Measurement of the Muon Antineutrino Charged Current Single Charged Pion Production Cross-Section on Water and Hydrocarbon with ND280

Liam O'Sullivan

Department of Physics and Astronomy,
The University of Sheffield

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy

July 4, 2021

Measurement of the $\bar{\nu}_\mu$CC1π^- Cross-Section on Water and Hydrocarbon with ND280

Liam O'Sullivan

Abstract

Tokai to Kamioka (T2K) is a long-baseline neutrino oscillation experiment based in Japan, performing measurements of neutrino oscillations and neutrino-nucleus interactions. Uncertainty in the modelling of neutrino-nucleus interactions on the order of GeV is a significant systematic uncertainty in T2K and other current oscillation experiments, and is set to become a limiting uncertainty in many measurements as new experiments are constructed. These uncertainties result from the difficulty in modelling the effects of the nuclear medium on the underlying neutrino interaction, where the kinematics and spatial distribution of nucleons is well understood for only the lightest nuclei. As such, a much improved understanding of nuclear effects on neutrino interactions is essential for neutrino interaction experiments of the near future.

This thesis describes the first measurement of the muon anti-neutrino charged current single pion production cross-sections on water and hydrocarbon using the passive water and active scintillator targets in the T2K ND280 near detector. The cross-section is measured as a function of the outgoing pion kinematics, which depend heavily on both the underlying neutrino interaction models and nuclear state model, and can be directly measured for each event by the detector, reducing dependence on the input simulation. A novel method of collapsing a double differential cross-section measurement into two correlated single differential cross-section measurements is presented and validated for this analysis. It is hoped that this measurement will provide additional constraints on neutrino-nucleus pion production models, enabling better understanding of the shortcomings in these models and where improvements are required.

Acknowledgements

The analysis that forms the basis for this doctoral thesis was performed using data taken by the T2K experiment, to which hundreds of scientists across the globe have contributed over many years. From the inception of T2K through to today, I am grateful to all who have contributed to the T2K experiment in some way, large or small. In particular I would like to thank Stephen, Lukas, Andrew, Sasha, Sophie, Ciro, Dan, Luke, Callum, Clarence, and Patrick, for numerous and often invaluable insights and contributions to my work over the years.

The greatest thanks to be given are naturally to my parents, Anton and Ger, who have tirelessly supported me through the good and the bad, and for raising me into the moderately-functional human that I have become; 'Twas far enough from pizza I was reared, never mind *neutrinos.*

To my supervisor Lee I give my heartfelt thanks and gratitude. Ever present but never over-bearing, I was given freedom but guidance and direction were available by the bucket when needed. I am very lucky to have been able to work under you.

Years of sitting at a PC in an office couldn't have been anywhere near as enjoyable as my time in Sheffield was without the great people I was fortunate to share an office with. Thanks to Ahmad, Andrew, Carl, Dave, Michael, Mike and Sam F. for helping make the office a good place to work, but also a good place to procrastinate on occasion. Thanks also to the rest of the neutrino/water-Cherenkov group at Sheffield. A special mention is deserved for (in chronological order) Patrick, Celeste, Sam J., and Jordan for making not only the office, but many trips abroad more enjoyable, and for helping me with various aspects of work and acclimatising to the complexities of English life (pasties, pies, etc.). Thank you to the staff in F10 for making the administrative tasks almost effortless.

Many thanks to those who made my time in Japan as enjoyable as it was. In particular to Matt and Trevor for making the UK flat a great place to live, and showing me many of Tokai's hidden (unpolished) gems. To the many natives who put up with my mediocre Japanese over the years and helped me improve, I am grateful. Without Oka-kun, Kusaka-san, Seika-kun, Koji-san, Tomo-san, Fukugawa-san, Onda-san, Amaya-kun, Iwaya-san, Iwamoto-san, and Chikuma-san my time in Japan would've been profoundly less exciting. I also owe a great many thanks to my contemporaries, in particular Artur, Will, Ed, and Yue, who made the downtime both on and off site as enjoyable as it was.

Finally, thanks to the university and department for funding my research.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Neutrinos are neutral fermionic leptons that interact with other matter only through gravitation and the weak nuclear force. The name ‘neutrino’ derives from Italian, as a diminutive of ‘neutral’, with the resulting word meaning ‘little neutral one’. In spite of their abundance in the known universe — second only to photons — and long history of study, many of their properties remain poorly understood.

1.1 A Brief History of Neutrinos

The first hint at the existence of a neutrino was in the continuous energy spectrum of electrons emitted during nuclear beta decay. As experiments observed only a single few-MeV electron from these nuclear decays, it appeared that this process was a two-body process, implying through energy and momentum conservation that the emitted particle should have a discreet energy. It was suggested that this was evidence of energy and momentum conservation not being strictly universal in response to this. In contrast, in 1930 Wolfgang Pauli postulated a new third particle to be present in this decay, allowing a straightforward explanation of the observations while maintaining energy and momentum conservation. An additional property of this proposed particle was that it was undetectable by the experiments of the time. In 1934 Enrico Fermi published his theory of beta decay, in which a neutron decayed into a proton, an electron, and the particle proposed previously by Pauli, which Fermi named the neutrino.

The first dedicated experiment to search for neutrinos was Project Poltergeist, started in 1951 by Frederick Reines and Clyde Cowan. The experiment proposed to search for neutrinos indirectly by measuring gamma rays associated with positron production in inverse beta-decay interactions. As the inverse of the reaction responsible for the previously mentioned neutrino production, a neutrino would interact with a proton to produce a neutron and a positron, where the positron would quickly annihilate with an electron in the surrounding matter to produce two gamma rays. To demonstrate the association of these gamma rays with neutrino interactions, it was necessary to compare rate measurements in the detector both with and without a neutrino source.

The first proposal for a neutrino source was an atomic bomb, which could provide a significant, prompt flux of neutrinos, allowing the potential for very high signal to noise detection. Although performing a physics measurement in close proximity to an uncontrolled nuclear explosion presents obvious challenges, Reines and Cowan proposed an experiment [1] to do just this. In this experiment, a detector was to be suspended in a long, vertical underground shaft under vacuum, with the atomic bomb located nearby above-ground. Moments after the detector would be released into freefall, the nuclear payload would detonate providing the neutrino flux, while the physics detector in freefall in a vacuum would be shielded from any associated shock wave. The bottom of the shaft was to be filled with feathers and rubber to provide a soft cushion, such that the detector would not suffer significant shock when decelerating upon impact. This proposal to use an atomic bomb for this experiment was approved by the director of Los Alamos National Laboratory. Before this experiment could be realised, it was proposed to instead use the neutrino flux from the reactor core at the nearby Savanah River nuclear reactor. This experimental setup consisted of a volume of liquid scintillator intsrumented with photomultiplier tuber (PMTs). To further improve signal purity, cadmium chloride was dissolved into the scintillator, as cadmium has a high neutron absorption cross section. After absorbing a neutron, the cadmium nucleus undergoes a nuclear de-excitation with the emission of a gamma ray, where this gamma ray is slightly delayed with respect to the gamma ray from positron annihilation, resulting in two pulses, as measured by the PMTs. [2]. This double pulse signal greatly reduced the background rate, and by 1956 the experiment had established the existence of the neutrino through inverse beta decay.

The discovery of a second, distinct type of a neutrino was reported in 1962 by Lederman, Schwarz, and Steinberger et al. using the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory [3]. While this experiment differed in almost all practical aspects to that of Reines and Cowan, the fundamental principles remained the same; neutrino production through particle decay and detection via their inverse reaction. The AGS accelerated protons to an energy of 15 GeV

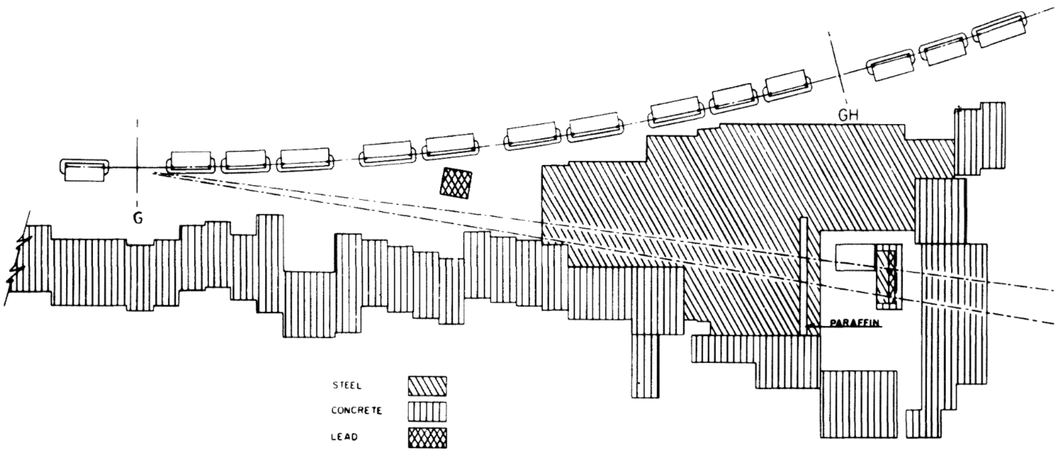

Figure 1.1: Overview of the AGS neutrino beamline setup, with the neutrino source shown on the left, and the detector in the cavity to the right. Figure from [3].

which then impinged on a beryllium target, producing a shower of primarily pions and kaons. The decay of charged pions in particular, $\pi^\pm \rightarrow \mu^\pm + \nu/\bar{\nu}$, provided the primary neutrino component. Located 7.5° off-axis from the initial proton beam direction some tens of metres away was placed a counting spark chamber surrounded by anti-coincidence shields to reject cosmic ray muons, with the neutrino path traversing a 13.5 m steel wall to sufficiently attenuate almost all muons from the pion decay. Through measuring the number of high energy muons appearing inside the physics detector and stringently rejecting many backgrounds, it was determined that the neutrinos produced in charged pion decay produced only muons through the inverse process, distinct from the electrons produced in the experiment by Reines and Cowan. This demonstrated not only the existence of two distinct flavours of neutrino, but that to each lepton — muon and electron at the time — was paired a neutrino. This experiment was also the first to produce neutrinos using a particle accelerator, a method widely used to this day.

Despite the discovery of the tau lepton in 1975 at Lawrence Berkeley National Laboratory, it would not be until 2000 that the corresponding tau neutrino was first measured. [4] The experimental setup greatly resembled that of the experiment to measure the muon neutrino, using the 800 GeV Tevatron accelerator and a tungsten target 36 m upstream of a particle detector, heavily shielded from all other particles. As the reaction of a tau neutrino produces a tau lepton, the ability to accurately identify tau leptons from other particles was key in this detection. The detector was comprised of alternating planes of iron and emulsion plates, where the iron

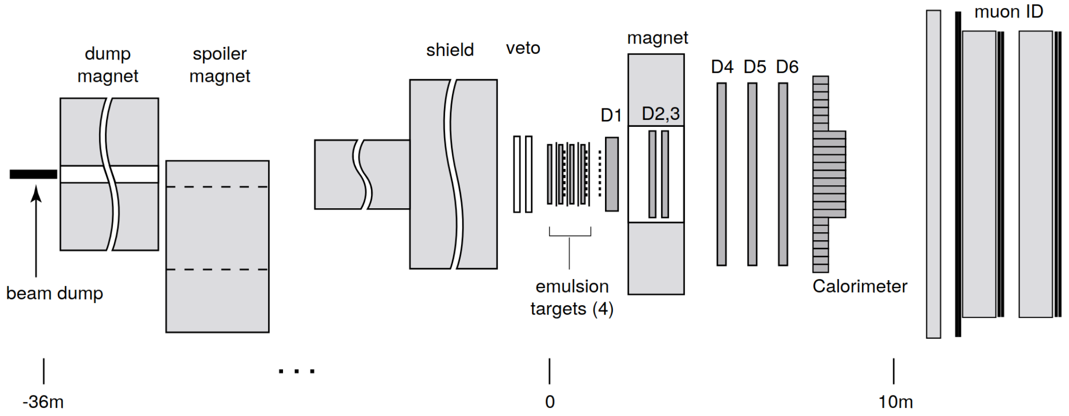

Figure 1.2: The experimental setup used by the DONUT collaboration to detect four tau neutrino interactions. Figure from [4].

provided a high density target for neutrino interactions, and the emulsion plates allowed precise 2D (within each plate) tracking and a measurement of energy loss. Due to the extremely short lifetime of tau leptons ($\widetilde{10}^{-13}$ s) and the majority decays resulting in a single charged particle, the signal was that of a track with a 'kink' — an abrupt change in direction — after traversing a few mm. The experiment measured 4 such events with a background of approximately 0.34 events, providing conclusive evidence of the existence of the third non-sterile neutrino flavour.

In 2005, the ALEPH collaboration published results determining there to be exactly 3 light ($m_\nu \leq 45$ GeV) neutrino flavours [5]. Using the LEP e^+e^- collider at cern, they measured the width of the Z^0 resonance for the process $e^+e^- \to f\bar{f}$. As the width of the Z^0 resonance depends directly on the number of possible final states to which it can decay, including $Z^0 \to \nu_\alpha\bar{\nu}_\alpha$ for each neutrino flavour α, a measurement of the width provides direct information about the number of such allowed decays.

Parallel to the experiments searching for the existence of these neutrinos, a number of other experiments aimed to measure specific properties of neutrinos, together with natural sources of neutrinos. While their extremely low interaction rate makes them challenging to detect, it also provides the potential to detect neutrinos through media that are opaque to photons. The first attempt at this was the Homestake experiment [6], which in the late 1960s aimed to measure the neutrino flux from the sun. At the time, stellar physics models were reasonably well understood, with an equation of state well approximated by simple first order differential equation. Moreover, the luminosity of the sun is determined directly by the rate of nuclear fusion in the core, and thus a measurement of the luminosity — a trivial astronomical measurement — provides a strong constraint on the flux of electron neutrinos from the sun. The experiment was performed in the Homestake mine in South Dakota, almost 1500 m underground. A 380 m^3 tank filled with a chlorine based cleaning agent comprised the detector. Chlorine was chosen as, when undergoing inverse beta decay, the nucleus is transformed into an argon nucleus. The argon can be separated from the chlorine and the amount counted, providing a measure of the rate of inverse beta decay. To the dismay of astrophysicists and particle physicists alike, the measured flux was approximately one third of the expectation. Despite efforts to find issues in the solar model, the inverse beta decay model, and other parts of the expected rate calculation, the issue persisted, known as the 'solar neutrino problem' [7].

The Soviet-American Gallium Experiment (SAGE) [8] began measurement of the solar neutrino flux in 1989 using a target mass of over 50 tonnes of liquid gallium metal, based in a mine in the Baksan Neutrino Observatory. Using the same concept as the Homestake experiment, gallium undergoing inverse beta decay produces a germanium nucleus, and the amount of germanium can be counted to determine the neutrino flux. Gallium has the advantage of a lower energy threshold for inverse beta decay compared to chlorine, allowing the experiment to access more of the solar flux, and thus a higher count rate. The same detection method was employed by the GALEX [9] (later GNO) experiment based in Laboratori Nazionali del Gran Sasso, which began data taking in 1991. Rather than liquid metal gallium, GALEX used gallium trichloride disolved in hydrochloride acid. Both SAGE and GALEX measured solar neutrino fluxes consistent with each other and with the value measured by the Homestake experiment, confirming the observed neutrino flux deficit.

The solar neutrino problem would remain experimentally unsolved until the Sudbury Neutrino Observatory (SNO) and Super-Kamiokande experiments published evidence of neutrino oscillations, explaining the solar flux deficit [10, 11]. In particular, SNO's sensitivity to neutral current interactions provided a powerful channel to measure the overall neutrino flux independent of flavour composition, where previous solar neutrino experiments measured charged current interactions for electron neutrinos only. This demonstrated that the total solar neutrino flux was as predicted by solar models, and that the flavour composition at Earth's distance was different to that produced in the Sun, a discovery which would earn both collaborations the 2015 Nobel Prize in physics. While this discovery solved one open problem in physics, in order for such oscillations to be possible, at least two neutrinos are required to be massive, at odds with the massless neutrinos of the standard model.

1.2 Neutrino Oscillations

Neutrino oscillation was formalised by Pontecorvo, Make, Nakagawa, and Sakata (PMNS) [12, 13] in 1962, providing a theoretical framework that accurately describes oscillation experiments to this day. Pontecorvo's initial theory drew a parallel to neutral kaon mixing ($\bar{K}_0 \rightarrow K_0$) which had been observed in the decade prior. Fundamental to this theory is the idea that neutrinos interact through the weak force via their flavour states, while they propagate through space and time as mass states. These mass and flavour states are superpositions of each other, and are related by a "mixing matrix". With the neutrino mass eigenstates represented by $|\nu_k\rangle, k = 1, 2, 3$ and the flavour eigenstates as $|\nu_\alpha\rangle, \alpha = e, \mu, \tau$, the relation between these states can be written as

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k}|\nu_k\rangle \tag{1.1}$$

where $U_{\alpha k}$ — the previously mentioned mixing matrix— is a unitary 3×3 matrix that describes the mixing between these two different states. The matrix being unitary prevents any change in the total probability of a neutrino being in any state. In its explicit form, the matrix is

$$U_{\alpha k} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \tag{1.2}$$

which closely resembles the Cabbibo-Kobayashi-Masakawa (CKM) mixing matrix for quarks. [14] In a more condensed notation

$$\begin{aligned} U_{\alpha k} &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\mathrm{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\mathrm{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{aligned} \tag{1.3}$$

where $c_{ij} = \cos\theta_{ij}, s_{ij} = \sin\theta_{ij}$ are functions of three "mixing angles" $\theta_{12}, \theta_{23}, \theta_{13}$, δ_{CP} is a Dirac charge-parity (CP) violating phase, and α_1 and α_2 are Majorana CP violating phases. In the case that neutrinos are not Majorana particles, the final matrix in equation 1.3 reduces to an identity matrix. Moreover, as the matrix contains only diagonal elements, these CP violating phases do not effect the overall oscillation probabilities. With natural units ($\hbar = c = 1$), neutrino mass states are eigenstates of the Hamiltonian

$$\begin{aligned} \mathcal{H}|\nu_k\rangle &= E_k|\nu_k\rangle, \\ E_k &= \sqrt{p^2 + m_k^2} \end{aligned} \tag{1.4}$$

where E_k are the energy eigenvalues. This also allows rewriting of the difference in energies from state to state,

$$E_k - E_j \approx \frac{\Delta m_{kj}^2}{2E} \tag{1.5}$$

where $\Delta m_{kj}^2 = m_k^2 - m_j^2$ is the squared mass difference, provided that masses m_k and m_j are small.

To determine the manner in which neutrino oscillations occur from this framework, consider first a neutrino created in an interaction with a known flavour state. From equation 1.1 this neutrino will be a superposition of mass states, for which the temporal evolution is given by the time-independent Schrödinger equation

$$|\nu_k(t)\rangle = e^{-iE_k t}|\nu_k\rangle \tag{1.6}$$

In the ultra-relativistic limit, this can be rewritten as

$$|\nu_k(L)\rangle = e^{-i\frac{m_k^2 L}{2E}}|\nu_k\rangle \tag{1.7}$$

where $L \approx t$ is the distance the neutrino has travelled, and $E_k = \sqrt{p_k^2 + m_k^2} \approx p_k + \frac{m_k^2}{2p_k} \approx E + \frac{m_k^2}{2E}$ with E being the total energy of the neutrino. These approximations are practically applicable to all detectable neutrinos, as the masses of all known neutrinos are so low that at any measurable energy $p_k >> m_k$ is satisfied.

Combining equations 1.7 and 1.1 allows the computation of the time evolution of state α,

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t}|\nu_k\rangle \tag{1.8}$$

which combined with equation 1.1 again for $|\nu_k\rangle$ yields

$$|\nu_\alpha(t)\rangle = \sum_\beta \sum_k U^*_{\alpha k} e^{-im_k^2 L/2E} U_{\beta k} |\nu_\beta\rangle. \tag{1.9}$$

This equation demonstrates that a neutrino starting at $t = 0$ with pure flavour state α evolves into a superposition of all flavour states as time evolves, based on the strength of the mixing encoded in the matrix U. The probability for this neutrino to be found in flavour state β is then given by

$$P_{\alpha\to\beta}(t) = |\langle\nu_\beta(t)|\nu_\alpha\rangle|^2 = \left|\sum_k U^*_{\alpha k} U_{\beta k} e^{-im_k^2 L/2E}\right|^2 \tag{1.10}$$

$$\begin{aligned} P_{\alpha\to\beta}(t) =& \delta_{\alpha\beta} \;-\; 4\sum_{k>j} \mathrm{Re}\left[U^*_{\alpha k} U_{\beta k} U^*_{\alpha j} U_{\beta j}\right] \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) \\ & \pm 2\sum_{k>j} \mathrm{Im}\left[U^*_{\alpha k} U_{\beta k} U^*_{\alpha j} U_{\beta j}\right] \sin^2\left(\frac{\Delta m^2_{kj} L}{2E}\right) \end{aligned} \tag{1.11}$$

where the positive and negative imaginary components are for neutrinos and anti-neutrinos respectively. In this form, the oscillation probability can be seen to depend only on the matrix elements — three angles and a CP violating phase — the squared mass splittings and the ratio of distance to neutrino energy. Due to its position in the mixing matrix, this CP violating phase can only alter oscillations if $\sin\theta_{13} \neq 0$. Modern neutrino oscillation experiments generally attempt to measure these mixing angles and δ_{CP} using a fixed distance L and a known neutrino energy spectrum.

1.3 Mass Hierarchy

One of the remaining unknowns of the neutrinos is the ordering of the mass states; that is whether the ν_3 mass eigenstate is lighter or heavier than the mass eigenstates of ν_1 and ν_2. While the effects of the mass eigenstates on neutrino-nucleus scattering are minimal, the ordering of the mass eigenstates has a significant impact on the results of oscillation experiments and double-beta decay experiments. A detailed overview of the observable differences in neutrino physics and the experiments built to measure them is presented in [17].

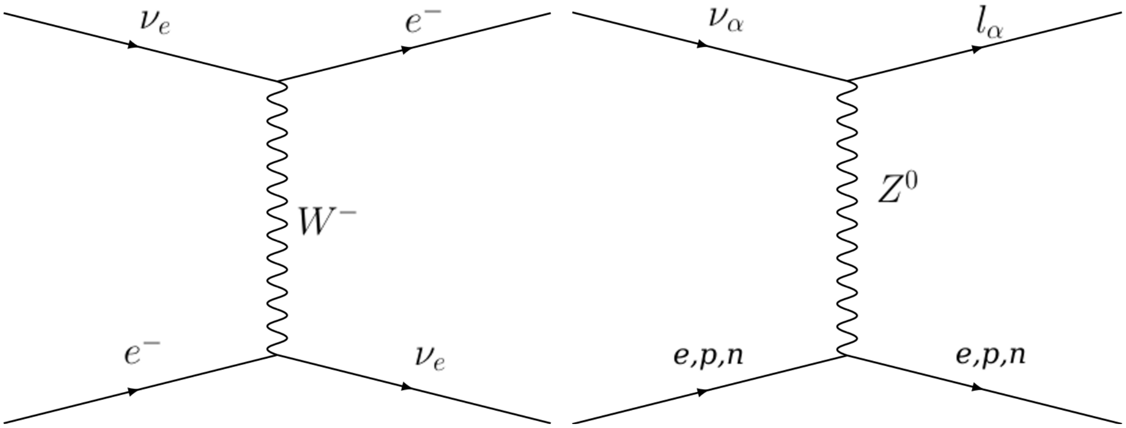

Figure 1.3: Diagrams of the two neutrino interactions that occur in matter that preserve the neutrino while altering oscillation probability.

1.4 Matter Effects on Neutrino Oscillation

The preceding derivation of the oscillation probability is for neutrinos in a vacuum. While traversing dense matter, neutrinos can undergo interactions that do not convert them to charged leptons, but affect the oscillation probability. [15] These interactions are the neutral current scattering off a proton, neutron, or electron, and the charged current scattering off an electron. Feynman diagrams for these processes are shown in figure 1.3.

For the case where this scattering occurs via a Z^0 boson, all neutrino flavours couple equally, resulting in an effect that is constant with neutrino flavour. For the $W^\pm$ interaction however, the presence of a charged lepton is required; as muons and tau leptons are generally not present in normal matter, the effect is only present for ν_e, causing an intrinsic asymmetry in how oscillations behave outside of a vacuum. Moreover, the $W^\pm$ interaction is not allowed for anti-neutrinos, due to the absence of positrons in matter. For experiments measuring neutrino CP violation involving $\nu_e, \bar{\nu}_e$, this matter effect must be accounted for, as it results in an asymmetry potentially indistinguishable from actual leptonic CP violation. These matter effects, unlike the oscillations described before, are sensitive to the signs of the mass splittings Δm^2_{kj}, providing a manner in which the mass hierarchy of neutrinos — whether $\Delta^2 m^2_{32} > 0$ or $\Delta^2 m^2_{32} < 0$ — can be determined experimentally. Solar neutrino observations have already determined that Δm^2_{21} is positive, due to the matter effect observed in the spectrum of electron neutrinos from the solar core. [16]

1.5 Modern Long-Baseline Neutrino Oscillation Experiments

It can be seen from equation 1.11 that other than the oscillation parameters, the oscillation probability depends on the neutrino energy and the distance travelled. Experiments measuring oscillations will generally use a two-detector setup, where a "near" detector measures the neutrino flux at or close to the source where oscillation effects are minimal, and a "far" detector located at a distance near the oscillation maximum for the relevant neutrino energies. The distance from the neutrino source to the far detector is referred to as the *baseline*. Modern neutrino oscillation experiments broadly fall into one of two categories: long-baseline, with baselines on the order of 100s of km and neutrino energies on the order of GeV, and short-baseline, with baselines on the order of 1 km and energies on the order of MeV. Short-baseline neutrino experiments are not discussed further in this thesis, but are comprehensively reviewed in [18].

The two currently running long-baseline experiments are T2K — which will be covered in more detail in chapter 3 — and NOνA [19]. Both experiments produce high-purity beams of muon (anti-)neutrinos using proton accelerators, and measure both the disappearance of muon (anti-)neutrinos and the appearance of electron (anti-)neutrinos at a distance a few 100s of km away. The neutrino rates measured at each detector, together with the neutrino flux, can be used to constrain the oscillation parameters. Measurement of muon (anti-)neutrino disappearance allows determination of $\sin^2\theta_{23}$ and Δm^2_{23}, while measurement of electron (anti-)neutrino appearance allows a measurement of $\sin^2\theta_{13}$ and δ_{CP}. Although conceptually simple, the process of making a rate prediction for such experiments is non-trivial. The observed neutrino rate at a detector depends on the overall flux, the neutrino interaction cross-section, the detector response to the neutrino events, and the oscillation parameters. The total number of predicted events in a detector — for a given set of model parameters $\vec{y}$ — can be expressed as

$$N_{\text{pred}}(\vec{y}) = \int \Phi(E_\nu) \times \sigma(E_\nu, \vec{y}) \times P_{\nu_k \to \nu_j}(E_\nu) \times R(E_\nu, \vec{y}) dE_\nu \tag{1.12}$$

where Φ is the neutrino flux, σ is the neutrino interaction cross section, and R is the detector response to the neutrino events. Experiments perform this task by using a Monte-Carlo simulation of the overall experimental setup, usually with specific Monte-Carlo simulations for each part of the experiment (detectors, beam, etc.). An important point to note is that in equation 1.12, all inputs are functions of neutrino energy, which due to the broad spectrum of the fluxes used in these experiments, is generally not known a-priori. As such, it is desirable to measure or reconstruct the neutrino energy for each interaction that occurs in each detector.

Figure 1.4 shows the reconstructed event rate for ν_μ-like and $\bar{\nu}_\mu$-like samples at SK, the T2K far detector for all available runs to date. A dip in both spectra can be seen at around 0.6 GeV, corresponding to the energy at which the probability for oscillation is maximal.

The reconstruction of neutrino energy is a challenging task, and the source of some of the largest systematic uncertainties in these experiments. To reconstruct the energy of a CC neutrino interaction to first order, all particles from the event should be detected and their energies reconstructed, with the sum of all this energy giving the neutrino energy. For real world detectors, there are generally energy thresholds for a particle to be detected, and some particles (e.g. neutrons) may not be detectable at all. Moreover, a significant amount of energy from the neutrino interaction may not even escape the nucleus to be detectable. In these cases a calorimetric approach is not feasible on its own, as some unknown amount of energy will always be unaccounted for. To overcome this limitation, a neutrino-nucleus interaction model is used, and the energy of each event inferred by the model based on the kinematics of some of the observed particles. This approach, while enabling neutrino energy reconstruction, is not without issues; the inner workings of the model may not accurately describe the true interaction physics, and the resulting predictions may bias the energy reconstruction in subtle but important ways. To minimise these effects, experiments often use only a subset of relatively well-understood interactions for their oscillation analysis.

Having taken data for a number of years, the primary uncertainties on the values of $N_{\text{pred}}(\vec{y})$ for both T2K and NOνA will soon no longer be statistically limited, instead limited by the systematic uncertainties on the various quantities in equation 1.12. In order to make efficient use of data being collected presently and in the future, better understanding of the overall modelling is required, in particular neutrino interaction and flux prediction models.

All neutrino oscillation mixing angles $\sin^2\theta_{kj}$ have been measured to reasonable precision by various experiments using different baselines, neutrino sources, and detector types. The current PDG values [14] for these quantities are derived from a best fit of all available constraints. For δ_{CP} however, precise measurement remains elusive due to the difficulty in accurately measuring oscillation of both neutrinos and

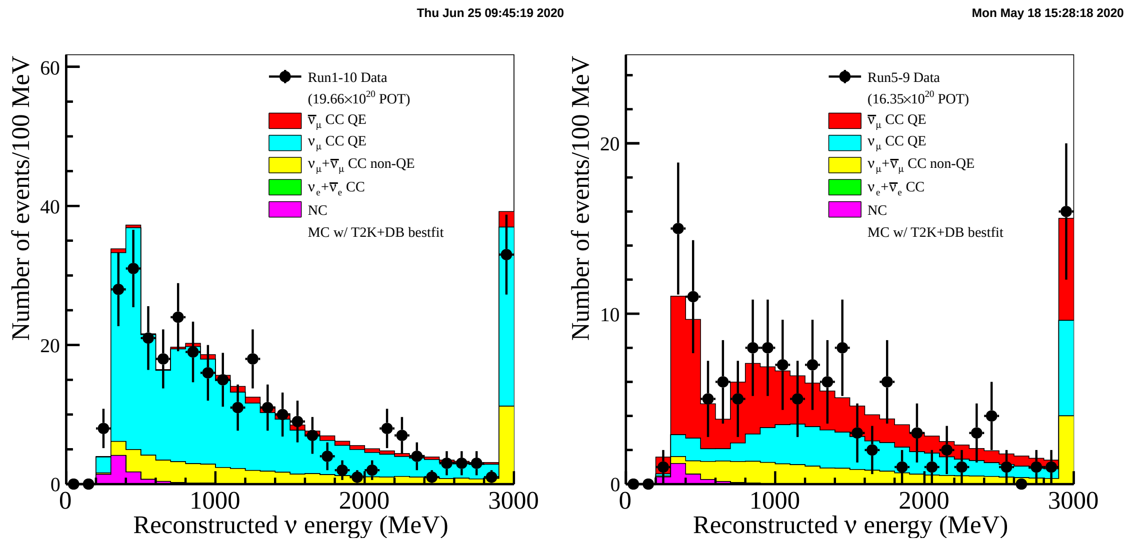

Figure 1.4: Reconstructed neutrino energy distributions at the T2K far detector for ν_μ enriched (*left*) and $\bar{\nu}_\mu$ enriched (*right*) samples, showing a dip in the spectrum at approximately 0.6 GeV, where some of the neutrino flux has oscillated. Figure from [20].

anti-neutrinos, combined with its degeneracy with other effects from cross-section models, matter effect, and other systematic effects. In July of 2020 T2K became the first neutrino experiment to exclude a value of $\delta_{\mathrm{CP}} = 0$ at 3σ significance [21], greatly restricting the credible interval for the parameter for both mass hierarchies. Figure 1.5 shows the constraints on δ_{CP} together with two PMNS phases.

