-out periods. In neutrino mode, the PØD contained water during partial periods of Runs 2 and 4.

5.1 Software Versions

The software chain has already been discussed in detail in Chapters 3 and 4. For completeness, a list of the specific versions used by this analysis is given in Table 5.2. With the exception of external neutrino generators NEUT and GENIE, the packages are given in terms of T2K internal

T2K Run	Data POT	MC POT
2 water	4.2896×10^{19}	1.2034×10^{21}
2 air	3.5509×10^{19}	9.2394×10^{20}
3c air	1.3482×10^{20}	3.0806×10^{21}
4 water	1.6282×10^{20}	3.4961×10^{21}
4 air	1.7625×10^{20}	3.4992×10^{21}
Water	2.0572×10^{20}	4.6995×10^{21}
Air	3.4658×10^{20}	7.5035×10^{21}
Total	5.5230×10^{20}	1.2203×10^{22}

Table 5.1: Detailed POT rates for the data taking periods used by this analysis.

versioning system.

5.2 Event Selection

The overall data sample contains many events not caused by $CC0\pi$ interactions in the PØD. In order to produce a measurement of the muon-neutrino cross section in this specific channel, it is necessary to first prune the sample of all reconstructed events into the ones relevant for this analysis. This procedure is simply referred to as the “event selection” and its purpose is to provide a filtered sample of events mostly occurring in the PØD fiducial volume and in the $CC0\pi$ interaction channel. The selection is based entirely on reconstructed information and therefore can be applied to both data and MC. As we performed a blind analysis, it was tested first using only MC samples. In addition to the $CC0\pi$ signal selection, two control samples, or “sidebands”, were used for constraining the two largest sources of background in the signal selection. This section describes both the signal $CC0\pi$ selection as well as the two sideband selections.

This analysis attempts to combine data from the PØD and Tracker into a coherent framework. As such, our event selection relies on previous work done by the PØD and Tracker groups. An inclusive, PØD-based charged-current selection was performed in [82] and serves as a basis for the initial inclusive selection. However, as mentioned in Section 3.3.3, the selection for this analysis relies on the global reconstruction instead of piecing together the individual subdetector reconstructions. Indeed, this is the first PØD-based analysis to use the ND280 global reconstruction and also the first

Software Package	Version	Purpose
NEUT	5.3.2	Generator
GENIE	2.8.0	Generator
nd280	v11r31	Simulation and reconstruction
highland2	v1r15	High level analysis framework
p0dNumuCCQEAnalysis	v1r13p1	Selection and detector systematics
T2KReWeight	v2015r1p5	Flux and theory systematics
RooUnfold	v1r7	Unfolding
xsTool	v2r3	Cross section extraction and error propagation

Table 5.2: Overview of software packages, versions, and a brief description of their purpose as used by this analysis. The interdependencies of these packages are rather complex, for example, `xsTool` relies on `RooUnfold`, `T2KReWeight`, and `highland2`. The `nd280` software is an overarching master package with many subpackages and dependencies used for production MC.

PØD+Tracker analysis. In addition, the tracker ν_μ charged-current inclusive measurement provided additional requirements on the tracker components of our sample [81].

5.2.1 Signal Selection




Figure 5.1: Schematic of an interaction occurring in the PØD and entering the Tracker to illustrate the signal selection. Due to the Tracker-going requirement, high-angle or low-momentum tracks do not make it into the selection.

The μ^- produced via CC interactions is a minimum ionizing particle. It is the final state particle most likely to travel through the PØD and into the tracker (Fig. 5.1), and we associate it with the highest momentum negative track in the bunch. As described earlier, the true signal definition is a CC0 π interaction consisting of a single muon, zero mesons, and any number of nucleons.

The $CC0\pi$ event selection is performed on a bunch-by-bunch basis as follows.

- (1) Beam and ND280 detector status both pass data quality requirements for real data
- (2) At least one reconstructed track in the TPC with interaction vertex in the PØD fiducial volume
- (3) Select the highest momentum negative (HMN) track in the bunch as the μ^- candidate
- (4) TPC track is composed of more than 18 reconstructed nodes (Section 3.3.3)
- (5) Only a single track reconstructed using the PØD per bunch

The PØD fiducial volume (FV), as shown in Fig. 2.11, is taken from [83] and defined to be entirely within the water target. The z-range is defined from the middle of the first water PØDule to the middle of the last water PØDule. The xy-range is approximately 25 cm from the edges of the water target (recall that the PØD is approximately a cube 2 m on each side). Although we select interactions occurring within the PØD, the requirement that the track enters the tracker and has at least 18 reconstructed TPC nodes allows for accurate momentum reconstruction, crucial for a double differential measurement. The last cut provides the $CC0\pi$ enhancement from all charged-current events.

The distribution of this selection is shown in Fig. 5.2. The major background comes from $CC1\pi^+$ as seen in Fig. 5.3. The selection **purity**, or the ratio of the number of selected events in a true category to the total number of selected events, and **efficiency**, or the ratio of the number of selected events in a true category to the total number of events in that category, can be calculated from MC. In the case of the PØD $CC0\pi$ selection, the total number of true events encompasses **all** true $CC0\pi$ interactions in the PØD FV (even those that are fully contained in the PØD). For water-in the overall $CC0\pi$ purity is 78 % and efficiency is 18 %. For water-out the purity is 82 % and efficiency is 18 %. The slight difference in the purity is due to the fact that the PØD water layers acts as a passive region that absorbs some tracks which would otherwise have been reconstructed in the active regions of the PØD.

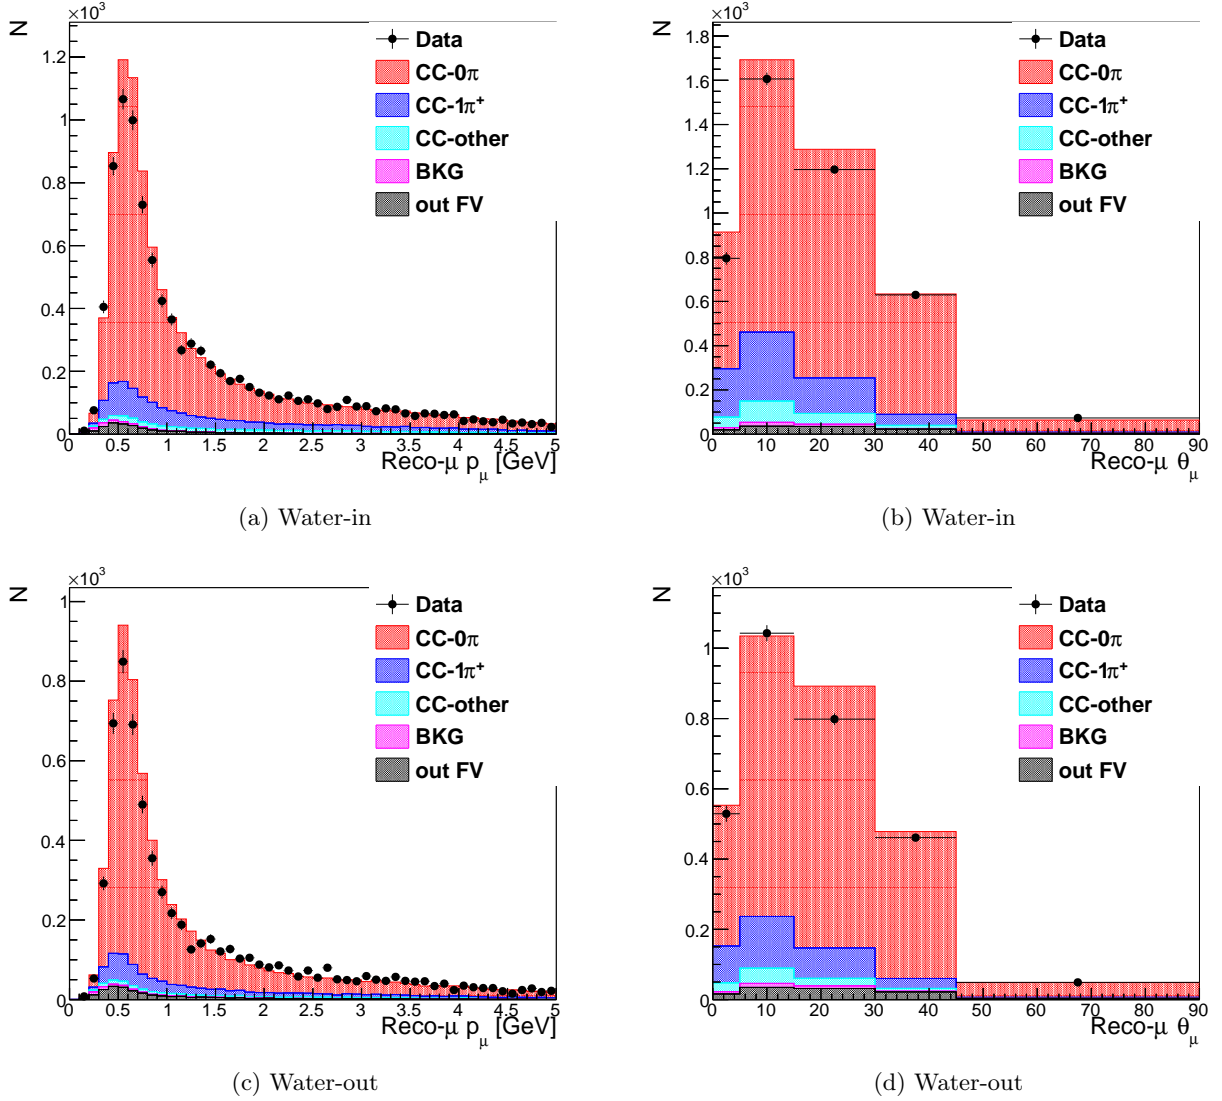


Figure 5.2: Data and MC distributions of the CC0 π selection for most of the water-in and water-out periods. The distributions are binned across the reconstructed muon momentum (p_μ) and angle (θ_μ) relative to the z-axis (Fig. 5.1). The NEUT MC has been normalized to data POT and is sorted into true observable topologies. The largest source of background comes from CC1 π^+ , followed by CCOther (mostly DIS interactions). Interactions without a muon in the final state are classified as “BKG” and interactions occurring outside the PØD fiducial volume are classified as “out FV” or “OOFV”

According to the MC, approximately 98 % of the tracks passing the selection criteria are truly due to μ^- s. A full break-down by particle, reaction, and topology is given in Tables 5.3 to 5.5. The phase space coverage of true CC0 π interactions with the given selection is shown in Fig. 5.4. The

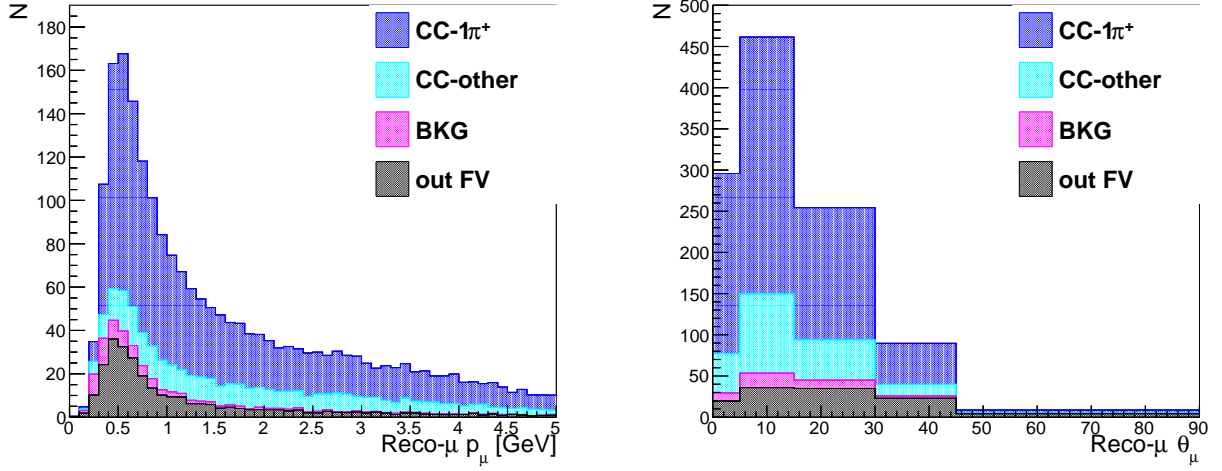


Figure 5.3: The total background distributions in p_μ and θ_μ with the CC0 π selection using water-in NEUT MC.

PØD to Tracker requirement means that we have minimal coverage in the backward-going bins. For comparison, the phase space of **all** true CC0 π interactions is shown in Fig. 5.5 and extends to the backward-going region. The unfolding method we use requires the inclusion of all true events, including events that fall outside the binned region (out-of-range bin). Thus, the full phase-space is used for the cross section extraction. However, the reported results will be presented only in the forward-going region using the binning described in Section 7.1. A few bins with zero efficiency are ignored during the cross section extraction.

In order to further reduce backgrounds, there is the possibility of rejecting spills containing Michel electrons in the PØD which occur after the primary interaction. Michel electrons are decay electrons from muons that come to rest in PØD. These lower energy muons are typically products of pion decay, so Michel electrons can indicate a non-CC0 π event and conversely, no Michel electrons can indicate a CC0 π event. The PØD subdetector reconstruction contains Michel tagging algorithms which allow for an addition 0-Michel cut after the single-track selection. A combined usage of global and local reconstruction makes this possible, and its effect on the CC0 π efficiency and purity is shown in Fig. 5.6. As expected, the 0-Michel cut improves purity at the expense of a slightly lower efficiency. One of the largest sources of uncertainty comes from data statistics, which gives some

Particle	Water-in [%]	Water-out [%]
μ^-	98.4	98.2
e^-	0.70	0.74
π^-	0.56	0.65
μ^+	0.08	0.11
e^+	0.02	0.03
π^+	0.11	0.15
p	0.09	0.10
other	0.00	0.01

Table 5.3: Composition of the CC0 π selection broken down by the true particle.




Figure 5.4: Two dimensional distribution in true p_μ and $\cos\theta_\mu$ space of CC0 π events in the selection. The inset shows the 2D distribution projected onto bins in a single dimension, starting from the lower-leftmost bin and moving right and up towards the upper-rightmost bin. Sample taken from NEUT water-in.

preference to a higher efficiency selection. Further, an additional 0-Michel cut introduces additional systematics. Since the original selection already is relatively high purity, we decided against imposing the 0-Michel cut in the signal selection, but it will be used as part of the sideband selection discussed in Section 5.2.2.

With regards to events with multiple reconstructed tracks in the PØD, the PØD reconstruction

Reaction	Water-in [%]	Water-out [%]
CCQE	62.3	67.3
2p2h	10.0	9.72
CCRES	18.9	14.4
CCDIS	2.87	1.90
CCCOH	2.29	1.89
OOFV	2.84	3.89

Table 5.4: Composition of the CC0 π selection broken down by the true reaction channel




Figure 5.5: Two dimensional distribution in true p_μ and $\cos\theta_\mu$ space of **all** CC0 π events. The inset shows the distribution projected onto bins in a single dimension, starting from the lower-leftmost bin and moving right and up towards the upper-rightmost bin. Sample taken from NEUT water-in.

Topology	Water-in [%]	Water-out [%]
CC0 π	78.1	81.8
CC1 π^+	14.1	10.6
CCOther	4.12	2.85
BKG	0.81	0.83
OOFV	2.84	3.89

Table 5.5: Composition of the CC0 π selection broken down by the true topology.

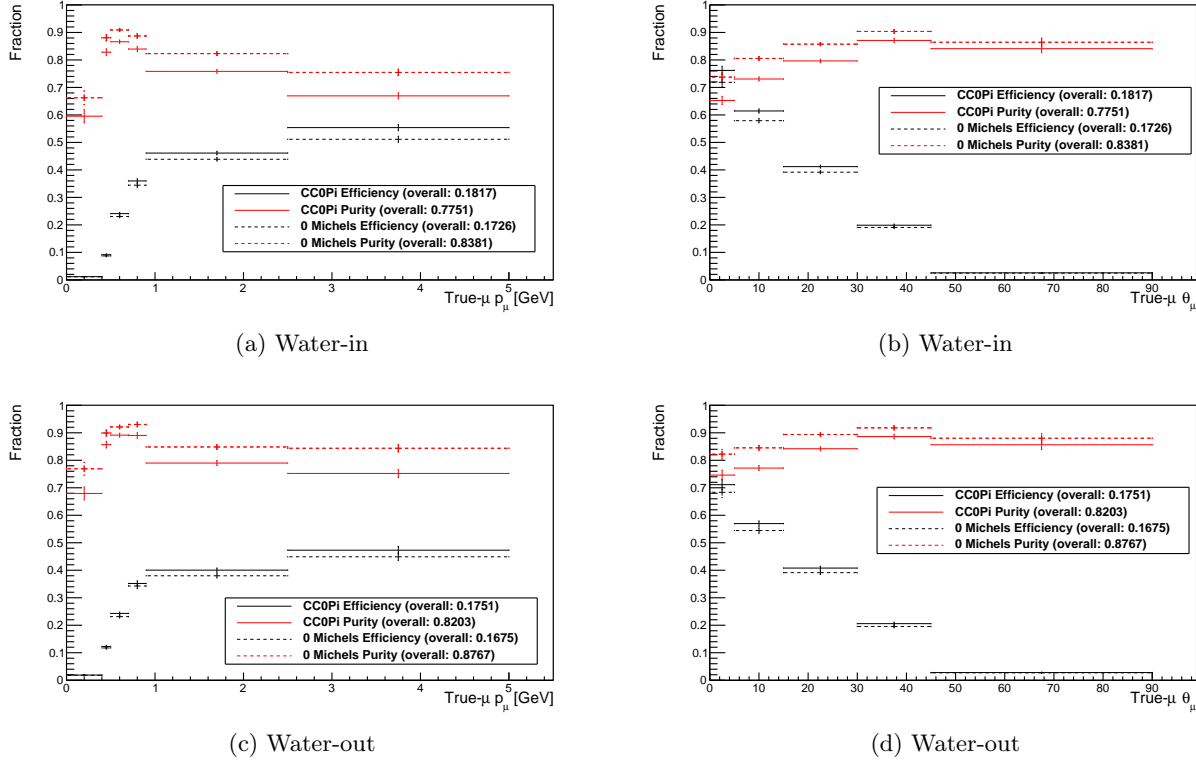



Figure 5.6: Comparisons of the efficiency and purity of the CC0 π selection used in this analysis (solid lines) against a selection with an additional 0-Michel cut (dashed lines). The common efficiency drop at low p_μ and high θ_μ is an artifact of requiring the track to enter the TPC as shown in Fig. 5.1.


tags all contained tracks as having PID `k0ther`. Because most secondary tracks will be entirely PØD-contained, the PØD PIDs are of limited usefulness for a 2-track selection. This makes it difficult to distinguish protons and charged pions without additional analysis of the energy deposition or Michel tagging, which is necessary for multi-track samples trying to distinguish CC0 π from CC $n\pi$. Michel electron tagging for a 2-track 0-Michel selection is possible, but decreases the CC0 π purity and introduces additional systematic uncertainties. Therefore, we decided to keep the selection simple and use multi-track samples as sidebands instead.

Figure 5.7 shows the water-in and water-out CC0 π signal selections in the global binning scheme used by this analysis (Section 7.1). The data is T2K Runs 2-4 and the MC is the production 6B NEUT sample with sideband constraints (Section 5.2.2) and all the corrections (Chapter 6) applied. The data errors are due to statistics. The MC errors are due to both statistics and

systematics (Chapter 6).



(a) Water-in



(b) Water-out

Figure 5.7: Comparison of the selection from data with a tuned and sideband constrained NEUT MC in the full binning scheme described in Section 7.1. The errors on the data selection are purely statistical. Errors on the MC selection are due to both statistics and systematics

5.2.2 Sidebands Selection

Sidebands are essentially background-enhanced selections that are used to estimate the background contribution to the signal region. To maximize the effectiveness of sidebands it is important to perform sideband selections with minimal signal ($CC0\pi$) contribution. Additionally, there must be no overlap between the signal selection and sideband selection, otherwise there could be events belonging to both. Multiple sidebands can be used, but they must also be non-overlapping and ideally describe different backgrounds. As the $CC0\pi$ selection is a restriction of the $CCInclusive$ selection to bunches with only a single reconstructed $P\bar{O}D$ object, sidebands for the background topologies are required to have multiple reconstructed $P\bar{O}D$ objects.

Two sidebands are used in this analysis, one for the $CC1\pi^+$ background and another for the $CCOther$ background. These are the two largest backgrounds in the signal selection. Sideband selection is identical to that described in Section 5.2.1 with the exception of the final cut, which is defined below.

$CC1\pi^+$ Sideband Exactly two $P\bar{O}D$ reconstructed objects in the bunch and nonzero number of $P\bar{O}D$ reconstructed Michel electrons.

$CCOther$ Sideband Greater than two $P\bar{O}D$ reconstructed in the bunch.

The reconstructed Michel electron requirement reduces the $CC0\pi$ component in the 2-track sample. This is important when selecting sidebands as they should ideally contain negligible contributions from the signal topology, in order to reduce correlations between the sideband and the signal that we are trying to measure. Both sideband samples are shown in Fig. 5.8 stacked by their true interaction topologies. It is difficult to isolate a highly pure $CC1\pi^+$ sample. However, both samples contain relatively small contributions from the signal topology. Discrepancies between data and MC exist, but are reduced when systematic errors are applied.

The selection for the $CC1\pi^+$ sideband requires an extra check for the number of Michel-tagged clusters. This is based on a $P\bar{O}D$ local reconstruction routine and may introduce additional systematics. A preliminary study using stopping cosmics to study data-MC discrepancies in the

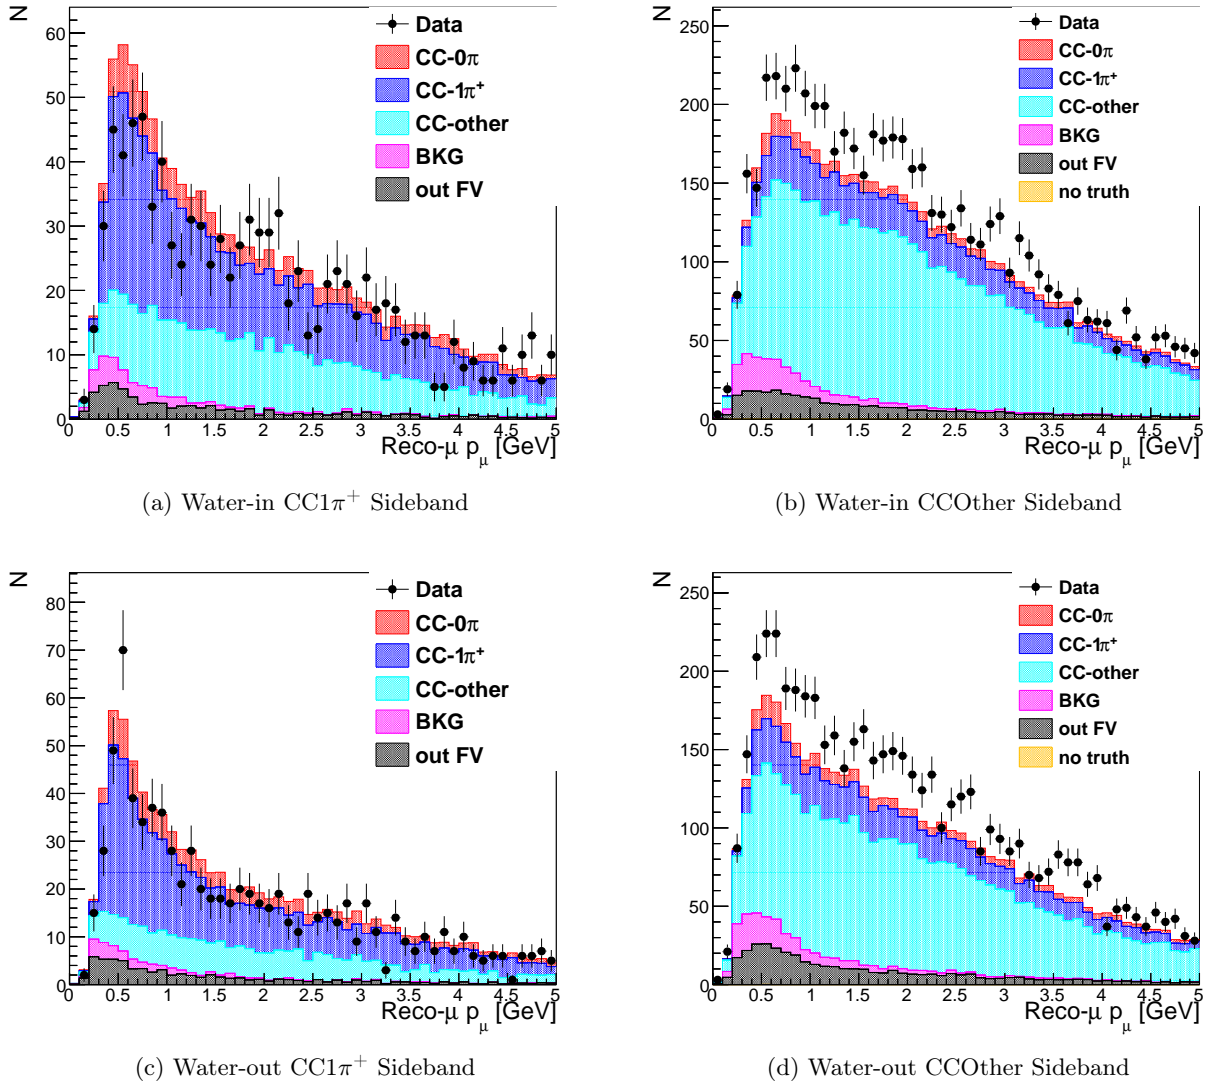


Figure 5.8: Data/MC comparison of the CC1 π^+ and CCOther sideband distributions. MC has been normalized to data POT. Error bars are statistical.

Michel reconstruction efficiency showed a disagreement on the order of 10%. However, this study is subject to less precision than a beam sample since cosmics trigger differently from the beam. This means that much of the difference may be attributed to the track timing and how tracks fall within the dead times of the readout cycles. A better approach is to use sand-muons, which follow the beam cycle timing. However, since this is only used as part of a sideband constraint on a component that is itself on the order of a 10% contribution to the total signal selection, the overall effect on the

signal should be no more 1 %. The Michel-tagging systematic is therefore negligible.



	χ^2/N_{DOF}	p-value
Water-in CC1 π^+	60.3/49	0.13
Water-in CCOther	45.2/49	0.63
Water-out CC1 π^+	94.0/49	0.00
Water-out CCOther	52.0/49	0.36

Table 5.6: The χ^2/N_{DOF} values obtained from comparison of the sideband data and MC distributions shown in Figs. 5.9 and 5.10.

	Water-in	Water-out
CC1 π^+	0.857	0.815
CCOther	1.129	1.127



Table 5.7: Overall sideband normalization scale factors obtained from a simple normalization rescaling of MC to data. These values were obtained with the corrected NEUT MC.

A comparison of the sideband MC selection to the sideband data is shown in Figs. 5.8 to 5.10. Figure 5.8 shows the Run 4 water-in and water-out distributions in a 1D p_μ binning and Figs. 5.9 and 5.10 shows a full comparison of the water-in sideband data and MC in the global binning scheme described in Section 7.1, which is used for the analysis. In the CC1 π^+ sideband, the MC tends to overestimate the data and in the CCOther sideband, the MC tends to underestimate the data. The errors shown in Fig. 5.9 come from both statistics and systematics. The data errors are purely statistical, while the MC errors have statistical uncertainties and systematic uncertainties added in quadrature. A χ^2 test gives the values listed in Table 5.6, and shows that the MC and data agree within errors. Detailed descriptions of the various sources of systematics are given in Chapter 6. The ratio between the overall data-MC normalizations give the scaling factors listed in Table 5.7. These can be used to constrain the background normalization in the signal selection as described in Section 7.3.2. Since we are using a simple normalization rescaling, shape differences between the sideband data and MC are not accounted for in the current framework. Studies have shown them to have little effect on the final result.

(a) Water-in $CC1\pi^+$ 

(b) Water-in CCOther

Figure 5.9: A comparison of the water-in data and MC sideband distributions in the binning scheme shown in Fig. 7.1. The MC has been scaled to match data POT. Error include all statistic and systematic error sources as described in the text.

(a) Water-out $CC1\pi^+$ 

(b) Water-out CCOther

Figure 5.10: A comparison of the water-out data and MC sideband distributions in the binning scheme described in Section 7.1. The MC has been scaled to match data POT. Error include all statistic and systematic error sources as described in the text.

Chapter 6

Systematic Errors and Corrections

The full ND280 MC simulation can be roughly divided into three stages (Chapter 3), beginning with the flux generation, moving through the neutrino interaction modeling, and ending with the ND280 detector simulation and track propagation through matter. In an ideal world, any uncertainty at any of these stages would be treated through a full re-simulation across the entire remaining MC chain. This is unfeasible in practice due to CPU limits and typically tweaks to underlying parameters are numerically propagated by event-by-event reweighting. A full list of all the underlying parameters in the first two stages of the simulation chain used by this analysis is given in Appendix C. Detector systematics are handled separately by the `highland2` and `psyche` framework as discussed in Section 4.2.2. Unlike flux and theory systematics, detector systematics are typically analysis and selection dependent since they are based on reconstructed variables.

Corrections can sometimes be confused with systematics. In general, corrections are applied to fix known problems with the MC and systematics are applied to address problems with no clear solution but are known to affect the MC. There are several corrections that are applied to the MC for this analysis. Each corresponds to a largely well understood difference between MC and reality. By correcting the MC ad-hoc, remaining differences between the MC and data can be attributed to the physics of interest (within uncertainties). Thus, corrections are applied to the nominal MC and affect the central value of the result, whereas systematics only affect the error as described in Section 7.4.

6.1 Flux

Flux corrections and uncertainties are parameterized in bins of true E_ν for each neutrino flavor at ND280 and SK [84]. Correlations between all bins, including cross-flavor and cross-detector correlations, are calculated to give the overall covariance matrix used by the oscillation and cross section analyses. The correlations between ND280 and SK allows for using near detector data to constrain flux uncertainties at SK.

6.1.1 Flux Correction

Hadron production is difficult to simulate correctly and T2K uses external data from NA61/SHINE [85, 86] to fine tune the nominal flux and constrain associated uncertainties. Measurements of pion and kaon production on two targets at NA61/SHINE have been used to tune the nominal JNUBEAM flux discussed in Section 3.1. The first dataset used corresponds to the 2007 data on a thin carbon target, called the thin-target data. The second data set corresponds to a 2009 release of both thin-target and a T2K-replica-target data [45]. The target can be removed to measure the out-of-target contamination. Additional constraints on pion rescattering multiplicities from the HARP experiment [87] have been recently included in the latest tunings [84]. Pion and kaon cross section measurements were performed and shown to cover most of the phase space relevant for T2K. From this a set of differential weights were generated from the ratio of data to both FLUKA (in-target) and GCALOR (out-of-target) models as a function of hadron kinematics [84]. These weights can then be applied to the nominal JNUBEAM flux, which contains the history of each neutrino produced.

Additional tuning due to measurements of proton beam parameters also are taken into account [84]. Proton beam parameters and errors are calculated based on measurements from the SSEMs, ESMs, and OTR monitors described in Section 2.2.2. Updated horn current measurements are also taken into account. The overall effect of tunings due to hadron interactions, beam profiles, and horn current is shown in Fig. 6.1.

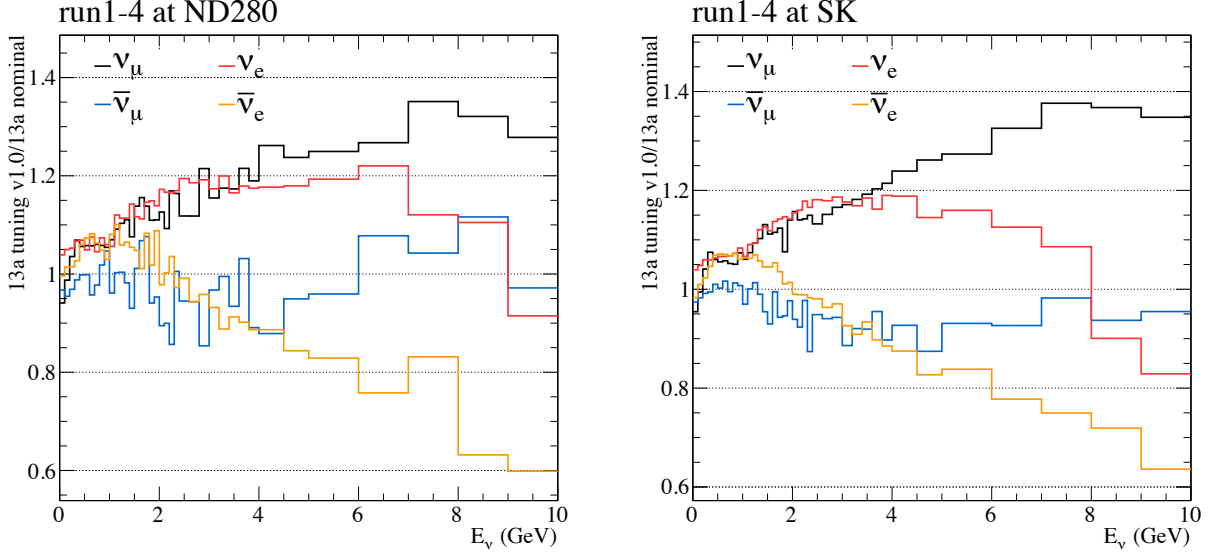


Figure 6.1: Flux tuning weights for Runs 1–4 at ND280 and SK. Weights are generated as a function of neutrino flavor, energy, and detector [84].

6.1.2 Flux Uncertainty

In T2K, the flux uncertainty tends to be the largest source of systematic uncertainty for cross section measurements. Various factors contribute to the flux uncertainty. The largest, as mentioned previously, arises from hadron production uncertainties. Even though external data is used to tune the nominal flux, there is still systematic uncertainty on the external measurements. Furthermore, uncertainties on scaling the data to different targets and incident particle kinematics contribute as well [47]. At low energies, uncertainties in pion and baryon kinematics and multiplicities dominate whereas at higher energies kaon production uncertainties dominate [84]. This is expected due to the preponderance of pion parents at low energies and kaon parents at high energies. The uncertainties from NA61/SHINE are propagated numerically to the nominal JNUBEAM flux using the same reweighting routine as was used for the tuning.

Other sources of flux uncertainty include beam alignment and monitoring, horn current and field, and MC statistics. All of these are shown stacked in Fig. 6.2 for the SK ν_μ flux. A proton beam normalization uncertainty is applied from the most downstream CT measurement (Section 2.2.2).




Figure 6.2: Fractional uncertainties from various sources on the ν_μ SK flux [45].

Proton beam profile uncertainties are dominated by the beam center and beam angle along the Y-axis. These are estimated by rerunning the JNUBEAM flux with a wide proton beam that covers the uncertainties of the measured beam profile [84]. The horn current uncertainty is estimated to be 2% or 5 kA. The effect of this uncertainty was studied by tuning JNUBEAM horn current values by this amount for each of the three horns [84]. The beam direction uncertainty, which translates to an error on the off-axis angle, is set by measurements from INGRID and their effects are propagated by moving the SK and ND280 detectors in JNUBEAM [47]. Additionally, target and horn alignment uncertainties were studied and propagated towards the overall flux uncertainty.

6.2 Interaction Modeling

Cross section and FSI uncertainties are parameterized either based on the internal NEUT parameterization or based on a template reweighting scheme from the T2K-internal Neutrino Interactions Working Group (NIWG) [88]. Several parameters were tuned away from the NEUT nominal values according to the recommended values obtained from fits to external neutrino scattering data [44]. Uncertainties on the parameters were propagated using part of the covariance matrix shown in Fig. 7.12. Since the PØD also contains brass, a tuning to account for the EMC effect (Section 6.2.2) is also performed for interactions on heavy targets.

6.2.1 Corrections to the Default NEUT MC

NEUT is the official MC used by T2K oscillation analyses and also for this cross section measurement. The ND280 MC relied on NEUT version 5.3.2, which was the first to use an implementation with the spectral function (SF) model for CCQE interactions on carbon and oxygen. Prior versions of NEUT used for MC production relied on a more simplistic Fermi gas (RFG) model. While the SF model is supposed to be a more accurate representation of nucleon momenta distributions within a nucleus than RFG, it does not take into account the Random Phase Approximation (RPA), which is a nuclear screening effect. Additionally, there are inconsistencies between the 1p1h (SF) and 2p2h calculations. As it currently stands, both SF and RFG are known to be imperfect models, but a choice had to be made on which model to use for obtaining the central values of the cross section.

A fit of three CCQE models,

- (1) SF+2p2h
- (2) RFG+relativistic RPA+2p2h
- (3) RFG+non-relativistic RPA+2p2h

to external MINERvA and MiniBooNE CCQE measurements was performed [88]. Results from these fits show that none of the models completely describe all of the data. The best fit model was the RFG+rel.RPA+2p2h and this is used for this cross section measurement [88]. In addition, a few of the cross section parameters listed in Appendix C were updated as part of the fit results. The tuned MC has an updated axial mass parameter $m_A = (1.15 \pm 0.03) \text{ GeV}/c^2$, 2p2h normalization parameter $2p2h \text{ norm.} = (27 \pm 12) \%$, and carbon Fermi momentum $p_F = (223 \pm 5) \text{ MeV}/c$.

6.2.2 Corrections due to the EMC Effect

In 1983 the European Muon Collaboration (EMC) discovered that the muon DIS cross section off nucleons bound in a nuclear medium exhibited dependence on the atomic number, A , as a function

of Bjorken x_{Bj} . Termed the EMC effect, much effort since its discovery has yielded no satisfactory answer as to its origin. As described in [89], the cross section ratio (to deuterium) σ_A/σ_D first gradually dips from approximately $0.3 < x_{Bj} < 0.7$ then sharply rises in the high x_{Bj} regime. The magnitude of this effect is seen to be larger for high A nuclei. Since the PØD water-target contains layers of brass (an approximately even mix of copper and zinc) originally intended to improve calorimetry, DIS interactions on these heavier targets was tuned using EMC effect data from [90].




Figure 6.3: Data points taken from [90]. In Fig. 6.3a the cross section ratios to deuterium are given for carbon and iron. σ_{Fe}/σ_C is shown in Fig. 6.3b and used to reweight interactions on Cu and Zn in the PØD.

As of this writing, the EMC effect has not been implemented in NEUT for any target. Indeed, it's not quite clear how to proceed in implementing some sort of EMC model for neutrino-nuclei interactions as naively reweighting all targets by σ_A/σ_D may conflict with other nuclear models. The most tested target in NEUT is carbon, so it was decided to reweight heavy target interactions relative to the carbon cross section. The original data in [90] was reported for discrete values of x_{Bj} . An approximation for the bin edges was assumed to be the midpoint between two adjacent x_{Bj} values as seen in Fig. 6.3. Due to missing data points, the σ_C/σ_D ratio is set to 1.0 for the three lowest and two highest bins. Furthermore, the original data did not report ratios for copper or zinc so the element closest in A , iron, was taken.

Currently, only DIS and multiple-pion interactions on copper and zinc are reweighted using Fig. 6.3. This includes both charged-current and neutral-current neutrino and antineutrino interactions. A comparison between cross section results with and without the EMC reweighting shows negligible differences because these specific interactions on heavy targets comprise a small percentage of the total number of events that pass our selection. Nevertheless, this correction is included for completeness, but no systematic uncertainty is assigned.

6.3 Detector

Detector systematics have a variety of sources. For this analysis, both PØD-based and Tracker-based systematics need to be included. The Tracker-based systematics have been studied intensively for Tracker selections and are implemented with modifications to make them work for interactions in the PØD fiducial volume [91]. PØD-based systematics have been studied independently and are covered individually in this section.

There are two types of detector systematic categories: efficiency and variation. Efficiency systematics are applied via event-by-event reweighting. V systematics are applied by altering the reconstructed kinematics of each event and rerunning the selection over the alternate reconstruction. In the `highland2` framework, each systematic can be either applied individually or in combination. Individually, their errors would be added in quadrature to obtain the total detector systematic. In combination, efficiency systematics have their weights multiplied together to get a total weight and variation systematics are applied in conjunction. For this analysis, the combination approach is used as it seems better suited to handle correlations between detector systematics. Detector systematics that affect this analysis are listed below.

- PØD momentum scale
- PØD momentum resolution
- PØD out of fiducial volume (OOFV)

- PØD fiducial volume mass
- PØD-TPC matching efficiency
- Tracker momentum scale
- Tracker momentum resolution
- Tracker B-field distortion
- Tracker charge efficiency
- Tracker TPC cluster efficiency
- Tracker TPC track efficiency

6.3.1 PØD Momentum Scale and Momentum Resolution

These two systematics stem from the energy loss of a track originating in the PØD and passing through to the tracker. The global reconstructed track can be separated into two parts, a PØD component and a tracker component. For each part, there is a momentum scale systematic and a momentum resolution systematic that alters the momentum as follows [92],

$$p' = p + (\Delta p)_{scale}^{tracker} + (\Delta p)_{res}^{tracker} + (\Delta p)_{scale}^{PD} + (\Delta p)_{res}^{PD}, \quad (6.1)$$

where,

$$(\Delta p)_{scale}^{PD} = (\mu_s + x_s \sigma_s)(p^{reco} - p_{TPC1}^{reco}) \quad (6.2)$$

$$(\Delta p)_{res}^{PD} = (\mu_r + x_r \sigma_r)(p^{reco} - p_{TPC1}^{reco} - p^{truth} + p_{TPC1}^{truth}). \quad (6.3)$$

Here $\mu_{s,r}$ are correction factors, $\sigma_{s,r}$ uncertainties, and $x_{s,r}$ are normally distributed variables thrown for systematic propagation.

A sample of backward-going stopping cosmics were used to study the effect of the PØD energy loss in data. Due to low MC cosmics statistics, particle gun MC was generated as a substitute and used for comparison. The stopping-cosmics sample is useful as it allows for a direct comparison of

Det	Vol	Water-out [%]	Water-in [%]
PØD	USECal	1.29	1.10
PØD	WT	20.2	26.9
PØD	CECal	78.5	72.0

Table 6.1: Percentage of PØD OOFV events in three different regions: The Upstream ECal ($z < -2969$ mm), Central ECal ($z > -1264$ mm), and Water Target (Section 2.4.1). The Central ECal is the most important part, followed by the water target. Again, these numbers are taken from the Run 4 ND280 MC dataset.

the reconstructed upstream TPC momentum to the reconstructed energy loss in the PØD. Since the tracks stop in the PØD, all its energy upon entering the PØD should be deposited in the detector. Differences between the data and MC TPC1 reconstructed momenta for tracks with the same length in the PØD can hint at differences between the GEANT4 detector model and the PØD as-built geometry. A total of eight variables (distinct sets of $\mu_{s,r}$ and $\sigma_{s,r}$ for water and air geometries) have been calculated and serve as the basis for propagating the scale and resolution systematics. An in-depth discussion can be found in [92].

6.3.2 PØD OOFV

In some cases, an event whose vertex occurred outside the fiducial volume may be mistakenly reconstructed within the fiducial volume. These are labeled Out of Fiducial Volume (OOFV) events. For this selection, OOFV events represent approximately 3% of the water-in sample and 4% of the water-out sample. Approximately 95% of these events occurred within the PØD but outside the fiducial volume. A small fraction also occur in the tracker ECal, with all other subdetectors representing less than 1% of all OOFV events.

Further investigation shows that the majority of events within the PØD occur in the PØD Central ECal (CECal), just downstream of the fiducial volume. The remainder largely occur within the water target (WT) region near the edges of the fiducial volume. These numbers are given in Table 6.1. Events in the CECal may be reconstructed in the fiducial volume due to backward-propagating tracks originating at the vertex. If these particles are not identified as separate tracks

from the muon candidate, the reconstruction will propagate the vertex backward to the most upstream hit. Events in the WT can also be reconstructed in the fiducial volume due to backward propagation of tracks. An additional topology where an event occurs outside the fiducial volume in non-active material such as water or brass but where the first hit in an active volume is in the fiducial volume is also possible.

In this analysis, OOFV events are separated into several classes based on their true vertex position:

- (1) PØD Central ECal
- (2) PØD Water Target
- (3) PØD Upstream ECal
- (4) FGDs
- (5) TPCs
- (6) ECals
- (7) SMRD
- (8) Other

Both the correction and systematic associated with OOFV interactions are applied as weight systematics where each OOFV event is given a weight $w = \mu + x\sigma$. Here, μ is a correction, σ is the uncertainty, and x is a normally distributed random number thrown separately for each category and for each systematic throw. The values for the TPC, ECal, SMRD, and Other categories are taken from the values used for tracker analyses. The FGDs are composed primarily of scintillator and, in the case of FGD2, water. These materials already have well-studied interaction uncertainties, so an additional rate uncertainty is unnecessary. The same is true for PØD water target events. Finally, the PØD central ECal events are due to backward migration, which may not be modeled well by the

generators used in the ND280 production MC or by the standard interaction uncertainties. This class of events is the only one assigned a reconstruction uncertainty.

The reconstruction systematic is taken from differences between data and MC results using the Run 4 dataset. We apply the signal selection (Section 5.2.1) on a different fiducial region to obtain a sample of events in the PØD central ECal. This alternative selection contains a mix of mainly carbon and lead interactions and is independent of the physics sample. The x scintillator layers within the PØD measure more events than the y layers due to events occurring in dead material between the PØDules. If the number of backward migrating tracks is being modeled improperly, the fraction of events being reconstructed in the x layers may differ between data and simulation. To estimate the possible magnitude of such a difference, we separate the MC samples into events that migrated backward at least one scintillator layer and events that did not migrate in such a way. We then extract the number of migrated events M_i per POT in each layer i and the number of non-migrated events N_i in that layer and compare that to the number of events per POT found in data, D_i . We construct a function $f(i|A,B) = AM_i + BN_i$ from the MC samples with two normalization factors A and B and fit the data to this function to extract best-fit scale factors for the migrated and non-migrated events. The results of these fits are plotted in Fig. 6.4. The fits do not perfectly match everywhere but are close to the data in all but a few bins, mostly in the first couple layers. This could be due to the transition between the PØD water-target region and the CECAL, which may be poorly modeled for backwards particles and thus lead to larger data-MC discrepancies that the fit cannot match within its constraints.