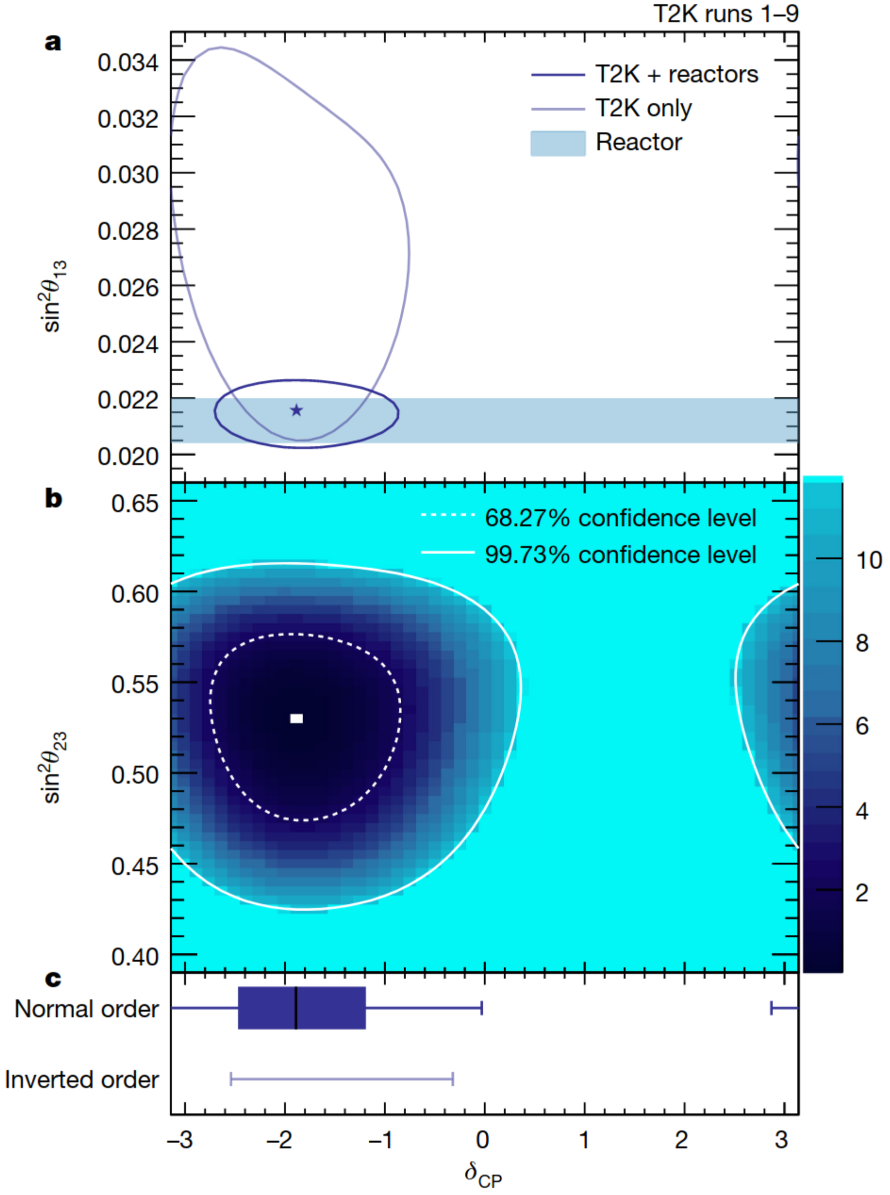

Figure 1.5: Joint constraints on δ_{CP} together with PMNS phases $\sin^2\theta_{13}$ (**a**) and $\sin^2\theta_{23}$ (**b**), with the effect of the neutrino mass hierarchy on the constraint shown in (**c**). The $\sin^2\theta_{13}$ result is combined with data from reactor-based neutrino oscillation experiments, as these provide a much stronger constraint on this parameter. Figure from [21].

Chapter 2

Neutrino-nucleus Interactions

In the previous chapter it has been shown that a solid understanding of the underlying neutrino interactions being observed is fundamental to a precise and robust oscillation analysis. In the context of future experiments, the uncertainties related to cross-section modelling — with the current generation of models — are expected to dominate the oscillation analysis of Hyper-Kamiokande [22]. This motivates both accurate measurements to understand the deficiencies in current models, alongside the development of improved models.

This chapter provides an overview of a number of charged-current neutrino interaction processes; quasi-elastic scattering, of most importance to the T2K oscillation analysis, resonant and coherent pion production, the interactions responsible for much of the signal in the analysis presented in this thesis, and multiple pion production and deep inelastic scattering, which produce much of the background events. Figure 2.1 shows the overall CC cross-section for neutrinos as a function of neutrino energy, with the contributions from each mode shown. Following the interactions, the effects of the nuclear medium on the particles produced are discussed, together with their effects in the context of cross-section and oscillation analyses.

2.1 Charged-current Quasi-elastic Interaction

The simplest charged-current process is charged-current quasi-elastic scattering (CCQE) on a free nucleon. The interaction is mediated by a charged W boson, with an incoming nucleon and neutrino being converted to a charged lepton and a nucleon with

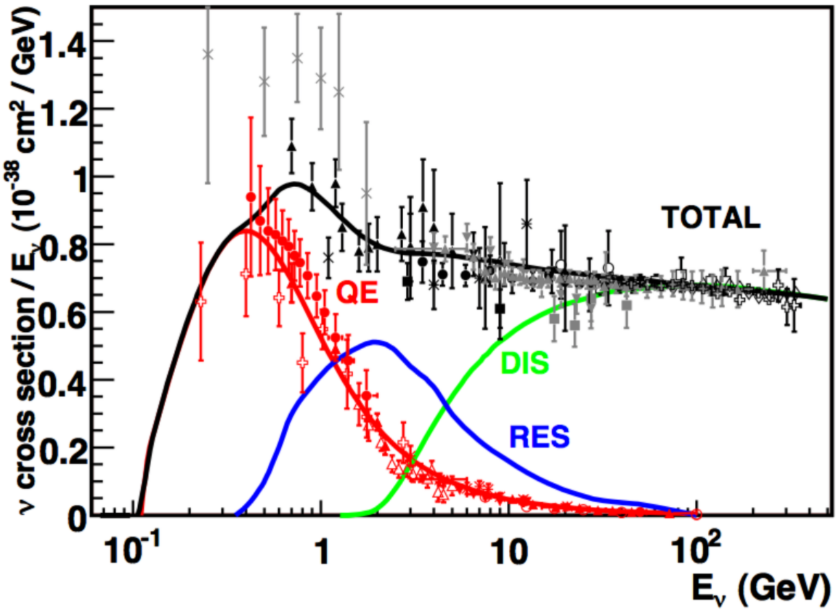

Figure 2.1: Total neutrino-nucleus charged-current scattering cross section as predicted by the NUANCE generator [23]. Shown is the cross section broken down by interaction mode together with the total cross-section, with available data (up to 2012) plotted for the total and quasi-elastic components. Figure from [24]

flipped iso-spin,

$$\begin{aligned} \nu_\alpha + n &\to l^- + p \\ \bar{\nu}_\alpha + p &\to l^+ + n \end{aligned} \tag{2.1}$$

where $\alpha = e, \mu, \tau$. Feynman diagrams for both ν and $\bar{\nu}$ CCQE scattering are shown in figure 2.2. As a nucleus is comprised of three quarks, the analytical computation of this cross section is not feasible. As an alternative to analytical calculation, the cross section is generally parameterised through the use of 'form factors', which can be determined through measuring various processes to constrain the free parameters in the model. The Llewellyn-Smith formalism [25] is commonly used in modern neutrino scattering theory, and parametrises the cross section as

$$\frac{d\sigma}{dQ^2} = \frac{m_N^2 G_F^2 |V_{ud}|^2}{8\pi E_\nu^2} \left[A(Q^2) \pm B(Q^2)\frac{s-u}{m_N^2} + C(Q^2)\frac{(s-u)^2}{m_N^4} \right], \tag{2.2}$$

where G_F is the Fermi coupling constant, m_N is the nucleon mass, V_{ud} is the up-down quark mixing element from the CKM matrix, E_ν is the neutrino energy, s and u are the Mandelstam variables, and the ± denotes neutrinos and antineutrinos

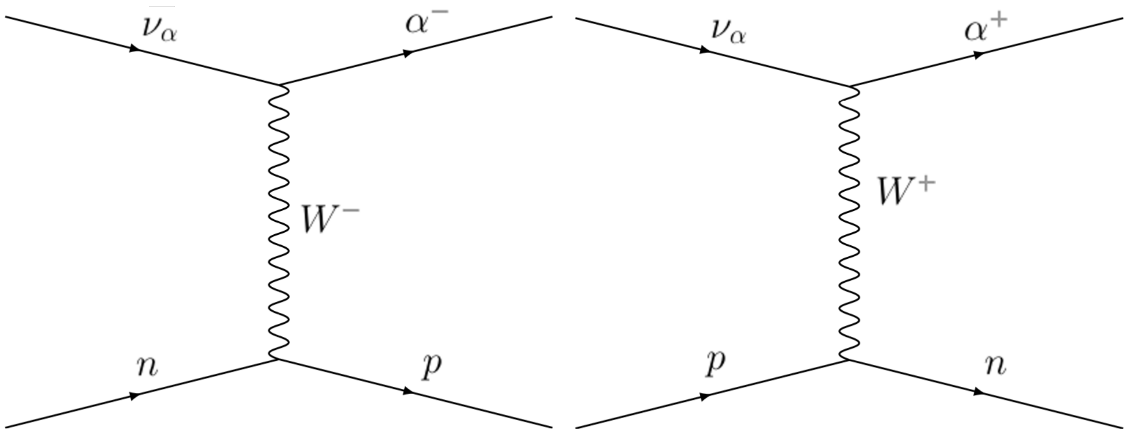

Figure 2.2: Feynman diagrams of CCQE scattering on free nuclei, where $\alpha = e, \mu, \tau$.

respectively. A, B, C are functions of the squared energy transfer Q^2, and the four form factors, F_1, F_2, F_A, and F_P. These form factors encompass the structure of the nucleon. F_1 and F_2 are vector form factors, encompassing the vector component of the weak interaction, while the axial and pseudoscalar components are described by F_A and F_P respectively. The vector form factors are present in all weak interactions, and as such can be constrained by electron-nucleon scattering experiments. The axial and pseudoscalar form factors in this formalism are only present for neutrino interactions, and so cannot be constrained by experiments using other weak interactions.

A similar formalism can be constructed for neutral-current elastic scattering, which differs primarily in the propagator, with neutral-current mediated by the neutral Z-boson. The resulting parameterisation differs primarily in the form factors employed, adding dependence on the strange quark contribution to the nucleon [26].

2.2 Charged-current Resonant Pion Production

If a neutrino nucleon interaction has centre of mass energy equal to or greater than the mass of a Δ baryon, a resonant state can be excited in the nucleon. For neutrino energies from about 1.5 to a few GeV, these resonances become especially relevant, as the cross sections for other processes are relatively low. These resonances decay extremely quickly — usually before the baryon has had time to exit a nucleus — into a nucleon and a pion, and can be induced through both charged and neutral-current scattering. The primary resonance for these processes is the Δ baryon, but various other resonant states with masses between 1 and 2 GeV can produce similar final states. A proper resonant pion production model must include not only the

ν_l	$\bar{\nu}_l$
$\nu_l + p \rightarrow l^- + p + \pi^+$	$\bar{\nu}_l + p \rightarrow l^+ + p + \pi^-$
$\nu_l + n \rightarrow l^- + n + \pi^+$	$\bar{\nu}_l + n \rightarrow l^+ + n + \pi^-$
$\nu_l + n \rightarrow l^- + p + \pi^0$	$\bar{\nu}_l + p \rightarrow l^+ + n + \pi^0$

Table 2.1: Primary pion production modes through Δ baryon decay.

Δ contributions, but all other relevant resonances, together with any interference effects between them.

As charged current interactions for ν_μ are mediated by W^+ and $\bar{\nu}_\mu$ by W^-, the possible pion production channels can be determined considering the possible permutations of target nucleon and decay pion, keeping charge conservation in mind. Table 2.1 shows the three primary pion production modes for both ν and $\bar{\nu}$ CC resonant interactions through the delta baryon.

The most commonly used theoretical model for this process is the Rein-Sehgal model [27], developed in the late 1970s to describe the single pion production data of the day. The model considers each possible final state as a superposition of all possible resonances at $W \leq 2$ GeV that can contribute to that final state, together with the effects of interference between resonances. The model also includes a non-resonant component, in which pions are produced without an intermediating resonance. To account for the extended nuclear target, nuclear form factors are employed in a manner similar to the Llewellyn-Smith model, although the form factors used differ. The Rein-Sehgal model introduces three constants that cannot be easily constrained experimentally; M_A^{RES}, a normalisation similar to the M_A in CCQE models, a normalisation constant for the Graczyk-Sobczyk form factor at $Q^2 = 0$, C_A^5, and a further normalisation constant of $I_{\frac{1}{2}}$ for the non-resonant background.

A more modern interpretation of this model is the MK model [28]. While based on the same underlying treatment used by Rein and Sehgal, this model includes effects such as the non-zero outgoing lepton mass for CC events — outgoing lepton mass is 0 for all NC events — which are neglected for simplicity in the Rein-Sehgal model. The MK model also includes updated helicity amplitudes for the resonances in the resonance region, together with the non-resonant background model described in [29]. The interference between the resonant and non-resonant contributions to the final state is considered more thoroughly. The MK model provides improved agreement with bubble chamber and neutrino scattering data over the Rein-Sehgal implementation in the NEUT generator.

2.3 Charged-current Coherent Pion Production

Coherent pion production occurs when a neutrino interacts with a nucleus as a singular entity. Such scattering is only possible at very low energy transfers, resulting in a very forward-going (relative to the neutrino direction) lepton. The models used in modern generators are largely based on Adler's partially conserved axial current (PCAC) theorem, which allows the coherent neutrino scattering cross section to be related to the pion-nucleus scattering cross section under the assumption of a massless lepton and very forward ($Q^2 \approx 0$) scattering angle. The Rein-Sehgal [30] coherent scattering model is based upon these assumptions, but includes a propagator term to account for the massive leptons present in CC interactions.

A more modern model of the same form is that of Berger-Sehgal [31], which improves upon the Rein-Sehgal implementation primarily by using more modern data for the total and differential pion-carbon scattering cross section. Relative to the Rein-Sehgal model, Berger-Sehgal predicts a much lower cross section at lower energy transfers, and provides a much better description of available CC coherent pion production data.

2.4 Deep Inelastic Scattering

As the energy of a neutrino increases beyond approximately 5 GeV, it is possible for the neutrino to resolve individual quarks within a nucleus. This scattering off the constituent quarks is known as Deep Inelastic Scattering (DIS). In CC DIS scattering, the interaction produces a charged lepton and a hadronic jet, with the target nucleon being destroyed in the process.

DIS modelling is done primarily by using a set of 'structure functions', which are form factors describing resonances that can be excited at lower energies, and a set for the higher energy continuum. These structure functions are expressed through Parton Distribution Functions that describe the nucleon substructure. At higher energies, only the continuum contribution are significant, and as there is a wealth of high energy neutrino scattering data available, models in this region are relatively well constrained [34]. At lower energies, particularly where nuclear resonances are still excited, there is large uncertainty in these models, as isolating these channels in a neutrino scattering experiment is difficult.

2.5 Shallow Inelastic Scattering and Multi-Pion Production

A particularly challenging region to model is the transition between resonant pion production and deep inelastic scattering. In this region, higher energy resonances can be excited from a nucleon that may decay into two or more pions. These resonances, combined with the low-energy tail of the underlying DIS continuum present a very challenging region to model, where a final state can have contributions from many different processes with significant interference. Notably, in this region, the primary two neutrino event generators used in this thesis have discontinuities in the cross section, resulting from where the resonance region is cut off in invariant hadronic mass. Within the rest of this thesis, any events where resonances produce multiple pions or other shallow scattering occurs will be referred to as DIS events for the purposes of simplicity.

2.6 Nucleon and Final State Interactions

Excepting coherent scattering, the aforementioned neutrino interactions are considered to occur on a single free nucleon. As modern experiments use nuclear targets, a model predicting the interactions of neutrinos with a nucleon must be combined with a model of the nuclear medium inside the nucleus.

The effects of nuclear models on the neutrino interaction can broadly be considered as two different effects, initial state effects and final state effects. As nucleons are bound inside a nucleus, the removal of a nucleon from the nucleus will require overcoming the binding energy, and the nucleon will have some non-zero initial momentum. Nucleons also interact with one another inside the nucleus, and as such neutrino interactions can occur on a pair (or more) of interacting nuclei, further altering the kinematics relative to the single free nucleon case. After the neutrino interaction, there is potential for final-state interactions (FSI) to occur, where outgoing hadrons re-interact inside the target nucleus before escaping, altering the outgoing kinematics or particle multiplicities.

2.6.1 Fermi Motion

The nucleons in the initial nuclear state will have a momentum distribution, known as Fermi motion, generally assumed to be isotropic. Since the initial momentum of the nucleon on which a neutrino interacted is not fixed, each neutrino event is affected by a boost in the lab-frame due to the nucleon momenta. As the momentum of a given nucleon during the interaction is unknowable, this results in a fundamental smearing of the outgoing kinematics of an event. As the number of nucleons in a nucleus increase, the difficulty of calculating this initial momentum distribution increases rapidly. Overall, this results in a significant smearing of the outgoing momentum and angle distributions of a neutrino interaction that is fundamental to scattering on heavy nuclei, prior to any smearing from detector effects.

The most common models used in neutrino interaction generators treat the nucleons inside a nucleus as a gas of non-interacting fermions inside some nuclear potential where they occupy momentum states from the lowest available up to some maximum momentum, the Fermi momentum. The most basic such model is the Relativistic Fermi Gas (RFG) model, which assumes the nuclear potential is constant for all nucleons, which results in a sharp cutoff in the nucleon momentum distribution. The Local Fermi Gas (LFG) model instead assumes that the potential is some function of the nuclear radius, and the local nuclear density at a given radius. A more sophisticated model for the nucleon momentum distribution is the Spectral Function (SF) [33], which considers shells of nucleons, where two and three body short range interactions can modify the orbitals. Both the NEUT and GENIE [35, 36] generators which are used in this thesis include the ability to use these models. An more detailed overview of these models and their implementations in event generators can be found in [32].

2.6.2 Final State Interactions

Neutrino interactions can occur anywhere within the target nucleus, and particles ejected or produced must propagate outward through the nuclear medium in order to be detected. Hadrons can re-interact in a number of ways, resulting in altered kinematics, hadrons being absorbed, or additional hadrons ejected. FSI presents a great challenge to neutrino scattering experiments using nuclear targets, as it is capable of decoupling the observed final state of an interaction from the fundamental interaction that occurred; for example in a resonant pion interaction the ejected pion may be absorbed within the nucleus, resulting in a final state identical to that expected from a CCQE interaction.

Modelling of FSI is performed by neutrino interaction generators, with the majority performing a cascade simulation. Such a simulation is performed by taking the particles produced in a neutrino interaction at some point in a nucleus, and propagating them through the nuclear medium in small steps. Step size is chosen usually as a fraction of the mean free path, and the probability of each possible re-interaction for each particle is computed based on the local nuclear density. Using this probability for re-interaction, it is decided whether the particle should re-interact, and if so, this interaction is simulated. This process continues until all particles have been absorbed or have left the nucleus. The probabilities of and models behind the hadronic re-interactions are tuned to external hadron scattering data to improve predictions. The details of FSI implementation for each neutrino event generator discussed in this thesis are presented in [35, 36, 23, 37] for NEUT, GENIE, NUANCE, and NuWro respectively. GiBUU is unique in its treatment of FSI; GiBUU solves a series of coupled hadronic transport equations for the particles [38].

2.6.3 Multi-Nucleon Correlations

Where a neutrino interacts with two or more correlated nucleons inside a nucleus, the event dynamics and kinematics can be greatly altered. At present, no models attempt to describe multi nucleon effects for the pion production processes relevant to this thesis, but descriptions of these effects for quasi-elastic scattering have been developed, such as that of Nieves et al. [39]. The Nieves model considers only correlations between pairs of nucleons — two-particle two-hole or 2p2h — where a virtual pion or other propagator connects the interaction vertex to another nucleon. The inclusion of this effect is of great significance in T2K, as 2p2h events exhibit significantly different kinematics from CCQE events while appearing almost identical in T2K's detectors. As such, these events form irreducible background that can significantly bias neutrino energy reconstruction, and thus oscillation measurements.

Chapter 3

The T2K Experiment

Tōkai to Kamioka (T2K) is a long-baseline neutrino oscillation experiment designed to measure the neutrino oscillation parameters. T2K produces a reasonably pure beam of muon (anti-)neutrinos in the J-PARC laboratory, which is measured by multiple near-detectors at a distance of 280 m, and then at 295 km from the beam source. The physics detector at 295 km from the beam is Super-Kamiokande, a 50 kt water Cherenkov detector located underground beneath Mt. Ikeno-Yama in Gifu Prefecture, Japan [40]. The beam is pointed 2.5° off-axis from the detectors, providing a much narrower neutrino energy spectrum in the beam, with an peak energy of 600 MeV. The difference in neutrino flavour composition and their energy spectra at each of these distances can then be used to determine the relative strength of the oscillations for a given $\frac{E}{L}$. T2K measures four channels of neutrino oscillation; the disappearance of muon neutrinos and anti-neutrinos, and the appearance of electron neutrinos and anti-neutrinos.

T2K's original design goals were to provide precision measurements of θ_{23}, Δm^2_{32}, together with improved constraints on other oscillation parameters. As T2K measures neutrino oscillation for both neutrinos and anti-neutrinos, the difference between these channels can be used to determine the charge-parity violating phase δ_{CP} [41, 42].

Alongside providing measurements of oscillation parameters, T2K regularly publishes cross section measurements from each of its detectors. Recent cross-section results from T2K include measurement of the inclusive ν_e and $\bar{\nu}_e$ cross-sections [43], and the double-differential ν_μ-CC0π cross-sections on both carbon and oxygen [44].

The remainder of this chapter provides a detailed overview of the ND280 tracker, and the T2K neutrino beam and flux prediction, both of which are used heavily in the analysis presented in this thesis.

3.1 The T2K Neutrino Beam

The T2K neutrino beam [46] is produced in the J-PARC laboratory in Tōkai-mura on the coast of Ibaraki prefecture, Japan. The beam is designed to produce a beam of either muon neutrinos or anti-neutrinos that is of high purity and with a well understood flux. The beamline is configured such that the neutrino beam used for oscillation measurements has a narrow energy distribution that is peaked at the oscillation maximum. The following sections provide an overview of the beamline itself, the manner in which the neutrino flux is predicted, and the effects of the off-axis beam configuration.

3.1.1 Neutrino Beamline

The J-PARC Main Ring proton accelerator [47] produces a 30 GeV beam of protons for use in a variety of applications, neutrino beam production among them. The T2K neutrino beamline is comprised of two sections, the so-called primary and secondary beamlines. From the Main Ring, protons are extracted and sent to the primary beamline where they impinge on a graphite target. The interactions between the protons and the target produce a cascade of hadrons, which are focused by magnetic horns to select either positive of negative hadrons, which then enter a decay volume where decaying mesons produce neutrinos. An overview of the beamline is shown in figure 3.1.

In the primary beamline, the extracted protons are prepared by a series of focusing magnets to both guide the protons to the target, and to ensure the beam is correctly oriented. At this stage, the beam is very closely monitored, and measurements of the beam and instrumentation are made during running to ensure that the beam can be accurately simulated and is running within specification.

In the secondary beamline, the focused protons enter the target station and impinge on the target. Particles produced in the target immediately enter the first magnetic horn, and subsequent two horns spaced slightly further apart to allow dispersed particles to leave the beam. These magnetic horns focus particles of one charge while causing the opposite charge to diverge, with the ultimate goal of selecting either positive or negative pions. The two modes used in the T2K beamline are Forward Horn-Current (FHC) and Reverse Horn-Current (RHC), preferentially selecting π^+ and π^- respectively. These focused hadrons then enter a 96 m long, helium filled decay volume. At the end of the decay volume is the beam dump, where most particles remaining — produced hadrons and their decay products —

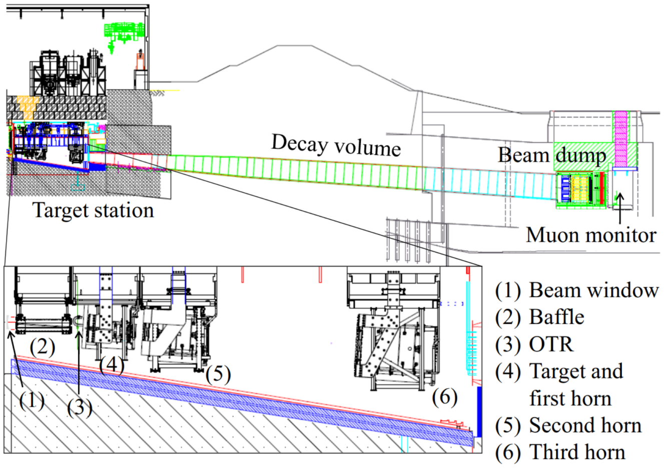

Figure 3.1: A labeled side view of the secondary beamline. Figure from [46].

are stopped, with neutrinos passing through towards the physics detectors. A small fraction of muons from the decay volume have sufficient kinetic energy (≥ 5 GeV) to pass through the beam dump, which are then detected by the muon monitor (MUMON). These muons provide a precise, real-time measurement of the beam direction and profile.

A detailed description of the entire beamline setup can be found in [46], with the MUMON detector described in [48].

3.1.2 Neutrino Production via Meson Decay

In the beamline decay volume, pions and kaons are allowed to decay to produce neutrinos. The primary modes for neutrino production in FHC are as follows:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \qquad \text{and} \qquad K^+ \rightarrow \mu^+ + \nu_\mu$$

Similarly, the primary decay modes in RHC are:

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \qquad \text{and} \qquad K^- \rightarrow \mu^- + \bar{\nu}_\mu$$

Unfortunately, the charge selection performed by the magnetic horns is imperfect, and some wrong-sign hadrons will enter the decay volume. This results in some level of wrong-sign neutrino — $\bar{\nu}_\mu$ in FHC, ν_μ in RHC — contamination in the beam. This wrong-sign contamination is significant in RHC, where a much larger ratio of wrong-sign mesons enter the decay volume.

The decay modes discussed above do not constitute 100% of the possible decays, and not all possible decays result in the same neutrino flavour being produced. The neutrino beam will have a relatively small ν_e and $\bar{\nu}_e$ contribution from decay modes such as

$$\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu \qquad \text{and} \qquad K^+ \to \pi^0 + e^+ + \nu_e$$

in FHC, and

$$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu \qquad \text{and} \qquad K^- \to \pi^0 + e^- + \bar{\nu}_e$$

in RHC. The overall production of ν_e and $\bar{\nu}_e$ is relatively small in both FHC and RHC due to the low branching fractions of these processes combined with the relatively low wrong-sign hadron contamination in the decay volume. For oscillation measurements however, such a low contamination obviously needs to be very well understood to search for $\nu_\mu \to \nu_e$ oscillations, and so such decay modes must be accurately accounted for.

3.1.3 Flux Prediction

Flux simulations of the T2K beam begin with simulating the interactions of the beam protons with the graphite target. FLUKA 2011 [49] is used to simulate the proton interactions within the target and baffles, which produce the secondary particles responsible for neutrino production. From this simulation, the information of each secondary particle — including kinematics and vertex position — is stored. The following stage in the chain is to simulate the propagation of these secondary particles through the secondary beamline. This simulation is performed with JNUBEAM, a Monte Carlo simulation package built upon Geant3 [51], where subsequent interactions of these particles are simulated with GCALOR [50] where they occur.

From the resulting simulation, the distribution of decay neutrinos is extracted, and neutrinos from this simulation are required to point towards a point within a detector plane (SK, ND280, etc.) to avoid simulating neutrinos that will not cross any detector. Simulated neutrinos are subsequently saved along with the likelihood of them travelling in a given direction, along with their energy when travelling that direction. This provides the basis for the final flux prediction for a given detector. The current T2K neutrino flux predictions are shown for forward and reverse horn currents respectively in figure 3.2, where the wrong-sign contamination can be seen to be far higher in reverse horn current than for forward horn current. This is due to a large excess of positive mesons produced in proton-nucleon collisions, with a

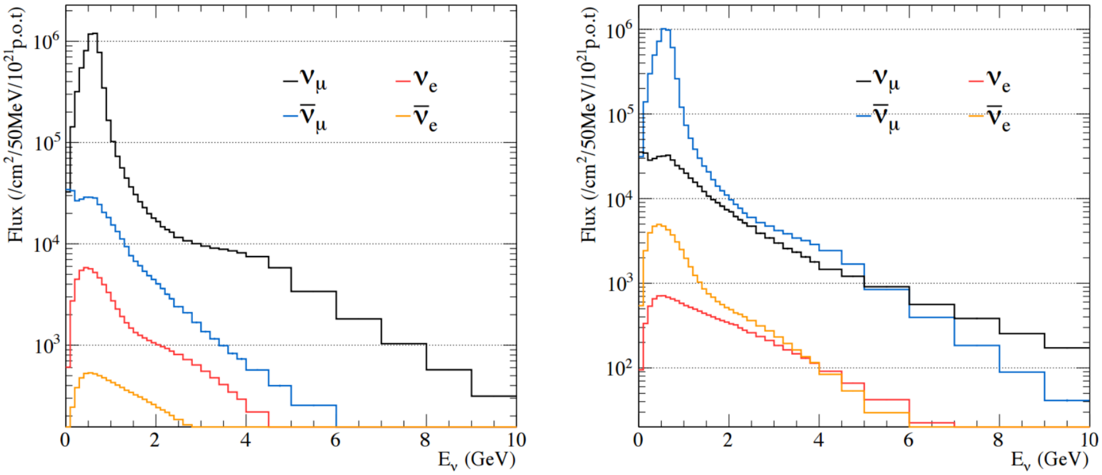

Figure 3.2: Neutrino flux predictions for forward (*left*) and reverse (*right*) horn currents split by neutrino flavour. The wrong-sign contamination is seen to be much lower in forward horn current. Figure from [46].

small but significant fraction entering the decay volume in RHC operation, and thus contributing a significant neutrino background at higher E_ν.

To improve the accuracy of the flux simulation, various quantities are measured while running, including measurements of the magnetic horn current and measurements of the incident proton beam. Further beam measurements are taken by the Muon Monitor (MUMON), and the on-axis neutrino detector INGRID. Measuring these quantities allows tuning of the flux simulation, reducing the uncertainty of the prediction. One of the primary sources of uncertainty in the flux simulation is the abundance of low energy strong interactions, both in the target and in the secondary beamline. To improve the modelling of these interactions, data from the NA61/SHINE experiment at CERN is used.

NA61/SHINE is an experiment dedicated to studying hadron production using a 31 GeV proton beam and two graphite targets; one a replica of the T2K beam target, the other a similar but thinner target [52]. The results from this experiment are used to tune the hadron production models in the flux simulation pipeline to better constrain the final flux.

Other external datasets are also used to constrain other parts of the simulation. Various external datasets for hadron production are used to constrain areas of phase-space not covered by the NA61/SHINE data. Data from the HARP experiment [53] are used to constrain the modelling of pion secondary interactions which are important in the secondary beamline simulation.

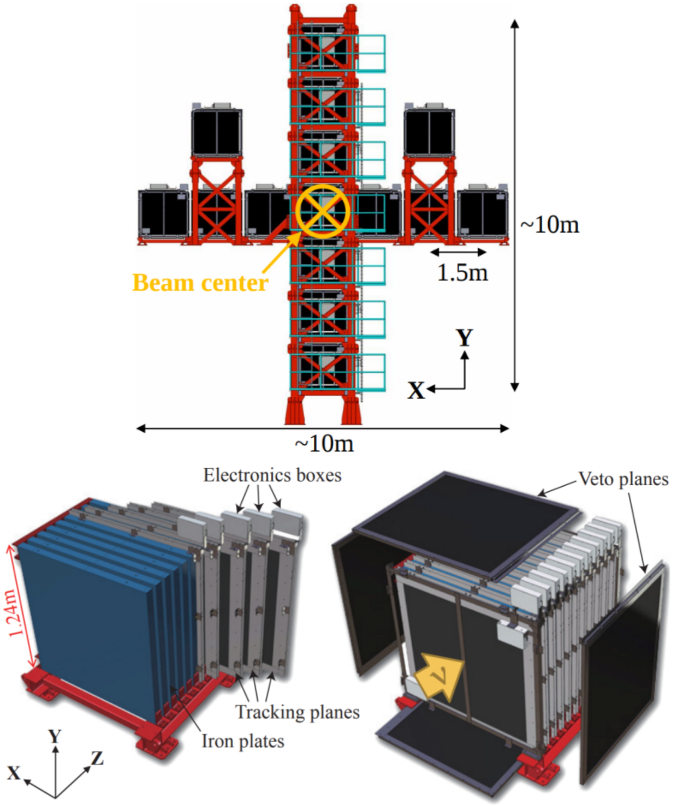

Figure 3.3: Configuration of the INGRID modules relative to the neutrino beam, with the detector structure illustrated in the bottom panels.

3.1.4 The Interactive Neutrino GRID Detector

The Interactive Neutrino GRID (INGRID) detector is located 280 m directly downstream of the graphite target, and is designed to provide precise, real-time measurements of the neutrino beam centre, luminosity, and profile. INGRID comprised 16 identical detector modules, each consisting of alternating layers of scintillator bars and iron, where the iron provides a large target mass to maximise neutrino interaction rates, and the alternating scintillating bars allow tracking of charged particles. The readout of the scintillating bars is largely the same as for the FGDs and ECals which are described in detail in sections 3.2.2 and 3.2.3.

Two special purpose modules — the proton module and the water module — are placed directly in line with the centre of the INGRID cross. The proton module is a fully active scintillator design where the iron is absent, and the water module contains water-filled panels in place of the iron. These modules allow measurement of the interaction rates on pure scintillator together with a simultaneous water and scintillator measurement, from which water cross-section data can be extracted.

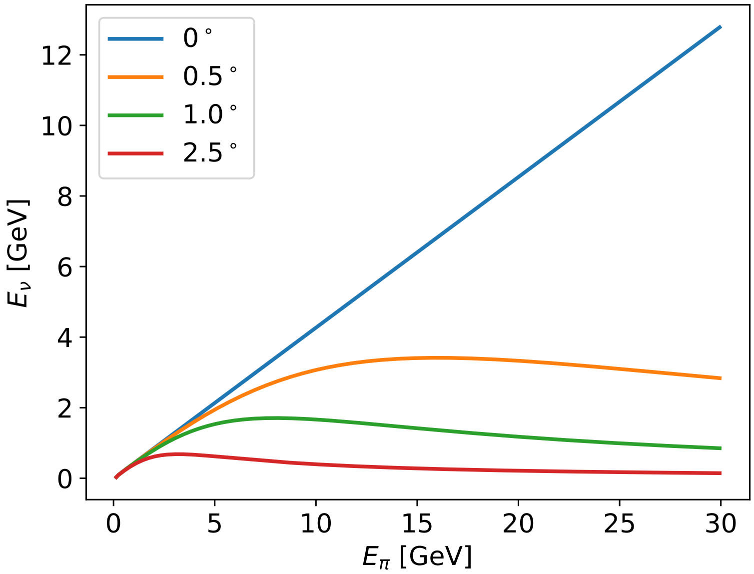

Figure 3.4: Outgoing neutrino energy at a given angle as a function of pion energy, for various angles. As the angle increases it can be seen that the neutrino energy becomes less dependent on the initial pion energy.

3.1.5 Off-Axis Beam

NA280 and SK are both positioned at 2.5° away from the beam axis. While the expected neutrino flux at this angle is significantly lower than that directly in the beam axis, the beam has a much narrower energy spectrum away from the axis, which can be beneficial for an oscillation experiment. The effects of off-axis angle on the beam energy can be well understood by considering the decay of pions in the secondary beamline, as pions are responsible for the majority of neutrinos produced. In the lab frame, the energy of the outgoing neutrino, E_ν, at an angle θ, can be expressed as a function of the parent pion energy and momentum, E_π and p_π , and the masses of the pion and decay muon, m_π and m_μ:

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2\left(E_\pi - p_\pi \cos(\theta)\right)} \tag{3.1}$$

Figure 3.4 shows the resulting neutrino energy at a given off-axis angle from the decay of a pion of a given energy. As the off-axis angle is increased, the resulting neutrino energy becomes increasingly decoupled from the parent pion energy, and as such makes a neutrino beam with significantly reduced energy spread.

Figure 3.5 shows the neutrino energy spectrum for different off-axis angles to-

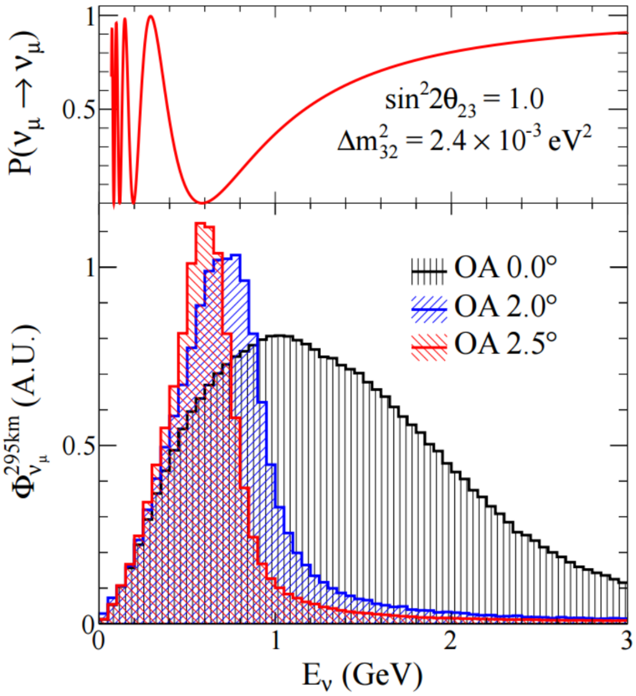

Figure 3.5: Neutrino flux prediction (arbitrarily normalised) for three off-axis angles as a function of neutrino energy. The value of 2.5° is used by T2K for oscillation measurements as it provides a narrow beam with a peak at the first oscillation maximum. Figure from [46].

gether with the oscillation probability. The result at 2.5° off-axis is a narrow-band neutrino beam peaked at the oscillation maximum, greatly simplifying the understanding of the oscillated spectrum.

3.2 The ND280 Detector

The Near Detector at 280 meters (ND280) is located 280 m downstream of the beam target, centered 2.5° away from the beam axis. The goal of the ND280 is to provide precise measurements of neutrino-nucleus interactions from the same beam composition that is seen at SK. An exploded view of the ND280 with its primary subdetectors labelled is shown in figure 3.6. The ND280 comprises five sub-detectors, the Time Projection Chambers (TPCs), the Fine-Grained Detectors (FGDs), the Electromagnetic Calorimeters (ECals), the Side-Muon Range Detectors (SMRDs), and the Pi-Zero Detector (PØD). The first three of these sub-detectors will be discussed in detail in the following subsections, as the event selection presented in chapter 5 makes use of these subdetectors exclusively. The function and operation of the PØD and SMRD can be found in [54] and [55] respectively.

The ND280 is a magnetised detector, being placed inside a magnet from the former UA1 experiment [56], which provides a magnetic field of 0.2 T in the detector, allowing momentum determination from track curvature, as well as charge determination from the direction of curvature.

3.2.1 Time Projection Chambers

ND280 is equipped with three gaseous argon Time Projection Chambers (TPCs) [58], designed to accurately measure the curvature of a particle in the magnetic field together with the energy loss as the particle traverses the TPC. Each TPC comprises an inner box containing the argon gas and a copper cathode plane along the centre of the box, and an outer box containing the inner box within a CO_2 atmosphere. Readout planes are located at the side of the TPC inner box, in the same plane as the central cathode. A cutaway view of a TPC is shown in figure 3.7. The electric field is oriented in the same direction as the magnetic field, and has a field strength of approximately 280 V/cm.

As charged particles pass through the gaseous argon ionise the gas, releasing electrons which drift along the electric field lines, with their path unaffected by the magnetic field as they move along the field lines. The readout plane measures

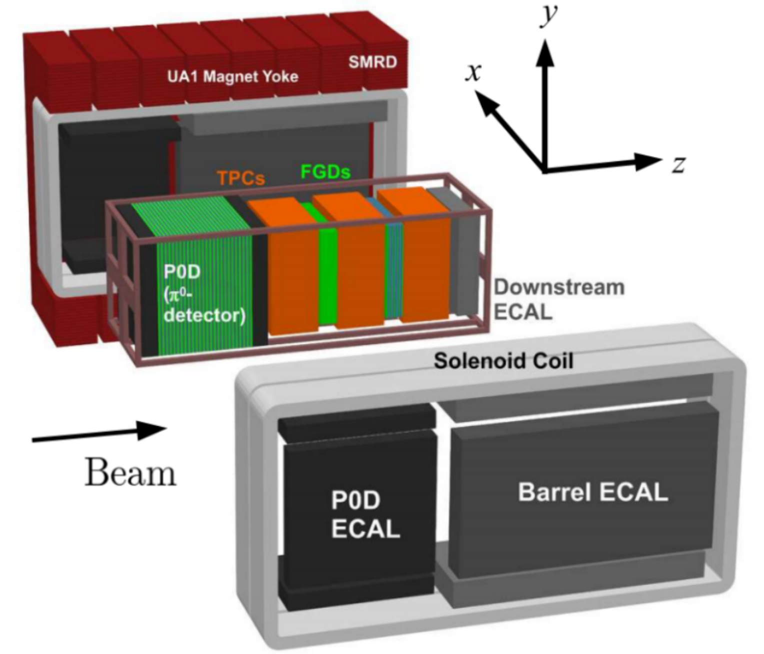

Figure 3.6: Cutaway view of the ND280 detector. Figure taken from [57].

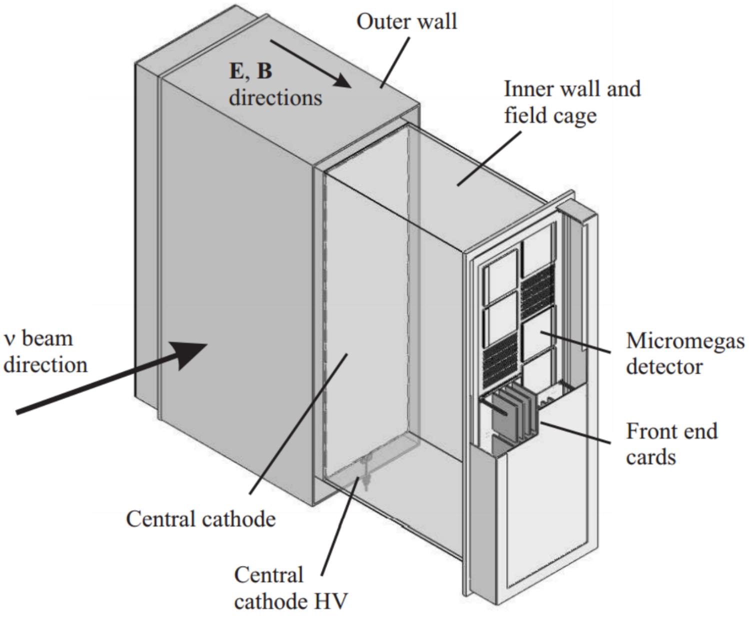

Figure 3.7: Cutaway view of an ND280 TPC, with key components marked. Figure taken from [57].

both charge and arrival time are measured at the readout plane by one of twelve 'micromegas' modules, with an overall spatial resolution of approximately 0.7 mm in the readout plane. At the readout plane electrons enter an amplification region with field strength 27 kV/cm to induce each electron to free more electrons from the gas, resulting in a large gain of order 1500. The micromegas output is digitised by 'AFTER' integrated circuit boards, with the digitised signal being sent to a front-end mezzanine board where it is condensed and prepared for storage.

The high spatial resolution afforded by the TPC, along with energy deposition measurements allow for very accurate characterisation of the energy loss and momentum of particles traversing the gas. How this information is used in particle identification and momentum determination are discussed in detail in section 5.1.3.

3.2.2 Fine Grained Detectors

The ND280 contains two FGDs [59] — FGD1 and FGD2 — which differ only in the presence of the passive water layers present in FGD2. The FGDs are designed to a massive target for neutrino interactions, while also providing accurate 3D tracking and vertex reconstruction. Each FGD has a TPC located both up and downstream to allow accurate measurement of most particles produced inside the FGDs. For particles that do not enter a TPC, the FGDs should be capable of constraining their properties to a reasonable degree. These design goals are achieved by constructing a detector where the target mass is itself an active detector medium. The geometry and properties of FGD1 are described first, followed by the differences in FGD2.

FGD1

FGD1 is a fully active detector, comprised of $5,760$ polystyrene scintillator bars measuring 9.61 mm × 9.61 mm × 18864.3 mm each, coated with a reflective TiO_2 coating. The bars are arranged into layers, where each layer comprises a plane of 30 parallel bars. These layers are arranged such that the bars in each layer are perpendicular to those in both neighbouring layers and the incoming neutrino beam direction. Every two neighbouring layers comprise an XY module.

Scintillation light produced as charged particles cross a bar is collected at both ends by a wavelength shifting optical fibre. The end of one fibre is connected to a Multi-Pixel Photon Counter (MPPC) [60] which measures the light yield from the bar together with the arrival time. The end of the other fibre is coated with a reflective aluminium coat to reflect light back, further improving light collection efficiency. The light collected by each MPPC is measured by counting the number of photoelectrons observed. The MPPC signal is then amplified with a gain of $\approx 5 \times 10^5$, and digitised by the same AFTER integrated circuit board used in the TPC, discussed in section 3.2.1. The entire FGD along with the MPPCs and fibre connections are contained within a dark box to prevent light from entering, to keep noise levels in the MPPCs low.

The alternating-layer design allows effective 3-dimensional track reconstruction when particles cross multiple layers, as each layer provides a strong positional constraint — to within the width of a single bar — in two dimensions. For particles in the FGD at a large angle from the neutrino beam axis, the situation can become somewhat complicated, as particles can travel a relatively long distance inside a single bar, resulting is tracks with very uneven energy deposits that are difficult to reconstruct. This results in a strong decrease in angular resolution as the angle away from the beam axis increases, which can be seen in figure 5.15 in section 5.4.

Due to the short overall length (~ 30 cm) of the FGD, particles will not usually leave a track that is sufficiently curved to establish its momentum. For particles that do not enter one of the TPCs, momentum can be determined by the distance travelled. This so-called 'momentum by range' depends on the species of the particle observed, and values are calculated for both the proton and muon hypotheses. The calculation is performed by summing the expected energy loss along the whole track in small sections, as the energy loss curves for both of these particles in the FGD scintillator are well known. A modified momentum by range calculation for the analysis presented in this thesis is detailed in section 5.4.1.

FGD2

FGD2 is composed of the same scintillating bars layered in alternating directions as FGD1, but between each set of XY modules is placed a layer of pure water. FGD2 comprises a total of 7 XY modules with 6 water panels interspersed. The water panels are constructed from corrugated polycarbonate panels of thickness 25 mm, connected to a water circulation system to allow for draining and filling of individual panels, as well as purification of the water. While the walls of the panels are thin, the presence of internal internal walls improves the structural integrity. The water panel is sized such that a single panel fits in the same volume that a single XY module would occupy, allowing the overall geometry of FGD2 to be kept the same as FGD1, which aids calibration and reconstruction efforts.

In order to make the overall elemental composition of the water panels as close to a mixture of pure water and polystyrene as possible, a sheet of polypropylene (CH_2) is glued to both ends of the water panel [59]. This ensures that the total elemental composition of the water panel matches, as closely as possible, an ideal composition of polystyrene and water, such that FGD2 and FGD1 differ only in the presence of water.

The water in FGD2 is continually circulated at a low rate. Water enters the panels through the bottom and is extracted at the top, with the entire water system maintained at a pressure slightly below atmospheric, such that any leaks that develop will result in air entering the system, and no water leaking into the detector. A small quantity of antimicrobial and anti-corrosive chemicals are added to the water to prevent biological growth.

During the draining of the water bags when shutting the ND280 down between runs, there have been two instances of water panels developing leaks. The first instance was before T2K Run 7 where the most upstream water panel developed a leak and was subsequently left empty. Most recently, while preparing for T2K Run 10, the second most downstream water panel developed a leak and had to be drained, remaining empty. The missing water is accounted for in detector simulations on a run-by-run basis. The result of this is a loss of one third of the water mass in FGD2 overall, but this also presents a small challenge for the analysis presented in this thesis, as the target mass changes over time. This effect is discussed further in the context of a cross section analysis in section 5.6.

3.2.3 Electromagnetic Calorimeter

The ND280 Electromagnetic Calorimeter (ECal) [61] is a lead-scintillator calorimeter that surrounds the primary sub-detectors of ND280. The ECal comprises three sections; the PØD ECal which surrounds the PØD, the Barrel ECal which surrounds the FGDs and TPCs, and the downstream ECal (DS ECal) directly downstream of the final TPC. The barrel ECal and DS ECal are often considered together and referred to as the 'Tracker ECal'. As the PØD ECal is designed to serve a somewhat different purpose to the tracker ECal and is not used in this analysis, this section describes only the tracker ECal in detail.

The design aim of the tracker ECal was to construct a subdetector to surround the main tracking region and allow measurement of any energy in electrons and photons that would otherwise escape the detector. The presence of lead results in a detector with a large number of radiation lengths, such that photons and electrons will not be able to exit the detector.

The tracker ECal is composed of alternating layers of scintillator bars and lead. Similarly to the FGD, each layer of scintillator bars is oriented perpendicular to the previous layer, enabling effective 3D track reconstruction. The scintillator bars used in the tracker ECal have cross-sectional dimensions 40×10 mm, with the overall bar length running the length of the subdetector in depending on orientation. Similar to the FGDs, the tracking efficiency for particles travelling at high angles to the scintillator planes suffers. However, for the detection of photons and electrons this is largely a non-issue, as the majority of energy deposit is not from the primary particle itself, but from secondary electrons and photons generated by bremsstrallung, etc.

The bars are read out by MPPCs connected by a wavelength shifting fibre in the same manner as for the FGD bars, and use the same reflective coating. Bars that are oriented parallel to the neutrino beam direction are read out to MPPCs at both ends due to the length of the bar, while bars perpendicular to the neutrino direction are read out on one side, with the opposite fibre being coated as in the FGD bars. The charge collected by an MPPC is collected by a capacitor to be read out every 580 ns. This time window is chosen such that beam events will arrive at roughly the centre of the window. The output from a group of 16 MPPCs is read by a single 'Trip-T' microprocessing chip. Groups of four Trip-T chips are housed in a 'Trip-T Front-end Board' (TFB), with the output of up to 48 TFBs being combined by a total of 12 Readout Merger Modules (RMMs). Data from the RMMs is then used in the local ECal reconstruction and the global detector reconstruction.

The PØD ECal serves a different function to the ECal in the tracker region. As the PØD itself is quite efficient at converting showers, the ECal is used primarily to determine whether a muon was present in the event, and contain any other particles such that minimal energy escapes the detector. As such, the PØD ECal is constructed slightly differently to the tracker ECal; The lead layers have thickness 4 cm, and only 6 scintillator layers are present. Moreover, all the scintillator bars are oriented parallel to the beam direction, as 3D track reconstruction is secondary to energy reconstruction.

3.2.4 Detector Physics Simulation

In order to understand the performance of event selections in terms of selection efficiency, purity, and the manner in which backgrounds enter a sample, it is vital to have a Monte Carlo simulation of the detector. This section details the manner in which such a simulation is prepared.

The ND280 event simulation chain begins with first determining the distribution of event vertices in the detector. The likelihood of an event at a given point is determined by the exact off-axis angle[1] and the composition of the detector at that point, with interactions being much more likely in high mass targets.

With the vertex placement determined, the NEUT [35] neutrino interaction simulation is run. NEUT predicts the cross sections of various neutrino-nucleus interactions, — and if produced inside a nucleus — propagates the particles through the nuclear medium using a cascade model to simulate final-state interactions (FSI). The probability of a specific reaction occurring at a vertex position is determined by the neutrino flux at that position, and the NEUT cross section predictions for that flux.

Following the prediction of particle production from neutrino-nucleus events, the behaviour of the emitted particles in the ND280 is to be simulated. All particles are propagated through ND280 using a GEANT4 [51] based simulation, simulating the appropriate particle energy loss and interactions with the detector medium. The output of this simulation is output to elecsim, ND280's electronics simulation package, where the detector response to each MC event is simulated. The results are stored in a format identical to real reconstructed data, but with all true information about the event stored alongside.

While this MC simulation with NEUT v5.4.1 is used as the default simulation, there are simulations produced in the same manner using different neutrino event generators. Genie [36] and NuWro [37] are used, along with older versions of NEUT to produce MC samples for validation purposes and comparison. Comparing the effects of a selection on multiple different models facilitates understanding to what level the performance of an event selection is susceptible to a specific model, and as such provides a way to reduce bias. The specific models used in each MC simulation, where relevant to this analysis, will be discussed in chapter 7.

3.2.5 ND280 Data Taking

T2K began data taking in 2009, and has been taking data regularly since. The data collected is generally quantified in terms of the number of 'protons on target' (POT); the number of protons incident on the neutrino beamline target. Since data taking began, the J-PARC proton beam has undergone various upgrades and improvements, increasing the beam power steadily over time. This has facilitated an increase in the rate of POT collection, and subsequently in the neutrino flux.

As this thesis deals only with $\bar{\nu}_\mu$ events, only T2K runs in RHC are used, those

---

[1] Due to the proximity to the beam source, ND280 spans a range of off-axis angles from 2° to almost 3°.

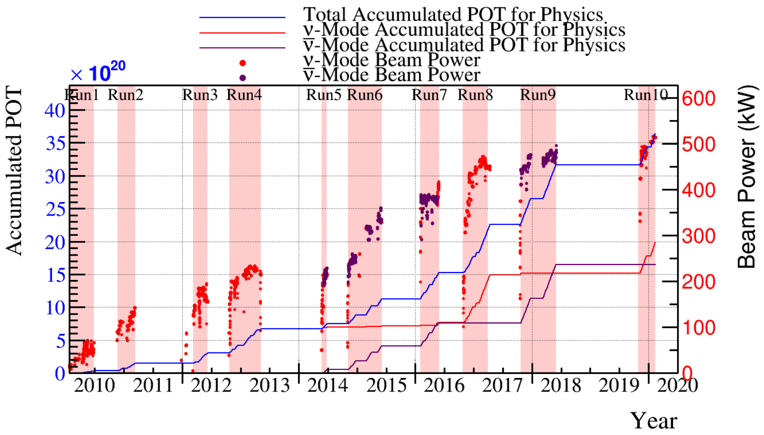

Figure 3.8: Accumulated POT over T2K's lifetime shown together with the J-PARC MR power.

being runs 5, 6, 7, and 9, with a total of 1.8×10^{21} POT, 9.4×10^{20} of which was taken with ND280 fully operational. While ND280 had no significant issues resulting in lost POT during runs 5, 6, and 7, during run 9 a number of issues surfaced resulting in an overall data taking efficiency of just over 40%.

Chapter 4

Analysis Strategy

4.1 Signal Definition

The signal is $\bar{\nu}_\mu$ Charged-Current (CC) events with a single, negatively charged pion in the final state, and no other mesons. Any number of nucleons are allowed in the final state, as the selection described in chapter 5 makes no attempt to reconstruct them. For reasons that will be discussed in detail in section 6.2, the phase space of of the signal is restricted such that only these events satisfying

- $p_{\mu^+} > 300\text{MeV}$
- $\cos\theta_{\mu^+} > 0.5$
- $p_{\pi^-} > 140\text{MeV}$
- $\cos\theta_{\pi^-} \in [-1, -0.35] \cup [0.35, 1]$

are considered to be signal. The angles θ_{μ^+} and θ_{π^-} are the angles that the outgoing μ^+ and π^- relative to the incident neutrino.

4.2 Flux-Integrated Cross Section Measurement Strategy

The aim of this work is the measurement of the single differential cross sections for $\bar{\nu}_\mu \text{CC}1\pi^-$ on water (H_20) and hydrocarbon (CH) as a function of the outgoing pion kinematics,

$$\frac{d\sigma^\tau}{dp_\pi} \qquad \text{and} \qquad \frac{d\sigma^\tau}{d\cos\theta_\pi},$$

where p_π and $\cos\theta_\pi$ are the momentum and outgoing angle[1] respectively, and τ denotes the target of the reaction, either H_20 or CH. The choice of these variables is motivated by the lack of previous measurements in these variables for single pion production channels. The vast majority of charged-current neutrino interactions are measured as a function of the outgoing lepton kinematics, or in other variables (e.g. neutrino energy) reconstructed from the lepton kinematics. As neutrino interaction generators often have parameters 'tuned' to better match experimental data, over time this can result in models over-fitting certain distributions. Presenting this result in variables to which models have not been tuned can provide a potentially much more valuable test of a model's ability to describe a more generalised data set.

In general, a measurement of the ratio of these two cross sections,

$$\frac{d\sigma^{\mathrm{H_20}}}{d\sigma^{\mathrm{CH}}}$$

can also be performed, but due to an unresolved issue in the cross-section propagation software such a ratio is not presented here. Such a cross section ratio has the advantage of being largely independent of any flux errors.

With limited statistics, there is a limit to how well the shape of the differential cross section can be resolved. As such, these quantities are measured in bins, such that the measurement in a kinematic bin is the integral of the cross section over that bin, i.e.

$$\left(\frac{d\sigma^\tau}{dx}\right)_i \approx \frac{1}{\Delta x_i}\int_{\Delta x_i}\frac{d\sigma^\tau}{dx}dx, \tag{4.1}$$

where x is the variable of interest $(p_\pi, \cos\theta_\pi, \dots)$, and i is the index of the kinematic bin, and Δx is the width of bin i.

The cross section depends on many more variables than the kinematic variable x, and these variables are integrated over implicitly. Among these variables is the neutrino energy, on which cross sections generally have a strong dependence. Denoting the integrated neutrino flux as Φ, it is clear that the integral of the cross section in a kinematic bin i will be equal to

$$\int_{\Delta x_i}\frac{d\sigma^\tau}{dx}dx = \frac{N_i^{\mathrm{signal},\tau}}{\epsilon_i^\tau \Phi T_\tau} \tag{4.2}$$

where ϵ_i^τ is the efficiency of detecting the signal in bin i, T_τ is the number of target molecules of τ, and $N_i^{\mathrm{signal},\tau}$ is the number of observed signal events in bin i, as this

---

[1] The outgoing angle is calculated with respect to the axis along which the neutrino was travelling.

is simply the interaction rate for the process in this bin, integrated over the runtime.

As such, the expression for the cross section becomes

$$\left(\frac{d\sigma^{\tau}}{dx}\right)_i = \frac{N_i^{\text{signal},\tau}}{\epsilon_i^{\tau}\Phi T_{\tau}\Delta x_i}, \tag{4.3}$$

with the ratio reducing to

$$\left(\frac{d\sigma^{\mathrm{H_2O}}}{d\sigma^{\mathrm{CH}}}\right)_i = \frac{N_i^{\text{signal},\mathrm{H_2O}}}{N_i^{\text{signal},\mathrm{CH}}} \times \frac{\epsilon_i^{\mathrm{CH}} T_{\mathrm{CH}}}{\epsilon_i^{\mathrm{H_2O}} T_{\mathrm{H_2O}}}. \tag{4.4}$$

To make this measurement, the values and errors on each of the quantities on the right hand side of this equation must be determined. Δx_i is simply the bin width, and is as such known. The number of molecules in the fiducial volume, T_τ is determined by the mass of the target and its molar mass. The integrated flux is discussed in the following subsection in more detail, with the determination of the efficiency and $N_i^{\text{signal},\tau}$ being described in detail in sections 4.3 and 4.5 respectively.

4.2.1 Flux

The T2K beam group perform detailed simulations [46] of the neutrino beamline. From their simulations they provide — in multiple binnings — a prediction of the neutrino flux and corresponding error in each neutrino energy bin, along with a correlation matrix giving the correlations between the bins. In this analysis the flux is allowed to be constrained by the data, and this constrained flux with its resulting errors and correlations are used to calculate the flux integral, Φ. More detail on how the flux is constrained is given in section 4.5.3.