The result is that the best fit is obtained when the number of backward-migrating events is increased by a factor of two. The results for the water and air geometries are very similar, as expected since the water lies upstream of the central ECal. The exact numbers with fit uncertainties are given in Table 6.2. Unfortunately, a very large uncertainty must be applied to this value. The PØD does not contain any scintillator-only regions, so it is difficult to separate the effects of mis-modeling of backward migration from data to MC discrepancies in the total cross section. These shape uncertainties could conceivably also be from a difference in the lead to carbon cross section ratio




Figure 6.4: Fits for the central ECal for water-out (left) and water-in (right) configurations of the PØD geometry. Data indicated by error bars, fit result by the solid line.

between data and MC. As a result, we apply a reconstruction uncertainty with a correction μ of 2 and an uncertainty σ of 0.5. We chose a conservative estimate of $\sigma = 0.5$ for the uncertainty as there may be other systematics that could cause some of the data-MC discrepancy so the correction may be overestimated. As such, we don't want to exclude the nominal $\mu = 1$ as a possibility, and having $\sigma = 0.5$ means the nominal still lies within the distribution. This being said, a large uncertainty on the OOFV systematic doesn't actually change our results too much since this particular background is shared by both the water and air samples and should be mostly cancelled out in a subtracted result on water only.

6.3.3 Mass of the PØD Fiducial Volume

The mass uncertainty is implemented by applying weight corrections and uncertainties to events interacting on various target nuclei in the PØD. The mass weight w for an event is $w = \mu(1 + x\sigma)$, where μ is the correction, σ is the fractional uncertainty, and x is a normally distributed variable that changes for each throw. These values are calculated based on [83], which is a dedicated study of the PØD fiducial volume mass uncertainty. The uncertainties for all elements except oxygen are treated as fully correlated, while oxygen is uncorrelated from the others. Due to the subtraction of water-in and water-out samples, most of the non-oxygen uncertainties on the fiducial volume mass

PØD Geom	A	B
Air	1.97 ± 0.11	0.518 ± 0.040
Water	1.92 ± 0.11	0.601 ± 0.041

Table 6.2: Summary of fit results for the migrating event scale factor (A) and the non-migrating event scale factor (B).

should cancel out. The numbers used are summarized in Table 6.3.

As discussed in Section 7.4, the PØD mass systematic is different from other detector systematics which are treated within `highland2`. This is necessary because, for a cross section measurement, the PØD mass systematic affects not only the selected events, but also all true $CC0\pi$ events and the number of targets. The `highland2` framework currently does not allow systematics to be propagated for all true events, so the PØD mass uncertainty was treated separately outside the `highland2` framework. Its uncertainty is reported as an individual, separate component and added in quadrature with the other error sources.

6.3.4 PØD-TPC Track Matching

Since this analysis uses multiple subdetectors in ND280, global reconstructed variables are used instead of the local detector reconstructions. Inefficiencies from PØD-TPC track mismatches may contribute to errors in the selection. To study this effect a sample of FGD triggered cosmics was used. The cosmics were selected with a base criteria as follows.

- (1) Single Global track with a component in the most upstream TPC (TPC1)

Element	μ	σ
Oxygen	0.9868	0.0084
Lead	0.9988	0.0218
Copper/Zinc	1.1448	0.0332
Other	0.9998	0.0093

Table 6.3: Corrections (μ) and fractional uncertainties (σ) applied to events on various target elements found in the PØD. The correction for brass (copper and zinc) differs significantly from unity due to the inclusion of dead material not otherwise modeled in the MC geometry.

- (2) Require the track have greater than 18 nodes in TPC1

These are further categorized as “PØD-going” with the additional requirements that

- (1) Track is present in the first couple layers of TPC1.
- (2) The linearly projected position into the PØD is within the fiducial volume.
- (3) The TPC1 reconstructed momentum is greater than 250 MeV.

Since cosmics trigger randomly, two issues arise. First, the hit times are not reconstructed relative to the beginning of an integration cycle, and second, hits do not always fall within the live time of PØD cycles. The first can be fixed by calculating an adjusted hit time,

$$t_{adj} = (t_{rec} + t_{MCMTSB} * 10) \mod 580, \quad (6.4)$$

where t_{MCMTSB} is the master clock module’s time since it last reset and each cycle has duration 580 ns. Before resetting the t_{MCMTSB} , the PØD bunch cycles are read out and the zeroth cycle is aligned with $t_{MCMTSB} = 0$. Regarding the second issue, the PØD cycle accommodates the beam spill such that the eight bunches extracted per spill all fall within the active regions of the first eight cycles registered by the PØD. However, for every bunch there is also a dead-time window of approximately 80ns during which the PØD electronics reset and no data can be collected. This forces an additional cut on the upstream TPC’s t_{adj} when selecting PØD-going cosmics. It is assumed that cosmics with the restricted TPC1 t_{adj} should have all hits falling within the PØD’s active-time window. By checking to see how often these active-time, PØD-going events actually have a globally reconstructed PØD object, it is possible to get a handle on the PØD reconstruction and PØD-TPC matching efficiency.

The active-time, PØD-going cosmics sample was divided into water-in and water-out. Both have less than 0.2 % overall matching inefficiencies. Further, a MC study using a sample of simulated cosmics shows 100 % MC efficiency. Therefore, this systematic is assumed to be negligible.

6.3.5 Tracker Systematics

Detailed studies of the ND280 Tracker systematics are covered in [91]. These systematics were originally implemented in the `highland2` framework for use by Tracker analyses, but many are applicable for tracks originating in the PØD by separating the tracks into separate PØD and Tracker components. The Tracker systematics that affect this analysis are listed at the beginning of this section. An overview is given here.

The TPC momentum uncertainties are applied in a similar manner as the PØD momentum uncertainties, but on the Tracker component of the global track. The TPC momentum scale uncertainty is obtained from uncertainties in the magnetic field measurements and set to be 0.57 % [91]. The TPC momentum resolution uncertainty was evaluated using tracks crossing multiple TPCs. This allows for the calculation of differences between the reconstructed momentum from two TPC segments of the same track. Energy loss in the intermediate FGD was corrected for prior to calculating the difference. The resulting distribution of the difference between the inverse transverse momentum to the magnetic field, $\Delta(1/p_t)$, is fitted to a Gaussian for both data and MC. Any difference between data and MC introduces a systematic error which is accounted for by an additional smearing factor in the error propagation [91].

The magnetic field distortions were mapped using Hall probe measurements of the magnetic field inside the ND280 basket. Results from this map are applied as a correction at the reconstruction level. A second empirical study was performed based on measurements from the TPC photoelectron calibration system [93]. This system pulses UV laser light onto aluminum dots in the cathode to produce photoelectrons. These then drift and are detected by the MicroMegas detectors (Section 2.3.2) and allow field mapping. This secondary, empirical study gave values to be used as the systematic uncertainty for the B-field distortions.

Charge misidentification probabilities are calculated based on the number of TPC segments, their individual charge reconstructions, and the curvature (inverse transverse momentum). Two probabilities are calculated [91]:

- probability of swapping the individual TPC's charge identification
- probability of the merged global track to swap the charge obtained from the best local TPC segment

and tracks are categorized by the number of TPC segments they passed through. Then, the charge misidentification systematic can be parametrized by comparing the differences between data and MC of these probabilities for the different track categories. It is propagated as a weight (or efficiency) systematic.

Differences between the data and MC of the TPC cluster reconstruction efficiency affect TPC track quality cut. The efficiency difference was calculated by fitting the MC distribution of the number of TPC nodes per track to the data using two different samples [91]. The first consisted of a CCInclusive sample originating in the FGD-1 and the second consisted of a sample of TripT-triggered cosmic rays. This difference is then propagated as the TPC cluster efficiency systematic.

Finally, the TPC track reconstruction efficiency, or the probability of a TPC successfully reconstructing a track from a passing particle, was evaluated for both data and MC for each of the three TPCs [91]. Control samples from both beam and cosmic datasets were used and the reconstruction efficiency was calculated using the upstream and downstream detectors around the tested TPC. If tracks exist in both the upstream and downstream detectors but not the tested TPC, that is a contribution to the inefficiency [91]. It turns out that the TPC track reconstruction efficiency is very high for all three TPCs, and so it was not deemed to depend on momentum or track length. The lowest value amongst all bins was taken as the overall track reconstruction efficiency separately for the three TPCs as given in Table 6.4. Uncertainties on the efficiency were used for the systematic propagation.

	TPC1	TPC2	TPC3
Data Efficiency	$99.9^{+0.1}_{-0.1}\%$	$99.7^{+0.2}_{-0.7}\%$	$99.3^{+0.5}_{-0.2}\%$
MC Efficiency	$99.6^{+0.2}_{-0.3}\%$	$99.5^{+0.3}_{-0.4}\%$	$99.8^{+0.1}_{-0.2}\%$

Table 6.4: The TPC track reconstruction efficiencies for all three TPCs [91].

6.4 Pile-up Correction

The ND280 beam MC does not simulate cosmic rays or neutrino interactions outside the ND280 magnet, backgrounds which may affect the selection. The MC selection as described in Section 5.2.1 does not factor into account pile-up events from cosmic rays falling within the beam bunches or neutrino induced muons produced outside the ND280 that then enter the detector (sand-muons). Pile-up is especially important for this analysis due to the single-track cut in the selection discussed in Section 5.2.1. Any sand or cosmic background that coincides with a beam bunch can therefore bias the MC selection and a pile-up correction factor was calculated from the data and applied to the MC.




Figure 6.5: Event rates for in-beam cosmic rays and sand-muons over the course of T2K runs 1–4. The sand-muons rate dominates and contributes to the pile-up correction factor.

In-beam cosmic rays were studied using events in the PØD that occurred prior to the first beam bunch. The event rate is shown in Fig. 6.5a with an overall rate of $3.4 \times 10^{-7} \text{ ns}^{-1}$ or $2 \times 10^{-4} \text{ bunch}^{-1}$. This gives an approximation for $\Pr(N_{\text{cosmics}}/\text{bunch} > 0) = 2 \times 10^{-4}$, which compared to the sand-muons background is negligible.

Sand-muons are the other data-only background and were studied with a selection of global tracks with hits in the three most upstream layers of the PØD. These hits are typically caused either

Run	p_{sand} [bunch ⁻¹]
1	0.0085
2 water	0.0144
2 air	0.0173
3b	0.0158
3c	0.0188
4 water	0.0217
4 air	0.0245

Table 6.5: Sand-muon pile-up rates for each run. These correspond to the average number of sand-muons per bunch in the specified run as measured from the event-rate difference between data and magnet MC, as explained in the text. MC bias is corrected for by applying a weight of $1 - p_{sand}$ to all selected events.

by interactions in the magnet in front of the PØD or from interactions outside the ND280. Since the MC does not simulate sand-muons, the event rate difference between data and MC can be attributed sand-muons. The total event rate per spill in data is shown in Fig. 6.5b. The general increasing trend corresponds to the increasing beam intensity. Overall background rates are given in Table 6.5 for each run. The single-track cut in the selection means that any time a sand-muon is reconstructed in the PØD the event is rejected. To correct for this in MC, the MC is reweighted by $1 - p_{sand}$ where p_{sand} is the rate of sand-muons per bunch as given in Table 6.5. While much higher than the in-beam cosmics rate, this corresponds to a less than 3% correction for all runs.

The uncertainty associated with this correction can be conservatively estimated using the total MC pile-up event rate, which is subtracted from the data in the calculation of the sand-muon pile-up rate. These are due to simulated interactions in the magnet and varies slightly depending on the MC beam setting. However, with the highest rate being $0.0034 \text{ bunch}^{-1}$, it is assumed to have a negligible effect on the final result.

Chapter 7

From Selection to Cross Section

We report a double differential cross section measurement in bins of p_μ and $\cos\theta_\mu$, the momentum and cosine of the angle relative to the z-axis of the outgoing μ^- immediately after the interaction. However, detector reconstructed kinematics are imperfect approximations to the μ^- 's true initial state. To correct for these we employ D'Agostini's Bayesian unfolding technique [77, 78]. The purpose of Bayesian unfolding is to remove imperfections in the reconstruction in order to achieve an accurate representation of how the μ^- emerged from the interaction.

To extract a measurement on water requires an additional subtraction step after unfolding. The PØD fiducial volume contains plastic scintillator layers and thin brass sheets sandwiched between layers of water [52]. The water layers act as a passive target making it difficult to know whether an interaction occurred on water or on some other target nucleus. To work around this, the PØD was designed to be drained and filled during different run periods. All things being equal except for the inclusion or exclusion of water, a subtraction of the true distributions should give the distribution of interactions on water. The idea for a measurement on water is thus to first separately unfold the reconstructed distribution for water-in and water-out to get an approximation of their true distributions, then subtract the post-unfolded results to get the distribution that should have occurred on water. The number of interactions on water is thus given as,

$$N_i^O = \frac{U_{ij}^w N_j^w}{\epsilon_i^w} - R \frac{U_{ij}^a N_j^a}{\epsilon_i^a}, \quad (7.1)$$

where the indexes i and j indicate true and reconstructed bins, respectively, N^w is the number of events measured during the water-in period, N^a during the water-out period, ϵ the selection efficiency,

and R the flux normalization factor between water-in and water-out periods. U_{ij} represents the unfolding matrix which will be discussed shortly. From this, the differential cross section on water can be expressed as,

$$d\sigma_i = \frac{N_i^O}{F^w N_n \Delta_i}, \quad (7.2)$$

where F^w is the integrated flux over the water-in period, N_n the number of neutrons, and Δ_i the bin area. Nominally, N_n is calculated to be 5.09×10^{29} based on the fiducial water mass reported in [83]. F^w is calculated based on the flux simulation discussed in Section 3.1. These are POT-normalized fluxes that may be slightly different between runs, so F^w is the POT-weighted sum of all water-in periods. Similarly, the total flux for the water-out periods, F^a , is calculated and used to scale the flux normalization ratio, $R = F^w / F^a$.

7.1 Binning

Since the cross section measurement stems from binned distributions across the muon kinematics, the binning used for unfolding is an important component of the analysis. This analysis uses the binning shown in Fig. 7.1 for both the true and reconstructed distributions. It is easiest to visualize plots and uncertainties in 1D so a global bin numbering scheme is used for the projection. Importantly, the bin-by-bin migration probabilities is dependent on the entire MC truth sample, which means out-of-range events must be included in their calculations. In the `xsTool`, these are handled by a single overflow bin corresponding to the 0th global bin.

The choice of this binning was motivated by previous analyses [81, 94] and a desire to have smaller bins in high statistics regions. Additional studies show the fractional statistical uncertainty to be relatively constant over the binning range and, after subtraction, comparable to systematic uncertainties. Ideally, we would have the fractional statistical uncertainty constant across all bins, but as seen in Fig. 7.2, there is a simultaneous dependence on p_μ and $\cos\theta_\mu$ that requires varying p_μ bins across different slices of $\cos\theta_\mu$. This is currently not possible under the constraints of the `xsTool` framework.




Figure 7.1: The binning used for this measurement with corresponding global bin number in red. Identical binnings are used for the reconstructed and truth distributions (i.e. before and after unfolding). Note that out-of-range events are included and treated as part of the unfolding procedure as a single overflow bin corresponding to the 0th global bin.




Figure 7.2: Unstacked fractional uncertainties from statistics and systematics. This was extracted using toy-MC with POT scaled to match actual data. Systematic sources are described in more detail in Chapter 6.

Figure 7.2 was generated using toy-MC but with similar POT as real data. The bin-by-bin uncertainties are calculated from the diagonals of the full covariance matrix whose construction is described in Section 7.4. Data statistics is a large source of uncertainty. Even though compared to FGD analyses there are relatively many interactions in the PØD, the subtraction increases the statistical uncertainty. The bin-by-bin fractional uncertainty from systematics and statistics varies in Fig. 7.2, but contribution from statistics is higher in most bins. However, from the uncertainties on the total cross section listed in Table 7.1, the contribution from systematics is higher than from statistics. This is expected since more bin-by-bin correlations exist for systematic errors, whereas statistical errors are generally independent across bins.

7.2 Bayesian Unfolding

Detailed descriptions of the Bayesian unfolding algorithm are given by D’agostini in [77] and [78]. An overview of the method is given here.

It is simplest to think about the unfolding in terms of two main categories, true kinematics (causes) and reconstructed kinematics (effects). Misreconstructions will smear the correspondence between cause and effect as shown in Fig. 7.3.

Source	Fractional Uncertainty [%]
Statistics (total)	7.9
Systematics (total)	10.7

Table 7.1: Contributions to the error on the total cross section due to statistics and systematics. Calculated using a full toy-MC study.




Figure 7.3: Cartoon diagram of smearing between a true distribution and reconstructed distribution.

In order to recapture the true distribution, a probabilistic approach is taken. First, $\Pr(E_j | C_i)$, the probability of an event in cause bin i migrating to effect bin j , is calculated from the MC. By Bayes' theorem the posterior can then be calculated as,

$$\Pr(C_i | E_j) \propto \Pr(E_j | C_i) \Pr(C_i) \quad (7.3)$$

to give the conditional probability of an event falling into true bin i given that it was measured in effect bin j . The prior probability $\Pr(C_i)$ is a choice, but is typically taken from MC truth. This is precisely what is needed to go from a measured reconstructed distribution to the true distribution. To continue the simple example from Fig. 7.3, the likelihood and posterior matrices can easily be seen to be,

$$\Pr(E_j | C_i) = \begin{pmatrix} \frac{2}{3} & \frac{1}{3} & 0 \\ 0 & \frac{2}{3} & \frac{1}{3} \\ 0 & 0 & 1 \end{pmatrix} \Rightarrow \Pr(C_i | E_j) = \begin{pmatrix} 1 & \frac{1}{3} & 0 \\ 0 & \frac{2}{3} & \frac{1}{3} \\ 0 & 0 & \frac{2}{3} \end{pmatrix}. \quad (7.4)$$

The unsmearing matrix can then be applied to a reconstructed distribution. Clearly, unsmearing the original reconstructed distribution used in the creation of the smearing matrix returns back the

original truth distribution. An example of the full unsmeared matrix used to unfold the water-in sample for this analysis is shown in Fig. 7.4.

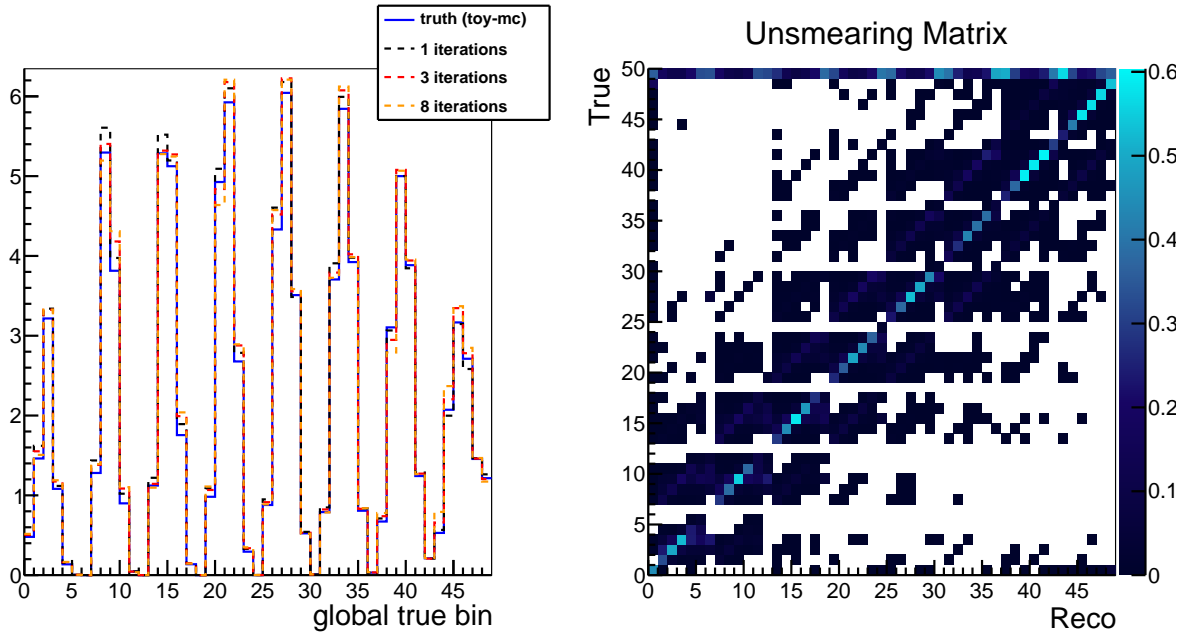


Figure 7.4: A toy-MC study of unfolded results across 1, 3, and 8 iterations of D'agostini's Bayesian unfolding algorithm. The unsmeared matrix is shown at right with the overall binning scheme explained in Fig. 7.1. The zeroth bin covers the out-of-range phase space. The last (extra) bin in truth covers all background contributions.

The `xsTool` package wraps an external unfolding framework, RooUnfold. Since the unfolding procedure can be rather complicated, it is important to ensure that the implementation works properly. To test this, two identical NEUT MC samples were used as both the MC prior and the toy-MC input. No reweighting or subtraction was applied to either sample to ensure a pure test on the unfolding procedure. The result is shown in Fig. 7.5. An exact agreement between the unfolded result and the truth is found, which indicates that the routine works properly.

7.2.1 Number of Iterations

This technique also allows for successive iterations with the prior updated after each iteration using the result from the previous iteration. However, a single iteration has typically been the




Figure 7.5: A test of the unfolding procedure using identical NEUT samples for the MC prior and toy-MC input. Exact agreement is found between the truth, which is the same for MC and toy-MC, and the unfolded result.

most favorable and used by other analyses [95]. This is believed to be due to starting with the MC prior which returns a reasonable result with smaller uncertainties after one iteration [96]. Multiple iterations increases the uncertainties on the result [97] and this is another limiting factor.




Figure 7.6: The fractional bias and fractional error vs the number of iterations set for the Bayesian unfolding engine. Both exhibit similar trends. Note that these were calculated using the total cross section, which is a simple proxy for the differential cross section. The iterations used are 1, 2, 4, 8, 16, 32, and 64.

For this analysis, an independent test was performed to check the bias and error as a function

of the number of iterations, starting with the MC prior. This test used the total cross section result as an approximation for calculating the bias and uncertainty over a range of iterations. The full MC statistics sample, separated into water-in and water-out datasets, was used. The fractional bias is defined as,

$$f_b = \frac{\sigma_{meas} - \sigma_{true}}{\sigma_{true}}, \quad (7.5)$$

and the fractional error (statistical only) is defined as,

$$f_e = \frac{\delta\sigma_{meas}}{\sigma_{meas}}. \quad (7.6)$$

As Fig. 7.6 shows, both f_b and f_e converge to some value at high iterations. While it is nice to see that the convergence is quick, f_e increases quickly in both cases, which is quite undesirable for a water-air subtraction.