As the cross section is computed as a function of the integrated flux, each kinematic bin can contain events contributed from all parts of the neutrino flux. Moreover, the cross section presented here will be specific to the integrated flux, and consequently to the ND280.

4.3 Efficiency Calculation and Correction

The so-called *efficiency correction* is the division by the efficiency to account for the number of events “missed” by the detector due to its imperfect efficiency. Simply, the efficiency in bin i will be the ratio of the number of observed signal events divided by the total number of signal events that occurred in the fiducial volume, i.e.

$$\epsilon_i^\tau = \frac{N_{\mathrm{obs.},i}^{\mathrm{signal},\tau}}{N_{\mathrm{true},i}^{\mathrm{signal},\tau}}. \tag{4.5}$$

While superficially simple, calculating the efficiency and associated error in a model-independent manner is non-trivial.

We define the *efficiency distribution* function, $\varepsilon(\vec{x})$, which gives the efficiency of detecting an event given by the kinematic variables[2] $\vec{x}$. More generally, the efficiency in a bin is given by the integral of the product of this efficiency distribution and the signal spectrum of the signal model, $f(\vec{x})$,

$$\epsilon_i = \int_{\mathrm{bin}\ i} f(\vec{x}) \times \varepsilon(\vec{x}) d\vec{x}. \tag{4.6}$$

In this form, is can be seen that when $\varepsilon(\vec{x})$ is a constant in bin i, ϵ_i will depend only on the total rate of the signal in the bin. When $\varepsilon(\vec{x})$ varies within the bin however, the value computed for ϵ_i will depend directly on the shape of the $\epsilon\,(x)$. From this it can be seen that unless the signal model exactly matches the signal distribution in nature, any bins where the efficiency varies will begin to acquire model dependence in the efficiency computed.

An obvious solution to the problems that arise from this is to bin finely, as the finer the binning, the less overall variation is expected in the efficiency in any given bin. Unfortunately with this method comes the need for high statistics — at the very least for the MC — to be certain that every bin has sufficient events to accurately characterise the event compositions, efficiencies, etc. in each bin. For analyses with a high signal rate where only a single particle is reconstructed (e.g. CC0π) the phase space is largely defined by two variables, and such a method is reasonable. In the case of multiple particles defining the final state, the dimensionality increases rapidly and the required MC statistics to bin similarly finely are increased exponentially.

4.4 Integrating Over Fine Bins

When using a fine binning for the purposes of a less model-dependent efficiency correction, it is still desirable to present a result where statistical fluctuations in each bin are low. Reducing statistical variations due to a low number of events and

---

[2] In the case of this analysis for example, the relevant variables would be $p_\mu, \cos\theta_\mu, p_\pi, \cos\theta_\pi, \cos\theta_{\mu-\pi}$.

increasing the number of bins are contradictory in nature, and in general these effects must be balanced. However, there is also the possibility of performing a cross section extraction in a fine binning, with each bin having a large statistical uncertainty and large correlation with nearby bins, and then integrating over subsets of bins such that the final integrated bins have appropriate statistical uncertainty.

In section 4.2 it was shown that the flux integrated cross section in a kinematic bin is given by equation 4.3, which generalises easily to multiple dimensions. For simplicity, here the two dimensional case is considered

$$\left(\frac{d^2\sigma^\tau}{dxdy}\right)_{i,j} = \frac{N_{i,j}^{\text{signal},\tau}}{\epsilon_{i,j}^\tau \Phi T_\tau \Delta x_i \Delta y_j}, \tag{4.7}$$

where x and y are the two variables of interest, and i and j are the bin indices in these variables respectively. This can be binned in both variables for the extraction and efficiency correction, but integrated over y such that the x-only binning will have lower statistical uncertainty. Starting with

$$\frac{d\sigma^\tau}{dx} = \int \frac{d^2\sigma^\tau}{dxdy} dy, \tag{4.8}$$

and using eq. 4.1 to obtain

$$\left(\frac{d\sigma^\tau}{dx}\right)_i = \frac{1}{\Delta x_i} \int_{\Delta x_i} \int_y \frac{d}{dx} \frac{d\sigma_{i,j}^\tau}{dy} dydx, \tag{4.9}$$

this can be rearranged to

$$\left(\frac{d\sigma^\tau}{dx}\right)_i = \frac{1}{\Delta x_i} \int_{\Delta x_i} \frac{d}{dx} \int_y \frac{d\sigma_{i,j}^\tau}{dy} dydx. \tag{4.10}$$

Substituting eq. 4.1 and using eq. 4.2 for the y bins turns the integral over y into a sum over j,

$$\left(\frac{d\sigma^\tau}{dx}\right)_i = \frac{1}{\Delta x_i} \int_{\Delta x_i} \frac{d}{dx} \sum_j \frac{N_{i,j}^{\text{signal},\tau}}{\epsilon_{i,j}^\tau \Phi T_\tau \Delta y_j} dx, \tag{4.11}$$

which reduces to[3]

$$\left(\frac{d\sigma^\tau}{dx}\right)_i = \frac{1}{\Delta x_i \Phi T_\tau} \sum_j \frac{N_{i,j}^{\text{signal},\tau}}{\epsilon_{i,j}^\tau \Delta y_j}. \tag{4.12}$$

This method of combining bins is functionally identical to that shown in [64],

---

[3] A similar expression can be derived for higher dimensions in the same manner.

where it was shown that combining bins in such a manner can provide the benefits of regularisation (e.g. smoothing the final distribution) without introducing bias towards a preferred model, and without sacrificing error coverage. While the analysis presented in this thesis lacks sufficient bins to fully demonstrate the regularisation-like benefits of this method, it is hoped that future analyses both on T2K and elsewhere can exploit this method to maximise their impact and usefulness.

4.5 Binned Likelihood Fit

To determine the number of selected true signal events in a given bin, a binned maximum likelihood method is used. The likelihood fit is performed by varying various fit parameters to find the most likely set values of these parameters for the given data. The software performing this fit is the Super-xsLLhFitter[4] [63], based on earlier such fitters used in T2K cross section analyses. The likelihood is given by

$$\mathcal{L} = \mathcal{L}_{\text{stat.}} \times \mathcal{L}_{\text{syst.}} \tag{4.13}$$

where $\mathcal{L}$ depends on the fit parameters and data. More useful to us is $-2\ln\mathcal{L}$, as the product above becomes a sum, and this quantity approaches the value of the χ^2 in the limit of an infinite sample [65]. The log-likelihood is then

$$\chi^2 = \chi^2_{\text{stat.}} + \chi^2_{\text{syst.}} \approx -2\ln\mathcal{L} = -2\ln\mathcal{L}_{\text{stat.}} - 2\ln\mathcal{L}_{\text{syst.}}. \tag{4.14}$$

With implicit dependence on the observed data and the free parameters.

The Poissonian log-likelihood is given by

$$\ln\mathcal{L}_{\text{stat.}} = \sum_{k}^{\text{reco. bins}} \left(\ln N_k^{\text{obs.}}! - N_k^{\text{exp.}} - N_k^{\text{obs.}} \ln N_k^{\text{exp.}}\right), \tag{4.15}$$

where $N^{\text{obs.}}$ and $N^{\text{exp.}}$ are the observed and predicted numbers of (total) events in each reconstructed bin. This expression can then be used to approximate the χ^2,

$$\chi^2_{\text{stat.}} \approx -2\ln\mathcal{L}_{\text{stat.}} = 2\sum_{k}^{\text{reco. bins}} \left(N_k^{\text{exp.}} - N_k^{\text{obs.}} + N_k^{\text{obs.}} \ln\frac{N_k^{\text{obs.}}}{N_k^{\text{exp.}}}\right), \tag{4.16}$$

---

[4] The Super-xsLLhFitter is hosted and documented at `https://gitlab.com/cuddandr/xsLLhFitter`

where $\ln(N_k^{\text{obs.}}!)$ is approximated by Stirling's approximation,

$$\ln(N_k^{\text{obs.}}) = N_k^{\text{obs.}} \ln(N_k^{\text{obs.}}) - N_k^{\text{obs.}} + \mathcal{O}(\ln N_k^{\text{obs.}}). \tag{4.17}$$

The systematic χ^2 contribution is given by

$$-2 \ln \mathcal{L}_{\text{syst.}} \approx \chi^2_{\text{syst.}} = \left(\vec{s} - \hat{\vec{s}}\right) \left(\mathbf{V}^{\text{cov.}}_{\text{prior}}\right)^{-1} \left(\vec{s} - \hat{\vec{s}}\right)^T, \tag{4.18}$$

where $\vec{s}$ is the vector of systematic parameters, $\hat{\vec{s}}$ is the vector of their prior values, and $\mathbf{V}^{\text{cov.}}_{\text{prior}}$ is the prior covariance matrix containing the uncertainties and pre-fit correlations on the systematic parameters. Implicit in this form of equation 4.18 is the assumption that the uncertainties on all systematic parameters can be described by a Gaussian.

4.5.1 Unfolding via Template Parameters

To calculate the cross section it is necessary to determine the true number of signal events observed, using the observed reconstructed events. The conversion from reconstructed to true variables in known as *unfolding.* Denoting the vector of reconstructed variables as $\vec{a}_{\text{rec}}$ and the vector of true variables as $\vec{a}_{\text{true}}$, the smearing matrix $\mathbf{S}$ is defined such that

$$\vec{a}_{\text{rec.}} = \mathbf{S}\vec{a}^T_{\text{true}} \tag{4.19}$$

and also the unfolding matrix $\mathbf{U}$ such that

$$\vec{a}_{\text{true}} = \mathbf{U}\vec{a}^T_{\text{rec.}}. \tag{4.20}$$

To unfold observed data, the simplest option is to calculate the matrix $\mathbf{S}$ and let $\mathbf{U} = \mathbf{S}^{-1}$, provided that $\mathbf{S}$ is invertible. $\mathbf{S}$ can generally be determined by simulating a large number of events in the detector and building the matrix from the relationships between the true and reconstructed kinematics.

Unfolding is performed in this analysis by assigning, for each target and each analysis bin, a free template parameter c_i^T that scales the signal rate in that bin, such that

$$N_i^{\text{signal}} = c_i N_i^{\text{MC signal}} \tag{4.21}$$

where $N_i^{\text{MC signal}}$ is the number of events predicted by the nominal MC. These parameters have no effect on background events, and in order to avoid biasing the result towards the signal model, the likelihood is not effected by the values of these template parameters. As the template parameters vary the rates of true events, the corresponding reconstructed event rates must also be scaled. The expression for the number of events in a reconstructed bin k then becomes

$$N_k^{\text{reco.}} = \sum_i \left(c_i N_i^{\text{MC signal}} + \sum_b^{\text{bkg.}} N_{ib}^{\text{MC bkg.}} \right) \mathbf{S}_{ik} \tag{4.22}$$

where b runs over all backgrounds.

4.5.2 Detector Systematic Uncertainty Parameters

In section 5.7 the systematic uncertainties associated with the event selection will be discussed and evaluated. The treatment of these uncertainties in the fit is done by adding a nuisance parameter d_k with prior uncertainty determined by the effects of the sum of all detector uncertainties in bin k. The correlations between the uncertainties in each bin are of great importance here, as the correlations are in general very strong between large numbers of these parameters.

This modifies the expression for the reconstructed event rate in bin k only slightly, as nuisance parameters are multiplicative, and the expression becomes

$$N_k^{\text{reco.}} = \sum_i \left(c_i N_i^{\text{MC signal}} + \sum_b^{\text{bkg.}} N_{ib}^{\text{MC bkg.}} \right) d_k \mathbf{S}_{ik}. \tag{4.23}$$

4.5.3 Flux Parameters and Errors

As discussed in section 4.2.1, the neutrino flux is parameterised in varying energy bins. The exact binning used is detailed in chapter 6. For each flux bin, a parameter is added to the likelihood fit. The prior uncertainty and correlations between these parameters are taken to be those calculated by the beam group.

To account for the effect of varying the flux on the reconstructed distributions, a matrix v_{ni} is defined which maps flux weights from flux bins f_n to weights in true bins i for both signal and background combined. Adding these terms to the expression for the number of events in a reconstructed bin gives

$$N_k^{\text{reco.}} = \sum_i \left(c_i N_i^{\text{MC signal}} + \sum_b^{\text{bkg.}} N_{ib}^{\text{MC bkg.}} \right) d_k \mathbf{S}_{ik} \sum_n^{E_\nu \text{ bins}} v_{ni} f_n. \tag{4.24}$$

4.5.4 Cross Section Model Uncertainty

The final effects to be included in the fit are those related to the cross section models used in the input MC. Each model parameter uncertainty is encoded into a response function $\omega_t(\vec{x})$ where t indexes the model uncertainty in question. The response function provides the appropriate weight to events based on their kinematics and other properties (e.g. interaction type), here given by $\vec{x}$. These weights are applied both to signal and background in general, while some parameters affect only one.

$$N_k^{\text{reco.}} = \sum_i \left[\left(c_i N_i^{\text{MC signal}} \prod_t^{\text{model}} \omega_t(\vec{x}) \right) + \left(\sum_b^{\text{bkg.}} N_{ib}^{\text{MC bkg.}} \prod_t^{\text{model}} \omega_t(\vec{x}) \right) \right] d_k \mathbf{S}_{ik} \prod_t^{\text{model}} \sum_n^{E_\nu \text{ bins}} v_{ni} f_n. \tag{4.25}$$

Equation 4.25 gives the final expression for the number of events predicted in reconstructed bin k as a function of all the parameters in the fit, allowing comparison between MC and data.

4.5.5 Fitter

The fit is performed using the Minuit2 [66] minimiser in combination with the MIGRAD algorithm to find the minimum of the log likelihood given in equation 4.14, where $\mathcal{L}_{\text{syst.}}$ and $\mathcal{L}_{\text{stat.}}$ are given by equations 4.18 and 4.16 respectively, with equation 4.25 substituted for $\mathcal{L}_{\text{stat.}}$. This minimum gives the best fit value for all fit parameters, and is computed to within a specified tolerance. The HESSE algorithm then computes the Hessian matrix (the matrix of second derivatives in parameter space) about the minimum using the method of finite differences. Under the assumption that the log likelihood surface in this region can be described by a multi-dimensional Gaussian, the Hessian matrix is then inverted to give the covariance matrix between the post-fit parameters.

4.6 Cross Section Calculation and Error Propagation

The fit described in section 4.5 returns the best fit values for all fit parameters, along with a post-fit covariance describing the uncertainties on each parameters and their correlations. These parameters allow the number of events in each true bin to be determined from equation 4.21.

To combine this into a measurement of a cross section, equation 4.3 is used, where the inputs are the output from the fit. Due to the large number of parameters and correlations between them, analytically propagating the error through this calculation is not feasible. In place of an analytical error propagation, a large number of toy Monte Carlos (henceforth toys) are thrown from the post-fit errors and covariance, and the cross section calculated for each toy. The distribution of these toys about the nominal cross sections allows the calculation of an uncertainty in each bin of the cross section, along with the correlations between them.

In order to generate toys with an appropriate distribution while accounting for correlations, the post-fit covariance matrix, $\mathbf{\Sigma}_{\vartheta}$, is first Cholesky decomposed into a lower-triangular matrix L and its conjugate transpose such that

$$\mathbf{L}\mathbf{L}^{*} = \mathbf{\Sigma}, \tag{4.26}$$

using the TDecompChol package in ROOT. This decomposed matrix $\mathbf{L}$ is multiplied with a vector $\vec{g}$ of random numbers from a Gaussian distribution with mean 0 and variance 1 to generate a single toy parameter vector $\vec{\vartheta}_{\text{toy}}$ such that

$$\vec{\vartheta}_{\text{toy}} = \vec{\vartheta}_{\text{best fit}} + \mathbf{L}\vec{g}. \tag{4.27}$$

For each parameter variation, the cross section is computed and the cross section covariance matrix is computed via

$$\mathbf{\Sigma}^{\sigma}_{i,j} = \sum_{t}^{\text{toys}} \left[\left(\frac{d\sigma}{dx}\right)_{i,t} - \left(\bar{\frac{d\sigma}{dx}}\right)_{i} \right] \left[\left(\frac{d\sigma}{dx}\right)_{j,t} - \left(\bar{\frac{d\sigma}{dx}}\right)_{j} \right] \tag{4.28}$$

In both equation 4.18 and in the use of the HESSE algorithm to determine the parameter covariance, an assumption has been made that the uncertainties on all fit parameters can be described by a Gaussian. For some specific systematic parameters — these will be discussed in section 5.8 — there is a defined range of values beyond which they cannot be used. Consequently, in cases where the best fit value approaches such a boundary to within a few sigma — as calculated by HESSE — these parameters will occasionally be thrown outside their defined range. To ensure that all throws are well defined for each toy, all parameters are required to be inside their defined range, and any parameters thrown outside this range are moved to the edge of that range. When parameters are pulled close to the edges of their defined ranges, this can become an significant issue and can bias the result through incorrect uncertainty propagation. A number of fake data studies are performed later to validate that such parameters stay well inside their defined range for credible data.

Integrating over dimensions as described in section 4.4 allows the computation of two distinct cross section measurements from a single likelihood fit result. More traditionally, single-differential cross section calculations have required a separate likelihood fit for each variable to be reported, as is the case in [72]. Propagating errors in manner described here, it is possible to compute multiple single-differential cross sections from a single likelihood fit result at the same time. Consequently, it is also possible to evaluate the correlations between the single-differential cross section bins in two different variables, i.e. the correlation between

$$\left(\frac{d\sigma^{\tau}}{dx}\right)_i \text{ and } \left(\frac{d\sigma^{\tau}}{dy}\right)_j . \tag{4.29}$$

Releasing multiple single-differential cross sections with the appropriate covariance matrix between all bins could allow all bins to be used to fit to a theoretical model without the risk of over-fitting to a given dataset in a global neutrino fit.

Chapter 5

Selection

The first step in measuring the rate of a process is to identify these events in a detector. In doing so one aims to understand how the signal behaves in the detector, and use this information to create a series of cuts to remove as large a fraction of background as possible whilst retaining high efficiency for signal events.

5.1 Existing T2K Selection

As part of the T2K oscillation analysis the collaboration developed a $\bar{\nu}_\mu$CC selection split by pion multiplicity. This selection first attempts to establish the presence of a μ^+ to find $\bar{\nu}_\mu$CC events. For events where a μ^+ is found, pions — charged and neutral — are searched for in the remaining tracks and hits in the detector in four distinct ways. Based on the number of each pions reconstructed, the events are divided into the following three branches:

- $\bar{\nu}_\mu$CC0π
    - No reconstructed pions are found. This branch is rich in CCQE and CCQE-like events.
- $\bar{\nu}_\mu$CC$1\pi^-$
    - A single negatively charged pion is found. These events are largely produced through resonant and coherent processes.
- $\bar{\nu}_\mu$CC-Other

- ○ Any events that do not fall into the two preceding branches. This branch is rich in DIS events, and other types of events that have many particles in the final state.

The cuts in this selection are the same as those in the T2K $\bar{\nu}_\mu$ selection with the exception that the μ^+ candidate is not required to be the highest momentum track in that event. The cuts are detailed in the following subsections.

5.1.1 Event Quality Cut

Various flags are attached to each ND280 subrun to indicate any issues with the detector, such as issues with inter-detector time calibration, particular sub-detectors being inoperational, etc. The first cut applied is to require that only data where there are no such issues is used, using these data quality flags. This cut also requires that the event be coincidental with the beam trigger.

5.1.2 Fiducial and Track Quality Cuts

A series of cuts is applied primarily to reject events where there are no useful reconstructed objects for this analysis. The first of these cuts requires that there is a reconstructed track crossing at least one of the TPCs, as the high momentum and $\frac{dE}{dx}$ resolution afforded by the TPCs is important in distinguishing μ^+ from other particles.

The second cut requires that an event have its vertex reconstructed inside the fiducial volume of an FGD, and that at least one positively charged track has ≥ 19 hits in a single TPC, as momentum and $\frac{dE}{dx}$ reconstruction is not well modelled below this number.

The following two cuts reject any events with their vertices reconstructed in the first and last layers (respectively) of an FGD. In some cases the track matching algorithm can fail to match a FGD track segment to a TPC track segment from a through-going particle, resulting in two tracks with a vertex at the edge of an FGD.

5.1.3 Anti-muon PID Cut

With at least one positively charged TPC track identified in an event, a $\bar{\nu}_\mu$CC candidate event is identified. However, there are many positive particles that can be produced in neutrino-nucleus interactions other than anti-muons, so it must be confirmed that the anti-muon candidate behaves as expected in the detector. Particle

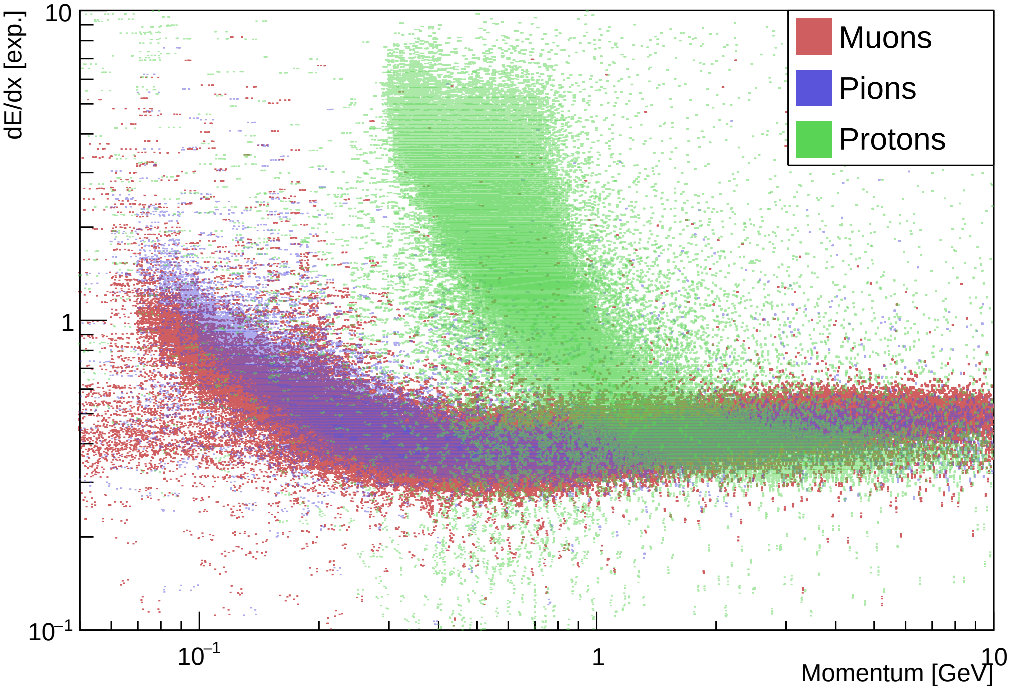

Figure 5.1: Distributions of $\frac{dE}{dx}$ against momentum in the ND280 TPCs for the three primary particle types selected as μ^+. Pions can be seen to be very similar to muons at most momenta, with the proton energy loss curve overlapping in the $1-2.5$ GeV regime.

identification in the TPCs is done primarily through the measure of both the momentum and the rate of energy loss in the gas, $\frac{dE}{dx}$.

Figure 5.1 shows how the $\frac{dE}{dx}$ distributions for the three main particles of interest in this selection vary with momentum. This plot makes it clear that there is no simple cut that can be placed on $\frac{dE}{dx}$ that would select most μ^+s while rejecting most other particles. Rather than uniquely consider the value of $\frac{dE}{dx}$, the *pull* is defined as the deviation from the expected energy loss for a given particle type, normalised by the standard deviation of the $\frac{dE}{dx}$ distribution at a given momentum for that particle, i.e.

$$\mathrm{Pull}_{\alpha}\left(p\right) = \frac{\left(\frac{dE}{dx}\right)_{\alpha}^{\mathrm{obs.}}\left(p\right) - \left(\frac{dE}{dx}\right)_{\alpha}^{\mathrm{exp.}}\left(p\right)}{\sigma_{\alpha}\left(p\right)} \tag{5.1}$$

where α denotes the particle hypothesis for which the pull is calculated, $\left(\frac{dE}{dx}\right)_{\alpha}^{\mathrm{obs.}}\left(p\right)$ and $\left(\frac{dE}{dx}\right)_{\alpha}^{\mathrm{exp.}}\left(p\right)$ are the observed and expected (theoretical) energy loss rates in a medium and σ_{α} is the standard deviation of the $\left(\frac{dE}{dx}\right)_{\alpha}^{\mathrm{exp.}}\left(p\right)$ distribution.

The cut placed to determine whether a particle is an anti-muon is as follows:

$$p_{\mu} > 500\mathrm{MeV} \qquad OR \quad \frac{\mathrm{Pull}_{\mu} + \mathrm{Pull}_{\pi}}{1 - \mathrm{Pull}_{\mathrm{Proton}}} > 0.9$$

and

$$\mathrm{Pull}_{\mu} > 0.1$$

5.1.4 Pion Identification

Having identified an anti-muon candidate, the next task is the identification of any pion candidates on which to base this selection. Pion candidates are tagged in four distinct and mutually exclusive ways;

1. All good quality TPC tracks other than the anti-muon candidate are checked for consistency with the charged pion hypothesis.
2. Tracks that are contained entirely within an FGD are checked for consistency with the charged pion hypothesis.
3. Stopping pions can decay, with the resulting muon decaying and emitting a so-called Michel electron. Michel electrons are searched for by looking for time-delayed hits in the same FGD as the event vertex.
4. Neutral pions are searched for by looking for shower-like objects in the ECals that are not associated with a track in the TPC or FGD.

Each of these methods is detailed further in the following paragraphs.

Pion PID in the TPCs

The cut placed on negative TPC tracks to select pion candidates requires that the track satisfy both of the following conditions:

$$p_{\pi} > 500\mathrm{MeV} \qquad OR \quad \frac{\mathcal{L}_{\mu} + \mathcal{L}_{\pi}}{1 - \mathcal{L}_{\mathrm{Proton}}} > 0.8$$

and

$$\mathcal{L}_{\pi} > 0.3$$

where $\mathcal{L}_{\alpha}$ is the likelihood for a track to result from a particle of species α. The

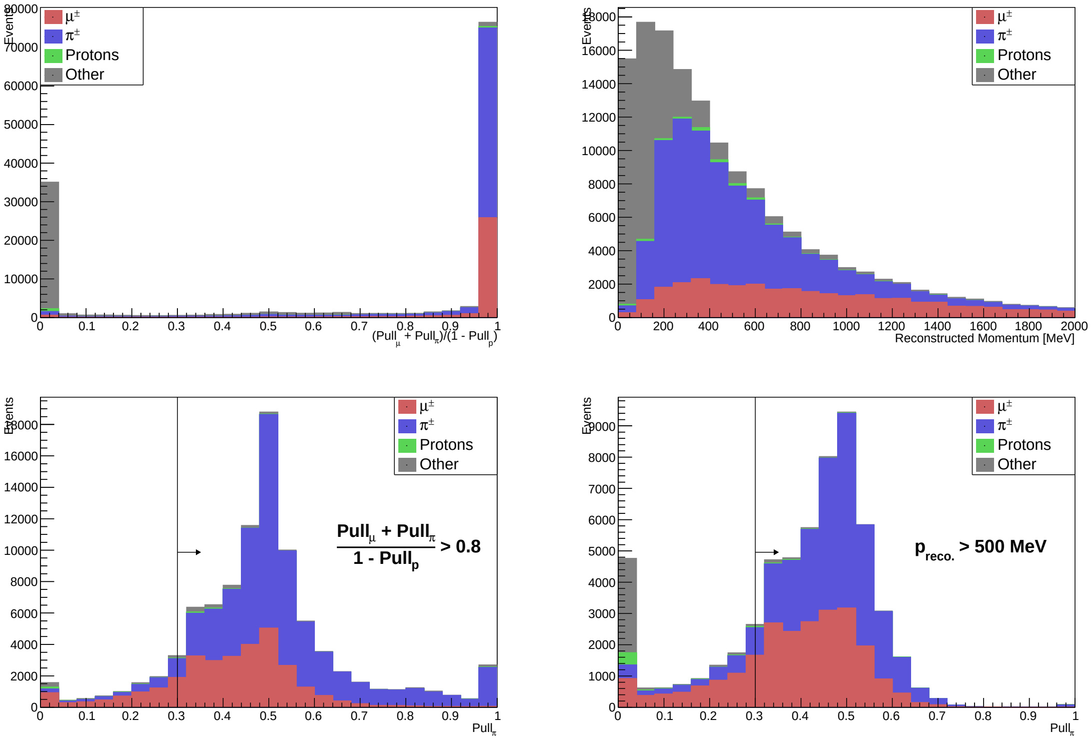

Figure 5.2: Distributions of particle type for all negatively charged TPC tracks for events passing the anti-muon PID cut. Particle composition is shown as a function of multiple pull variables (*top left*) and momentum (*top right*), with the cuts shown in Pull_π for both cut cases in the bottom plots.

likelihood is defined as

$$\mathcal{L}_\alpha = \frac{\text{Pull}_\alpha}{\text{Pull}_\mu + \text{Pull}_\pi + \text{Pull}_\text{e} + \text{Pull}_\text{Proton}}. \tag{5.2}$$

Figure 5.2 shows both cases of the cut individually, as an event needs only pass one of two potential requirements.

Pion PID in the FGDs

Pions contained in the FGDs are selected, as in the TPCs, by cutting on Pull variables rather than purely on $\frac{dE}{dx}$. Unlike the TPCs, the FGDs do not provide a similarly accurate measurement of momentum. As a result of the low momentum resolution, the discriminating power of the Pull variable suffers, being unable to accurately determine the expected value for the energy loss. As a proxy for momentum, particle range is used instead, as discussed in section 3.2.2. The FGD contained reconstruction is discussed in more detail in section 5.4, and analysis-level improvements to the

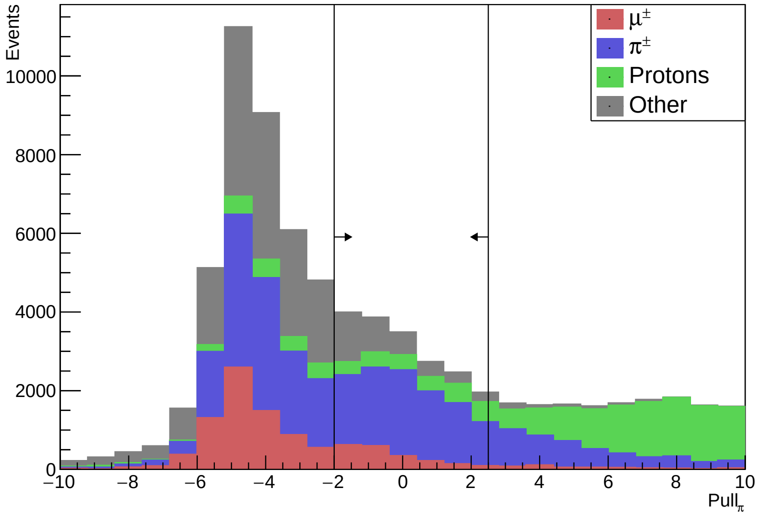

Figure 5.3: Cut on Pull$_\pi$ used to select pion-like isolated FGD tracks.

reconstructed kinematics are presented. These improved kinematics — in particular momentum — have the potential to improve the discriminating power of the Pull variables, but as these variables cannot currently be recalculated without reproducing the detector simulation, they are not used at this time.