The interplay between the response matrix (likelihood in (7.3)) and prior that determines the unfolding matrix (posterior) and ultimately the unfolded result is tricky to intuitively understand. In essence increasing the number of iterations is akin to a feedback loop, which hints at why f_e tends to increase. The response matrix and efficiency remain unchanged through all iterations. Only the prior, normalized to unity, is updated after each iteration. This has the effect of altering the unfolding matrix. If the background is added as an additional cause bin (c.f. Appendix B) this will in essence change the purity, along with the signal unsmearing matrix, after each iteration.

In the case of unfolding two identical samples as in Fig. 7.5, only the MC truth prior gives an exact agreement between the truth and the result since both the response matrix and the prior were generated from that MC sample. When unfolding non-identical samples, the response matrix remains entirely from the MC while the prior is updated. Even if, after a certain number of iterations, the prior is in perfect agreement with the data (toy-MC) truth, unless the response matrix is exactly identical for both samples, the next iteration will not give a perfect result. This explains why there can be no absolute convergence unless the response matrices are identical.

While it seems that multiple iterations allows for some convergence, the limiting factor is really the increase in uncertainty. Both statistic and systematic uncertainties increase. With realistic




Figure 7.7: A toy-MC study of the double differential result (without subtraction) for one and four iterations. NEUT was used as MC and GENIE as toy-MC. Errors are statistical.

MC, there is only slight gains in bias-reduction with increased iterations. Additionally, the total cross section is a simple proxy for the behavior of the double differential result. A full phase space comparison of the unfolded result is shown in Fig. 7.7. With four iterations, the statistical uncertainty increases dramatically for perhaps a slight improvement in the central value across a few bins. Therefore, this analysis proceeds with a single iteration.

7.3 Background Treatment

The largest source of background in the $CC0\pi$ selection is $CC1\pi^+$ as shown in Fig. 5.3. Secondary contributions from the $CCOther$ topology consist of primarily multi-pion events from DIS. Rounding out the background categories, events migrating in from outside the PØD fiducial volume are classified as the out-of-fiducial-volume (OOFV) background, and events without a muon in the final state are classified as general backgrounds (BKG). Several approaches to handling backgrounds were studied for this analysis,

- (1) Purity correction

(2) Background subtraction

(3) Simultaneous unfolding

and the relative merits of each, along with the effect of using sidebands, will be discussed in this section. As discussed in Appendix B, under certain circumstances purity correction is actually identical to the default background treatment in `RooUnfold`, whereby backgrounds are incorporated as an additional cause bin and unfolded along with the signal.

Depending on the background treatment method, the result may be different, the errors on the result may be different, and the effect of sidebands may be different. As the addition of sidebands can complicate things, the first part of this section will focus on the effect of different background treatments without sidebands. Following this, we discuss how sidebands can affect the background. In Section 7.4 we discuss how backgrounds can affect errors. All error bars shown are statistical, unless otherwise noted.

7.3.1 Without Sidebands

First consider the case of no sidebands. In a binned analysis, purity correction simply adjusts the data bin-by-bin using a scaling factor extracted from the MC. Background subtraction subtracts the MC background directly from the data. Simultaneous unfolding reduces to unfolding a single sample and is identical to purity correction.

The difference between background subtraction and purity correction affects only N_j^w and N_j^a of Eq. (7.1). The unfolding matrix itself is constructed using the true signal component of the MC selection. Thus, we can focus on how backgrounds affect the number of events. Ignoring the water-air subtraction, for a single bin the cross section behaves approximately as,

$$\sigma \propto N_d - B \quad (\text{Background subtraction})$$

$$\sigma \propto pN_d. \quad (\text{Purity correction})$$

Here, N_d is the number of measured events in the data, B is the MC predicted background, and

$p = 1 - B/N_m$. Thus, for purity correction the cross section can be rewritten as,

$$\sigma \propto N_d - B \frac{N_d}{N_m}. \quad (7.7)$$

Since B comes entirely from the MC, it's independent of N_d . On the other hand, p is simply a scaling factor that is applied directly to N_d . This means that the effective subtracted background in the case of purity correction is dependent on N_d . In other words, with purity correction, the data background estimate depends on the measured data.

7.3.2 With Sidebands

Now consider the case with sidebands. Sidebands work by constraining the MC background using an additional data sample. For this analysis, the two sidebands mentioned in Section 5.2.2 are used. The importance of sidebands is illustrated in Fig. 7.8. In this study, the toy-MC had its $CC1\pi^+$ (primary background) component reweighted by a factor of ten. The sample is then unfolded using the default NEUT MC. As the MC does not model the toy-MC's background correctly the background is not treated properly without sidebands, and this leads to results vastly different from truth. Sidebands provide a data-driven method for estimating the backgrounds.

The binning used for the sidebands described in Section 5.2.2 is the same as the signal (Section 7.1). In the case of a sideband comprised of only a single background category (i.e. a pure sideband), the background is constrained through an overall normalization as,

$$B' = B \frac{S_d}{S_m}, \quad (7.8)$$

where B' denotes the sideband constrained background, B is the original background in the MC signal selection, and S_d and S_m symbolizes the measured sidebands from data and MC respectively. Here B and B' represents any bin in a binned analysis.

In the case of background subtraction, B' is then directly subtracted from the total data N_d . Since there is no constraint on S_d the value of B' may be greater than N_d and thus lead to negative bins after background subtraction. This effect has been seen in other analyses, specifically in [94].




Figure 7.8: Here, the toy-MC sample has all its $\text{CC1}\pi^+$ background reweighted by a factor of ten but the MC is not sideband constrained. As can be expected, the background treatment is incorrect for all methods and the results deviate from truth by a large amount. Background subtraction tends to be higher than purity correction because, as the toy-MC increases, the purity factor reduces data by a larger amount than the background subtraction, which relies on the constant MC absolute background.

We also see this when using pure sidebands to extract the cross section with the $\text{CC1}\pi^+$ -reweighted toy-MC sample in Fig. 7.9. In a more direct measurement, perhaps it is fine to quote a statistically insignificant signal in bins where the background estimate tends to be higher than the measured data. However, unfolding introduces additional complexities as those negative reconstructed bins will affect neighboring bins through the unsmeared matrix. In theory, this can also occur without sideband constraints. This type of behavior is quite undesirable, especially for a blind analysis. If the method can lead to unphysical values with toy-MC, then it can lead to unphysical values with real data. In addition, as discussed in Section 7.4.1, the statistical errors on the result using background subtraction tend to be larger than those from purity correction.

In the case of purity correction, instead of subtracting the scaled background directly, an

adjusted purity factor is calculated as,

$$p' = 1 - \frac{B'}{N'_m}, \quad (7.9)$$

where $N'_m = N_m - B + B'$. This is equivalent to subtracting an effective background B_{eff} which can be seen by setting,

$$N_d - B_{eff} \equiv p' N_d \quad (7.10)$$

$$B_{eff} = N_d - p' N_d \quad (7.11)$$

$$= B' \frac{N_d}{N'_m}. \quad (7.12)$$

Since $B' \leq N'_m$, we must have $B_{eff} \leq N_d$. Thus, no negative values are allowed in the case of purity correction. In the case of multiple sidebands, this result holds still as long as N'_m is taken to be the final overall number of MC events after all sidebands have been taken into account.

The last method considered is simultaneous unfolding, which by definition affects the result only if sidebands are used. The idea is to expand the unfolding to include both reconstructed signal and reconstructed sideband and unfold those simultaneously into the true signal and background. As realistic sidebands are not 100 % pure in the background they are meant to constrain, the signal and sideband regions can affect each other through the unsmeared matrix. In principle, the sideband data should constrain the background after successive iterations of the unfolding, leading to increasingly accurate MC priors even if the MC originally fails to model the data correctly. This behavior can be seen in Fig. 7.10.

As already mentioned, to test the behavior of these different treatment methods, we reweighted $CC1\pi^+$ background events in toy-MC by a factor of ten. This meant directly using the MC would produce very inaccurate results as the MC completely underestimated the data background. Comparisons of results using the three different methods without sidebands are shown in Fig. 7.8. The first 0–48 bins correspond to the $CC0\pi$ selection, followed by the $CC1\pi^+$ sideband, and ending with the $CCOther$ sideband. The purity scaling factor has a larger effect than the subtraction of absolute backgrounds since it is correcting the increased, reweighted data. This explains why




Figure 7.9: Here, the toy-MC sample has all its $CC1\pi^+$ events reweighted by a factor of ten and the MC is sideband constrained with artificial sidebands forced to be 100 % pure. Bins 1–48 correspond to the $CC0\pi$ signal, 49–96 the $CC1\pi^+$ sideband, and 97–144 the $CCOther$ sideband. The background subtraction exhibits negative bin behavior. The simultaneous unfolding shows poor results after one iteration as the reweighted sideband data feeds into the unfolded signal which leads to an increase in the unfolded result.

background subtraction is higher than purity correction for this test without sidebands. On the other hand, a 100 % pure sideband sample does a good job of constraining the MC background using purity correction or background subtraction after one iteration as seen in Fig. 7.9. However, because background subtraction allows for negative bins it is not used for this analysis. The simultaneous unfolding performs poorly after a single iteration. After successive iterations, the simultaneous unfolding can be seen to converge to the truth as seen in Figs. 7.10 and 7.11. However, as discussed in Section 7.2.1, multiple iterations cause larger statistical uncertainties and a single iteration is used for this analysis. Coupled to the fact that with realistic sidebands simultaneous unfolding requires upwards of 20 iterations and still doesn't seem to perform better than purity correction, the simultaneous unfolding method was also rejected. This analysis therefore uses sidebands and

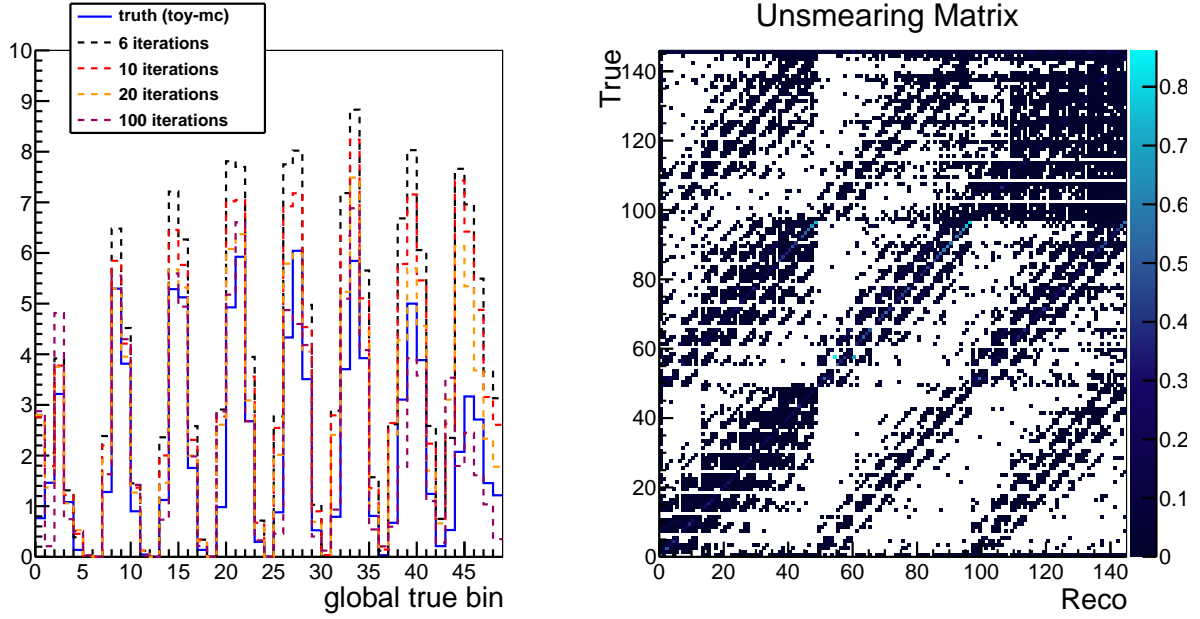


Figure 7.10: A study of the simultaneous unfolding method. True $\text{CC1}\pi^+$ events are reweighted by a factor of ten in the toy-MC. The MC is then sideband constrained using the sidebands defined in Section 5.2.2. Only after successive iterations does the result approach the toy-MC truth. The axes of the unsmearing matrix is separated into signal, sideband 1 and sideband 2, each of which is then binned across the physical (p_μ, θ_μ) -space. Statistical errors are not shown but increase with the number of iterations.

treats backgrounds using the purity correction method. There are caveats with purity correction as well that stem from the fact the MC purity may not match reality. An extreme example is shown towards the end of Section 7.5.2, with the conclusion being that a realistic MC with reasonable model uncertainties is required to properly perform any such cross section measurement.

7.4 Propagation of Uncertainties

Uncertainties can be roughly categorized into five independent categories. Each category can affect different components of Eq. (7.1) and Eq. (7.2). The five categories are summarized as follows.

Statistical treated by `xsTool` using Poisson throws.

Detector (PØD mass) treated by by `xsTool` via stand-alone implementation




Figure 7.11: Identical configuration as Fig. 7.9 but with six iterations instead of one. The simultaneously unfolded result can be seen converging across both signal and sideband regions. However, statistical uncertainties increase as a side effect.

Detector (other) treated by `highland2` via variation and reweighting

Flux treated by `xsTool` and `T2KReWeight`

Cross section treated by `xsTool` and `T2KReWeight`

FSI treated by `xsTool` and `T2KReWeight`

All uncertainties are numerically calculated using pseudo-experiments or “throws” that are each run through the unfolding framework. Besides the data statistical uncertainty which alters the actual data selection, all the other error sources affect only the MC and the overall normalization. Each throw corresponds to an alternative hypothesis. This means that the unfolding matrix, efficiency, purity, integrated flux, and number of targets, may change depending on which error source is being propagated. Each throw will therefore give a different cross section result and a covariance matrix is built throw by throw across the true bins in which the cross section measurement is made. A sample

covariance matrix given by

$$\text{cov}(X_i, X_j) = \frac{1}{T-1} \sum_{t=1}^T (x_{ti} - \overline{x_i})(x_{tj} - \overline{x_j}), \quad (7.13)$$

where X_i and X_j symbolize random variables corresponding to the result in bins i and j , T is the total number of throws, x_{ti} and x_{tj} are the measured cross section values in those bins for throw t , and the overline denotes the sample mean across all throws, can be calculated for each error source to form a bin-by-bin covariance across differential bins in the cross section result. In this way, a covariance matrix is built for each of the five categories enumerated at the beginning of this section. The five covariances are then added to form the final overall covariance matrix which is then used to extract the overall uncertainty for each bin. The diagonals of the covariance matrix gives the variance for each bin. As discussed in Section 7.4.5, correlated and uncorrelated errors are taken into account across the water-air subtraction. Concrete examples of how the throws feed through the cross section calculation are given in Section 7.4.6.

For this analysis 2000 statistical throws were used to generate the statistical errors. Propagation of systematics is much more CPU intensive, so a smaller number of throws was used. In order to determine an optimal number of throws, we required that the variance on the sample variance correspond to a 5% fractional error. For a Gaussian distributed sample, the variance of the sample variance behaves approximately as [98],

$$\text{Var}[s^2] = \frac{2\sigma^4}{N-1}, \quad (7.14)$$

where $\sigma^2 = \text{E}[s^2]$. This corresponds to a fractional error on the variance that goes as $f_{s^2} = \sqrt{2}/\sqrt{N}$, and requires $N = 800$ for $f_{s^2} = 0.05$. Thus, the final result is obtained with 800 throws for each systematic.

7.4.1 Propagation of Statistical Errors

Statistical errors affect both MC and data in the same way. Both are thrown assuming Poisson statistics. Unlike the other systematic uncertainties, the statistical throws are not performed

event-by-event but instead on a bin-by-bin basis. To handle MC statistics, the MC reconstructed and truth distributions are thrown to calculate a new unfolding matrix, efficiency, and background. To handle data statistics the data reconstructed distribution is thrown. All the bins are assumed to be independent of each other. The water-in and water-out samples are assumed to be statistically uncorrelated.

As discussed in Section 7.3, the background in the case of background subtraction is independent of the measured data. This means that when propagating statistical errors, the fractional uncertainty is larger on the background-subtracted result. For example, when propagating statistical uncertainty on N_d , the background from MC, B , is held constant. Thus,

$$\text{Var}[N_d - B] = \text{Var}[N_d], \quad (7.15)$$

which means the fractional error on $N_d - B$ will be larger than σ_{N_d}/N_d . Using a purity correction N_d is simply rescaled, which preserves the fractional error.

7.4.2 Propagation of Detector Errors

Detector uncertainties are treated either by event-by-event reweighting or redoing the selection based on perturbed kinematic variables. In **highland2**, these are respectively called “weight systematics” and “variation systematics”. This analysis contains detector systematics in both categories. Both types of systematics can affect the MC by altering the unfolding matrix U_{ij} and the efficiency ϵ_i in Eq. (7.1), as well as the background which alters $N_j^{w,a}$.

In **highland2**, all systematics can be thrown simultaneously and thus allow detector systematics to be correlated amongst themselves. The $P\bar{O}D$ mass systematic is a special case that affects not only the selection tree, but the truth tree and the overall $N_{targets}$. In this way, the mass uncertainty is similar to the flux systematic and is handled through a stand alone package in **xsTool**. All other systematics only affect the reconstruction. A detailed discussion of individual systematics is given in Section 6.3.

The average of throws from the **highland2** systematic tree is used as a reference from which

the covariance associated with detector uncertainties is evaluated. This reference may be different from the default `highland2` nominal distribution. Additionally, event-by-event weights from detector systematics are evaluated without theory corrections (e.g. flux tuning and NIWG tuning) applied. However, these theory corrections are included in the final cross section result. The fractional error as evaluated on the detector systematics is preserved and transferred onto the final, tuned nominal cross section result.

7.4.3 Propagation of Theory Errors




Figure 7.12: Comparison of covariance built from 5000 throws in `xsTool` of the theory parameters and the original covariance matrix used as the input. The xy-axes correspond to the index of the parameters listed in Appendix C, which are grouped into flux, FSI, and cross section categories. The two covariances agree with a $\chi^2/N_{DOF} = 0.88/2116 = 0.00042$.

The final three groups of errors consist of the flux, cross section, and FSI systematics. These rely on `T2KReWeight`, a framework that presents users with a set of underlying parameters, or dials, that can be tweaked to some value different from the nominal. A reweighting engine is then

run to generate a weight for each event. Correlations between parameters are taken into account with a covariance matrix across all parameters. The `xsTool` framework performs throws using this covariance matrix across the three systematic groups, flux, cross section, and FSI. For each group, the parameters within that group are simultaneously varied while the parameters outside that group are kept at their nominal values.

To ensure that the throwing procedure is implemented properly, the `xsTool` provides a validation routine to check that the covariance matrix produced from its own throws matches the input covariance across the parameters (Fig. 7.12). A chi-squared test is also run to ensure the two are quantitatively equivalent. In addition, the central values of all parameters are ensured to match up exactly with the input nominal values.

7.4.4 Sidebands and Error Propagation

The importance of sidebands and how they can be used to constrain backgrounds through a normalization rescaling was discussed in Section 7.3.2. However, their use can complicate the error propagation treatment.

As an example, let's take the flux uncertainty. This uncertainty affects both the numerator and denominator in the cross section measurement, which can be simplified as,

$$\sigma_B \propto \frac{N_d - B}{F} \quad (7.16)$$

$$\sigma_p \propto \frac{pN_d}{F} \quad (7.17)$$

for background subtraction, σ_B , and purity correction, σ_p . As the flux uncertainty affects F , B and p the numerator and denominator are coupled throughout the flux error propagation. The total flux uncertainty as calculated using the `xsTool` throwing mechanism is approximately 8.5%. A test of the overall uncertainties on the total **cross section** due to the flux error is given in Table 7.2 for different combinations of sidebands/background treatment methods. Clearly, there are some significant differences.

The cases without sidebands are quite simple to understand. For purity correction, p remains

Sidebands	Subtraction [%]	Purity [%]
No	11.22	8.33
Yes	9.53	7.21

Table 7.2: Fractional uncertainties on the total cross section due to the flux are given for different combinations of the background treatment method and sideband usage. These numbers were obtained using toy-MC.

relatively constant since the signal and background change in a correlated manner when the flux is varied. For background subtraction, increasing F will also increase B . Thus, without sidebands the numerator and denominator are uncorrelated (anticorrelated) in the case of purity correction (background subtraction). We can see this behavior in Fig. 7.13, which shows the throw-by-throw distributions of the total background B and the purity p vs integrated flux F . These were created over many Gaussian throws of the flux from the covariance matrix shown in Fig. 7.12. This means that the fractional error on the **cross section** increases for background subtraction but is preserved for purity correction.




Figure 7.13: Distributions of the total background, B , and purity, p , vs integrated flux, F , over many throws of the flux. B is correlated with F while p is uncorrelated.

With sidebands, the throw proceeds as normal before the sideband constraint is applied. For

the flux, the MC and integrated flux are again both reweighted, then the sideband constraint is imposed. In this way, each throw corresponds to a perturbation of the flux input parameters and sidebands are applied on altered MC in the same way as on the nominal MC. This ensures that all throws are treated on equal footing, each corresponding to a result obtained from different input parameters. Thus, for each throw, we obtain a constrained background and purity.




Figure 7.14: Distributions of the total sideband constrained background, B' , and purity, p' , vs integrated flux, F , over many throws of the flux. B' is less correlated with F than B while p' is more correlated than p .

In terms of the flux systematic, this has the effect of reducing the uncertainty on the cross section. In the case of background subtraction, B' (Eq. (7.8)) remains fixed as the flux varies since B and S_m change in a correlated manner. Now, the numerator becomes uncorrelated with the denominator in Eq. (7.16). In the case of purity correction, p' becomes correlated with F . This is because as B' remains fixed, N'_m still changes with the flux due to its signal component. For a larger F , N'_m increases even with sidebands, which forces $p' = 1 - B'/N'_m$ to increase as well. Both of these results are shown in Fig. 7.14. A summary of the approximate behavior of propagated flux uncertainty is given in Table 7.3.

The physical interpretation for why this occurs lies in how sidebands act as an additional

Sidebands	Subtraction	Purity
No	increases	preserved
Yes	preserved	decreases

Table 7.3: Summary of the behavior of the propagated flux error.

modification to the event rate **after** the parameter is thrown. In a sense, sidebands counteract the change to the background event rate due to the flux throw by constraining all the other components that go into the event rate (interaction cross section, detector reconstruction, etc.). This is especially apparent with the flux (and mass) systematics since for those systematics both the numerator and denominator in the cross section formula is allowed to change, and so extra correlations can appear.