The cut placed to select pions contained in the FGDs is

$$-2 < \text{Pull}_\pi < 2.5,$$

with the effect of the cut shown in figure 5.3.

Michel Electrons

As low momentum charged pions stop in the FGD, they decay into a (anti-) muon and muon (anti-) neutrino, where the (anti-) muon will in general have low momentum, and will also quickly decay into a so-called ‘Michel electron’. Where detected, these Michel electrons can indicate the presence of a charged pion in an event, which is used in this selection to further tag pions. Michel electrons are searched for by looking for delayed hits in the FGD in the time after an event is reconstructed, where the hits are consistent with an electron-like particle. As there are large populations of both π^+ and π^- in the events relevant to this selection, and the energy of the electrons is too low to reliably determine charge, the presence of a Michel electron cannot reliably be used to determine the charge of the primary pion in general. Due to the much higher absorption cross-section for π^- on nuclei however, it is much more likely for a stopping pion in an FGD to be π^+, and so where a Michel electron is found, a π^+ is assumed to have been present in the event.

ECal Neutral Pion Detection

With alternating layers of active scintillator, high energy photons and electrons will generally initiate an electromagnetic cascade in the ECals, depositing their energy in a shower-like pattern. Minimum ionising particles will generally penetrate the lead layers more easily while still leaving more concentrated hits in the scintillator layers, resulting in a more track-like energy deposit. PID in the ECal is based primarily on the energy deposit per unit track-length, shower-like appearance, and whether the particle appears to have stopped in the ECal. On the global reconstruction level, three ECal specific PID variables are computed via a multivariate method, with the low-level hit information from the ECal local reconstruction as inputs. These three variables are largely degenerate with one another, and in general, use of more than one provides little benefit. The ECal also provides a measure of the energy deposited in a reconstructed track/shower, as well as a measure of the energy of the initial particle under the assumption that the particle was an electron or photon.

Neutral pions will decay quickly into two photons with a branching ratio ≥ 0.98. Due to the relatively low density of the FGD scintillator, the very low density of the TPC gas, and the presence of lead layers in the ECals, photons are most likely to interact in the ECals. Moreover, when a decay photon converts in an ECal, the photon should have left no or minimal traces in the tracker, and it is expected that the ECal hits will not be associated with any other tracks in the detector. It is also expected that, due to the high density, photon conversion should occur in the first few layers traversed.

To search for neutral pions, the most energetic ECal track not associated to a global track in each event is checked. For each of these tracks, if the electromagnetic energy deposit is above 30 MeV[1] the value of the Log-Likelihood variable PIDMipEm is greater than 0, and the track begins within the first five layers of the ECal, the track is marked as a decay photon from a neutral pion. The cut value is located at 0 as the relative systematic uncertainty for a cut at this value has been calculated for previous measurements with the ND280. While some additional discrimination power or greater efficiency can be achieved by moving this cut value, any improvement is for

---

[1] ECal reconstruction at very low energies is not well modelled, and as such no attempt is made to extract analysis-level information from such tracks.

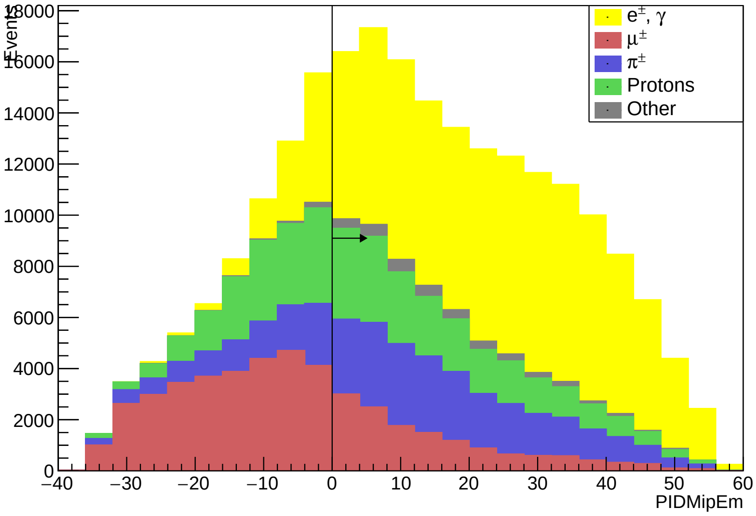

Figure 5.4: Distribution of isolated ECal tracks in the PIDMipEm variable for suitable decay photon candidates, broken down by particle incident on the ECal. All events in this plot have passed the anti-muon PID cut.

this measurement is slight, while there is an increased systematic uncertainty penalty, as data and simulation agreement are poorer towards the edges of the distribution.

The true particle distribution for all isolated ECal tracks with electromagnetic energy deposit above 30 MeV that start within the first five ECal layers is shown as a function of the PIDMipEm variable in figure 5.4 for all events passing the anti-muon PID cut. The effect of this reconstruction on the selection as a whole is shown and discussed in section 5.2.4.

5.2 Signal Selection

The selection described up to this point provides a good point from which to begin. For the purposes of extracting a cross section, it is beneficial to improve the purity and constrain the dominant backgrounds with some control regions.

As discussed in the preceding section, a large source of contamination is positive pions and protons being incorrectly reconstructed as anti-muons. In attempting to remove this contamination it is important to consider only the effects on the ability to correctly select anti-muons. Optimising such a cut based on its effect on the signal only (which includes a pion) could conceivably introduce more model dependent effects in the selection than are necessary. Keeping this in mind, the cuts described in the following sections have been determined considering only their effects on the

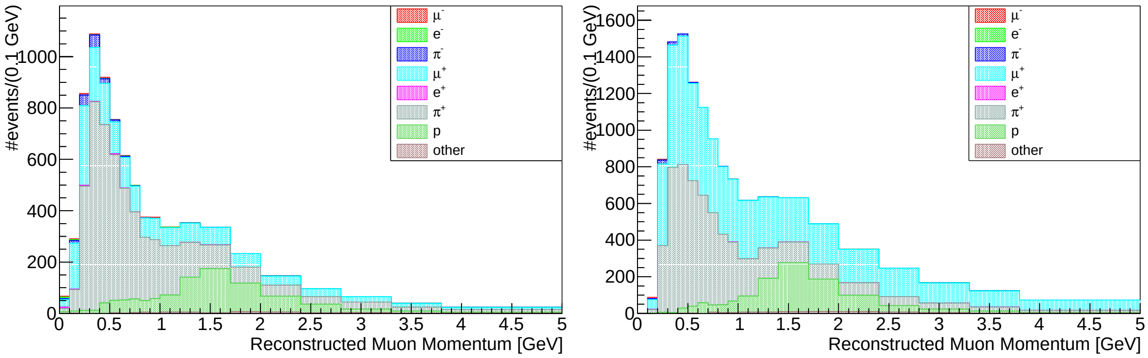

Figure 5.5: Distributions of MC events in anti-muon candidate momentum with no ECal segment (*left*) and with one or two ECal segments (*right*).

anti-muon selection purity.

5.2.1 Negatively Charged Pion Cut

As the signal should have a single negative pion in the final state, exactly one reconstructed pion is required in either the TPC or FGD. The signal selection is at this point branched in two, split by the subdetector in which the pion is tagged. Two further branches are created for the two control regions described in sections 5.3.1 and 5.3.2. Events which have more than one negatively charged pion, or one negatively charged pion and at least one other charged pion are placed in the second of these control regions.

5.2.2 Tracks with ECal Segments

As the anti-muon candidate track is the highest reconstructed momentum track the majority of the time, it is natural to assume that this track should generally interact with more of the detector than the pion candidate. The anti-muon candidate track is also required to have a segment in at least one TPC. At this point in the selection the only detector that can reasonably be expected to give significant information that is not already in use is the ECal (both barrel and downstream).

It is clear that low momentum anti-muons can not reasonably be expected to propagate to the ECals, but above a few hundred MeV the majority of the anti-muon candidates that do not have a track segment in any ECal are actually particles other than an anti-muon. Figure 5.5 shows the reconstructed momentum distributions of anti-muon candidate tracks with and without segments in the ECals.

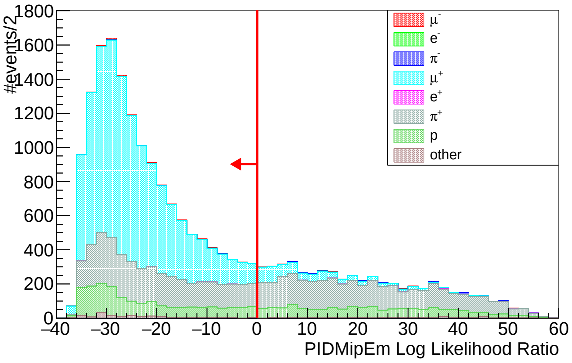

Figure 5.6: Distribution of MC events in the PIDMipEm Log-Likelihood variable, for all signal region events with an ECal segment.

A cut is applied requiring that any anti-muon candidate track with momentum over 220 MeV have a segment in an ECal.[2]

5.2.3 ECal PID

For all tracks remaining in the selection with an ECal segment, a cut is made on the variable PIDMipEm from the ECal reconstruction. PIDMipEm is a multivariate-based Log-Likelihood ratio where values below 0 are MIP like and values above are shower-like. Figure 5.6 shows the distribution of events for this variable, with events left of the cut line passing.

This cut can be seen to remove approximately half of the incorrectly identified muon candidates while removing relatively few true anti-muons, allowing both a significant increase in purity and providing — together with events failing the previous cut — a sample rich in background from misidentified anti-muons, which is discussed further in section 5.3.1.

5.2.4 ECal π^0 Veto

Neutral pions are tagged via their decay photons as described in section 5.1.4. If one or more neutral pions are tagged in the event, it is excluded from the signal selection and classified as a CC-Other event and placed into the CC-π^-+X control sample described in section 5.3.2. Figure 5.7 shows the events passing and failing this cut

---

[2] Tracks with two ECal segments are also accepted, but tracks with more than two segments are rejected. The reason for this is that two ECal segments in a track is somewhat natural if the anti-muon is incident at a high angle on the edge of the DS ECal, as the barrel ECals overlap the sides of the DS ECal. These events are very rare.

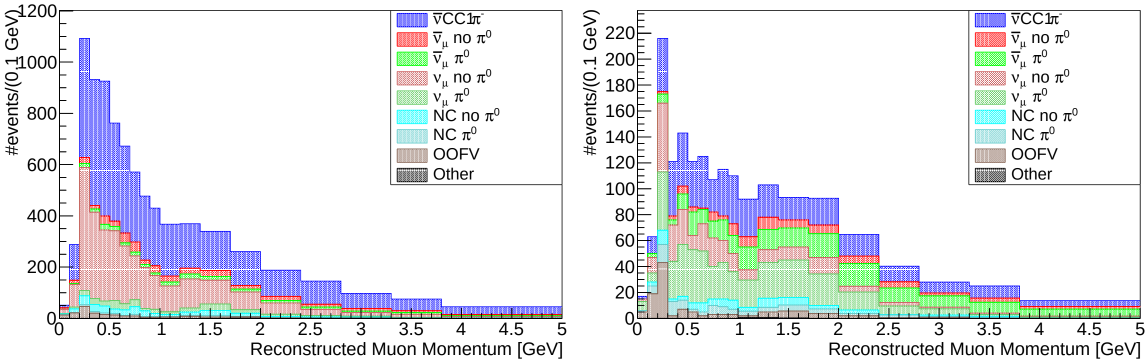

Figure 5.7: Distributions of MC events passing (*left*) and failing (*right*) the ECal π^0 veto cut for both FGDs.

for both FGDs combined.

5.2.5 Signal Selection Summary

The predicted kinematic event distributions are shown in figure 5.8 and 5.9, in the reconstructed momentum and cosine of both the pion and muon candidates. The TPC pion sample is the most pure due to the ability to accurately measure charge, momentum, and $\frac{dE}{dx}$ for both particles, where the measurements of these quantities are relatively limited when the pion is contained in the FGD. Overall signal selection efficiency within the signal samples is approximately 18%, with a purity of approximately 63%.

5.3 Control Regions

As discussed in section 5.1, the two primary types of background stem from particles being incorrectly identified as anti-muons, and pions (both charged and neutral) not being well reconstructed. To attempt to constrain these backgrounds and validate their modelling, a number of selection samples are introduced with relatively high purity in these backgrounds, as well as kinematics representative of those events in the signal sample.

5.3.1 Failing ECal PID/Segment Cut

Events passing the 1π cut but failing either of the subsequent cuts described in sections 5.2.2 and 5.2.3 are placed in this sample. This sample is rich in events

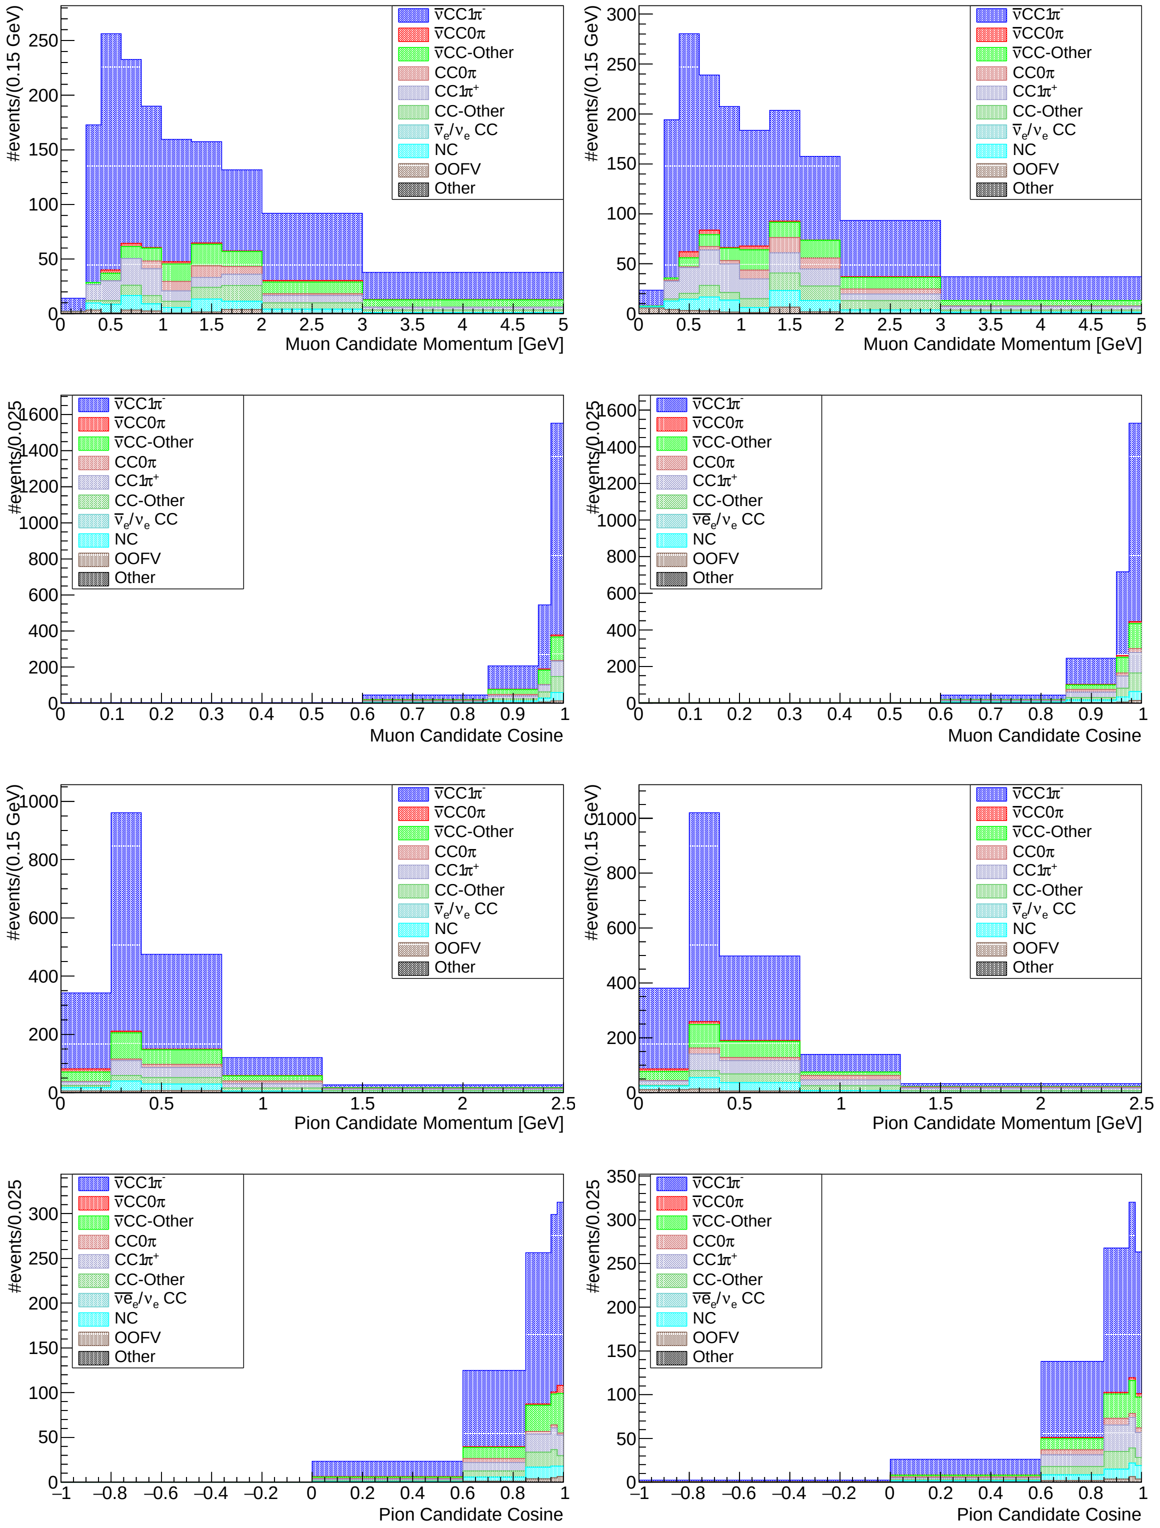

Figure 5.8: Distributions of MC with the pion candidate in the TPC events passing the signal selection for FGD1 (*left*) and FGD2 (*right*).

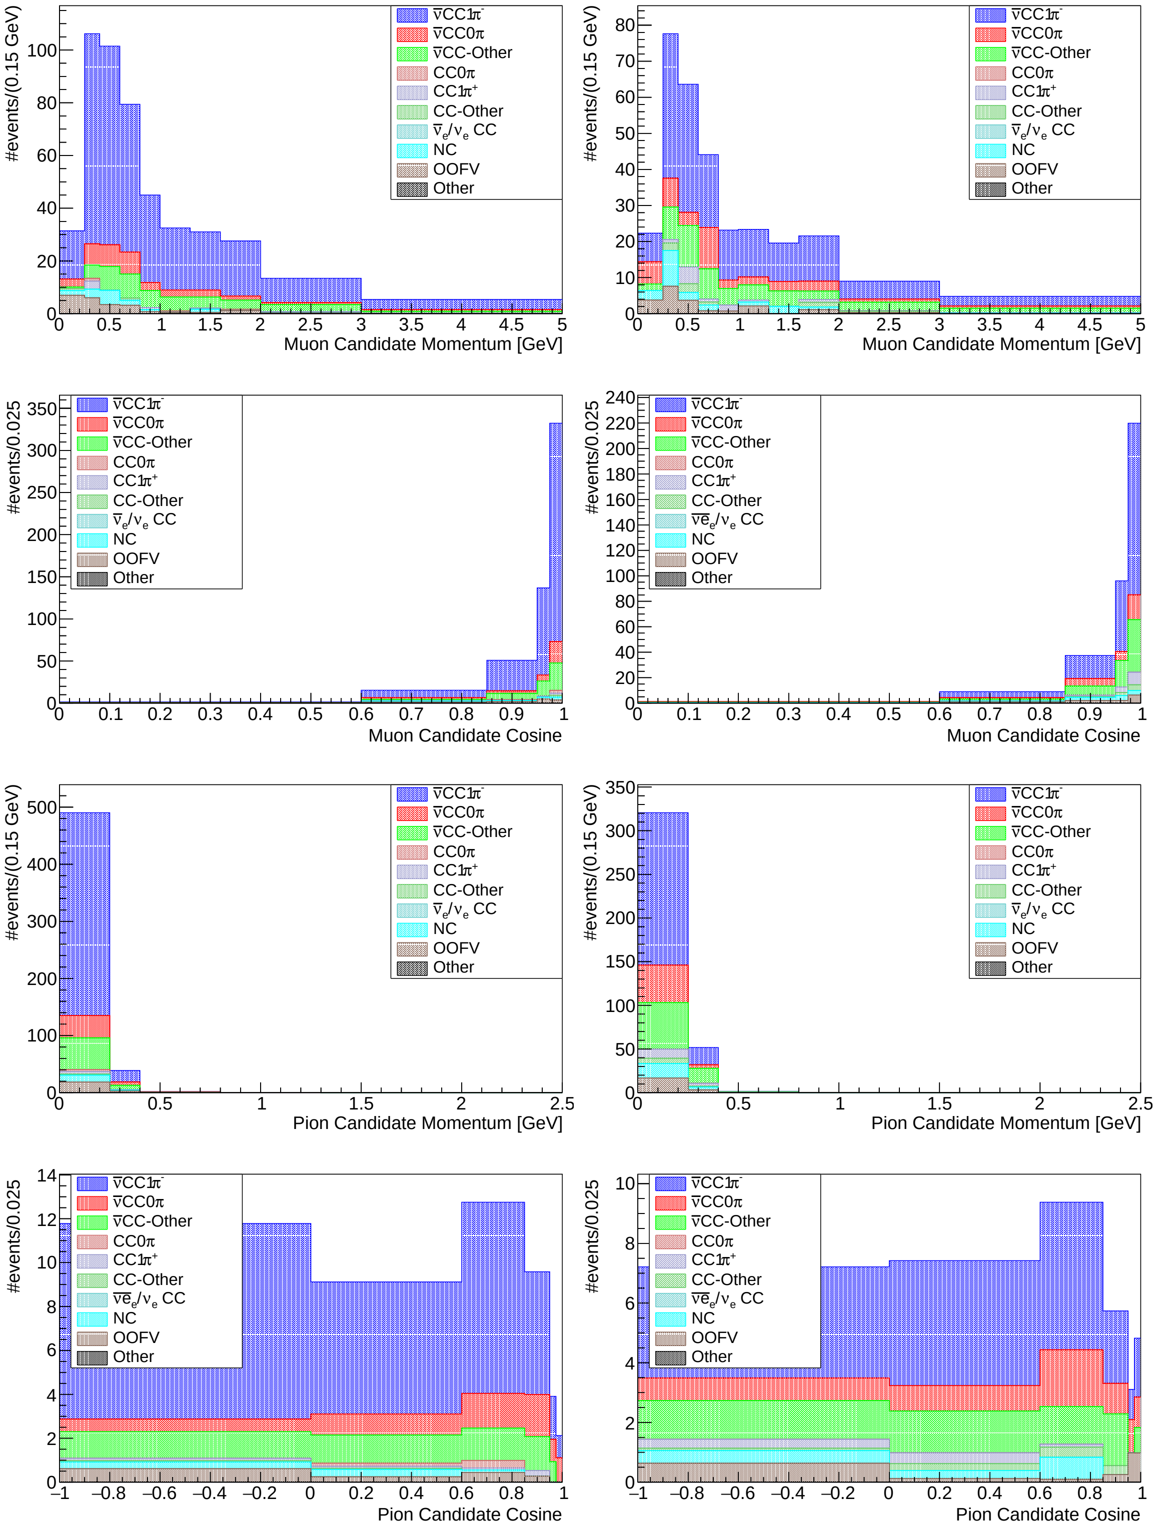

Figure 5.9: Distributions of MC with the pion candidate in the FGD events passing the signal selection for FGD1 (*left*) and FGD2 (*right*).

that pass the anti-muon TPC PID cut, but are not true anti-muons, while being largely kinematically representative of the background that passes these two cuts. The kinematic distributions broken down by event topology can be seen in figure 5.10.

The motivation behind using these events in the final fit is to better constrain the large background from events where the anti-muon candidate is a proton or pion. Large discrepancies in this region could indicate that the rate or distribution of these events is not accurately modelled, or could be the result of the cut behaving differently than predicted by the detector MC. If these events were simply discarded, such issues would be entirely missed if present in the data.

The cases where both positive pions and protons are identified as anti-muons can be seen to share very similar reconstructed kinematic distributions in figures 5.11 and 5.12, for pions and protons respectively.

5.3.2 CC-π^-+X

This control region contains any reconstructed CC events with a reconstructed π^-, and at least one other reconstructed charged or neutral pion. Despite the selection criteria, this control region is perhaps surprisingly deficient in $\bar{\nu}_\mu$CC-Other events, due to the difficulty of accurately identifying a muon, particularly in high-multiplicity events. The event distributions for this control region are shown in figure 5.13. The kinematic distribution of both ν_μ and $\bar{\nu}_\mu$ CC-Other events is also markedly different from their distributions in the signal samples, but importantly there are no regions in these kinematics that are devoid of events, suggesting that this control region should still be capable of providing some constraint for these backgrounds over the entire signal region.

5.4 FGD Contained Tracks

Due to the T2K near-detector fit framework fitting — at the time of writing — only to the muon candidate momentum and angle, the kinematics of other particles are rendered irrelevant to the analysis, and consequently very little work has been performed on accurately reconstructing these quantities. In the context of the analysis presented in this thesis however, the kinematics of both the pion and anti-muon are of interest.

Momentum determination in the ND280 tracker is performed by fitting a helix —

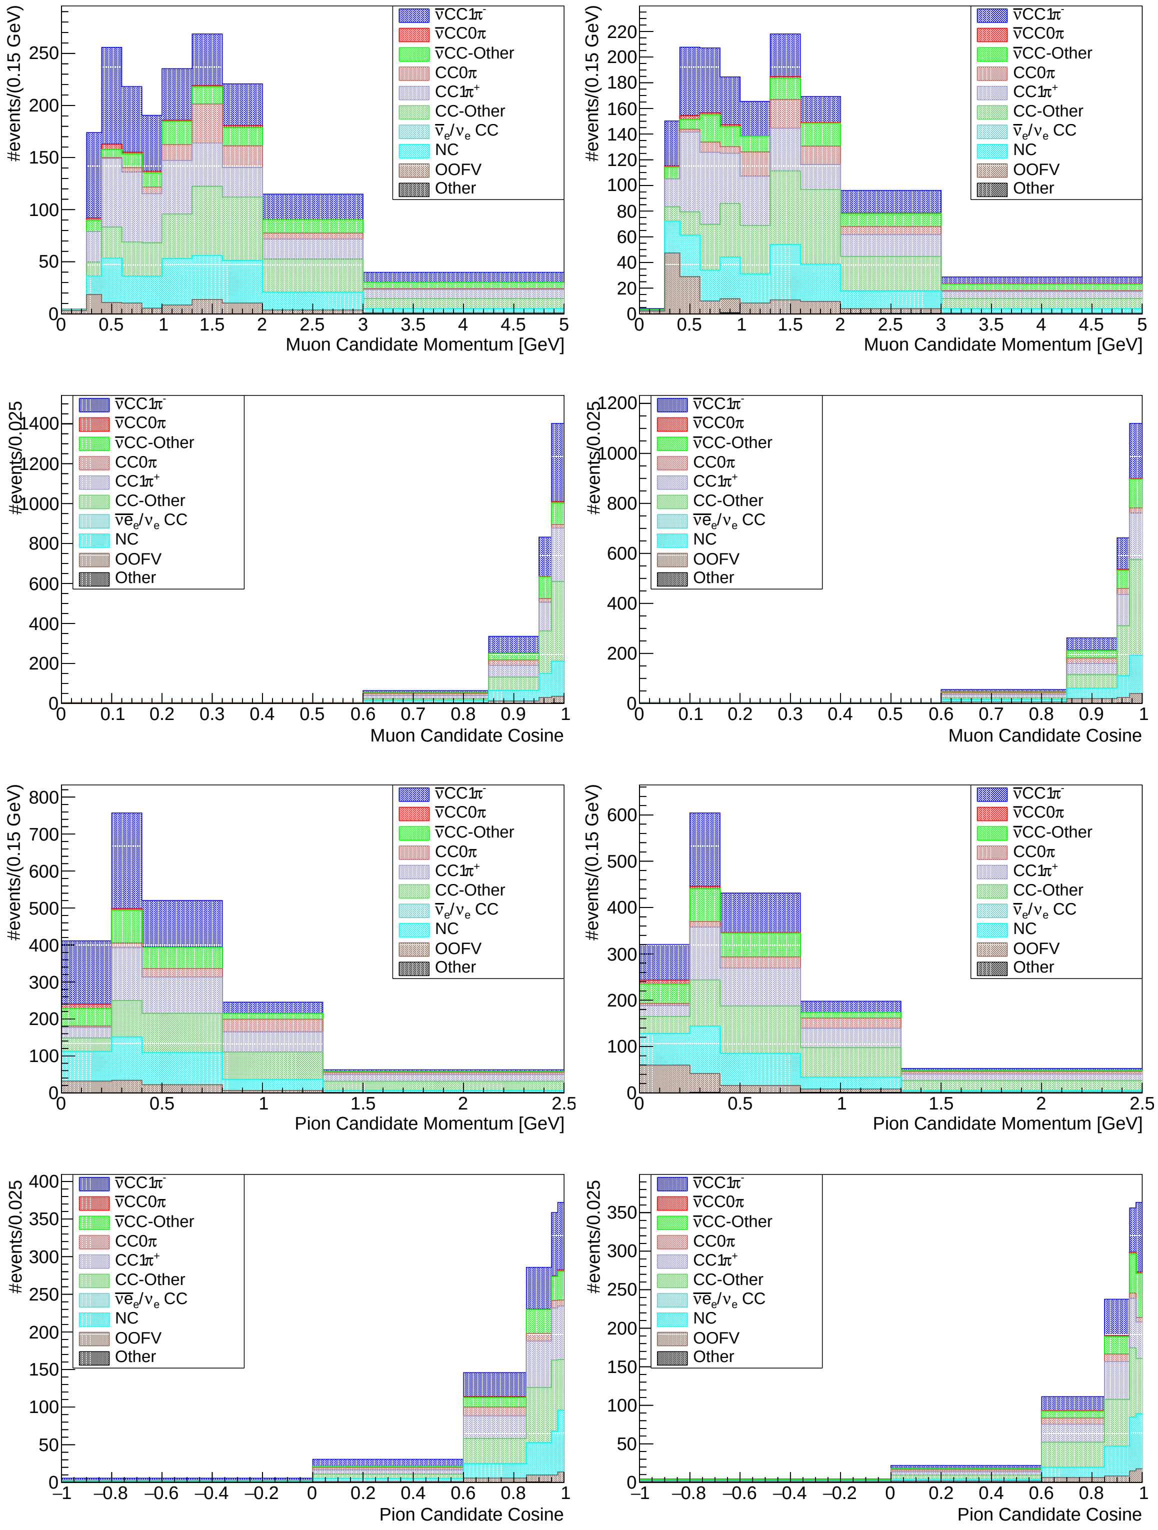

Figure 5.10: Distributions of MC events in muon and pion candidate kinematic variables in the ECal PID/Segment Control Regions for FGD1 (*left*) and FGD2 (*right*).

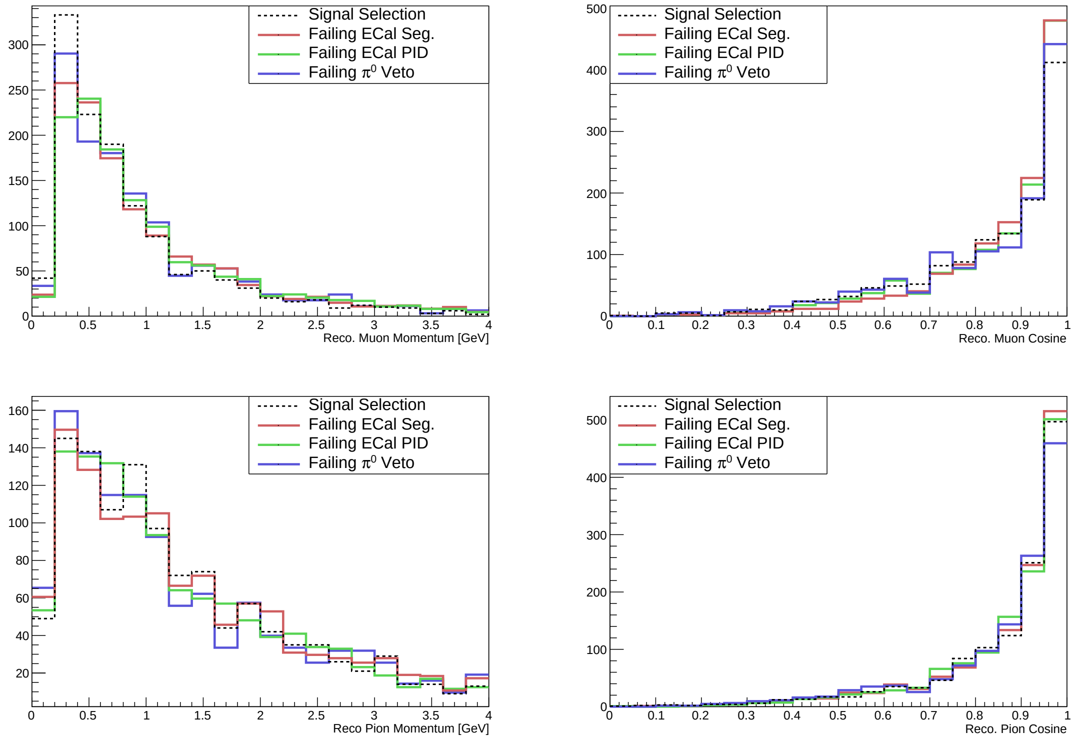

Figure 5.11: Distributions of events with true pions misidentified as anti-muons by failed cut for both FGDs combined. Distributions shown in the four primary reconstructed kinematic variables. All histograms are normalised.

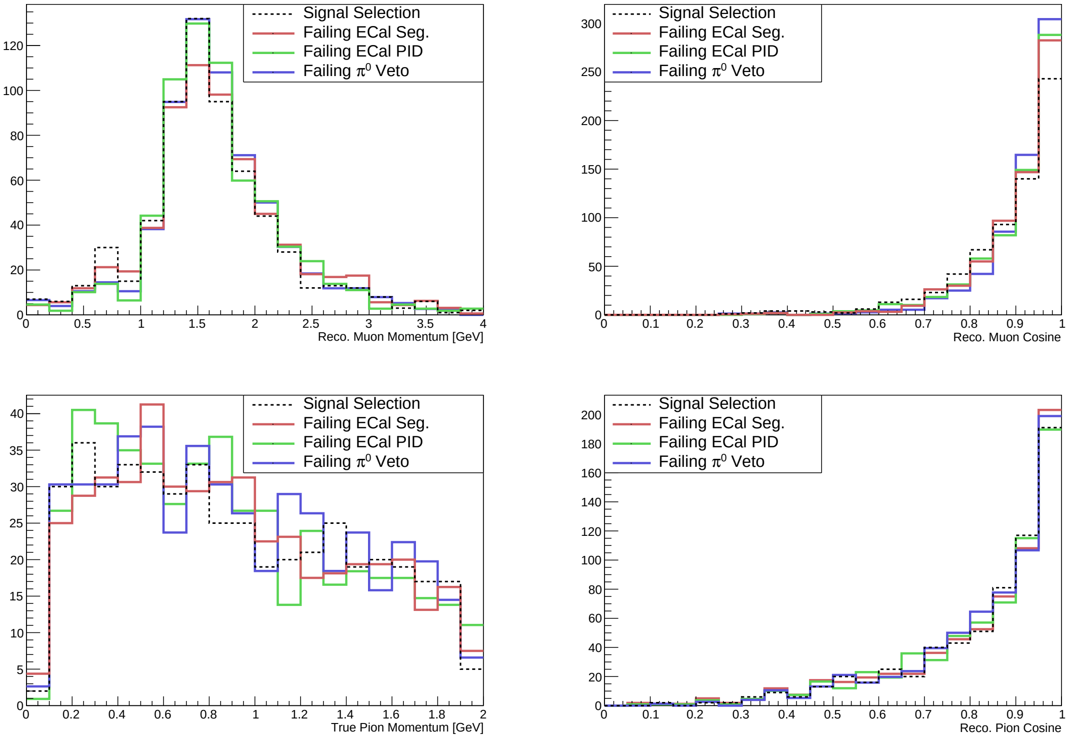

Figure 5.12: Distributions of events with true protons misidentified as anti-muons by failed cut for both FGDs combined. Distributions shown in the four primary reconstructed kinematic variables. All histograms are normalised.

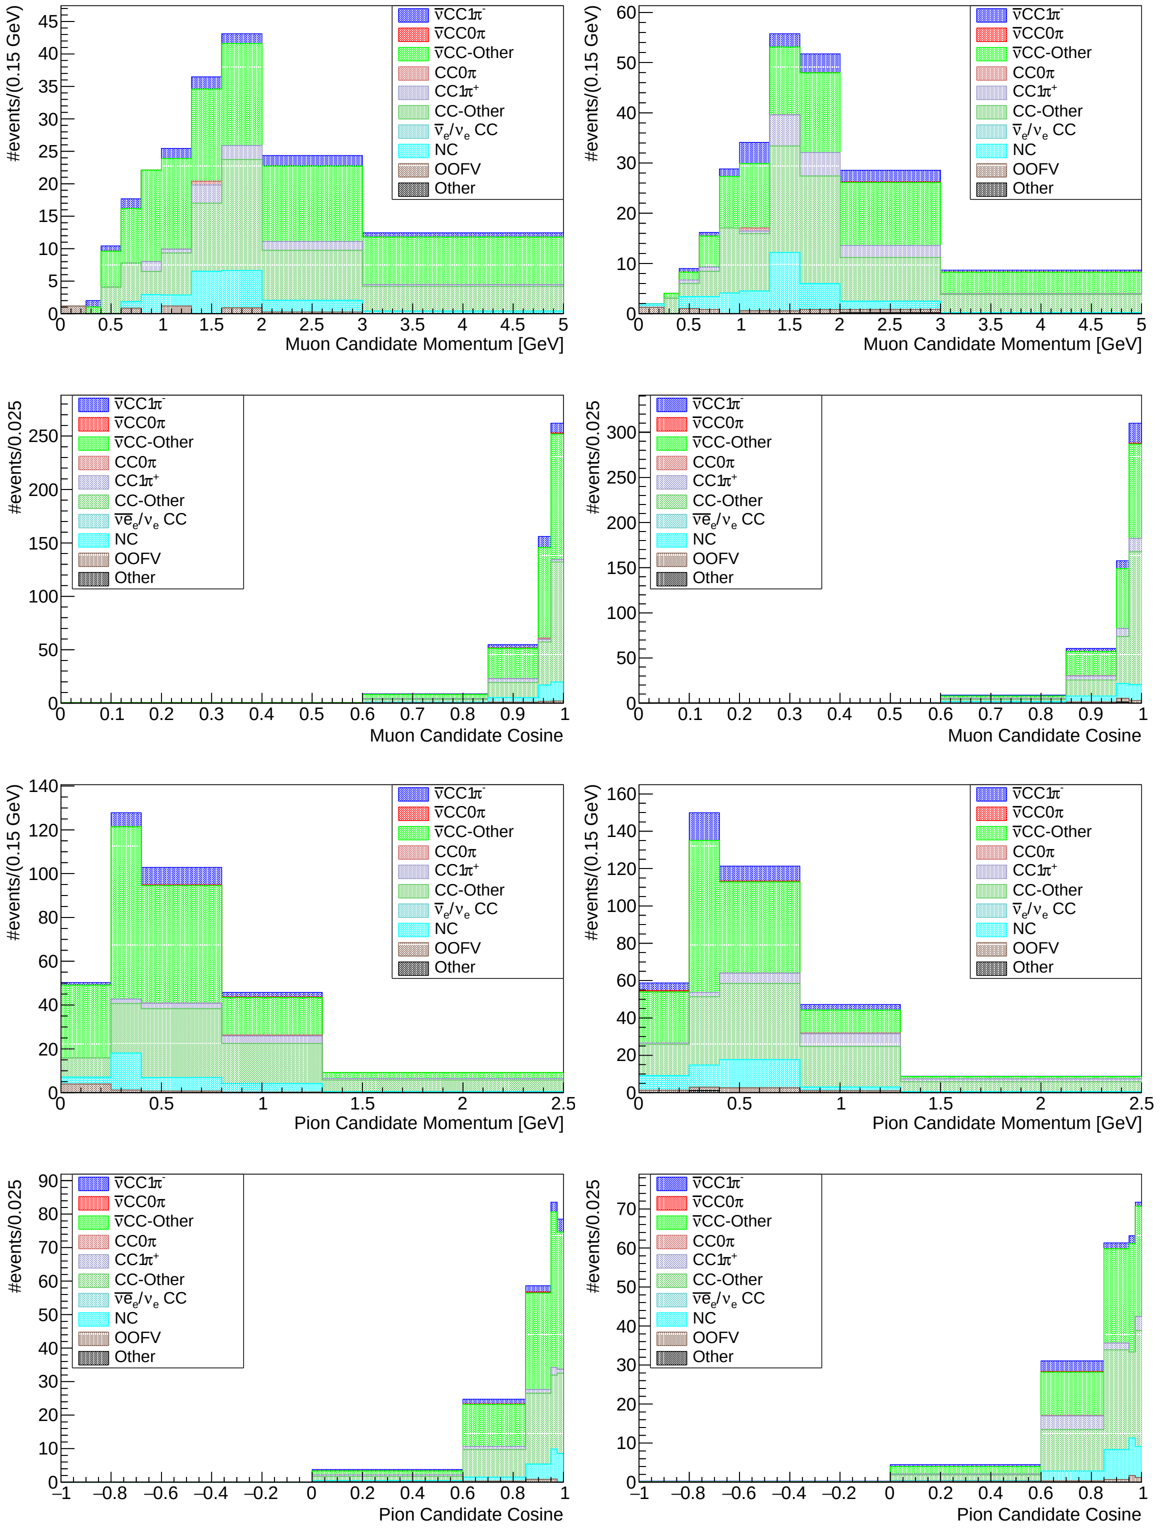

Figure 5.13: Distributions of MC events in muon and pion candidate kinematic variables in the CC-π^-+X Control Regions for FGD1 (*left*) and FGD2 (*right*).

the expected trajectory of a charged particle in a magnetic field — to the global track in question. When a particle enters a TPC, this reconstruction is straightforward as the TPCs provide an accurate measure of position and have a relatively high spatial resolution. When a particle is reconstructed in the FGDs however, momentum reconstruction by fitting a helix to the track is complicated by the short length of these tracks and the relatively low spatial resolution of the FGDs compared to that of the TPCs. The helical fit will often fail to converge for very straight tracks[3] and return a value for the momentum close to the prior value of 1 GeV.

Another way of determining the momentum of a particle is by the range which it travels in a medium. As the energy loss of a given particle in the materials of ND280 is generally known, the range which it is expected to travel can be computed as a function of initial momentum. This momentum by range is well established in the FGDs for muons in previous T2K CC0π analyses [44], but the existence of a TPC track complicates the situation. In the global reconstruction, TPC tracks are used to seed the reconstruction of global tracks first. In the case where two particles originate from a common vertex in an FGD, if one particle enters a TPC and the other does not, all hits in the vertex layer will be assigned to the TPC track object in the reconstruction, resulting in the second track being reconstructed as shorter than it really is. This obviously can present a problem if one wishes to compute the momentum of the track based on the range that the particle travelled. The second issue is that of track sense; the FGD internal timing is not sufficiently precise to determine the sense of tracks without the use of external detectors. Without timing information to inform the reconstruction of track sense, the reconstruction makes the assumption that all tracks are forward going; i.e. that the most upstream end is the start position.

In this analysis — and in the more general case where a particle enters a TPC — the vertex of the event is identified as the start of the anti-muon candidate. For both cosine and momentum reconstruction this information plays a key role, as it gives crucial information about the direction and origin of the particles being reconstructed. The relevant information is that the true direction of the contained track should generally be away from the event vertex, which can be used to correct the track sense if it is incorrect. The true range of the particle should then be from the vertex to the end of the contained track.

---

[3] This issue is compounded by the FGD contained track reconstruction algorithm assuming that tracks should appear close to a straight line.

5.4.1 Momentum Reconstruction

To evaluate the momentum of these tracks, the RECPACK [45] reconstruction toolkit is used. A new wrapper function was written to pass RECPACK the muon candidate vertex position and the end position of the contained track for momentum fitting. The result of this correction is a much improved correlation between reconstructed and true momentum, which can be seen in figure 5.14.

There is a notable number of events seen scattered above the main population in figure 5.14. As the events are primarily above the diagonal, the length of the FGD contained track segment is shorter than one would expect for a pion with such a momentum. The two primary causes for this effect are:

- Pions leaving the FGD where the global reconstruction fails to associate the FGD segment with the other segments of the particle's track.
- Pions scattering and/or absorbing on nuclei in the FGD.

Uncertainties in the relative rate of these effects are included in the detector systematic uncertainty evaluation in section 5.7.

While the number of events with reconstructed momenta being significantly overestimated is rather low, there is at least one reconstruction effect that can cause this. In the case where the muon candidate and a low momentum particle (e.g. a proton) are emitted approximately back-to-back, the hits from the backward-going particle are likely to be assigned to the global track seeded from the muon candidate in the TPC. This has the effect of shifting the reconstructed event vertex away from the true vertex position, and as such artificially increasing the distance from the vertex to the end of the pion candidate track which results in an increased reconstructed momentum. This effect is similarly capable of biasing the momentum of the muon candidate, although this effect is much smaller due to the presence of a TPC track that can be fit with a helix to determine the momentum.

This momentum reconstruction is an important addition to this analysis. In later sections it will be shown that the efficiency depends strongly on pion momentum at low momenta, and the majority of events in the FGD pion sample have pions in the kinematic region where the efficiency varies rapidly.

5.4.2 Cosine Reconstruction

Correcting the sense of the track in this case is relatively easy. As mentioned previously, the true start position of the track should in general be the end closest to the

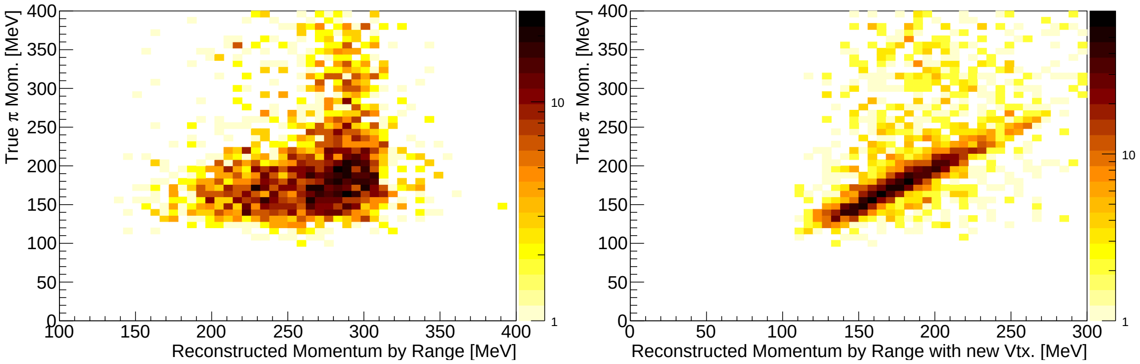

Figure 5.14: Momentum reconstruction using only the global reconstruction track information (*left*) versus improved momentum reconstruction accounting for the event vertex (*right*).

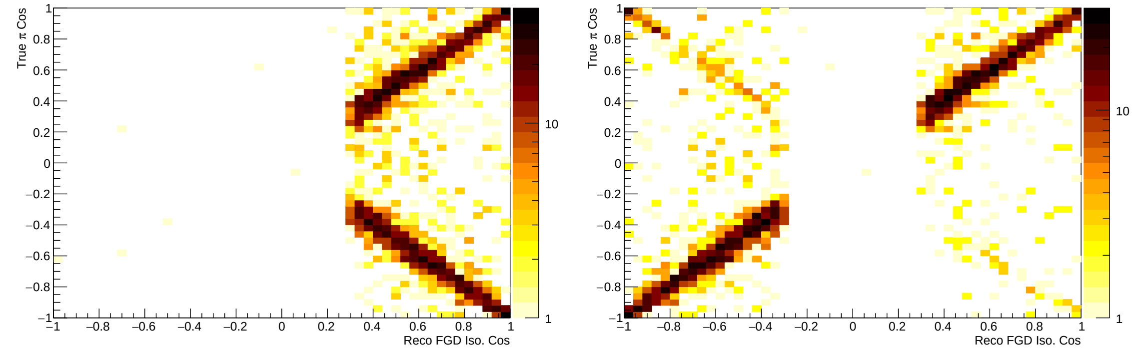

Figure 5.15: Cosine reconstruction using only the global reconstruction track information (*left*) versus improved reconstruction accounting for the event vertex (*right*).

vertex. The value of the cosine from the reconstruction is generally quite good and is difficult to improve upon even with the extra information, so only the sign of the cosine is affected. The improvement can be seen in figure 5.15.

5.5 Sand Muons

Neutrino interactions from beam neutrinos are not restricted to the detector, and a large number of neutrinos will interact in the walls and other surrounding material in the detector complex. A significant number of these interactions will produce (anti-) muons[4] capable of travelling a significant distance from their interaction vertex.

---

[4] Many species of particle are produced in these reactions, but the high rate of (anti-)muon production from a (anti-)muon neutrino beam combined with the high penetration of (anti-)muons results in the vast majority of these particles that enter the detector being (anti-)muons.

When these (anti-)muons propagate into the detector, they are referred to as *sand muons*.

On an analysis level these particles can be problematic as they will appear within the beam window. In the simplest case, a sand muon that enters the ND280 from the outside and stops inside the fiducial volume will leave a track that — excepting time of arrival in each subdetector — can appear identical to a muon of the opposite charge originating in the FGD fiducial volume and propagating out of the detector.

All branches in this selection requires a minimum of two tracks (particles), and this greatly reduces the acceptance of sand muons. The primary way for sand muons to enter this analysis is for the muon stopping in the FGD to decay, with the resulting Michel electron being tagged as a FGD contained pion track. As Michel electrons are often delayed in time slightly from when the (anti-)muon stops and electrons are unlikely to leave hits consistent with a pion track, the rate with which these events enter this selection is very low. For events where the Michel electron is correctly identified as a Michel electron, this can be taken to indicate the presence of a positive pion as discussed in section 5.1.4.

The anti-neutrino beam induced sand muon simulation for the ND280 was run through the selection, with a total of 3.3×10^6 events. Only two of these events passed the selection, with neither event being reconstructed in any signal samples. The expected number of sand events in each sample over the entire data to be used is on the order of 1, and as such they are not considered any further, and no systematic uncertainties on sand muons are applied.

5.6 Number of Targets

Of the six water layers in the FGD2, two have been decommissioned as discussed in section 3.2.2, with one having been drained since the end of run 6 and the most recent failure occurring during the shutdown following run 9. As the data collected since run 9 has been with the beam in FHC mode, the latter panel failure has no impact on this analysis. As the first panel was decommissioned between runs 6 and 7, runs 5 and 6 have data collected with 6 water layers while runs 7 and 9 have data collected from 5 layers. With all other factors being equal, runs 7 and 9 can be expected to have an event rate for water interactions equal to $\frac{5}{6}$ of that of runs 5 and 6.

In the computation of a cross-section value using equation 4.3, the number of targets T is used alongside the integrated neutrino flux Φ. The product $T\Phi$ is the

Run	Total POT	Good POT	ND280 Efficiency
5	7.61×10^{19}	6.88×10^{19}	90.47%
6	3.79×10^{20}	3.59×10^{20}	94.76%
7	4.02×10^{20}	2.83×10^{20}	70.58%
9	9.09×10^{20}	2.30×10^{20}	25.34%

Table 5.1: Total ND280 POT and POT where there are no data quality issues for the runs used in this analysis.

‘luminosity’, L, of the experiment. Where T and the instantaneous neutrino flux are functions of time, their product must be considered as a function of time, such that

$$L = \int T(t)\phi(t)dt. \tag{5.3}$$

As T is a constant within each T2K run, the luminosity of a single run is simply $T_i\Phi_i$, and the total luminosity can be written as

$$L = \sum_{i}^{\text{runs}} T_i\Phi_i = T_{5,6}\Phi_{5,6} + T_{7,9}\Phi_{7,9} \tag{5.4}$$

where T_i and Φ_i are the number of targets and integrated flux for run i respectively. The individual POT for each run, together with the data collection efficiency is shown in table 5.1. The nominal numbers of CH and H_2O targets are:

$$T_{\text{CH}} = (7.439 \pm 0.036) \times 10^{29}$$

$$T_{\text{H}_2\text{O}} = (2.581 \pm 0.019) \times 10^{29}$$

5.7 Detector Systematic Uncertainties

As the majority of the cuts and actions performed in the selection are imported from existing selections, most of the relevant systematics are already applied in the selection. The cuts added to improve purity — ECal Segment requirement, ECal PID, ECal π^0 veto — require extra systematics not included in the standard ND280 selection. Table 5.2 lists the detector systematics enabled in HighLAND2 for this analysis.

The ECal related systematics will be discussed here, with the remaining systematics shown and discussed in Appendix A.1. The effects of individual systematic

Systematic	**Reference**	**Used in TN-273**
TPC-ECal Matching	T2K-TN-279	No
ECal PID	T2K-TN-279	No
ECal π^0 Veto Pileup	T2K-TN-270	No
TPC Cluster Eff.	T2K-TN-212	Yes
TPC Momentum Resolution	T2K-TN-212	Yes
TPC Momentum Scale	T2K-TN-212	Yes
TPC PID	T2K-TN-212	Yes
TPC Tracking Eff.	T2K-TN-212	Yes
TPC Charge ID Eff.	T2K-TN-212	Yes
B-Field Distortions	T2K-TN-212	Yes
TPC-FGD Matching	T2K-TN-212	Yes
FGD PID	T2K-TN-212	Yes
FGD Hybrid Tracking Eff.	T2K-TN-212	Yes
OOFV Background	T2K-TN-212	Yes
Pile-up	T2K-TN-212	Yes
Proton Secondary Interactions	T2K-TN-212	Yes
Pion Secondary Interactions	No Technote	Yes

Table 5.2: Table of detector systematics used in HighLAND2. Corresponding T2K technical notes are [67], [68], and [69]. T2K Technical Note 279 forms the basis for the majority of T2K's current CC neutrino analyses, and as such systematics included in it are well understood in this context.

uncertainties on a selection can be simulated in HighLAND2 [70]. For a given event binning, the distribution of event weights or event number can be used to determine the uncertainty on a bin's event rate, under the assumption that the uncertainty in a bin is Gaussian. In the following subsections and in Appendix A.1, the relative uncertainty is plotted against reconstructed event kinematics in an arbitrary binning.

5.7.1 TPC-ECal Matching Efficiency

The efficiency of matching TPC tracks to their corresponding ECal segments becomes an issue if the ECal information is to be used, and all the more so if a cut is made on the presence or absence of a segment. This systematic has previously been evaluated in TN-279 [67]. To evaluate this uncertainty, events coincident with the beam are selected where a TPC track with at least 18 nodes is present. The particles satisfying this selection are split into three groups based on their pulls, forming three samples that are rich in muons, electrons, and protons respectively. For each sample, each particle is tracked through the TPC to find if it exits the TPC, and if so, whether or not it is incident on an ECal. For particles determined to be incident on an ECal, any ECal hits are present in the expected incident region are assumed to correspond to the particle. As such, a ratio of such particles with ECal segments to the total selected particles provides a measure of the efficiency. This measure includes the inherent efficiency of reconstructing ECal tracks (whether or not they are matched). The efficiency is evaluated for both MC and data, and a systematic uncertainty applied based on the agreement and uncertainty. Figure 5.16 shows the uncertainty from this systematic as a function of reconstructed muon candidate momentum for the four selected samples for FGD2.

5.7.2 ECal PID

As with the TPC-ECal matching efficiency systematic, high-purity samples of muons, protons, and electrons reconstructed in the ECal are used. Unlike the previous case however, each track must have a track segment in an ECal and a TPC, such that the TPC PID provides a benchmark for that of the ECal. Cosmic muons are also selected to increase both the event rate and angular coverage of the muon samples. For each PID likelihood variable, the efficiency to pass the cut (at 0 for each PID variable) is evaluated as the difference between the data and MC performance, accounting for statistical error. The error from this uncertainty is shown in figure 5.17 against anti-muon candidate momentum for the four FGD2 samples.

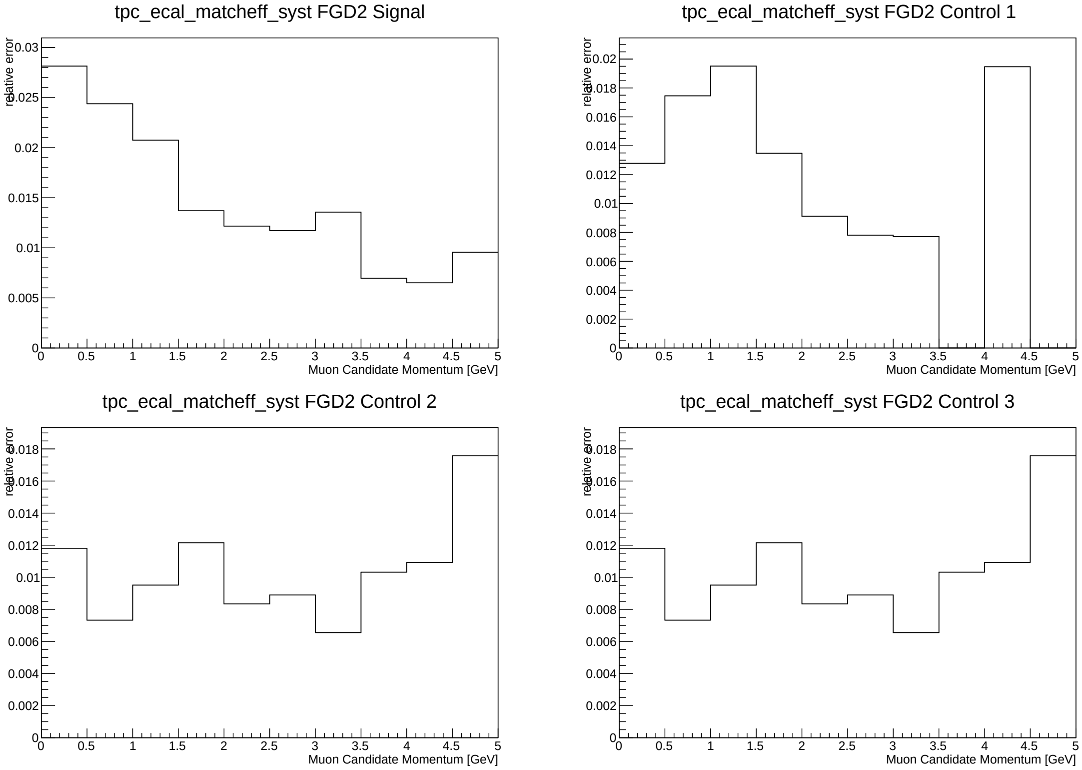

Figure 5.16: Relative error from TPC-ECal matching uncertainties as a function of reconstructed μ^+ candidate momentum.

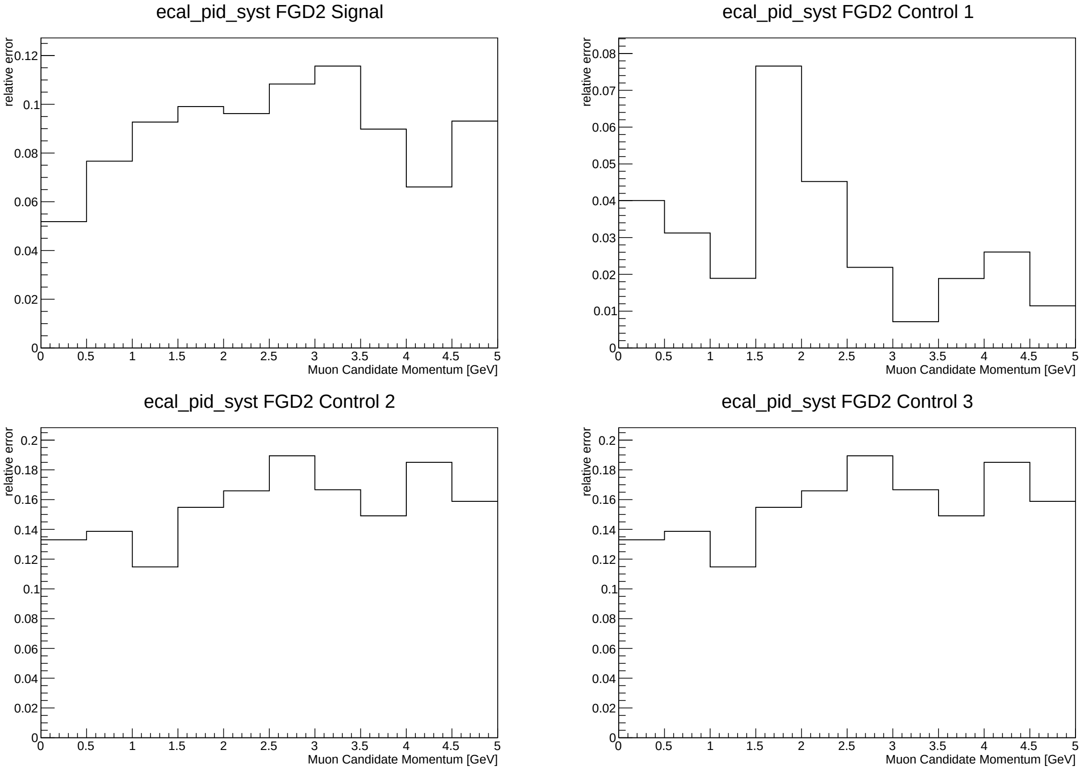

Figure 5.17: Relative error from ECal PID uncertainties as a function of reconstructed μ^+ candidate momentum.