In order to preserve the propagated uncertainty on the cross section measurement a conservative approach is used. When propagating the flux and mass systematics, the purity correction factor is kept fixed. Other sources of uncertainty did not exhibit much dependence on the enabling or disabling of sidebands. Thus, they are propagated normally, with the sidebands allowed to alter the background for each throw.

7.4.5 Subtraction and Error Propagation

In the following discussion, the term “subtraction” is used to refer to a subtraction of the water-in and water-out samples after unfolding and correcting the measured distributions. This should not be confused with the background treatment method of background subtraction, discussed in Section 7.3, in which the backgrounds are subtracted from the selection prior to unfolding.

The PØD was designed to be run with the water targets either filled or drained. These two detector configurations should be identical with the exception of the water, thus allowing for a measurement on water to be performed by subtraction. Essentially, performing a subtraction across the true $CC0\pi$ distributions should give the correct distribution on water. This is seen to be true within statistical uncertainties in Fig. 7.15.

Practically speaking, the detector acceptance will be different between water-in and water-out.




Figure 7.15: Two independent GENIE samples were used, one for the true oxygen measurement and the other for the subtracted cross section. Normalized residuals are drawn on the lower canvas and exhibit no outliers. The results agree within statistical uncertainties, thus validating the subtraction methodology in the ideal case.

This means that we cannot directly subtract the raw selected distributions. The subtraction is taken after the unfolding and efficiency correction as illustrated in Eq. (7.1). The idea is that this serves as an estimate for the underlying true distribution.

The PØD water configuration up to run 4 is shown in Table 5.1. Assuming the water-in and water-out samples are statistically independent, the variance on their difference is the sum of their individual variances. Thus, the statistical uncertainties on the subtracted result will contribute a larger fraction than the standard $1/\sqrt{N}$. An analysis of the statistical uncertainties shows that

$$\frac{\sigma_{N_O}}{N_O} \approx \frac{1}{\sqrt{N_O}} \sqrt{1 + (1 + R) \frac{N_{nonO}}{N_O}}, \quad (7.18)$$

where N_{nonO} indicates the number of interactions that occurred on non-oxygen targets. This is an undesirable aspect of the subtraction technique, and causes the error due to data statistics to be a large source of uncertainty. Appendix A discusses a possible approach to minimize the statistical uncertainty with additional data-taking.

On the other hand, correlated errors will lead to approximately unchanged fractional uncertainty after subtraction. To see this is indeed the case, assume absolute correlation between water-in and water-out samples. Then the fractional uncertainty on the subtraction, f_{w-a} , is,

$$f_{w-a} = \frac{\sigma_w - \sigma_a}{N_w - N_a} \quad (7.19)$$

$$= \frac{\frac{\sigma_w}{N_w N_a} - \frac{\sigma_a}{N_w N_a}}{\frac{N_w - N_a}{N_w N_a}} \quad (7.20)$$

$$= \frac{\frac{f_w}{N_a} - \frac{f_a}{N_w}}{\frac{1}{N_a} - \frac{1}{N_w}} \quad (7.21)$$

$$= f_w = f_a, \quad (7.22)$$

where the last equality holds if $f_w = f_a$. Flux, cross section, and FSI systematics are prime examples of correlated errors. Uncertainties due to the detector simulation and reconstruction are a bit trickier to interpret, but typically they are correlated to the extent that they affect each sample simultaneously. For example, an uncertainty on the water mass only affects the water-in sample, and so will not subtract out, but other masses may. The description and evaluation of these uncertainties is discussed in Section 7.4.

The `xsTool` framework was used to perform the unfolding, subtraction, and error propagation for all error sources. To test the subtraction of correlated and uncorrelated sources, several tests were performed. The first test used an identical sample for water-in and water-out. When subtracted, all errors besides statistical uncertainties should cancel exactly as shown in Fig. 7.16. The second tests performed an air-air subtraction with two statistically independent water-out samples as shown in Fig. 7.17. Both these tests validate the error propagation across the subtraction.

7.4.6 Examples of Single Parameter Variations

To further illustrate how the error propagation works in the `xsTool` and to give some intuition with how the results can change for specific parameters, a check was performed that compared components feeding into Eq. (7.1) over variations of a single parameter. We chose to look at a few individual cross section parameters as they are typically more uncorrelated than other theory




Figure 7.16: Subtraction of two identical samples. The subtracted result should be exactly zero with correlated errors such as the detector (orange), flux (green), cross section (yellow), and FSI (red) errors canceling out completely.




Figure 7.17: Subtraction of two independent water-out samples. The central value tends towards zero with some slight fluctuations, but the dominant error source is again statistical (turquoise and cyan).

systematics. These parameters affect different interaction channels and tweaking them individually should have a different effect on the components of Eq. (7.1). Hopefully, these examples can help give some intuitive understanding of how the errors are calculated on the final cross section result.

Figures 7.18 to 7.23 show the effect of varying a single parameter from $[-2, 2]\sigma$, where σ was calculated from the square root of the diagonal element corresponding to the specified parameter in the covariance matrix shown in Fig. 7.12. Since these are throws on the cross section parameters, only the MC is affected in the error propagation. The distributions affect the unfolding matrix U_{ij} and are explained as follows.

The “reconstructed” distribution corresponds to the overall MC selection, which includes both signal and backgrounds. The “signal” distribution is a subset of the selection that is flagged as true $CC0\pi$. The “purity” and “efficiency” are defined in Section 5.2.1 and correspond to how those values can change as the parameter is varied. As sidebands were included in the generation of these throws, the “purity” shown here corresponds to p' discussed in Section 7.4.4. The “result” distribution is the actual cross section result after putting everything (unfolding, efficiency correction, and normalization) together for a single water-in sample. The “subtracted result” distribution then is the same thing but including the subtraction of the water-out distribution. These distributions were generated using a small subsample of approximately 1×10^{20} POT of the NEUT MC sample. Most of the large relative deviations, seen primarily in the subtracted result, occur in low statistics bins.

Two points are immediately evident. The first is that the different cross section parameters have different effects on the distributions. This is expected as the parameters affect different interaction channels that are present to different degrees in the selection. The reaction breakdown for our $CC0\pi$ selection is given in Table 5.4 and the reader is encouraged to refer to that when looking at the plots of the individual parameter throws. We note that the MaQE parameter affects both the signal and reconstructed distribution. The CA5RES parameter seems to affect the reconstructed distribution a bit more than the signal, probably due to the fact that there are more CCRES interactions in the selection than the $CC0\pi$ component of the selection. The CCCoh parameter has no effect on the signal or the efficiency but has a slight effect on the reconstructed.

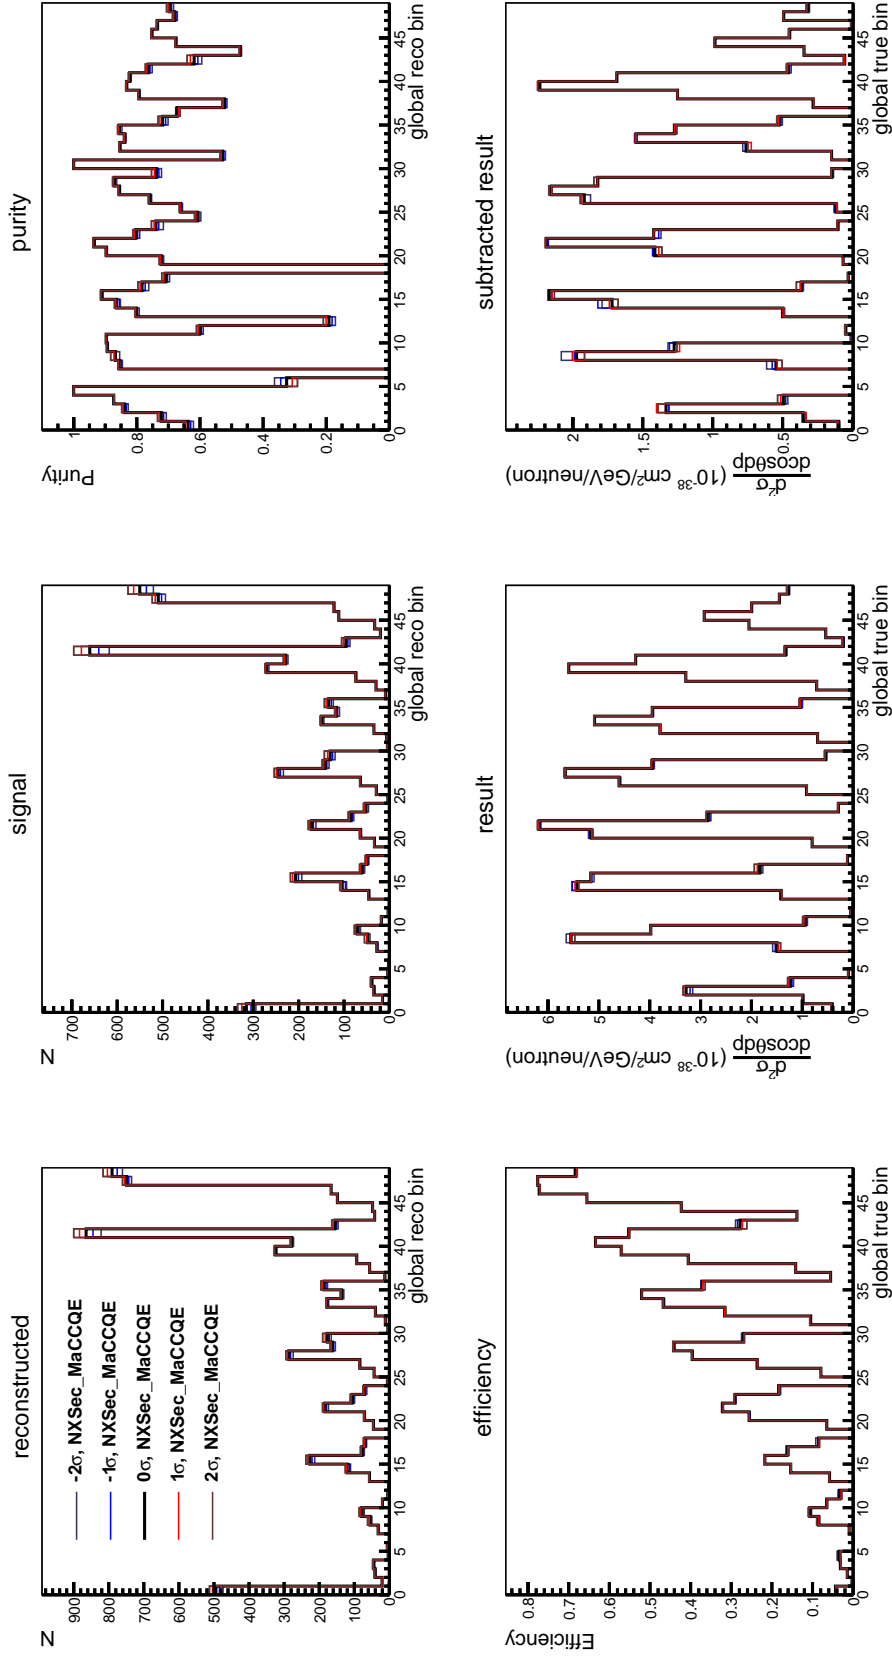


Figure 7.18: Individual throws on the NXSec_MaCCQE parameter for various distributions used in the cross section calculation. An explanation of each distribution is given in the text.

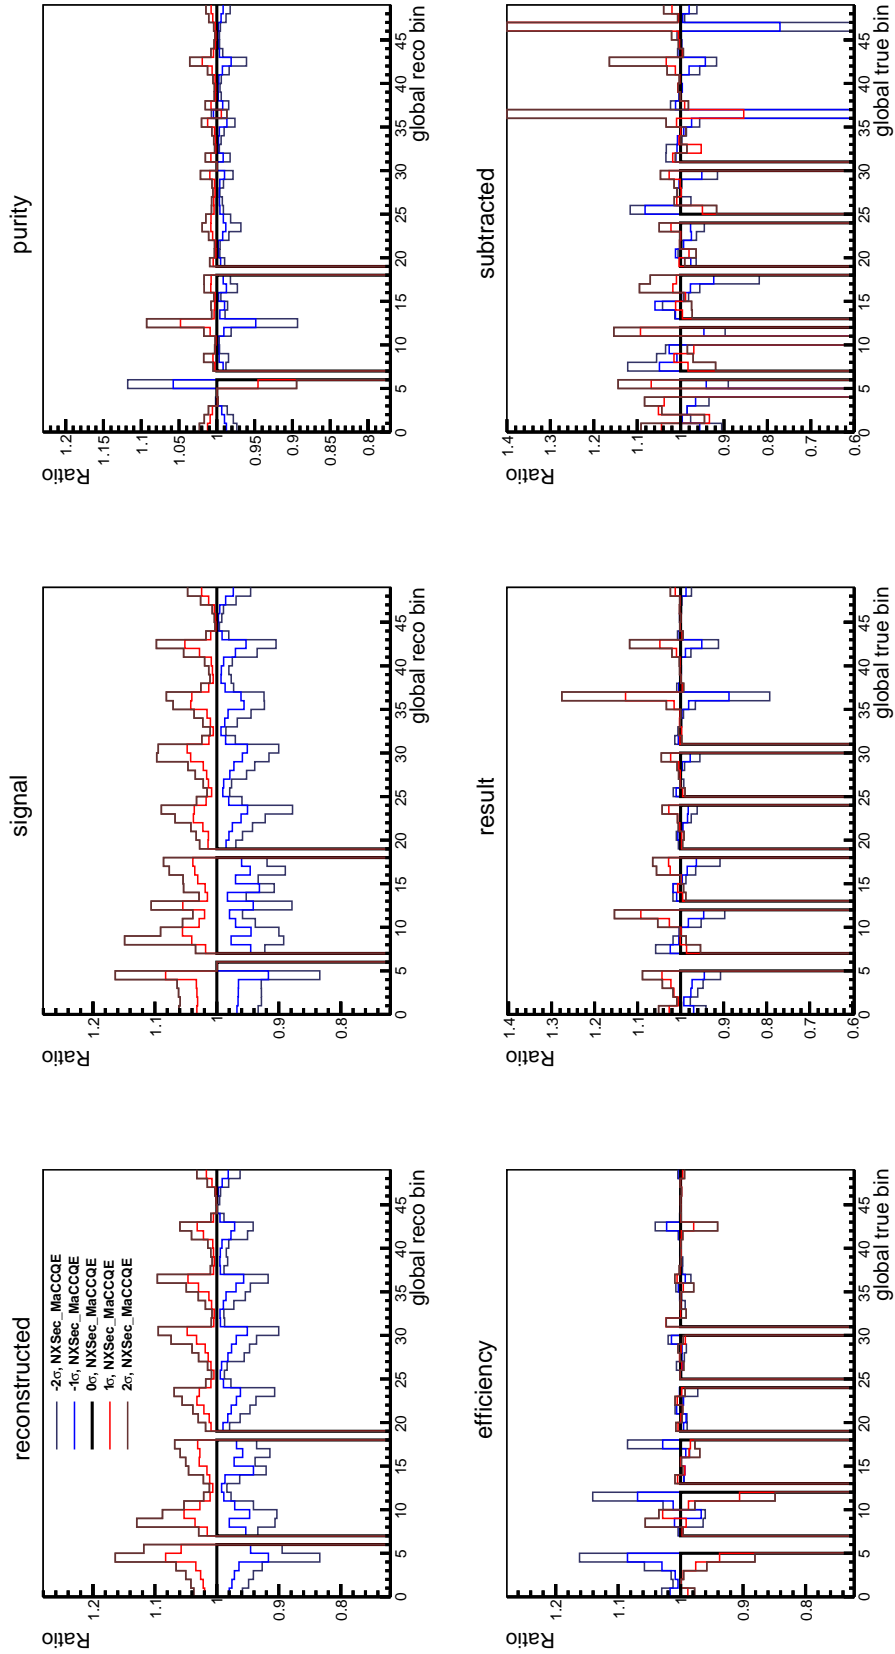


Figure 7.19: Ratios to the nominal of throws on NXSec_MaCCQE.

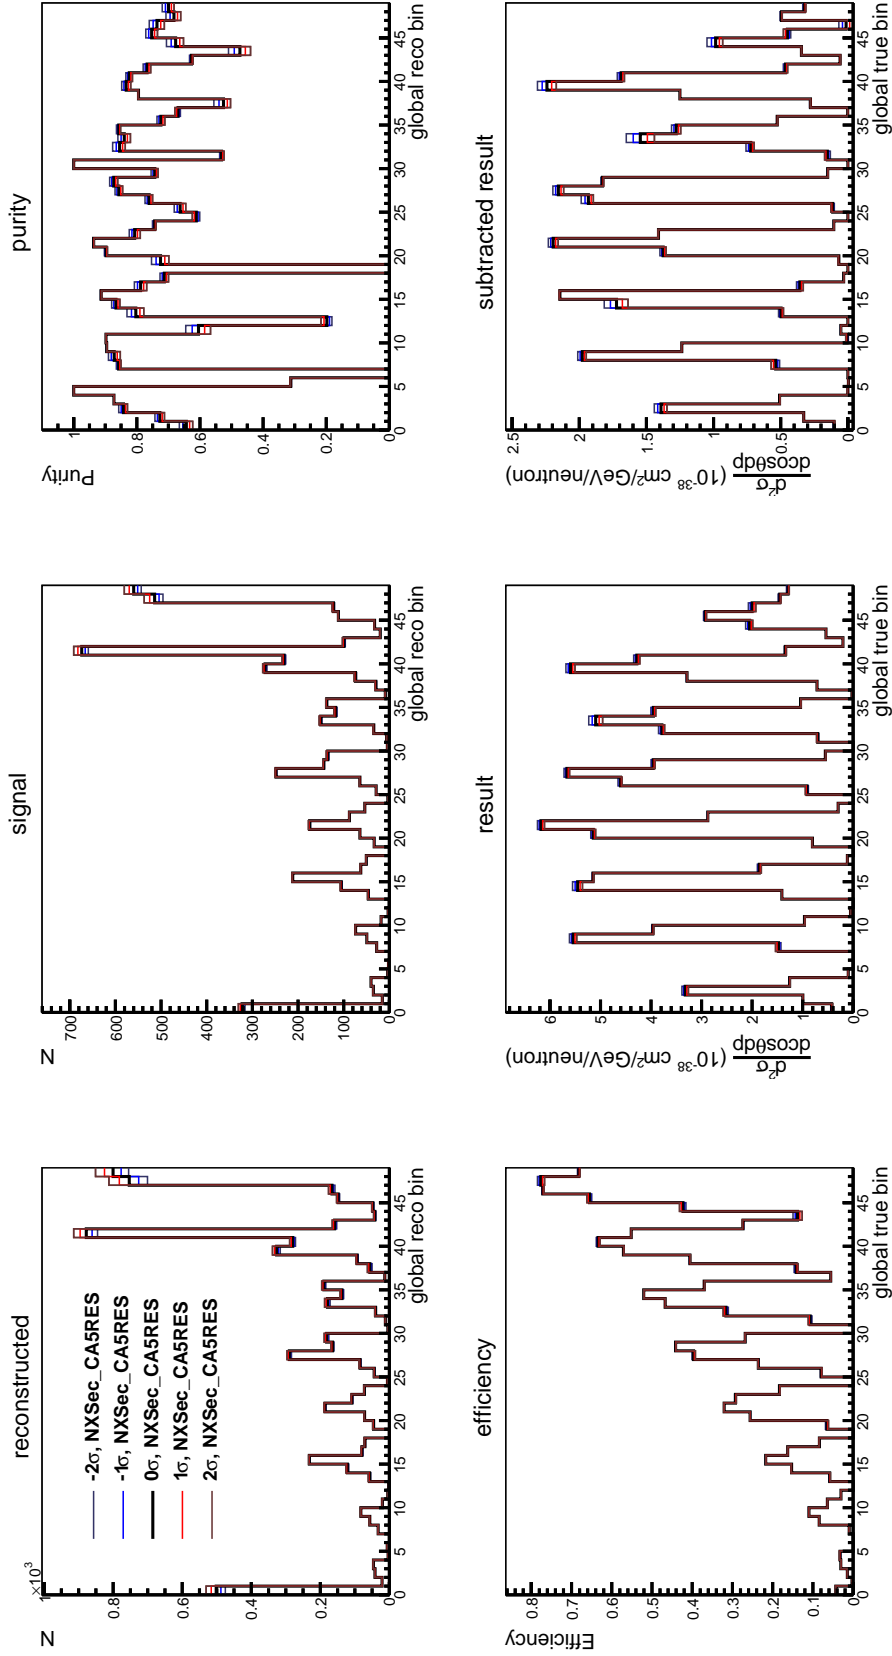


Figure 7.20: Individual throws on the NXSec_CA5RES parameter for various distributions used in the cross section calculation. An explanation of each distribution is given in the text.

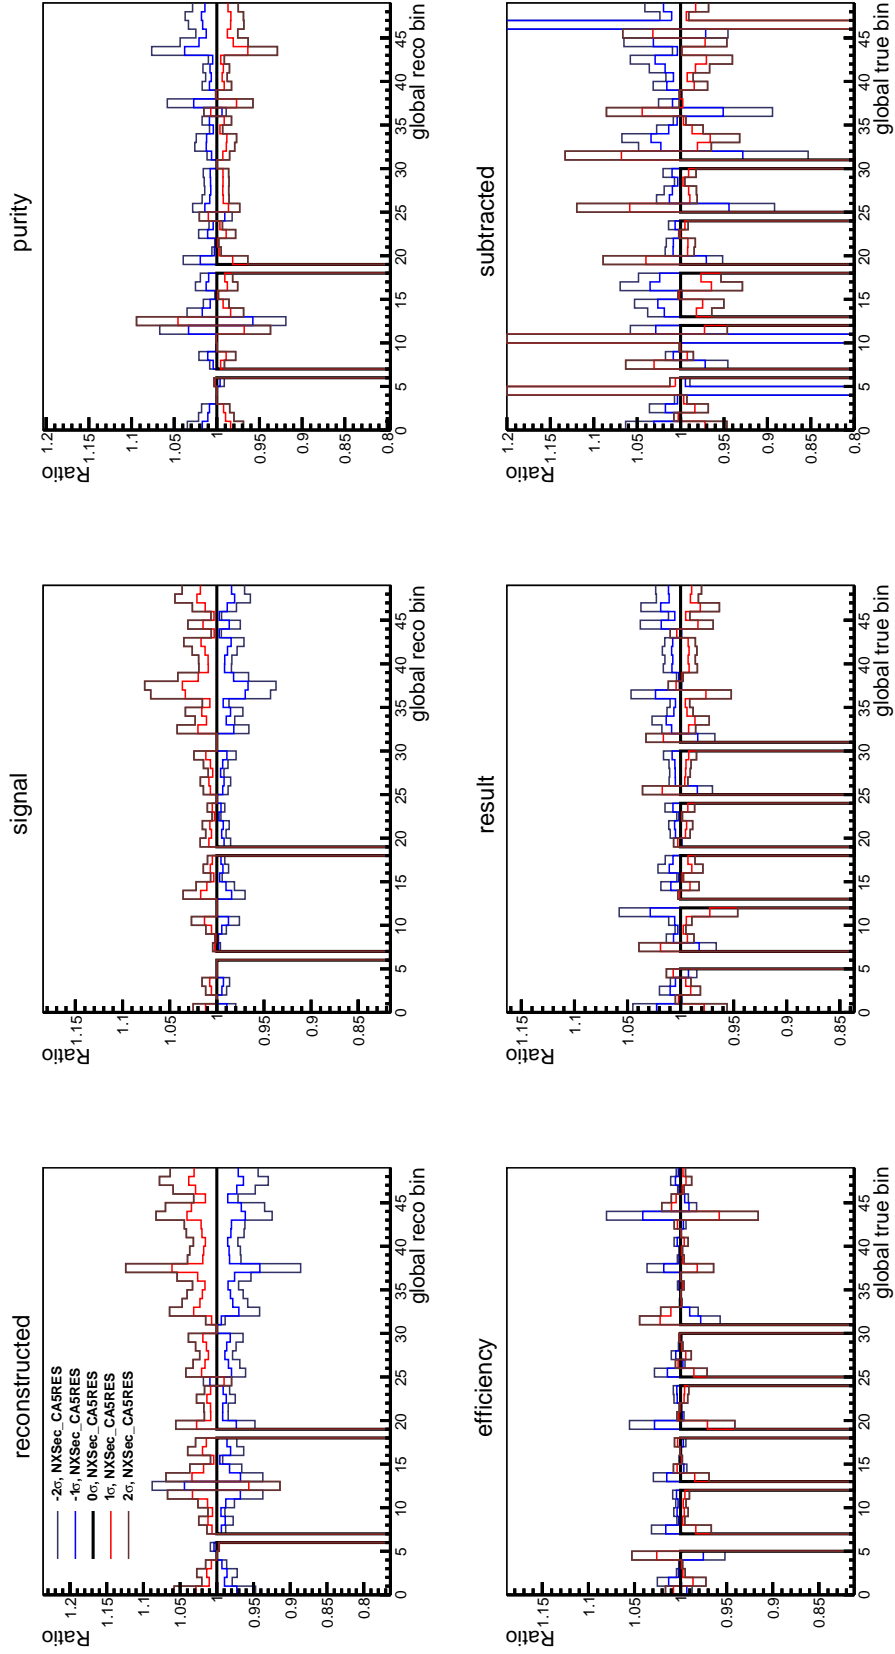


Figure 7.21: Ratios to the nominal of throws on NXSec_CA5RES.

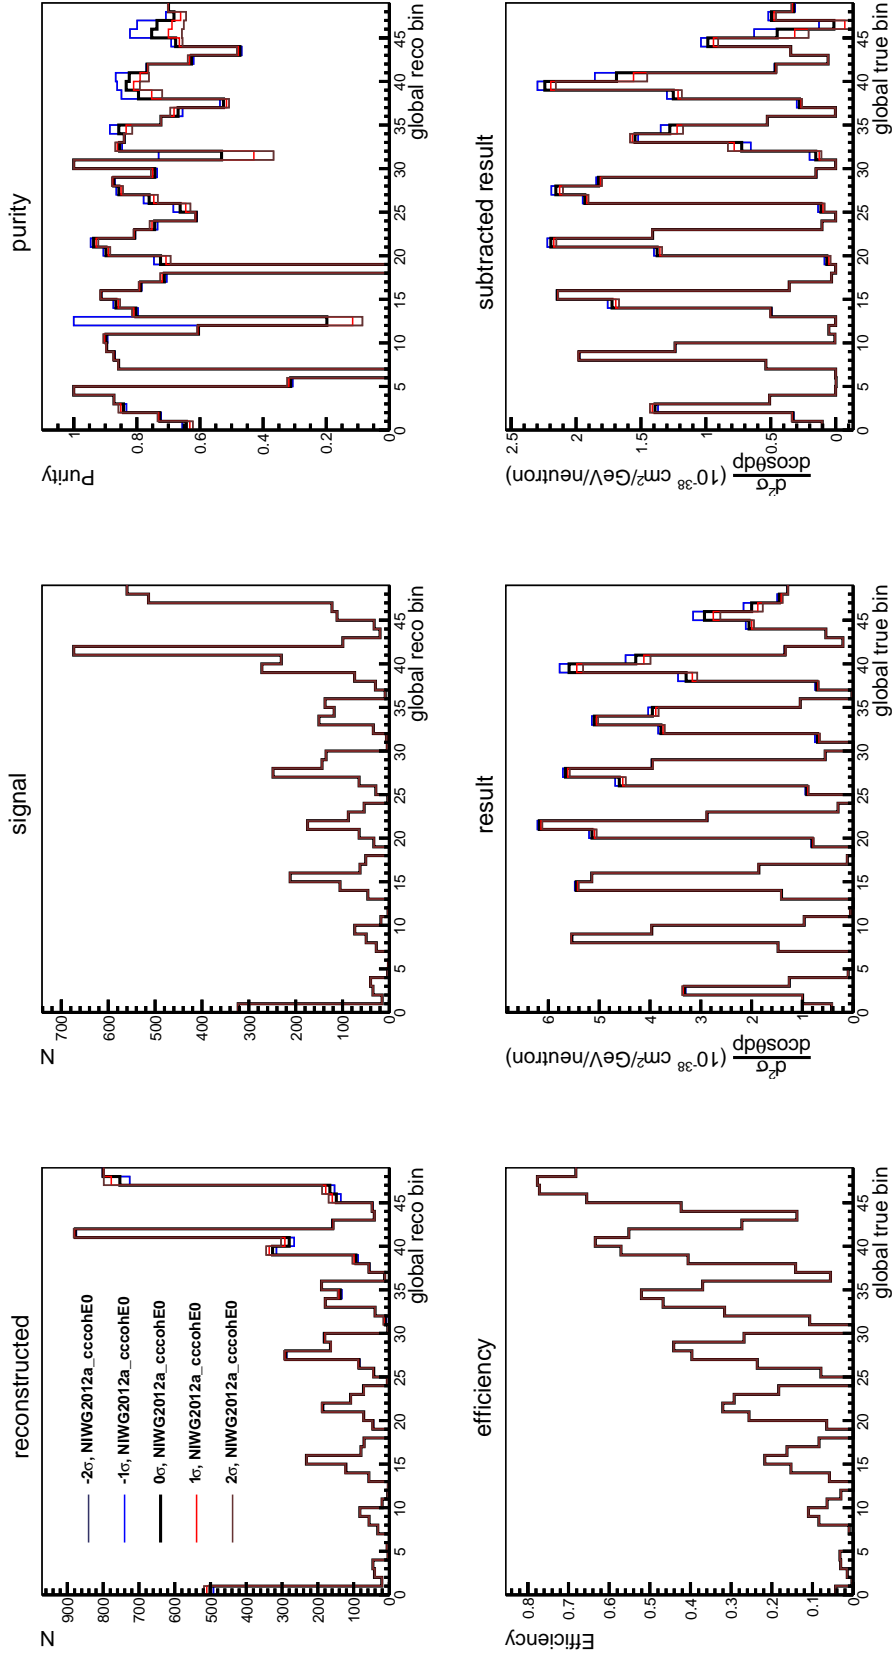


Figure 7.22: Individual throws on the NIWG_cccohE0 parameter for various distributions used in the cross section calculation. An explanation of each distribution is given in the text.

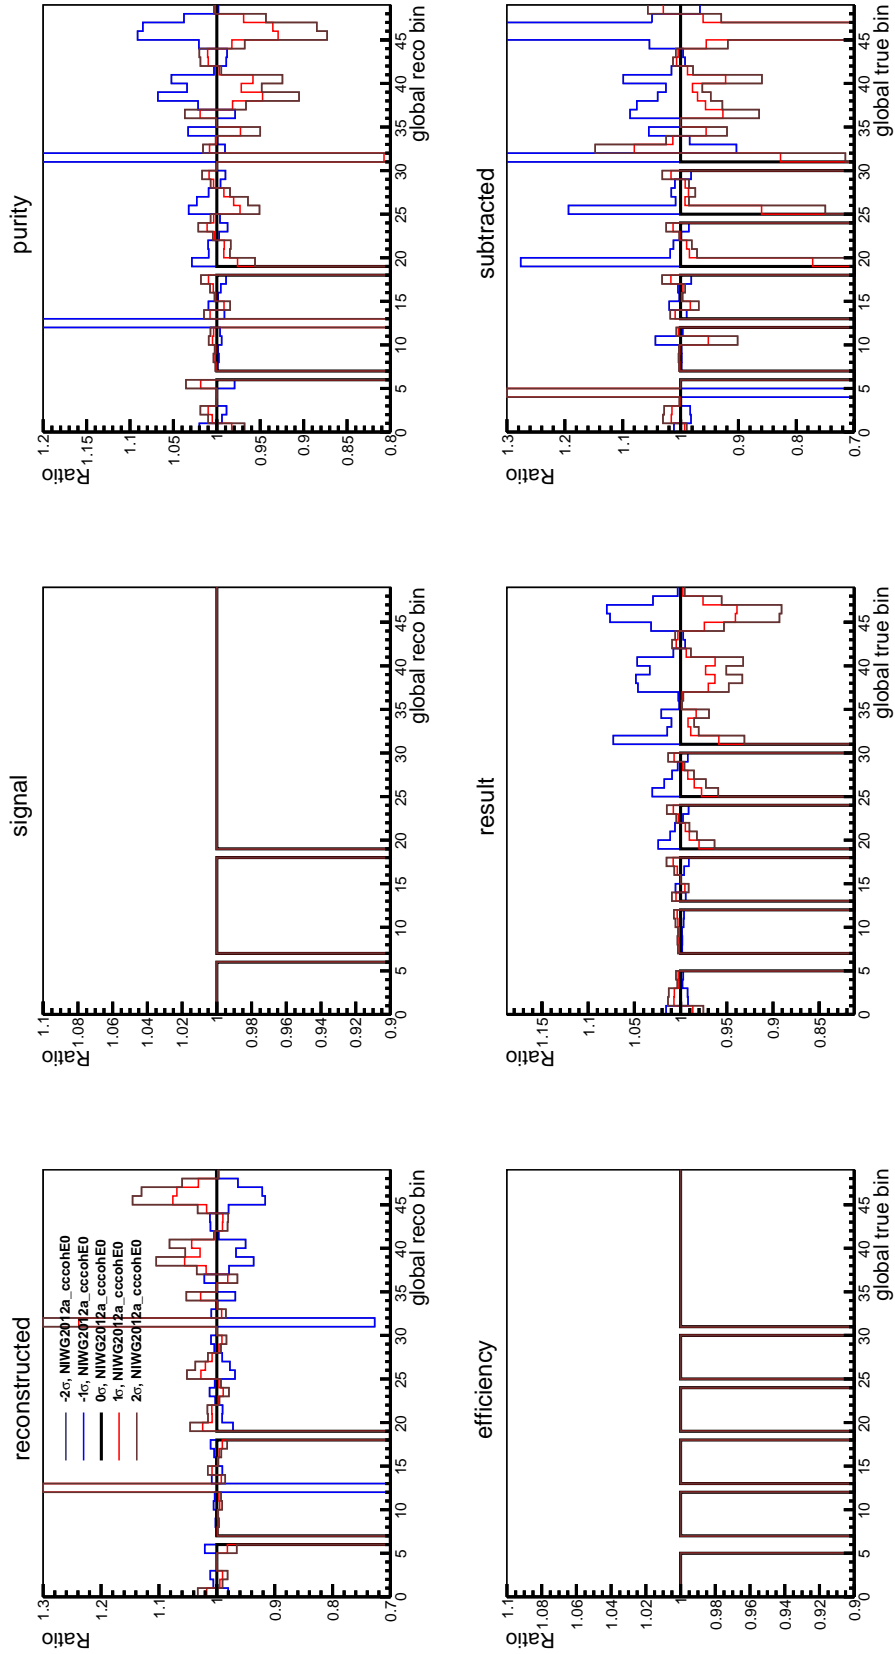


Figure 7.23: Ratios to the nominal of throws on NIWG_cccohe0.

The second point is that correlations between raw distributions can affect compound distributions such as the purity or the efficiency. For example, throws of the MaQE parameter have a similar effect on the signal, reconstructed, and truth (not shown) distributions, leading to a cancellation in the overall effect on the purity and efficiency. This goes back to the discussion given in Section 7.4.4 regarding correlations between the numerator and denominator of Eq. (7.2) that can affect the final result. These correlations need to be accounted for when numerically propagating the uncertainties.

7.5 Model Dependence and Bias

In Eq. (7.1), U_{ij} and ϵ_i are constructed from MC truth. Further, implicit in that equation is the removal of backgrounds, which also depends on the MC truth. It is therefore important to address the question of model dependence, or how much the MC may bias the result.

An intuitive way to think about model biases is to imagine a situation in which the data selection perfectly matched the MC selection. This would be identical to using the MC to unfold itself and the result would exactly match the MC truth (Fig. 7.5), leading us to believe that the real life cross section is exactly reproduced by the MC. However, it could be the case that it is the cross section in combination with some other effect that caused the reconstructed data and MC to agree perfectly, and that in reality the cross section is quite different from the MC truth. Nevertheless, in such a hypothetical scenario, the technique described here can only return the MC truth as the result. This naïve example illustrates that model dependence is inescapable unless a full description of all physics separate from the processes of interest (e.g. detector reconstruction, flux) is modeled perfectly by the MC. In practice, perfect descriptions are impossible and systematic uncertainties attempt to encapsulate our imperfect knowledge. While complete model independence may be impossible, results should be as unbiased as possible within uncertainties. In this section, several tests of the MC bias on unfolded and subtracted results are discussed.

7.5.1 Dependence on Background Treatment

Before discussing the bias tests, it is important to address the sometimes controversial topic of background treatment and model dependence. One issue that is often raised against purity correction is that p depends on the MC signal and background, whereas with a background subtraction B depends only on the MC background. Therefore, it might seem that that purity correction may lead to a more MC-biased result.

However, there are strong assumptions made in the case of background subtraction that can introduce model dependence as well. First and foremost is the assumption that the **absolute** MC background is representative of the data background. Background subtraction places more emphasis on the MC background, which for this analysis is probably generally less understood than the $CC0\pi$ signal. Another feature of background subtraction is that B is allowed to vary independently from the data. This can lead to several issues, including larger uncertainties as discussed in Section 7.4.1 and negative bins. It seems physically and statistically incorrect to allow the estimate of the background in the measured data to be larger than the measured data itself.

A purity correction approach can be thought of as a probabilistic treatment of the background. Its assumption requires that one trusts the ratio of signal to background, but its benefit is that the estimate of the background in the measured data, B , is dependent on N_d . This is no more different than asking for the number of tails that occurred in N_d flips of an unfair coin for which the probability of heads can be estimated. In comparison to background subtraction, the chance to get a completely unphysical results is much smaller using purity correction.

Additionally, a treatment scheme where the background is subtracted prior to unfolding does not fall inline with the mindset of Bayesian unfolding. In fact, such an approach would no longer be a Bayesian treatment of the background and conflicts with the unfolding methodology. On the other hand, as discussed in Appendix B, purity correction is simply allowing for an extra cause bin associated with all backgrounds. This type of treatment fits within the mentality of Bayesian unfolding. If the signal unfolding can be thought of as a probabilistic event migration between

kinematic bins, then purity correction can be thought of as a probabilistic event migration between the selection and background. Of course, the MC model is required to behave reasonably, but there is no way around this with either treatment method.




Figure 7.24: Comparison of results from a toy-MC study using two background treatment methods. The full NEUT MC statistics was used and the GENIE toy-MC had its POT scaled to match data POT. Error bars are statistical.

Practically speaking, there is little difference between the central values of the two results when using a NEUT MC sample to unfold GENIE toy-MC. A direct comparison is shown in Fig. 7.24. The central values agree well within errors. However, as mentioned several times already, the errors obtained using background subtraction tends to be larger than those obtained using purity correction.

7.5.2 Tests of Model Dependence

NEUT and GENIE are two different MC generators useful for tests of model dependence. The CC0 π water cross sections from truth (via subtraction) as reported by these two generators are shown in Fig. 7.25 and indicate significant differences. The selection efficiency in the denominator of

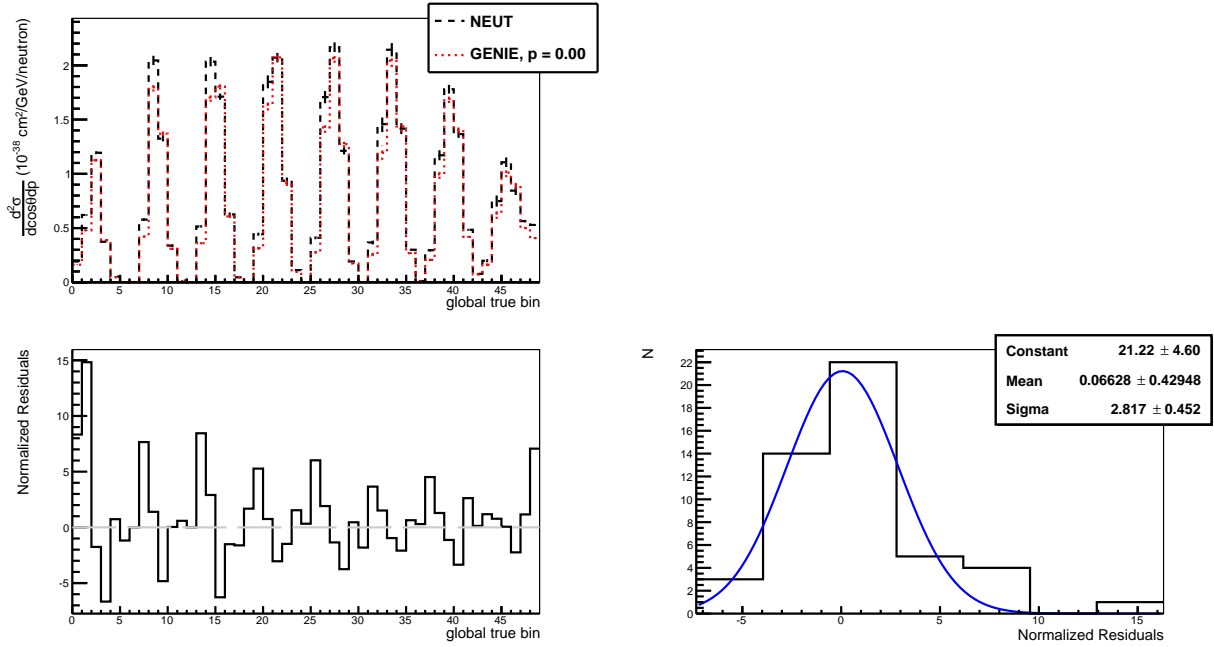


Figure 7.25: Comparison of **true** CC0 π cross section on water for NEUT and GENIE (top-left). Normalized residuals (bottom-left) and their distribution (bottom-right) are included and indicate significant differences.

Eq. (7.1) is affected through a combination of the cuts and interaction models. A comparison of the efficiencies calculated using NEUT and GENIE is shown in Fig. 7.26. Specific selection cuts, such as the requirement for a track to go from the PØD into the TPC and the single PØD reconstructed object cut can cause the efficiency to become model dependent.

As used in this analysis, Bayesian unfolding is expected to mainly unsmear detector effects. Unfolding is performed across directly reconstructable kinematics and has no dependence on target, neutrino energy, Q^2 , or any other pre-FSI parameter. In building the unsmearing matrix, the cross section model primarily affects the prior $\text{Pr}(C_i)$ in Eq. (7.3) and the background as discussed in Section 7.3. A test of the model dependence was performed by comparing the unfolded result obtained with identical samples for MC and toy-MC against the unfolded result using the same samples but with the MC CCQE cross section on oxygen reweighted by a factor of two. This comparison is shown in Fig. 7.27. Additionally, Fig. 7.28 shows a comparison of the unfolded result from the reweighted MC, plotted against the true MC cross sections (nominal and reweighted by a

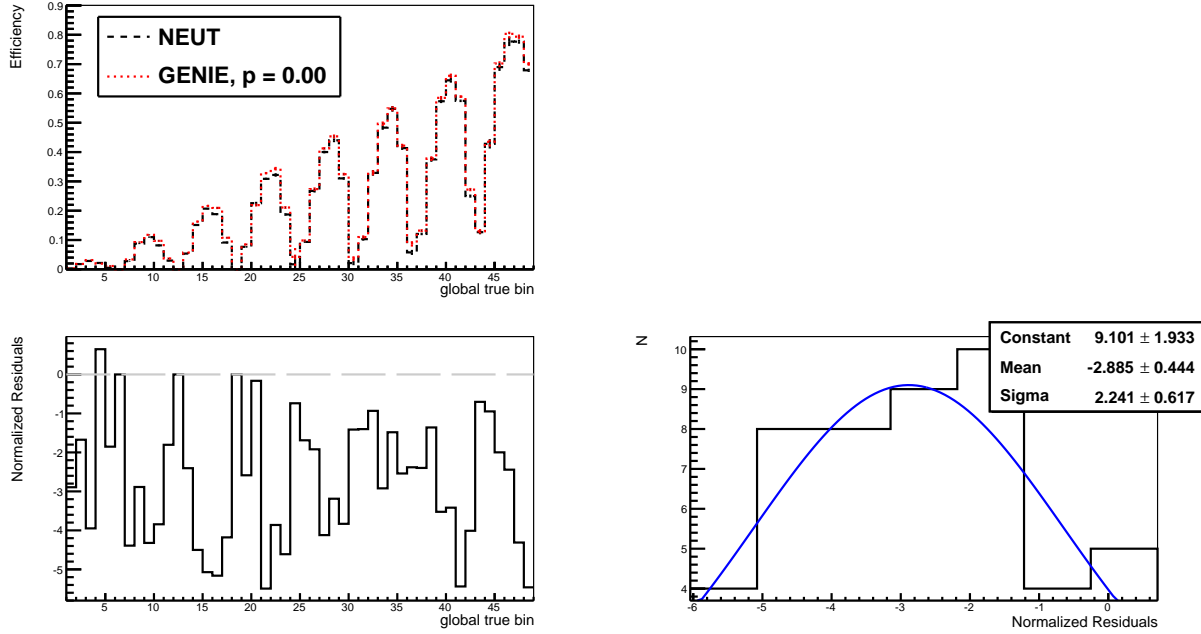


Figure 7.26: Comparison of selection efficiencies for NEUT and GENIE (top-left). Normalized residuals (bottom-left) and their distribution (bottom-right) are included and indicate differences between the two generators.

factor of two). As expected, the reweighted oxygen cross section is exactly twice the nominal cross section. The unfolded result is mostly unaffected by the reweighting in the MC.

These plots show that the Bayesian unfolding technique is relatively insensitive to incorrect cross section normalizations in the MC. Such differences in the MC normalization affect both the truth and reconstructed samples, leaving the efficiency and detector response mostly unchanged. The model dependence is certainly nonzero, as evidenced by the slight bias in the result towards a higher measured cross section. Nevertheless, these tests show that a 100 % (artificial) bias in the MC seems to still give rather reasonable results.