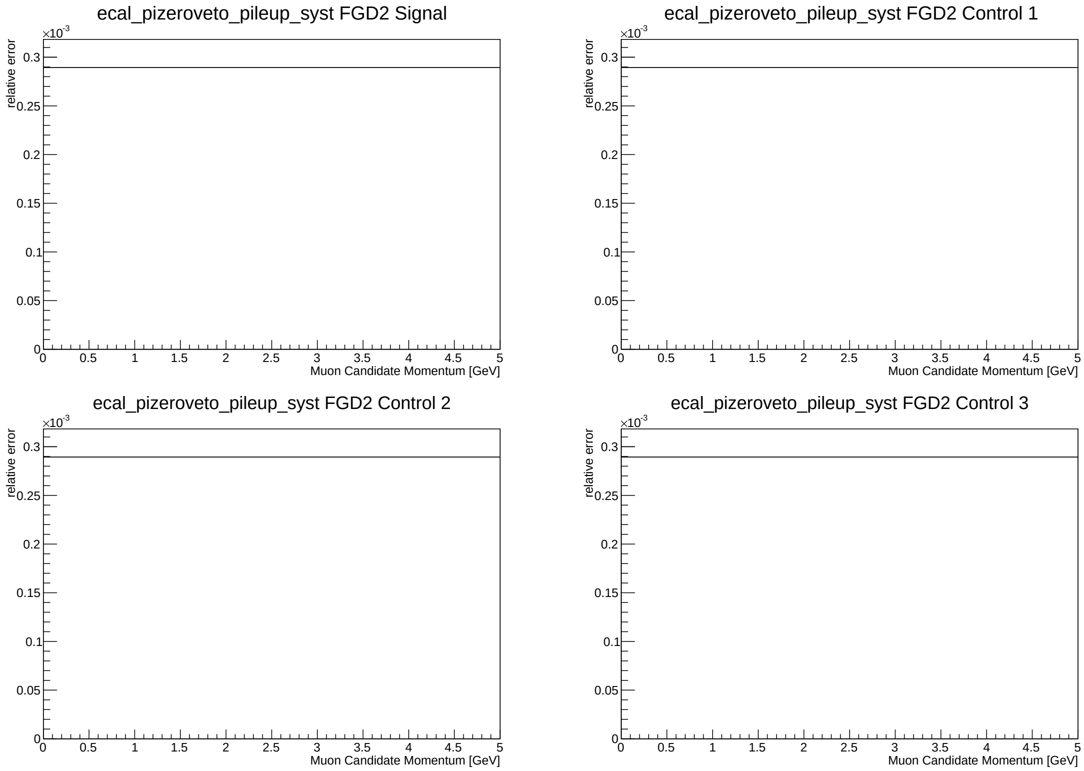

ecal_pizeroveto_pileup_syst FGD2 Signal

ecal_pizeroveto_pileup_syst FGD2 Control 1

ecal_pizeroveto_pileup_syst FGD2 Control 2

ecal_pizeroveto_pileup_syst FGD2 Control 3

Figure 5.18: Relative error from ECal π^0 reconstruction uncertainties as a function of reconstructed μ^+ candidate momentum.

5.7.3 ECal π^0 Veto

As π^0s are inferred through the conversion of photons, inter-detector matching is not a factor, and the reconstruction efficiency for these events is purely the efficiency to reconstruct a π^0 decay photon in the ECals. As ECal PID variables and the reconstructed energy of the shower are also cut on, the cumulative effects of the cut are analysed, with the pass/fail efficiencies evaluated in data and MC for the most common particles. The error from this uncertainty is shown in figure 5.18 against anti-muon candidate momentum for the four FGD2 samples. This uncertainty can be seen to be independent of anti-muon candidate momentum, which is largely expected as this uncertainty results from backgrounds unrelated to the event being considered.

5.7.4 Total Detector Systematic Uncertainty

With all sources of detector error considered simultaneously, a full detector error matrix can be constructed with correlations. Figure 5.19 shows projections of this total detector systematic uncertainty in reconstructed muon candidate momentum

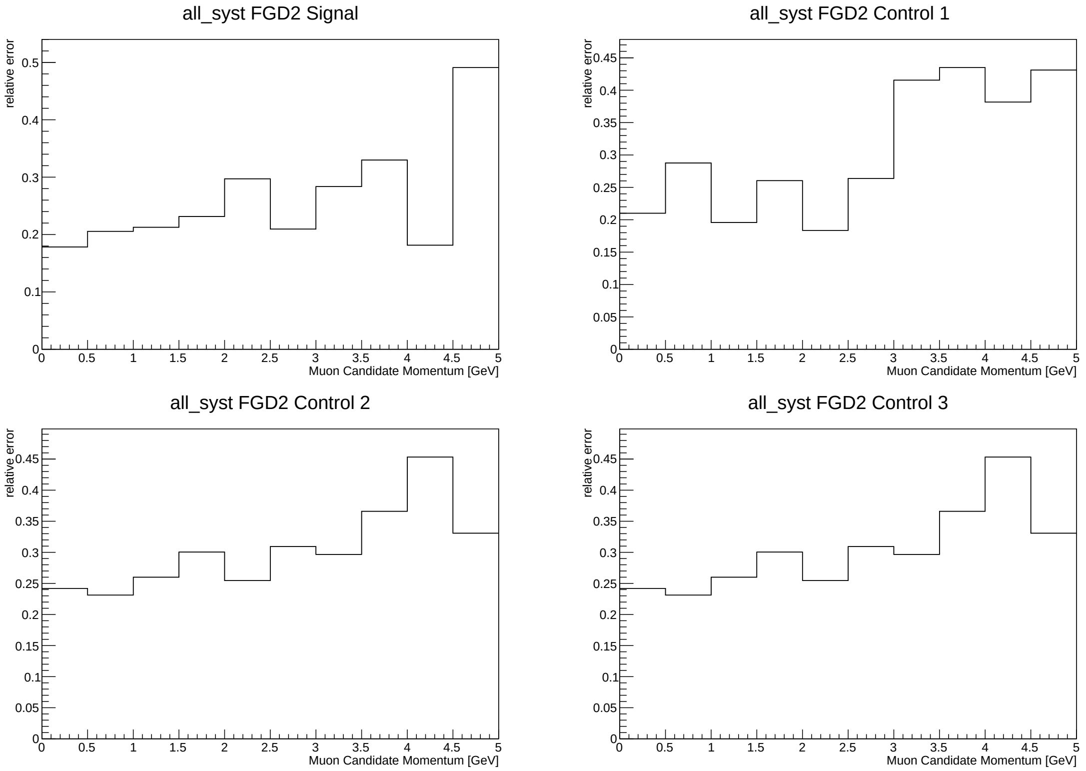

Figure 5.19: Relative error from all detector uncertainties combined as a function of reconstructed μ^+ candidate momentum.

for the four FGD2 selection samples.

5.8 Model Systematic Uncertainties

The neutrino-nucleus interaction models used in the detector simulation are known to be imperfect, and uncertainties for many of the simulation parameters have been determined from fits to sets of neutrino-nucleus scattering data. The majority of free parameters in the models used for the MC simulation can be ‘reweighted’ to. Examples of effects which cannot be trivially reweighted are particle momenta and related effects, as small changed in particle momentum can drastically change the appearance of an event in a detector, or cross a cut threshold. Reweighting is the application of weights to individual events to reflect a change in their probability of occurring as the model parameter varies. Where a model parameter variation can be reweighted, the T2KReWeight software facilitates this reweighting for ND280 event simulations. Each uncertainty is varied through a “dial”, where a dial generally scales the underlying model parameters. Most dials in T2KReWeight have their nominal at

Parameter Name	**Type**	**Prior**	**Error**
C_A^5	Signal	1.01	0.12
M_A^{RES}	Signal	0.95	0.15
$I = \frac{1}{2}$RES	Signal	1.3	0.2
CC-Coherent	Signal	1.0	0.15
CC Multi-Pi	Background Norm.	1.0	0.25
DIS Multi-Pi	Background Shape	1.0	0.4
DIS Norm. ν_μ	Background Norm.	$1,0$	0.25
DIS Norm. $\bar{\nu}_\mu$	Background Norm.	$1,0$	0.35
Bodek-Yang DIS	Background Shape	0.0	1.0
AGKY DIS	Background Shape	0.0	1.0
AGKY Multi-Pi	Background Shape	0.0	1.0
M_A^{QE}	Background Shape	1.21	0.2
FSI Pion Prod.	FSI	1.0	0.5
FSI Pion Abs.	FSI	1.0	0.41
FSI Inel. Low	FSI	1.0	0.41
FSI Inel. High	FSI	1.8	0.34
FSI Chrg. Exchg. Low	FSI	1.0	0.57
FSI Chrg. Exchg. High	FSI	1.8	0.28

Table 5.3: List of T2KReWeight dials used in this analysis, with corresponding priors and errors.

1, with a associated uncertainty, usually implemented as a fractional uncertainty. In chapters 7 and 8, all T2KReWeight dial values are normalised when shown grapically, such that all uncertainties appear fractional for clarity. The T2KReWeight dials used in this analysis are shown in table 5.3.

5.8.1 Signal Dials

The C_A^5, $M_A^{textRES}$, and $I = \frac{1}{2}$RES dials effect the resonant pion production model used in the NEUT simulation, which provides the majority of signal events. C_A^5 is a normalisation factor for a poorly constrained form-factor under the Graczyk-Sobczyk parameterisation which effects both the normalisation and shape of the model prediction. M_A^{RES} is the axial mass, which also scales the model normalisation, but also has modest impact the shape. $I = \frac{1}{2}$RES is a normalisation only dial which directly scales the non-resonant pion production background used in the model. The CC-Coherent dial is a simple normalisation applied to all coherent CC events.

5.8.2 Background Dials

The CC Multi-Pi and DIS Normalisation dials for ν_μ and $\bar{\nu}_\mu$ all apply a normalisation to these events. Here the distinction between multi-pion and DIS is made at 2 GeV, where DIS is above this threshold. The DIS Multi-Pi dial varies the shape of the predicted distribution by applying a weighting based on the energy of the neutrino. M_A^{QE} is the axial mass in the prediction for CCQE events, which varies both the shape and normalisation of the predicted distribution.

5.8.3 Final State Interaction Dials

NEUT propagates pions through the nuclear medium using a cascade model. For a given event, the start position within the nucleus is randomly chosen, with weighting based on the predicted nuclear density profile. From this start position, each pion is propagated outwards in small time steps, with a probability for interactions at each step based on the pion kinematics and local nuclear density. When an interaction occurs, the kinematics of the pion can be altered, or the pion can be absorbed, or another produced. The simulation propagates until all remaining pions have left the nucleus or stopped. Pion interaction probabilities are based on the model in [71] for pions with $p_\pi < 500$ MeV, where for pions above this threshold interaction probability is tuned to pion-proton and pion-deuteron scattering data from the PDG [14]. In the region of 500 MeV these models are modified slightly such that there is no discontinuity in interaction probability.

Uncertainties related to this treatment are parameterised in the six FSI dials in table 5.3. The dials affect the probabilities for different interactions possible within the simulation. The first two dials directly modify the probability for a pion to be produced and absorbed respectively. The third and fourth dials scale the probability for a pion to inelastically scatter at low and high energies respectively, with the fifth and sixth dials scaling the probability for a pion to undergo charge exchange at low and high energies.

Chapter 6

Binning and Efficiency

6.1 Efficiency

Considering a single particle from a given vertex at a given momentum in a detector, it should either be detected or not. As such, efficiencies — when plotted against against the particle's kinematics — should be step-functions, with a turn-on when the detector becomes sensitive to the particle. Efficiency in this sense depends — in some parts of phase space — quite heavily on variables such as vertex position in the FGDs (i.e. x, y, z), as well as the specific direction in which the particles travel. As such, when considering the efficiency against some kinematic variable, it is implicitly integrated over the detector fiducial volume and all other kinematics, resulting in a curve that is distinctly not that of a step function.

As a neutrino interaction generator will not predict detector-level variables, this information must be taken into consideration for the comparison of data and theory. In this analysis there are two particles, and consequently a phase space that can be largely decomposed into 5 dimensions; the momenta and angle of the two particles, along with the angle between their vectors. The efficiencies in each of these variables is shown for FGD1 in figure 6.4 and for FGD2 in figure 6.3. In each of these plots, the phase-space restrictions discussed in section 6.2 are applied for all variables other than those being plotted, and as such the efficiencies shown appear slightly different to those in section 6.2.

6.2 Signal Phase-Space Restriction

While the goal is to measure the cross-section of $\bar{\nu}_\mu$-CC1π^-, the detector determines what parts of the phase-space can be measured. It is obviously counterproductive to attempt to measure the cross section in parts of the phase-space where a detector has no or very limited efficiency. Beyond the reasoning that it is difficult to do so, events in these regions will be subject to a large, possibly heavily model-dependent efficiency correction. A more detailed discussion of the dangers of correcting for poorly understood or rapidly changing efficiencies can be found in [62].

In the angular efficiency plots in figures 6.1 and 6.2, it can be seen that the efficiency is very low for both muons and pions below $\cos\theta = 0.3$, with the efficiency for pions returning somewhat below $\cos\theta = -0.3$. In the case of pions, the majority at high angles are reconstructed as FGD isolated tracks, and systematic errors for pions at high angles ($-0.3 < \cos\theta < 0.3$) are not well understood due to the lack of a control sample. In momentum space, both pions and muons show a sharp turn-on in their efficiency curves, becoming somewhat stable at around 250 MeV. While both NEUT and GENIE show largely similar behaviour across all variables, there is discrepancy observed at low muon momenta, with GENIE predicting a higher efficiency. As this region of phase-space is outside of the signal region in this analysis, this is not considered problematic, as no efficiency correction is applied in this region. These criteria form the signal definition presented in chapter 4 in the restricted phase-space, and the signal now becomes events with one μ^+ and one π^-, with no other mesons in the final state, where

- $p_{\mu^+} > 300$ MeV
- $\cos\theta_{\mu^+} > 0.5$.
- $p_{\pi^-} > 140$ MeV
- $\cos\theta_{\pi^-} \in [-1, -0.35] \cup [0.35, 1]$.

In the angle between the two particles, the efficiency can be seen to vary rather slowly for the majority of the range. There is a drop at very small angles due to the reconstruction of near-collinear particles being difficult. When the efficiency is plotted later with these phase-space constraints accounted for in figures 6.3 and 6.4, it is flatter.

6.3 Efficiency for Restricted Phase-Space Signal

The efficiencies in each of the 5 variables is shown for FGD2 in figure 6.3 and for FGD1 in figure 6.4. In each of these plots, the phase-space restrictions discussed in section 6.2 are applied for all variables other than those being plotted, and as such

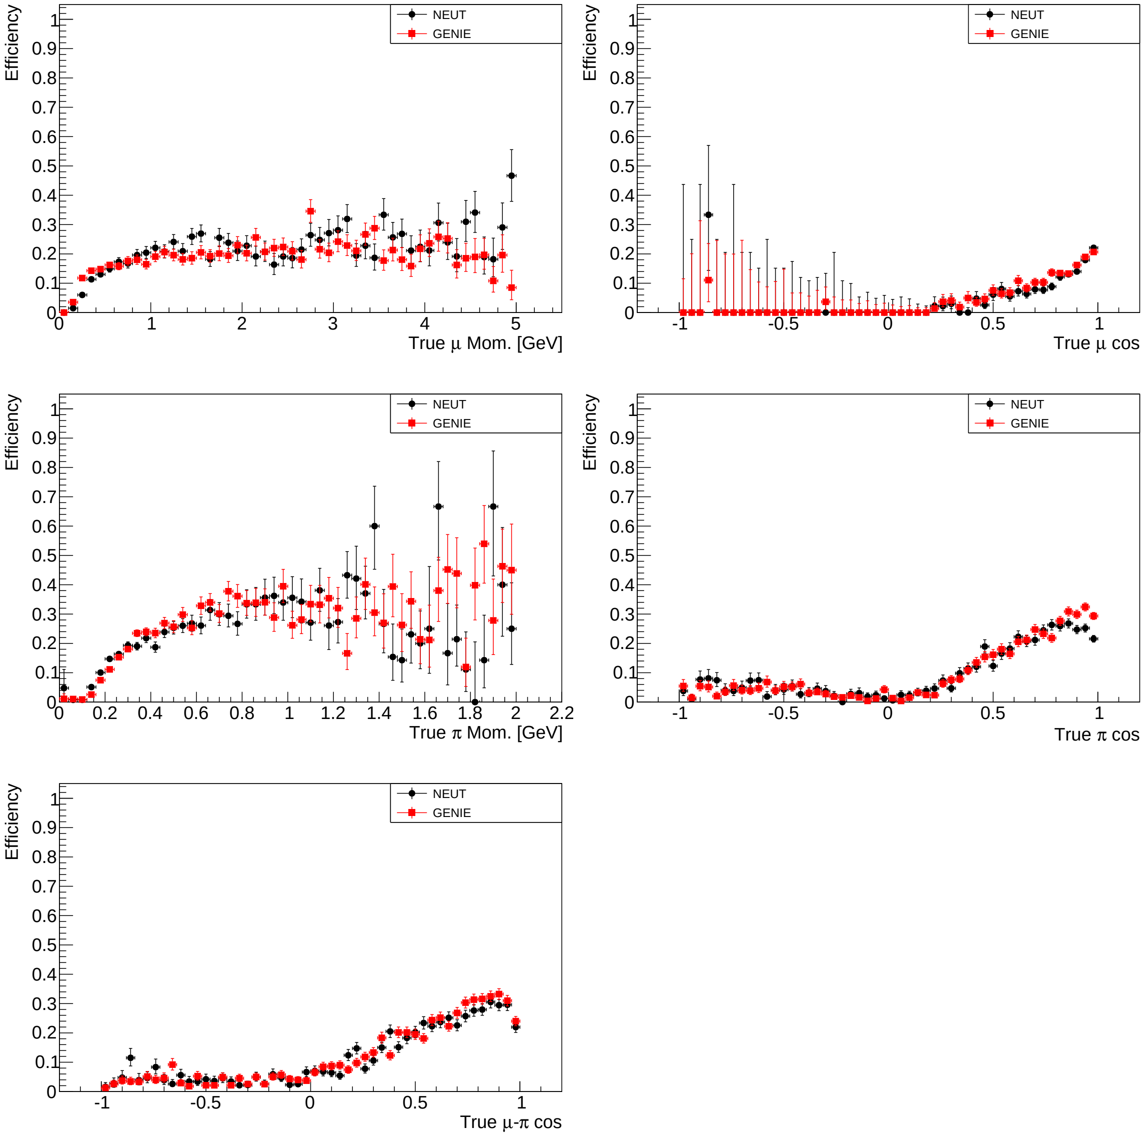

Figure 6.1: Efficiencies in the 5 kinematic variables for the FGD2 selection, for both NEUT and GENIE MC. Error bars are statistical only.

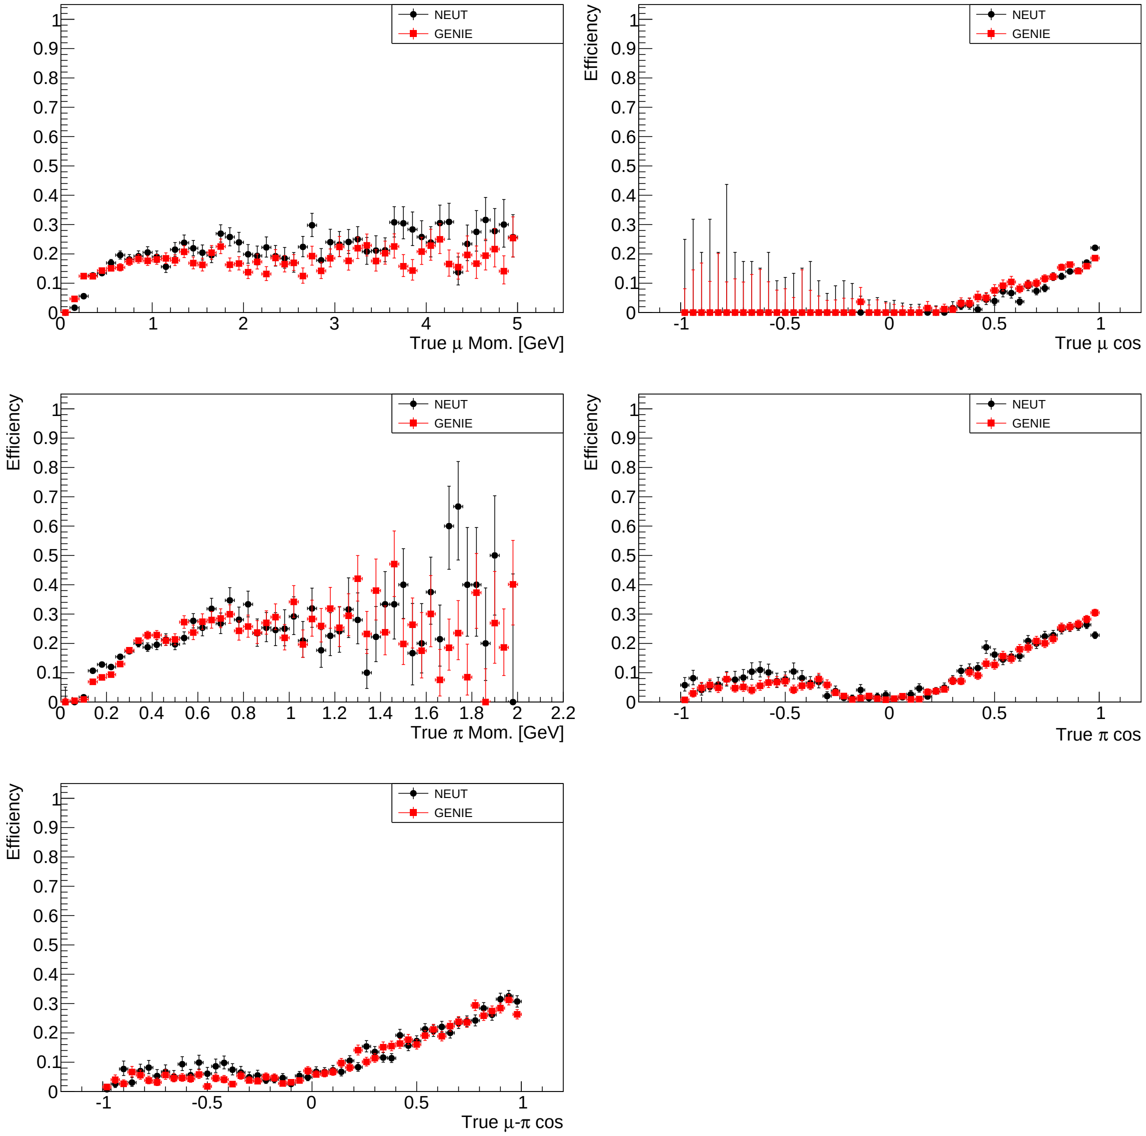

Figure 6.2: Efficiencies in the 5 kinematic variables for the FGD1 selection, for both NEUT and GENIE MC. Error bars are statistical only.

the efficiencies shown are different to those in that section. Compared to figures 6.1 and 6.2, some of the features have been amplified. In particular, the momentum turn-on becomes much more prominent, and the overall shape of the $\mu - \pi$ angular distribution becomes much smoother.

6.4 Binning

The selection of bins in an analysis is one of compromise. To maximise the usefulness of the result, finer bins are desirable to resolve the shape of the cross-section and distinguish various models. The primary limit is often the number of events that can be collected with a detector, and a binning should be coarse enough that statistical fluctuations are minimised, or at least a sub-dominant error.

In choosing a binning in which to present a cross-section result, the following factors should be taken into consideration:

- To mitigate the effects of statistical fluctuations on the final cross-section, a number of reconstructed events on the order of at least 100 per bin brings the statistical error to a similar level as the other errors (e.g. flux, detector systematics, etc.).

- The transfer matrix between true and reconstructed bins should be as diagonal as possible. Significant migration of events can cause difficulty in constraining the true kinematics with the reconstructed information. When direct observables (e.g. momentum) are being measured, having the same binning in true and reconstructed variables generally returns a near-diagonal matrix.

- The bin width should always generally be greater than the detector's resolution in that bin. Ignoring this requirement could result in large anti-correlations between neighbouring bins.

- Efficiency should not vary greatly across a bin, as this will make the efficiency correction sensitive to the distribution of true events across the bin, and thus model dependent.

With the available POT from runs 5–9, approximately 400 events are expected in the FGD2 signal selection with vertices reconstructed in an x-layer, motivating a 4 or 5 bin result. The range of possible values of all variables should be accounted for, so that no events are unintentionally left out of the final binned fit. As all angular variables will have their cosine within the interval $[-1, 1]$, there is little else

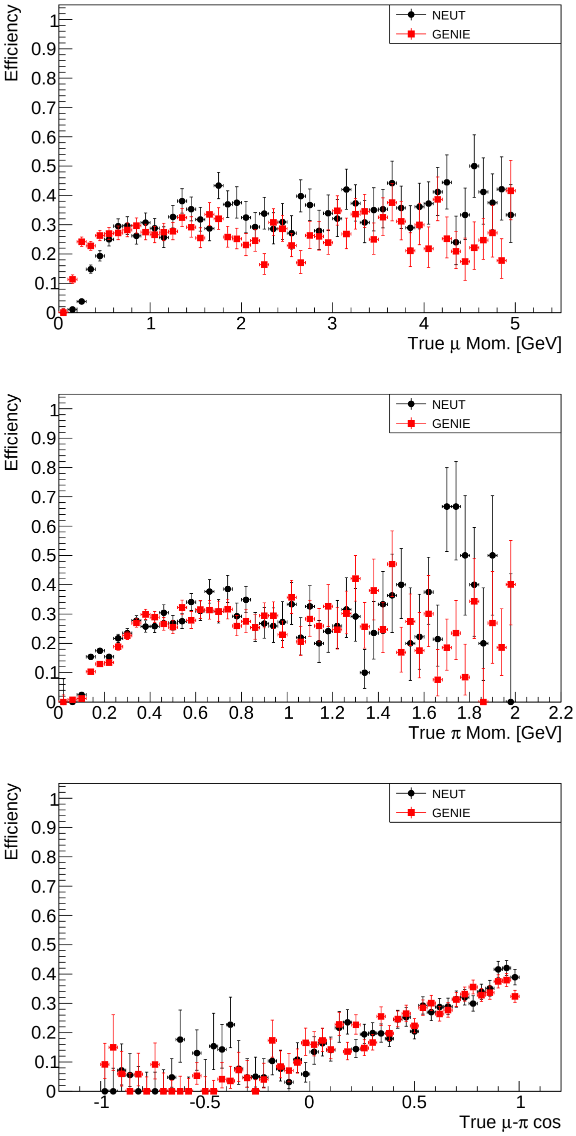

Figure 6.3: Efficiencies in the 5 kinematic variables for the FGD2 selection, for both NEUT and GENIE MC. Phase-space restrictions are applied for all variables other than that being plotted. Error bars represent statistical errors only.

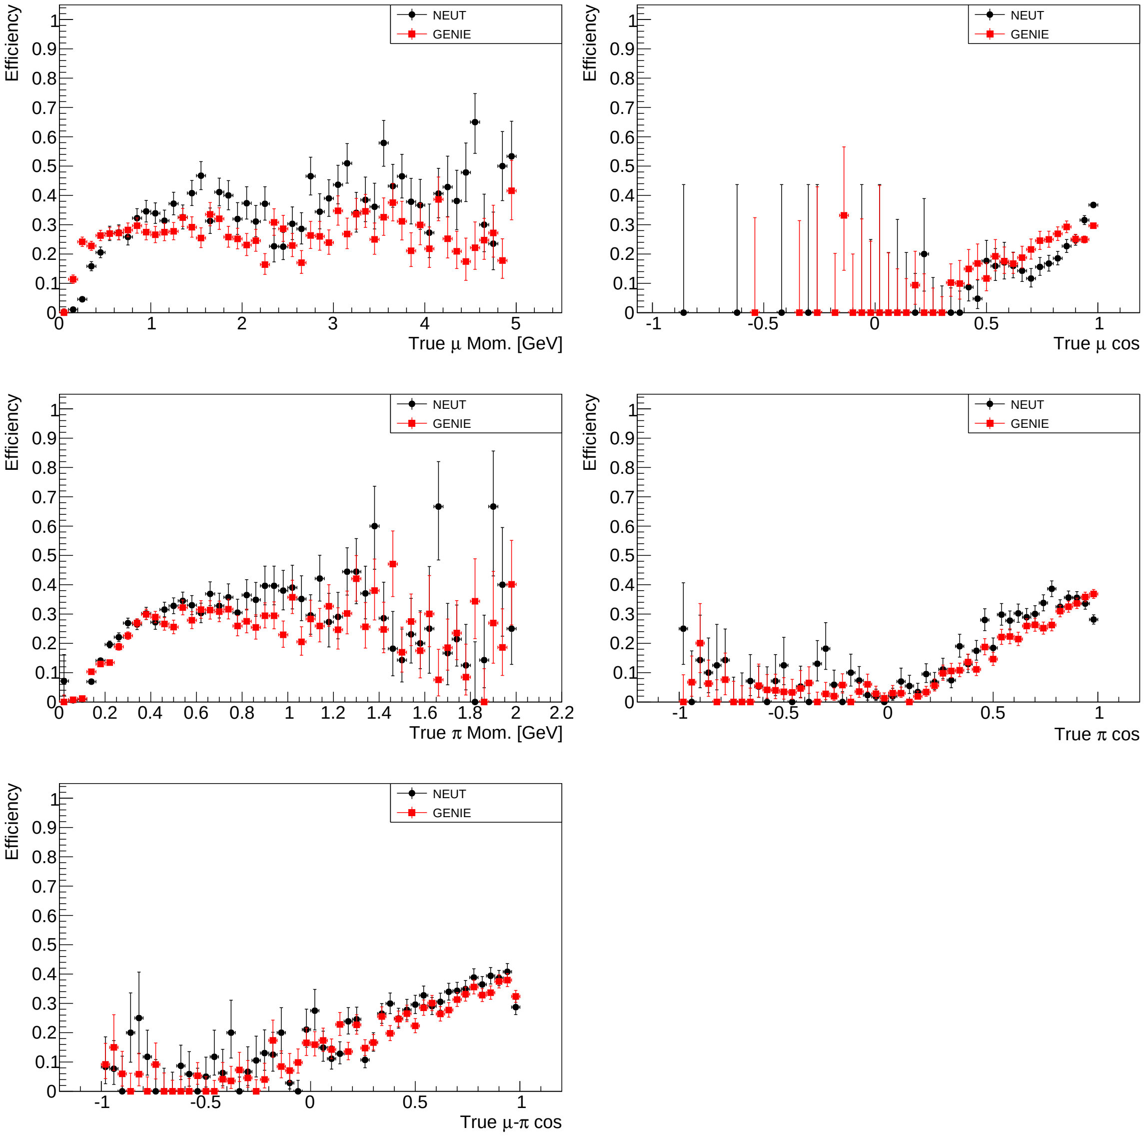

Figure 6.4: Efficiencies in the 5 kinematic variables for the FGD1 selection, for both NEUT and GENIE MC. Phase-space restrictions are applied for all variables other than that being plotted. Error bars represent statistical errors only.