These results are reinforced by studies using NEUT as MC and GENIE as toy-MC and vice versa. In both cases, the extracted cross section measurement has better agreement with the toy-MC truth than the MC truth. This is seen in Fig. 7.29 and is especially noticeable in bins where the two truth cross sections differ significantly. With the inclusion of model systematics, a chi-squared test produces a reasonable p-value with the toy-MC truth as seen in Fig. 7.30. Also included are




Figure 7.27: The black line shows the cross section obtained using identical samples for MC and toy-MC. The red line shows the cross section obtained with the same samples, but with MC CCQE interactions on oxygen reweighted by a factor of two. Although shifted slightly higher, the bias is still within the statistical uncertainties. The slight differences between truth and nominal are due to the truth being calculated from true interactions on oxygen instead of a subtraction.






Figure 7.28: Similar to Fig. 7.27 but with the artificially biased MC truth plotted (solid blue). The result in black was obtained using identical samples for MC and toy-MC, but with MC CCQE interactions on oxygen reweighted by a factor of two. This shows that the result is in better agreement with the toy-MC truth than the artificially biased MC.

plots of the normalized residuals as calculated via ROOT's chi-squared test routine for comparison of two weighted histograms. The residuals should be distributed according to a standard normal distribution. This is seen in Fig. 7.30c.



(a) NEUT MC and GENIE toy-MC




(b) GENIE MC and NEUT toy-MC


Figure 7.29: Comparison of extracted cross section against the MC prior and toy-MC truth. In both cases, the extracted cross section agrees better with the toy-MC truth (dashed line) than MC prior (solid line).

7.6 Alternative Methods


Bayesian unfolding is not the only unfolding procedure, and unfolding is not the only cross section extraction method. Many analyses rely on template likelihood fits which vary the true cross



(a) Result Comparison



(b) Normalized Residuals



(c) Distribution of Normalized Residuals

Figure 7.30: Unfolded result with cross section errors and statistical uncertainties propagated. This plot was made using NEUT to unfold GENIE toy-MC. In (a) model uncertainties can be seen to cover additional differences between the result and toy-MC truth. The normalized residuals are calculated with ROOT's chi-squared test routine for weighted histograms. Its distribution can be seen to be approximately standard normal.

section to fit the reconstructed data. In a previous measurement, Bayesian unfolding was performed in parallel, but independently, with a likelihood fit, and the two methods gave results that agreed within errors [95]. As for other unfolding procedures, the `xsTool` includes several different unfolding engines that use different routines. The different algorithms used for each will not be discussed here, but results from different engines have been validated and compared against each other using the framework's automated testing framework [80].

Chapter 8

Double Differential $\text{CC0}\pi$ Cross Section

In Chapter 7, we detailed the procedure used to extract the muon-neutrino double-differential $\text{CC0}\pi$ cross section on water from the selection described in Chapter 5. Namely, an unfolding method is first applied to the selection to remove detector effects and allow for the extraction of an estimate for the true interaction distributions. The PØD water-in and water-out datasets are unfolded separately so that a subtraction can be performed to obtain a cross section on water.

This analysis was conducted using a data-blind technique. All the steps towards a cross section measurement were studied with toy-MC data and approved by an internal committee before they were applied on real data. These results were obtained following the method described in Chapter 7 and applying the corrections and systematics listed in Chapter 6 onto the selection described in Chapter 5. Statistic and systematic errors are categorized by their source. Errors are separated in this manner since the various sources are treated independently prior to addition in quadrature.

8.1 Results from T2K

The result shown in Fig. 8.1 uses data from T2K Runs 2–4. It is reported as a double differential cross section in the outgoing muon kinematics, $p_\mu - \cos\theta_\mu$. The colorized error bars show the cumulative contributions from various sources, starting with the data statistics and moving up to the detector systematics. Since errors from the various sources are added in quadrature, the colors just show the contribution of the additional source to the total cross section uncertainty. The black data points shows the double differential result with all errors. Tables containing the full result and

errors can be found in Appendix D.

The fractional error contribution from each source of uncertainty is shown in Fig. 8.2, binned and plotted in the same scheme as Fig. 8.1. In most regions, the statistical error from the data is the single most dominant source of bin-by-bin uncertainty. As discussed in Appendix A, statistical uncertainties are minimized if there is more water-in POT than water-out POT, but unfortunately this is the opposite in beam data. Aside from the really low sensitivity bins, the fractional bin-by-bin errors lie on the order of 10–20 %.

Model	χ^2/N_{DOF}	p-value
NEUT (tuned)	49.2/47	0.388
GENIE	72.4/46	0.007

Table 8.1: Results from a χ^2 test across all bins in the double differential measurement.

A χ^2 test was performed against the tuned NEUT and default GENIE MC predictions. The result of the test across all bins is given in Table 8.1, while p-values for individual slices of $\cos\theta_\mu$ are shown in Fig. 8.1. Overall, it appears that the tuned NEUT prediction is favored, but certain sliced regions exhibit better agreement with GENIE. These simple goodness-of-fit tests complement the final results, but should not be taken as a rigorous assessment of model preference.

The total cross section, in the full phase space, was calculated by integrating over the double differential cross section to be

$$\sigma_{\nu_\mu \text{H}_2\text{O}}^{\text{CC}0\pi} = (1.15 \pm 0.08(\text{stat}) \pm 0.13(\text{syst})) \times 10^{-38} \text{ cm}^2/\text{n}. \quad (8.1)$$

The fractional uncertainty on the total cross section for each source is given in Table 8.2. The largest single contributor is the flux uncertainty, followed by the data statistics.

8.2 Concluding Remarks

This thesis presents a double differential measurement of the muon-neutrino CC0 π cross section on water using data collected by the near detector of T2K. This result is important as both the interaction channel and the target play large roles in neutrino oscillation experiments. Although

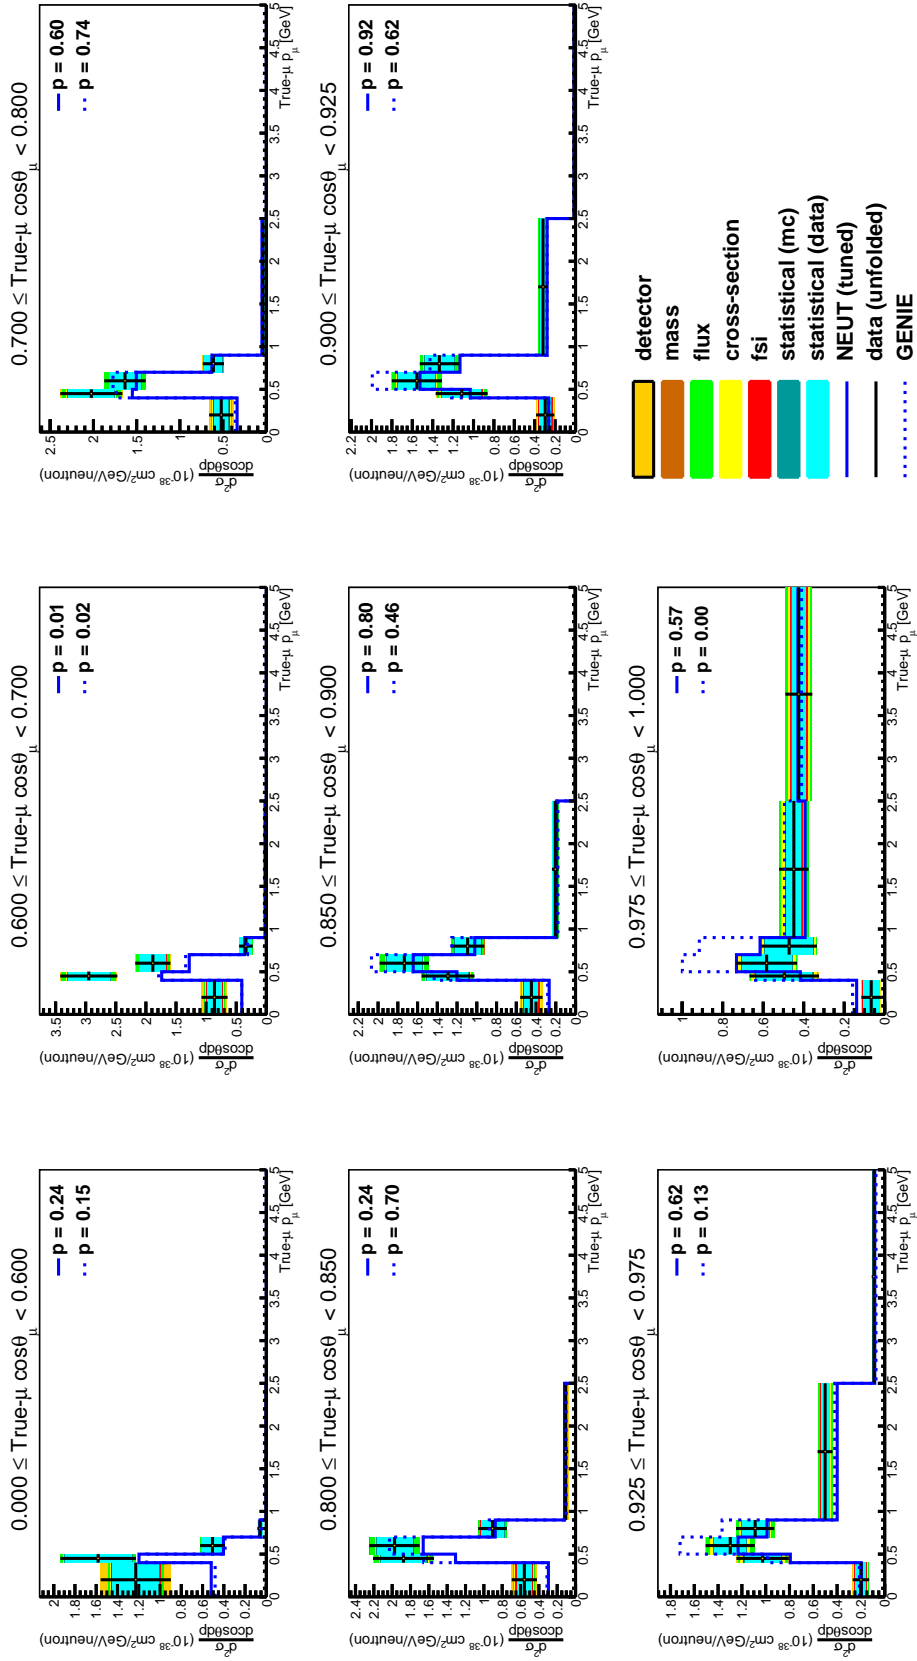


Figure 8.1: The double differential ν_μ CC0 π cross section on water. The p-values listed are calculated with respect to the two MC predictions for each slice in $\cos \theta_\mu$.

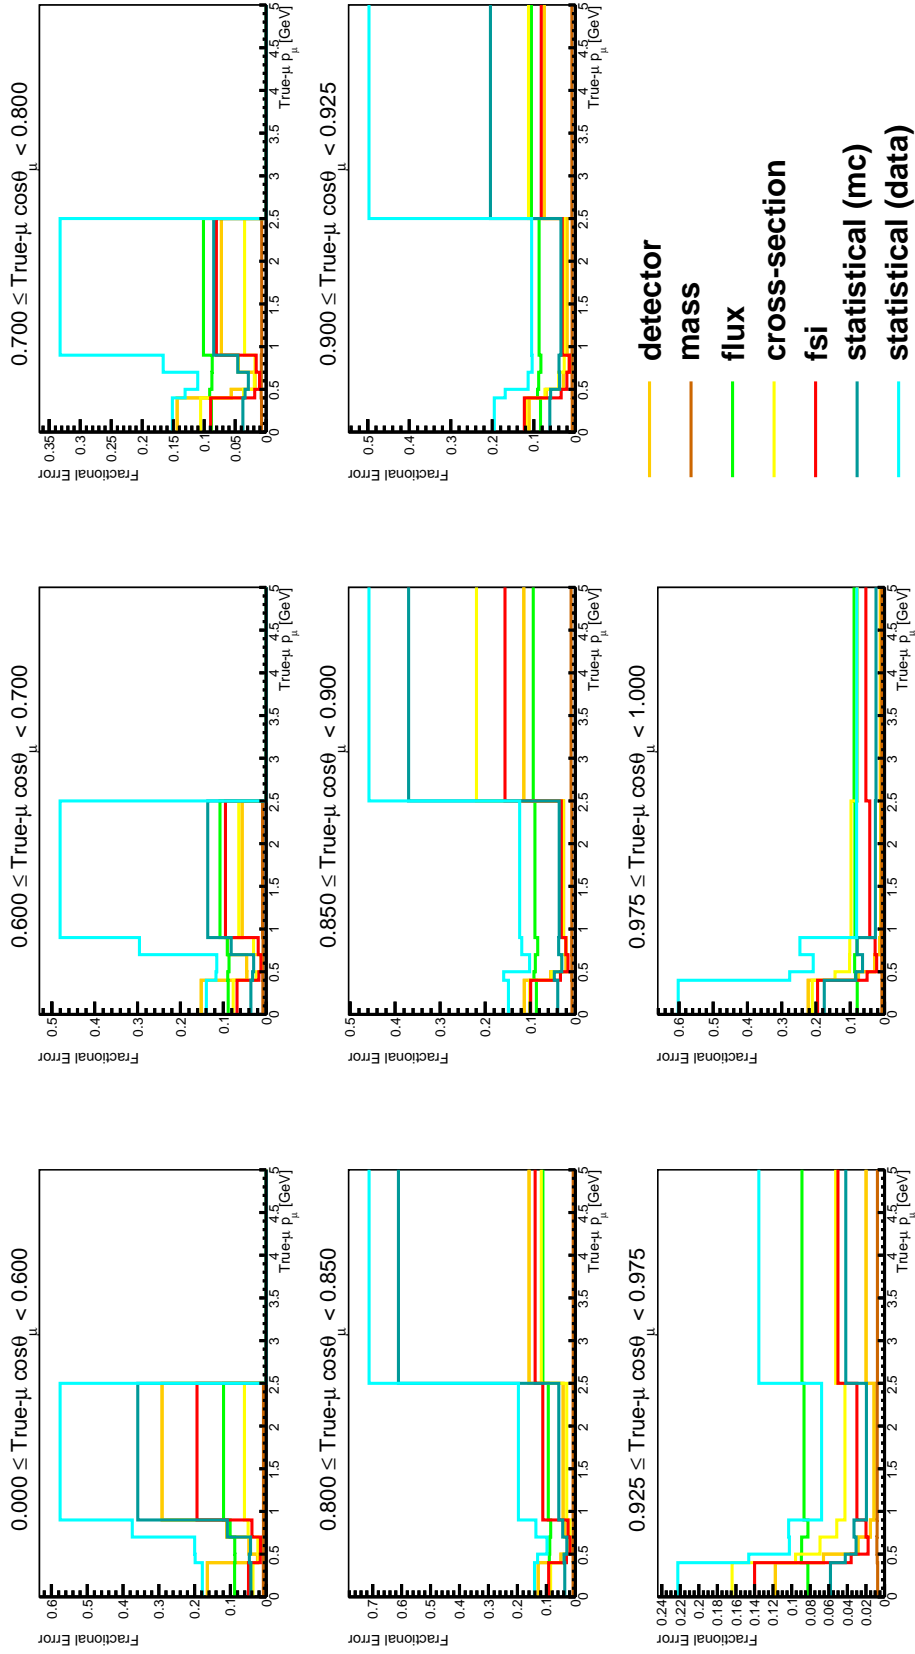


Figure 8.2: The fractional error on the double differential ν_μ CC0 π n cross section on water.

Source	Fractional Uncertainty [%]
Statistics (data)	6.6
Statistics (mc)	1.6
Detector (mass)	0.8
Detector (other)	5.3
FSI	3.2
Cross section	3.2
Flux	8.4

Table 8.2: Fractional uncertainties on the total cross section using T2K Runs 2–4. These sources are considered independent and the overall errors are obtained by adding them in quadrature.

the measurement is not independent of the T2K flux, it can be used in conjunction to constrain neutrino interaction models.

One question that may arise is why was the cross section not measured as a function of the neutrino energy. Such a measurement, while perhaps allowing for a more direct comparison with predictions, ends up largely dependent on the interaction model used in its extraction. Since the neutrino energy is not directly detectable, it must be reconstructed based on observable kinematics and an assumed interaction channel. These, as discussed in Section 1.4, are complicated by a couple of confounding factors. The first is a practical point that detectors imperfectly reconstructs observable particles, and sometimes misses them entirely. The second is a theoretical limitation due to intranuclear effects. Detector misreconstructions are independent of the interaction model and can be corrected, but intranuclear effects cannot be separated from the model. Thus, the neutrino energy cannot be disentangled from the model guiding its reconstruction and is less useful as a result independent of model assumptions.

On the other hand, a measurement in double differential lepton kinematics is based on observable quantities. The only additional step needed is to unsmeared from reconstructed to true kinematics as described in Section 7.2. Such a procedure ought to be less model dependent, and much discussion on this subject is provided in Section 7.5. However, there is a technical limitation as the smearing and unsmearing matrices are not inverses of each other. While Fig. 7.5 shows that using the same MC and toy-MC samples returns exact agreement with the MC truth, this is

a specific instance where the vector of reconstructed and true bin contents are eigenvectors with eigenvalue one of the product of the unsmeared and smearing matrices. In general, this will not be the case for an arbitrary reconstructed (or true) distribution. That is, smearing the unsmeared data does not, in general, return exactly the original data. There is thus an ambiguity as to which true distribution is less biased, the truth corresponding to the unsmeared reconstructed data, or the truth that, when smeared, best fits the reconstructed data. In general these two will not be identical, but both methods were used in [95] and shown to return results in agreement within uncertainties. In terms of the results reported in this thesis, the errors should cover the true unsmeared matrix, and thus the true result.

Our knowledge of neutrino physics has grown considerably in the last two decades. Throughout, there has been a constant give and take between the theoretical and experimental aspects of the field. There is no better example of this interplay than the study of neutrino interactions and cross sections. This result is itself the product of many previous studies. Perhaps, one day, it can find similar usage.

Bibliography

- [1] C. Giunti and K. C. Wook, Fundamentals of Neutrino Physics and Astrophysics (Oxford Univ., Oxford, 2007).
- [2] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, *Science* **124**, 103 (1956), <http://www.sciencemag.org/content/124/3212/103.full.pdf>.
- [3] E. Sudarshan and R. Marshak, *Phys. Rev.* **109**, 1860 (1958).
- [4] E. C. G. Sudarshan and R. E. Marshak, Origin of the universal V-A theory, in 50 Years of Weak Interactions, edited by D. Cline, DOE-ER-40200-023. IMSC-TP-85-001. VPI-HEP-84-8, p. 27 p, 1984.
- [5] R. Feynman and M. Gell-Mann, *Phys. Rev.* **109**, 193 (1958).
- [6] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Phys. Rev.* **105**, 1413 (1957).
- [7] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).
- [8] M. Goldhaber, L. Grodzins, and A. W. Sunyar, *Phys. Rev.* **109**, 1015 (1958).
- [9] G. Danby et al., *Physical Review Letters* **9**, 36 (1962).
- [10] F. Hasert et al., *Physics Letters B* **46**, 138 (1973).
- [11] J. Bahcall, W. A. Fowler, I. Iben Jr, and R. Sears, *The Astrophysical Journal* **137**, 344 (1963).
- [12] J. Bahcall, N. Bahcall, and G. Shaviv, *Phys. Rev. Lett.* **20**, 1209 (1968).
- [13] Z. Maki, M. Nakagawa, and S. Sakata, *Progress of Theoretical Physics* **28**, 870 (1962).
- [14] B. Pontecorvo, *Sov. Phys. JETP* **26**, 40 (1968).
- [15] K. Hirata et al., *Phys. Rev. Lett.* **58**, 1490 (1987).
- [16] R. Bionta et al., (IMB Collaboration), *Phys. Rev. Lett.* **58**, 1494 (1987).
- [17] K. Kodama et al., (DONUT Collaboration), *Physics Letters B* **504**, 218 (2001).
- [18] Y. Fukuda et al., (Super-Kamiokande Collaboration), *Physical Review Letters* **81**, 1562 (1998).
- [19] Q. Ahmad et al., (SNO Collaboration), *Physical Review Letters* **89**, 011301 (2002).

- [20] F. An et al., (Daya Bay Collaboration), Physical Review Letters **108**, 171803 (2012).
- [21] K. Abe et al., (T2K Collaboration), Phys. Rev. Lett. **112**, 181801 (2014).
- [22] N. Said, E. Di Valentino, and M. Gerbino, Physical Review D **88**, 023513 (2013).
- [23] E. Giusarma, R. de Putter, S. Ho, and O. Mena, Physical Review D **88**, 063515 (2013).
- [24] K. Olive et al., (Particle Data Group), Chin.Phys. **C38**, 090001 (2014).
- [25] B. Kayser, eConf **C040802**, L004 (2004), hep-ph/0506165.
- [26] S. Chatrchyan et al., Physics Letters B **716**, 30 (2012).
- [27] G. Aad et al., Physics Letters B **716**, 1 (2012).
- [28] F. Halzen and A. D. Martin, Quarks and Leptons: An Introductory Course in Modern Particle Physics (John Wiley and Sons, Inc, 1984).
- [29] J. A. Formaggio and G. P. Zeller, Rev. Mod. Phys. **84**, 1307 (2012).
- [30] D. Casper, Nuclear Physics B - Proceedings Supplements **112**, 161 (2002).
- [31] V. Bernard, L. Elouadrhiri, and U.-G. Meißner, Journal of Physics G: Nuclear and Particle Physics **28**, R1 (2002).
- [32] D. Rein and L. M. Sehgal, Annals of Physics **133**, 79 (1981).
- [33] C. Andreopoulos et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **614**, 87 (2010).
- [34] E. J. Moniz et al., Phys. Rev. Lett. **26**, 445 (1971).
- [35] C. Andreopoulos et al., arXiv preprint (2015), 1510.05494.
- [36] J. Nieves, I. R. Simo, and M. J. V. Vacas, Journal of Physics: Conference Series **408**, 012040 (2013).
- [37] T. Katori, (2013), 1304.6014.
- [38] J. Spitz, (ArgoNeuT Collaboration), Journal of Physics: Conference Series **203**, 012108 (2010).
- [39] K. S. McFarland, (MINERvA Collaboration), Nucl. Phys. B, Proc. Suppl. **159**, 107 (2006).
- [40] A. O. Bazarko, Nuclear Physics B - Proceedings Supplements **91**, 210 (2001), Neutrino 2000.
- [41] K. Mahn, Nuclear Physics B - Proceedings Supplements **159**, 237 (2006), Proceedings of the 4th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region Proceedings of the 4th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region.
- [42] Y. Oyama, arXiv preprint (1998), hep-ex/9803014.
- [43] Y. Hayato, (T2K Collaboration), Nuclear Physics B - Proceedings Supplements **143**, 269 (2005), {NEUTRINO} 2004 Proceedings of the {XXIst} International Conference on Neutrino Physics and Astrophysics.

- [44] C. Wilkinson *et al.*, Phys. Rev. D **93**, 072010 (2016).
- [45] K. Abe *et al.*, T2K Collaboration, Phys. Rev. D **91**, 072010 (2015).
- [46] K. Abe *et al.*, (T2K Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **659**, 106 (2011).
- [47] K. Abe *et al.*, (T2K Collaboration), Phys. Rev. D **87**, 012001 (2013).
- [48] N. Abgrall *et al.*, (NA61/SHINE Collaboration), Journal of Instrumentation **9**, P06005 (2014).
- [49] K. Abe *et al.*, The T2K Collaboration, Phys. Rev. D **91**, 112002 (2015).
- [50] P.-A. Amaudruz *et al.*, (T2K ND280 FGD Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **696**, 1 (2012).
- [51] N. Abgrall *et al.*, (T2K ND280 TPC Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **637**, 25 (2011).
- [52] S. Assylbekov *et al.*, (T2K ND280 PØD Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **686**, 48 (2012).
- [53] S. Fukuda *et al.*, (Super-Kamiokande Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **501**, 418 (2003).
- [54] A. Rodriguez *et al.*, (K2K Collaboration), Phys. Rev. D **78**, 032003 (2008).
- [55] G. P. Zeller, Journal of Physics: Conference Series **136**, 022028 (2008).
- [56] J. L. Alcaraz-Aunion and J. Walding, AIP Conference Proceedings **1189**, 145 (2009).
- [57] A. A. Aguilar-Arevalo *et al.*, (MiniBooNE Collaboration), Phys. Rev. D **81**, 092005 (2010).
- [58] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, (2005), SLAC-R-773.
- [59] R. Brun, F. Bruyant, M. Maire, A. McPherson, and P. Zancarini, CERN Report No. DD-EE-84-1, 1987 (unpublished).
- [60] Y. Hayato, Nuclear Physics B - Proceedings Supplements **112**, 171 (2002).
- [61] T. Golan, J. Sobczyk, and J. Żmuda, Nuclear Physics B - Proceedings Supplements **229–232**, 499 (2012), Neutrino 2010.
- [62] O. Lalakulich, K. Gallmeister, and U. Mosel, Journal of Physics: Conference Series **408**, 012053 (2013).
- [63] T. Dealtry, Using genie and neut with nd280control, T2K internal presentation, 2011.
- [64] D. J. Gross and C. L. Smith, Nuclear Physics B **14**, 337 (1969).

- [65] L. A. Ahrens et al., Phys. Rev. D **35**, 785 (1987).
- [66] A. Bodek and U. K. Yang, Journal of Physics G: Nuclear and Particle Physics **29**, 1899 (2003).
- [67] A. Finch, ND280 Software Workbook, <http://www.hep.lancs.ac.uk/nd280Doc/devel/invariant/nd280Doc/workbook/>, Accessed: 2016-02-25.
- [68] S. Agostinelli et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250 (2003).
- [69] I. J. Taylor, Development of T2K 280m Near Detector Software for Muon and Photon Reconstruction, PhD thesis, Imperial Coll., London, 2009.
- [70] T. Lindner and C. McGrew, oaEvent Manual, Accessed: 2016-02-25, 2010.
- [71] A. Hillairet et al., Multiple Institutions Report No. T2K-TN-072-v1.0, 2011 (unpublished).
- [72] J. Illingworth and J. Kittler, Computer Vision, Graphics, and Image Processing **44**, 87 (1988).
- [73] A. Cervera-Villanueva, J. J. Gomez-Cadenas, and J. A. Hernando, Nucl. Instrum. Meth. **A534**, 180 (2004).
- [74] J. A. Hernando, Recpack: A reconstruction toolkit, Accessed: 2016-02-27, 2003.
- [75] A. Cervera, The highland2 framework, <http://hep.lancs.ac.uk/nd280Doc/stable/invariant/nd280Highland2/>, Accessed: 2016-03-04.
- [76] T. Adye, arXiv preprint (2011), 1105.1160.
- [77] G. D'Agostini, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **362**, 487 (1995).
- [78] G. D'Agostini, arXiv preprint (2010), 1010.0632.
- [79] C. Andreopoulos, J. Dobson, and S. Dytman, Multiple Institutions Report No. T2K-TN-007-v1.0, 2009 (unpublished).
- [80] M. Hierholzer, The xsTool, <http://www.hep.lancs.ac.uk/nd280Doc/xsTool/stable/>, 2015, Accessed: 2016-03-11.
- [81] K. Abe et al., T2K Collaboration, Phys. Rev. D **87**, 092003 (2013).
- [82] T. Campbell et al., Colorado State University, University of Colorado Boulder Report No. T2K-TN-080-v4, 2014 (unpublished).
- [83] K. Gilje, Stony Brook University Report No. T2K-TN-073-v3.1, 2012 (unpublished).
- [84] M. Friend et al., Multiple institutions Report No. T2K-TN-217-v5, 2015 (unpublished).
- [85] M. Gazdzicki, (NA61/SHINE Collaboration), Journal of Physics G: Nuclear and Particle Physics **36**, 064039 (2009).
- [86] N. Abgrall et al., (NA61/SHINE Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **701**, 99 (2013).

- [87] M. Catanesi et al., (HARP Collaboration), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **571**, 527 (2007).
- [88] A. Bercellie et al., Multiple institutions Report No. T2K-TN-192-v2.2, 2015 (unpublished).
- [89] K. Rith, Present status of the emc effect, Lecture notes prepared for Proceedings of the International School of Subnuclear Physics, 2013.
- [90] J. Gomez et al., Phys. Rev. D **49**, 4348 (1994).
- [91] P. Bartet et al., Multiple institutions Report No. T2K-TN-212-v2.1.1, 2015 (unpublished).
- [92] J. P. Lopez, University of Colorado Boulder Report No. T2K-TN-275, 2015 (unpublished).
- [93] C. Bojecho et al., Multiple Institutions Report No. T2K-TN-061-v2.0, 2013 (unpublished).
- [94] A. P. Furmanski et al., University of Warwick Report No. T2K-TN-215-v2.4, 2014 (unpublished).
- [95] K. Abe et al., T2K, Phys. Rev. **D93**, 112012 (2016), 1602.03652.
- [96] A. Longhin, A quick test on xstool unfolding using the original code from g. d'agostini, T2K internal presentation, 2015.
- [97] A. Longhin, Anti-numu inclusive cc cross section in fgd1, T2K internal presentation, 2016.
- [98] E. W. Weinsstein, Sample variance distribution, <http://mathworld.wolfram.com/SampleVarianceDistribution.html>, 2016, Accessed: 2016-05-05.

Appendix A

Ideal Water-in POT for Reduction of Statistical Errors

As data statistics is the second largest source of uncertainty, it is important to ask what how best it can be reduced. Using coarser bins should decrease the fractional error due to statistics, but is not ideal as the results become averaged across more regions of the muon phase space. The alternative is to take more data, but due to the subtraction of water-in and water-out it is inappropriate to take data entirely in one mode or the other. The question then is whether there is an ideal fraction of water-in POT that will minimize the statistical uncertainty on the subtracted result.




Figure A.1: Plot of ξ_{\min} , the fraction of water-in POT that minimizes the statistical uncertainty on the subtracted result, as a function of η . For the PØD, $\eta \approx 0.6$.

To see how the statistical uncertainty behaves as a function of the water-in to water-out POT ratio, note that N_w and N_a in Eq. (7.1) are dependent on the POT. For simplicity, ignore the unfolding and efficiency correction and assume that the measured number of events is proportional to the flux with $N_w = \alpha_w F^w$ and $N_a = \alpha_a F^a$. Then, with $\eta \equiv \alpha_a/\alpha_w$, the fractional error squared goes as,

$$\frac{\sigma_{N_O}^2}{N_O^2} = \frac{N_w + R^2 N_a}{N_O^2} \quad (\text{A.1})$$

$$= \frac{N_w + R^2 N_a}{(N - R N_a)^2} \quad (\text{A.2})$$

$$= \frac{1 + \eta R}{(1 - \eta)^2 N_w} \quad (\text{A.3})$$

$$= \frac{1 + \eta R}{(1 - \eta)^2 \alpha_w F^w} \quad (\text{A.4})$$

$$= C \frac{(1 + \eta R)(1 + R)}{(1 - \eta)^2 R}, \quad (\text{A.5})$$

where,

$$C = \frac{1}{(F^w + F^a)\alpha_w} \quad (\text{A.6})$$

is a constant. Thus, the fractional error on the subtraction depends on both η and R . It can be rewritten in terms of the fraction of water-in POT, $\xi \equiv F^w/(F^w + F^a) = R/(R + 1)$,

$$\frac{\sigma_{N_O}^2}{N_O^2} = C \frac{\xi - \eta\xi - 1}{(1 - \eta)^2 \xi(\xi - 1)}. \quad (\text{A.7})$$

The values of R and ξ that minimize the subtracted fractional uncertainty are,

$$R_{min} = \frac{1}{\sqrt{\eta}} \quad (\text{A.8})$$

$$\xi_{min} = \frac{1}{\sqrt{\eta} + 1}. \quad (\text{A.9})$$

ξ_{min} is plotted as a function of η in Fig. A.1. η can be thought of as the ratio of PØD water-out FV mass to the water-in FV mass. We see that for all reasonable values of η , ξ_{min} lies above 0.5. Therefore, ideally we would like to have more water-in POT than water-out. In reality, the water-in POT is slightly less than the water-out POT over Runs 2–4, as shown in Table 5.1, making the statistical uncertainty more dominant in the subtracted result.

Appendix B

Purity Correction and Unfolding

`RooUnfold`, the unfolding framework used by `xsTool`, treats backgrounds, or what it calls “fakes”, as an additional cause bin, alongside the true kinematic bins towards which we are unfolding. This is seen in the additional truth bin in Fig. 7.4, corresponding to $\Pr(C_b | E_j)$, the probability of finding a background event in effect bin j . The total number of backgrounds, calculated from the difference between selection and signal histograms, alters the prior $\Pr(C_i)$. It can be shown that when the MC prior is used, $1 - \Pr(C_b | E_j)$ is equivalent to the purity for effect bin j and that unfolding the overall (signal+background) reconstructed distribution with C_b included is equivalent to unfolding the purity-corrected reconstructed distribution using a signal-only unfolding matrix.

To see this, first note that the likelihood is calculated from the MC as $\Pr(E_j | C_i) = N(E_j \cap C_i)/N(C_i)$, where $N(A)$ denotes the number of events in category A . Further, the MC prior assumption gives $\Pr(C_i) = N(C_i)/N(\bigcup_k C_k) = N(C_i)/N(C)$. The normalization, Z , required for calculating the posterior, can be calculated from the requirement that $\sum_i \Pr(C_i | E_j) = 1$, which gives,

$$Z = \sum_i \Pr(E_j | C_i) \Pr(C_i) \tag{B.1}$$

$$= \sum_i \frac{N(E_j \cap C_i)}{N(C_i)} \cdot \frac{N(C_i)}{N(C)} \tag{B.2}$$

$$= \sum_i \frac{N(E_j \cap C_i)}{N(C)} \tag{B.3}$$

$$= \frac{N(E_j)}{N(C)}. \tag{B.4}$$

This is true regardless of whether a background cause bin exists. Thus, the posterior can be written Eq. (7.3),

$$\Pr(C_i | E_j) = \frac{N(E_j \cap C_i)}{N(E_j)}, \quad (\text{B.5})$$

and it is easy to see that $1 - \Pr(C_b | E_j)$ gives the purity for E_j .

Next we show that the unfolded result is the same with or without a background cause bin. In other words, the result from unfolding with a purity-corrected, signal-only sample is the same as including C_b and unfolding the default reconstructed distribution, which is what `RooUnfold` actually implements. Given a measured distribution $\tilde{\mathbf{N}}(\mathbf{E})$, the unfolded and efficiency corrected true distribution is given by,

$$\tilde{N}(C_i) = \frac{1}{\epsilon_i} \sum_j \Pr(C_i | E_j) \tilde{N}(E_j) \quad (\text{B.6})$$

$$= \frac{1}{\epsilon_i} \sum_j N(E_j \cap C_i) \frac{\tilde{N}(E_j)}{N(E_j)}. \quad (\text{B.7})$$

In a purity-corrected sample, both N and \tilde{N} are scaled by the same factor. Therefore, the ratio $\tilde{N}(E_j)/N(E_j)$ is the same before and after correction. Further, for all $i \neq b$, ϵ_i is the true selection efficiency and remains unchanged, and $N(E_j \cap C_i)$ restricts itself to signal events falling into E_j which also has no dependence on C_b . This means results with or without C_b are identical.

A simple test was performed to verify that unfolding default reconstructed distributions with an additional background cause bin gave the same results as unfolding purity-corrected distributions with a signal-only unfolding matrix. Figure B.1 shows exact agreement between the two methods.




Figure B.1: Comparison of results from the `RooUnfold` Bayesian unfolding framework with the default background correction vs a manually applied purity correction whereby the selection histograms were corrected prior to being passed to `RooUnfold`. The two results are identical.

Appendix C

Parameters for Flux and Theory Systematics

The following is a list of all parameters used for propagation of flux, FSI, and interaction modeling uncertainties. The reweighting is performed using the **T2KReWeight** package described in [79]. The indexes correspond to the axes in Fig. 7.12.

C.1 Flux Parameters

- (0) JEnu2013a_nd5numu0: ν_μ flux at ND280 magnet (0.0–0.4 GeV)
- (1) JEnu2013a_nd5numu1: ν_μ flux at ND280 magnet (0.4–0.5 GeV)
- (2) JEnu2013a_nd5numu2: ν_μ flux at ND280 magnet (0.5–0.6 GeV)
- (3) JEnu2013a_nd5numu3: ν_μ flux at ND280 magnet (0.6–0.7 GeV)
- (4) JEnu2013a_nd5numu4: ν_μ flux at ND280 magnet (0.7–1.0 GeV)
- (5) JEnu2013a_nd5numu5: ν_μ flux at ND280 magnet (1.0–1.5 GeV)
- (6) JEnu2013a_nd5numu6: ν_μ flux at ND280 magnet (1.5–2.5 GeV)
- (7) JEnu2013a_nd5numu7: ν_μ flux at ND280 magnet (2.5–3.5 GeV)
- (8) JEnu2013a_nd5numu8: ν_μ flux at ND280 magnet (3.5–5.0 GeV)
- (9) JEnu2013a_nd5numu9: ν_μ flux at ND280 magnet (5.0–7.0 GeV)
- (10) JEnu2013a_nd5numu10: ν_μ flux at ND280 magnet (7.0–30.0 GeV)

- (11) JEnu2013a_nd5numub0: $\bar{\nu}_\mu$ flux at ND280 magnet (0.0–0.7 GeV)
- (12) JEnu2013a_nd5numub1: $\bar{\nu}_\mu$ flux at ND280 magnet (0.7–1.0 GeV)
- (13) JEnu2013a_nd5numub2: $\bar{\nu}_\mu$ flux at ND280 magnet (1.0–1.5 GeV)
- (14) JEnu2013a_nd5numub3: $\bar{\nu}_\mu$ flux at ND280 magnet (1.5–2.5 GeV)
- (15) JEnu2013a_nd5numub4: $\bar{\nu}_\mu$ flux at ND280 magnet (2.5–30.0 GeV)
- (16) JEnu2013a_nd5nue0: ν_e flux at ND280 magnet (0.0–0.5 GeV)
- (17) JEnu2013a_nd5nue1: ν_e flux at ND280 magnet (0.5–0.7 GeV)
- (18) JEnu2013a_nd5nue2: ν_e flux at ND280 magnet (0.7–0.8 GeV)
- (19) JEnu2013a_nd5nue3: ν_e flux at ND280 magnet (0.8–1.5 GeV)
- (20) JEnu2013a_nd5nue4: ν_e flux at ND280 magnet (1.5–2.5 GeV)
- (21) JEnu2013a_nd5nue5: ν_e flux at ND280 magnet (2.5–4.0 GeV)
- (22) JEnu2013a_nd5nue6: ν_e flux at ND280 magnet (4.0–30.0 GeV)
- (23) JEnu2013a_nd5nueb0: $\bar{\nu}_e$ flux at ND280 magnet (0.0–2.5 GeV)
- (24) JEnu2013a_nd5nueb1: $\bar{\nu}_e$ flux at ND280 magnet (2.5–30.0 GeV)

C.2 FSI Parameters

- (25) NCasc_FrInelLow_pi: inelastic rescattering probability for low energy pions
- (26) NCasc_FrInelHigh_pi: inelastic rescattering probability for high energy pions
- (27) NCasc_FrPiProd_pi: pion production probability
- (28) NCasc_FrAbs_pi: pion absorption probability
- (29) NCasc_FrCExLow_pi: charge exchange probability for low energy pions

- (30) NCasc_FrCEXHigh_pi: charge exchange probability for high energy pions

C.3 Cross Section Parameters

- (31) NXSec_MaCCQE: MaCCQE
- (32) NIWG2014a_pF_C12: carbon Fermi momentum
- (33) NIWGMEC_Norm_C12: carbon 2p2h normalization
- (34) NIWG2014a_Eb_C12: carbon binding energy
- (35) NIWG2014a_pF_016: oxygen Fermi momentum
- (36) NIWGMEC_Norm_016: oxygen 2p2h normalization
- (37) NIWG2014a_Eb_016: oxygen binding energy
- (38) NXSec_CA5RES: CA5 for single pion interaction
- (39) NXSec_MaFFRES: MaRES
- (40) NXSec_BgSc1RES: $I = 1/2$ background scale factor for single pions
- (41) NIWG2012a_ccnueE0: CC-nue normalization
- (42) NIWG2012a_dismpishp: DIS multi-pi normalization
- (43) NIWG2012a_cccohE0: CC-Coherent normalization
- (44) NIWG2012a_nccohE0: NC-Coherent normalization
- (45) NIWG2012a_ncotherE0: NC-Other normalization

Appendix D

Detailed Cross Section Results

$0.00 \leq \cos \theta_\mu < 0.60$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	1.2281	0.2415	0.2253
$0.4 \leq p_\mu < 0.5$	1.5820	0.1498	0.3227
$0.5 \leq p_\mu < 0.7$	0.5065	0.0529	0.1036
$0.7 \leq p_\mu < 0.9$	0.0595	0.0097	0.0232
$0.9 \leq p_\mu < 2.5$	0.0017	0.0006	0.0011
$2.5 \leq p_\mu < 5.0$	0.0000	0.0000	0.0000

Table D.1: The ν_μ CC0 π cross section on water in the first $\cos \theta_\mu$ slice.

$0.60 \leq \cos \theta_\mu < 0.70$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.8617	0.1753	0.1241
$0.4 \leq p_\mu < 0.5$	2.9489	0.3089	0.3576
$0.5 \leq p_\mu < 0.7$	1.8845	0.1880	0.2240
$0.7 \leq p_\mu < 0.9$	0.3385	0.0449	0.1038
$0.9 \leq p_\mu < 2.5$	0.0166	0.0028	0.0083
$2.5 \leq p_\mu < 5.0$	0.0000	0.0000	0.0000

Table D.2: The ν_μ CC0 π cross section on water in the second $\cos \theta_\mu$ slice.

$0.70 \leq \cos \theta_\mu < 0.80$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.5217	0.1140	0.0814
$0.4 \leq p_\mu < 0.5$	2.0243	0.2328	0.2737
$0.5 \leq p_\mu < 0.7$	1.6348	0.1511	0.1866
$0.7 \leq p_\mu < 0.9$	0.6126	0.0620	0.1056
$0.9 \leq p_\mu < 2.5$	0.0380	0.0058	0.0130
$2.5 \leq p_\mu < 5.0$	0.0000	0.0000	0.0000

Table D.3: The ν_μ CC0 π cross section on water in the third $\cos \theta_\mu$ slice.

$0.80 \leq \cos \theta_\mu < 0.85$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.5581	0.1114	0.0813
$0.4 \leq p_\mu < 0.5$	1.8765	0.2129	0.2553
$0.5 \leq p_\mu < 0.7$	1.9741	0.1852	0.2016
$0.7 \leq p_\mu < 0.9$	0.9055	0.0887	0.1295
$0.9 \leq p_\mu < 2.5$	0.1073	0.0167	0.0220
$2.5 \leq p_\mu < 5.0$	0.0004	0.0001	0.0004

Table D.4: The ν_μ CC0 π cross section on water in the fourth $\cos \theta_\mu$ slice.

$0.85 \leq \cos \theta_\mu < 0.90$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.4470	0.0899	0.0686
$0.4 \leq p_\mu < 0.5$	1.2922	0.1598	0.2141
$0.5 \leq p_\mu < 0.7$	1.7339	0.1697	0.1836
$0.7 \leq p_\mu < 0.9$	1.0930	0.1000	0.1361
$0.9 \leq p_\mu < 2.5$	0.2050	0.0210	0.0264
$2.5 \leq p_\mu < 5.0$	0.0030	0.0009	0.0018

Table D.5: The ν_μ CC0 π cross section on water in the fifth $\cos \theta_\mu$ slice.

$0.90 \leq \cos \theta_\mu < 0.93$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.2969	0.0656	0.0608
$0.4 \leq p_\mu < 0.5$	1.1157	0.1570	0.1998
$0.5 \leq p_\mu < 0.7$	1.5550	0.1610	0.1864
$0.7 \leq p_\mu < 0.9$	1.3349	0.1201	0.1485
$0.9 \leq p_\mu < 2.5$	0.3184	0.0315	0.0353
$2.5 \leq p_\mu < 5.0$	0.0132	0.0025	0.0071

Table D.6: The ν_μ CC0 π cross section on water in the sixth $\cos \theta_\mu$ slice.

$0.93 \leq \cos \theta_\mu < 0.97$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.1996	0.0517	0.0459
$0.4 \leq p_\mu < 0.5$	1.0259	0.1551	0.1562
$0.5 \leq p_\mu < 0.7$	1.2979	0.1535	0.1388
$0.7 \leq p_\mu < 0.9$	1.0886	0.1096	0.1180
$0.9 \leq p_\mu < 2.5$	0.5005	0.0511	0.0353
$2.5 \leq p_\mu < 5.0$	0.0911	0.0106	0.0129

Table D.7: The ν_μ CC0 π cross section on water in the seventh $\cos \theta_\mu$ slice.

$0.97 \leq \cos \theta_\mu < 1.0$			
Bin (GeV)	$\frac{d^2\sigma}{d(\cos \theta_\mu)dp_\mu}$ (10^{-38} cm ² /GeV/n)	Systematics (10^{-38} cm ² /GeV/n)	Statistics (10^{-38} cm ² /GeV/n)
$0.0 \leq p_\mu < 0.4$	0.0665	0.0248	0.0417
$0.4 \leq p_\mu < 0.5$	0.4959	0.0947	0.1434
$0.5 \leq p_\mu < 0.7$	0.5830	0.0813	0.1269
$0.7 \leq p_\mu < 0.9$	0.4719	0.0628	0.1225
$0.9 \leq p_\mu < 2.5$	0.4486	0.0619	0.0384
$2.5 \leq p_\mu < 5.0$	0.4242	0.0557	0.0357

Table D.8: The ν_μ CC0 π cross section on water in the eighth $\cos \theta_\mu$ slice.