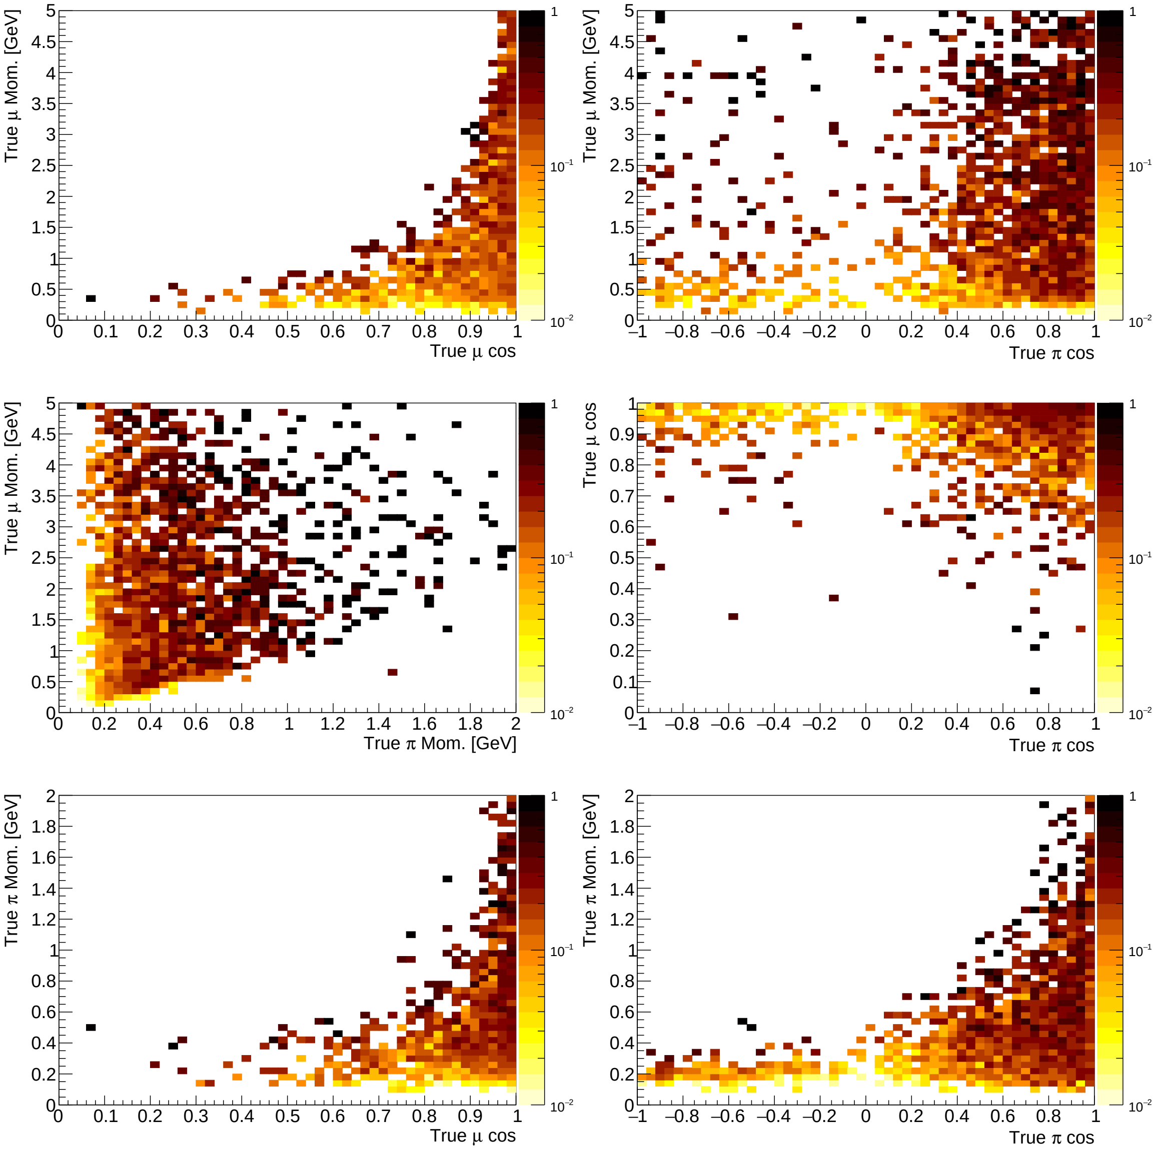

Figure 6.5: 2D Efficiencies in the 4 main kinematic variables for the FGD2 selection. Phase-space restrictions are applied for all variables other than that being plotted.

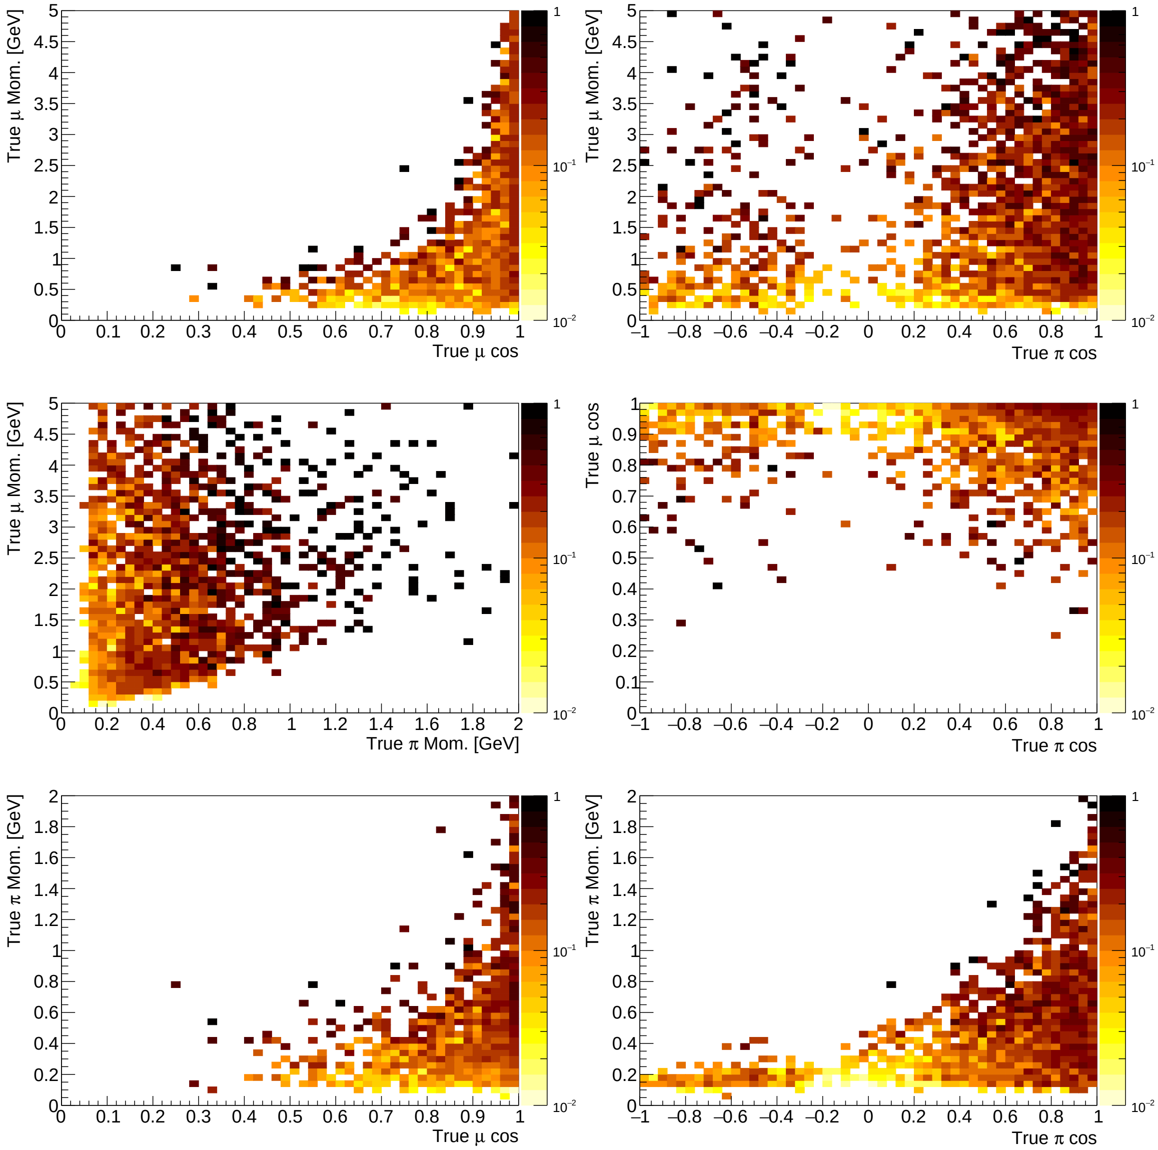

Figure 6.6: 2D Efficiencies in the 4 main kinematic variables for the FGD1 selection. Phase-space restrictions are applied for all variables other than that being plotted.

to do than make sure all such values fall into one and only one bin. In the case of particle momenta, there is slightly more to consider. For the true momenta of the particles in the signal, 30 GeV is the maximum possible momentum. The choice of 30 GeV is motivated by the energy of the protons used to generate the neutrino beam. Reconstructed momenta are determined by fitting a helix to the reconstructed global track, which can conceivably return any numerical value allowed by the function. A value of 5×10^{40} MeV was chosen for the upper bin edges in momentum, as this is slightly above the highest reconstructed value for the momentum of a particle in the selected MC. In the following sections describing the binning schemes, 'max' denotes this value.

6.4.1 Binning Definitions

Table 6.1 shows the bins in each variable for the analysis. For each kinematic bin, there are two corresponding bins, one for water-enhanced FGD2 layers, and one for FGD2 scintillator layers and all FGD1 layers. With these bin edges, a pseudo 4-dimensional rectilinear binning can be formed in which the cross-section fit is performed. After cross-section fitting, three axes will be collapsed to form a 1-dimensional result. The upper limit for the momentum of a single particle is set at 30 GeV in this analysis, as particles with momenta greater than this cannot be produced by neutrinos generated from a 30 GeV proton beam.

Of these bins, a number are in regions of very low event rates, and as such some of these bins are to be merged. The first bin in each variable is the Outside Phase-Space (OPS) bin, where the efficiency is too low or varies too rapidly for the efficiency correction to be reasonable. The fit result in these bins will not be reported in the final cross-section result, and the inclusion of these events is only to give these events a template parameter in the fit, and to ensure that signal events both inside and outside of this restricted phase space are treated the same. As the efficiency is low in these areas, the event rates are very low in the OPS bins, and all bins are merged with one variable in the OPS region into one of four bins with adequate event numbers for the fitter.

The bin merging is done in the following order:

- Events with true p_{π^-} below 140 MeV are placed in the first OPS bin
- Remaining events with true $\cos\theta_{\mu^+}$ below 0.5 are placed in the second OPS bin
- Remaining events with true $\cos\theta_{\pi^-}$ from $-0.35, 0.35$ are placed in the third

Bin Index	p_{μ^+}[MeV]	$\cos\theta_{\mu^+}$	p_{π^-}[MeV]	$\cos\theta_{\pi^-}$
1	0.0–30000.0	−1.0–1.0	0.0–140.0	−2.0–1.0
2	0.0–30000.0	−1.0–0.5	140.0–30000.0	−2.0–1.0
3	0.0–30000.0	0.5–1.0	140.0–30000.0	−2.0–−1.0
4	0.0–300.0	0.5–1.0	140.0–30000.0	−1.0–1.0
5	300.0–30000.0	0.5–1.0	140.0–220	−1.0–0.6
6	300.0–30000.0	0.5–1.0	140.0–220	0.6–0.85
7	300.0–30000.0	0.5–1.0	140.0–220	0.85–0.93
8	300.0–30000.0	0.5–1.0	140.0–220	0.93–1.0
9	300.0–30000.0	0.5–1.0	220–325.0	−1.0–0.6
10	300.0–30000.0	0.5–1.0	220–325.0	0.6–0.85
11	300.0–30000.0	0.5–1.0	220–325.0	0.85–0.93
12	300.0–30000.0	0.5–1.0	220–325.0	0.93–1.0
13	300.0–30000.0	0.5–1.0	325.0–425.0	−1.0–0.6
14	300.0–30000.0	0.5–1.0	325.0–425.0	0.6–0.85
15	300.0–30000.0	0.5–1.0	325.0–425.0	0.85–0.93
16	300.0–30000.0	0.5–1.0	325.0–425.0	0.93–1.0
17	300.0–30000.0	0.5–1.0	425.0–650.0	−1.0–0.6
18	300.0–30000.0	0.5–1.0	425.0–650.0	0.6–0.85
19	300.0–30000.0	0.5–1.0	425.0–650.0	0.85–0.93
20	300.0–30000.0	0.5–1.0	425.0–650.0	0.93–1.0
21	300.0–30000.0	0.5–1.0	650.0–30000.0	−1.0–0.6
22	300.0–30000.0	0.5–1.0	650.0–30000.0	0.6–0.85
23	300.0–30000.0	0.5–1.0	650.0–30000.0	0.85–0.93
24	300.0–30000.0	0.5–1.0	650.0–30000.0	0.93–1.0

Table 6.1: Full kinematic definition for each analysis bin, same binning is applied for each target.

OPS bin

- Remaining events with true p_{μ^+} below 300 MeV are placed in the fourth OPS bin

6.4.2 Detector Level Binning

For the likelihood fitter to determine the best fit parameters for the model, a direct comparison between the model prediction and measured data is needed. This is done by comparing the MC and data in reconstructed variables directly from the detector, without any unfolding performed. As described in section 4.5.1, the unfolding is performed by the inclusion of template parameters in the fit, which add the necessary freedom for the fit in reconstructed variables to constrain true quantities. To assist in the unfolding, analysis and reconstructed variables that correlate strongly should be chosen, with binning schemes that complement one another. The reconstructed binning should also be finer in general than the analysis binning, as unfolding is an ill-posed problem.

The binning for each detector sample is considered based on the purpose of that sample in the fit. For the signal samples (TPC π, FGD π), the primary goal is to constrain the signal parameters, and so the binning for these samples should correlate with the analysis bins. As the purpose of the control regions is to determine the background strength and shape, it is not necessary to have a direct relationship to the analysis bins.

As the analysis is performed using the variables $p_{\mu^+}, p_{\pi^-}, \cos\theta_{\mu^+}, \cos\theta_{\pi^-}$ and target, the reconstructed samples should inform the fit about these quantities. As such, the reconstructed quantities $p_{\mu^+}^{\text{reco}}, p_{\pi^-}^{\text{reco}}, \cos\theta_{\mu^+}^{\text{reco}}, \cos\theta_{\pi^-}^{\text{reco}}$ are used as these variables correlate heavily with the true quantities. In lieu of a direct way of determining target, the FGD layer information is passed to the fitter, with interactions being reconstructed in either FGD1, FGD2x, or FGD2y layers. The water component of the signal should be primarily constrained by FGD2x and FGD2y layers, with the hydrocarbon component constrained primarily by FGD2y and FGD1 layers.

For consistency the same binning scheme is used for each FGD layer selection subsample, for example all TPC pion selected events will have the same nominal binning regardless of the FGD layer in which they occurred. In some cases, this criterion results in bins with low entries, for which evaluating the likelihood and estimating the uncertainty become difficult. Wherever a bin with 4 or fewer entries presented itself, the offending bin was merged with a neighbouring bin. Bins with low numbers of events present an issue for this analysis in two distinct ways. A limited number of events in a bin reduces the accuracy of the correlations computed between it and the true bins. The other potential issue is in the evaluation of uncertainties; as the fit machinery assumes all errors are Gaussian at multiple points in the fit, any deviation from a Gaussian has the potential to introduce bias. As the uncertainty in counting individual events is Poissonian in nature, for very low event numbers the assumed Gaussian uncertainty will deviate from the true uncertainty. The influence of these effects will be evaluated on the analysis in the next chapter.

Pion Cosine Phase Space

Three of the four variables on which phase space restrictions are imposed have a single value at which the cut is made. For $\cos\theta_{\pi^-}$ however, events within the range $(-0.35, 0.35)$ are OPS. In general, it is desirable to have reconstructed bins capable of constraining the corresponding true OPS bins. To enable easy binning in this reconstructed variable for both OPS and normal signal, events with reconstructed $\cos\theta_{\pi^-}$ in the range $(-0.35, 0.35)$ will have -1.5 added to their value, such that these events are now equivalently distributed in the range $(-1.85, -1.15)$. This ensures that any reconstructed bin spanning the normal OPS range will have no reconstructed OPS events in the bin, and that all of these events can be found in $(-2, -1)$. A number of bins can then be chosen to accept the events with $\cos\theta_{\pi^-}$ in this new range.

TPC Pion Samples

The majority of the events selected for this analysis are selected with both the pion and muon in a TPC, and the majority of the kinematic phase space for this analysis is measured with the TPC. As this comprises the bulk of the signal information the fitter is to receive, this binning scheme should be at least as fine as the analysis binning for the majority of the phase space. As the analysis binning is primarily in $p_{\pi^-}^{\text{true}}$ and $\cos\theta_{\pi^-}^{\text{true}}$ the majority of the information is encoded in $p_{\pi^-}^{\text{reco}}$ and $\cos\theta_{\pi^-}^{\text{reco}}$. The TPC pion bins are shown in table 6.2.

FGD Pion Samples

Pions detected in the FGD are generally of lower momenta and/or high angle, and as such the binning of these events should be adapted to the relevant kinematic region. As previously, the majority of the information is encoded in p_{π^-} and $\cos\theta_{\pi^-}$. The

Bin Index	$p_{\mu^+}^{\mathrm{reco}}$[MeV]	$\cos\theta_{\mu^+}^{\mathrm{reco}}$	$p_{\pi^-}^{\mathrm{reco}}$[MeV]	$\cos\theta_{\pi^-}^{\mathrm{reco}}$
1	0.0–max	−1.0–1.0	0.0–140.0	−2.0–1.0
2	0.0–max	−1.0–0.5	140.0–max	−2.0–1.0
3	0.0–300.0	0.5–1.0	140.0–max	−2.0–1.0
4	300.0–max	0.5–1.0	140.0–200.0	−1.0–0.4
5	300.0–max	0.5–1.0	140.0–200.0	0.4–0.7
6	300.0–max	0.5–1.0	140.0–200.0	0.7–0.8
7	300.0–max	0.5–1.0	140.0–200.0	0.8–0.93
8	300.0–max	0.5–1.0	140.0–200.0	0.93–1.0
9	300.0–max	0.5–1.0	200.0–275.0	−1.0–0.4
10	300.0–max	0.5–1.0	200.0–275.0	0.4–0.7
11	300.0–max	0.5–1.0	200.0–275.0	0.7–0.8
12	300.0–max	0.5–1.0	200.0–275.0	0.8–0.93
13	300.0–max	0.5–1.0	200.0–275.0	0.93–1.0
14	300.0–max	0.5–1.0	275.0–350.0	−1.0–0.4
15	300.0–max	0.5–1.0	275.0–350.0	0.4–0.7
16	300.0–max	0.5–1.0	275.0–350.0	0.7–0.8
17	300.0–max	0.5–1.0	275.0–350.0	0.8–0.93
18	300.0–max	0.5–1.0	275.0–350.0	0.93–1.0
19	300.0–max	0.5–1.0	350.0–425.0	−1.0–0.4
20	300.0–max	0.5–1.0	350.0–425.0	0.4–0.7
21	300.0–max	0.5–1.0	350.0–425.0	0.7–0.8
22	300.0–max	0.5–1.0	350.0–425.0	0.8–0.93
23	300.0–max	0.5–1.0	350.0–425.0	0.93–1.0
24	300.0–max	0.5–1.0	425.0–650.0	−1.0–0.4
25	300.0–max	0.5–1.0	425.0–650.0	0.4–0.7
26	300.0–max	0.5–1.0	425.0–650.0	0.7–0.8
27	300.0–max	0.5–1.0	425.0–650.0	0.8–0.93
28	300.0–max	0.5–1.0	425.0–650.0	0.93–1.0
29	300.0–max	0.5–1.0	650.0–max	−1.0–0.4
30	300.0–max	0.5–1.0	650.0–max	0.4–0.7
31	300.0–max	0.5–1.0	650.0–max	0.7–0.8
32	300.0–max	0.5–1.0	650.0–max	0.8–0.93
33	300.0–max	0.5–1.0	650.0–max	0.93–1.0

Table 6.2: Reconstructed binning for the TPC pion samples. One such binning exists for each FGD layer group. For the FGD2x sample bins 29+30 are merged.

Bin Index	$p_{\mu^+}^{\text{reco}}$[MeV]	$\cos\theta_{\mu^+}^{\text{reco}}$	$p_{\pi^-}^{\text{reco}}$[MeV]	$\cos\theta_{\pi^-}^{\text{reco}}$
1	0.0 –max	-1.0–1.0	0.0– 140.0	-2.0 –1.0
2	0.0 –max	-1.0–0.5	140.0– max	-2.0– 1.0
3	0.0 –300.0	0.5 –1.0	140.0– max	-2.0– 1.0
4	300.0 –max	0.5– 1.0	140.0– 180.0	-1.0– 0.4
5	300.0 –max	0.5– 1.0	140.0– 180.0	0.4 –0.6
6	300.0 –max	0.5– 1.0	140.0– 180.0	0.6 –0.8
7	300.0 –max	0.5– 1.0	140.0– 180.0	0.8 –1.0
8	300.0 –max	0.5– 1.0	180.0– 220.0	-1.0– 0.4
9	300.0 –max	0.5– 1.0	180.0– 220.0	0.4– 0.6
10	300.0 –max	0.5– 1.0	180.0– 220.0	0.6– 0.8
11	300.0 –max	0.5– 1.0	180.0– 220.0	0.8– 1.0
12	300.0 –max	0.5– 1.0	220.0– max	-1.0– 0.4
13	300.0 –max	0.5– 1.0	220.0– max	0.4 –0.6
14	300.0 –max	0.5– 1.0	220.0– max	0.6 –1.0

Table 6.3: Reconstructed binning for the FGD pion samples. One such binning exists for each FGD layer group. For the FGD2x sample bins the following bins are merged: 4+5, 8+9.

relatively low event rate in this sample restricts how fine the binning can be, as can be seen relative to the TPC pion sample. The FGD pion sample bins are shown in table 6.3.

Control Samples

These samples are designed to contain minimal signal, and as such should have little direct impact on the signal parameters in the fit. As such, there is little reason for these binning schemes to correlate with the signal parameters, and this is reflected in the schemes chosen for these samples in this thesis. The schemes chosen are the same for both control regions; regular binning in all variables but $p_{\mu^+}^{\text{reco}}$.

6.5 Cross-Section Model Uncertainties

The cross-section model related uncertainties discussed in section 5.8 must be parameterised in such a way as to be usable in the fit. As detailed in section 4.5.4, this is done by encoding the model variations into a set of response functions. For a given model dial variation, events will be reweighted based on their true kinematics and properties. The reweighting of these events is applied to the reconstructed level data/MC comparison, modifying the model event rate in each bin affected by the

Bin Index	$p_{\mu^+}^{\text{reco}}$[MeV]	$\cos\theta_{\mu^+}^{\text{reco}}$	$p_{\pi^-}^{\text{reco}}$[MeV]	$\cos\theta_{\pi^-}^{\text{reco}}$
1	0.0 –max	-1.0– 0.89	0.0 –325.0	-2.0 –0.75
2	0.0 –max	-1.0– 0.89	0.0 –325.0	0.75 –1.0
3	0.0 –max	-1.0– 0.89	325.0– 425.0	-2.0– 0.75
4	0.0 –max	-1.0– 0.89	325.0– 425.0	0.75– 1.0
5	0.0 –max	-1.0– 0.89	425.0– 650.0	-2.0– 0.75
6	0.0 –max	-1.0– 0.89	425.0– 650.0	0.75– 1.0
7	0.0 –max	-1.0– 0.89	650.0– max	-2.0– 0.75
8	0.0 –max	-1.0– 0.89	650.0– max	0.75– 1.0
9	0.0 –max	0.89– 0.96	0.0– 325.0	-2.0 –0.75
10	0.0 –max	0.89– 0.96	0.0– 325.0	0.75 –1.0
11	0.0 –max	0.89– 0.96	325.0- 425.0	-2.0– 0.75
12	0.0 –max	0.89– 0.96	325.0- 425.0	0.75– 1.0
13	0.0 –max	0.89– 0.96	425.0- 650.0	-2.0– 0.75
14	0.0 –max	0.89– 0.96	425.0- 650.0	0.75– 1.0
15	0.0 –max	0.89– 0.96	650.0- max	-2.0– 0.75
16	0.0 –max	0.89– 0.96	650.0- max	0.75– 1.0
17	0.0 –max	0.96– 1.0	0.0 –325.0	-2.0 –0.75
18	0.0 –max	0.96– 1.0	0.0 –325.0	0.75 –1.0
19	0.0 –max	0.96– 1.0	325.0 –425.0	-2.0– 0.75
20	0.0 –max	0.96– 1.0	325.0 –425.0	0.75– 1.0
21	0.0 –max	0.96– 1.0	425.0 –650.0	-2.0– 0.75
22	0.0 –max	0.96– 1.0	425.0 –650.0	0.75– 1.0
23	0.0 –max	0.96– 1.0	650.0 –max	-2.0 –0.75
24	0.0 –max	0.96– 1.0	650.0 –max	0.75 –1.0

Table 6.4: Reconstructed binning for the first control sample. One such binning exists for each FGD layer group. For the FGD1 sample bins 1+2 are merged. For the FGD2y sample the following bins are merged: 3+4, 5+6, 21+22, 23+24.

Target Code	Events
0	Events occurring on water
1	Events occurring on hydrocarbon
2	Other targets
3	Events occurring outside the FGD fiducial volume

Table 6.5: Modified target codes output by HighLAND2 for use this analysis.

dial. While it is possible to evaluate the weights corresponding to each dial variation on an event-by-event basis, this offers relatively little benefit in resolving the finer effects of a dial while adding significant computational overhead. As the dials used in this analysis are generally quite coarse in their effects, this motivates binning events together in a manner which groups similar events, but allows reasonably accurate parameterisation of the shape effects of individual dials in the context of this analysis.

The primary variables in which model uncertainties are generally parameterised are Q^2, the energy transfer squared, and W, the invariant hadronic mass. Furthermore, other variables in the detector (such as target) play a significant role in how uncertainties are applied. Target is accounted for by splitting events in four, according to the target code shown in table 6.5. These sets of events are further split according to the sample the event was measured in (FGD1 TPC pion, etc.). In lieu of further binning in W directly, a reaction code output by HighLAND2 is used. Table 6.6 details the reaction codes as implemented in HighLAND2. From this table it can be understood that — within the context of the events most important for this analysis — each reaction code broadly corresponds to a specific region in W. Within each of these categories there is room to further bin in Q^2, and so the binning scheme is shown in table 6.7.

The resulting binning scheme splits all selected events by target (4), reaction (7), selected sample (8), and Q^2 (11) for a total of 2464 distinct bins for the events to occupy. For each cross-section dial a number of dial variations are evaluated, with the effects on individual events in a bin averaged and stored. Performing this for a number of dial variations about the nominal (e.g. from -5σ–$+5\sigma$ in 1σ steps) allows characterisation of the general dial effect for a bin, which is then fit with a spline to allow determination of a bin weight for all values of the dial. These splines are used as the ω_i as described in section 4.5.4.

Reaction Code	Events
0	CCQE and CCQE-like (2p2h, etc.) events
1	CC Resonant pion production events
2	CC DIS ($W > 2$ GeV) events
3	CC Coherent pion production events
4	$\bar{\nu}_e$ and ν_e events
5	All neutral current scattering events
6	All other events

Table 6.6: Modified reaction codes output by HighLAND2 for use this analysis.

Bin Index	Range Q^2 [GeV]
0	$0.0, 0.1$
1	$0.1, 0.2$
2	$0.2, 0.3$
3	$0.3, 0.4$
4	$0.4, 0.6$
5	$0.6, 0.8$
6	$0.8, 1.0$
7	$1.0, 1.3$
8	$1.3, 1.7$
9	$1.7, 2.5$
10	$2.5, \mathrm{max}$

Table 6.7: Kinematic spline binning in Q^2.

Chapter 7

Fit and Analysis Validation

This chapter explores the cross section fit and extraction detailed in chapter 4. A number of tests are performed to verify that the fit behaves as expected, and that the extracted cross section values are reasonable given the input. Each test is performed by providing the fitter with ‘fake data’ — colloquially fake data is a set of data that was not measured experimentally — generated for the purpose of testing a specific aspect of the analysis. Aspects of the fit to be validated include the dependence of the result on the model priors and uncertainties, fit stability and convergence, and sensitivity to signal and background.

7.1 Summary of the Fit

The input model has four main categories of parameter to be fit:

- Template parameters
    - One for each kinematic bin for each target (Water, hydrocarbon) for a total of 48. No prior value or uncertainty to prevent bias to the input model.
- Flux parameters
    - One for each bin in neutrino energy for the flux model, for both $\bar{\nu}_\mu$ and ν_μ fluxes. Prior value, uncertainty, and correlations given by the ND280 beam group.
- Cross section model parameters

  - Discussed in section 5.8, each model parameter has a free parameter assigned in the fit.

- Detector nuisance parameters
  - One for each reconstructed kinematic bin for each sample, as described in section 6.4.2. Each parameter is a simple weight for all events in the corresponding bin with a prior of 1, and uncertainty and correlations between these parameters set by the detector simulation and HighLAND2, in section 5.7.

The fitter determines the goodness of fit through the likelihood functions detailed in section 4.5. As discussed in section 4.5.5, the best fit is found by minimising the likelihood as a function of the model parameters using Minuit, with the MIGRAD algorithm used to evaluate the uncertainty on each parameter.

7.2 Asimov Fits

An Asimov fit is one in which the data used is the same as the input simulation, and is commonly used as a basic test of fit machinery. As the fitter begins to explore the parameter space, the movement of parameters away from the nominal will result in a decrease in the log-likelihood, allowing the fit to converge reasonably quickly. The expected result for a correctly functioning Asimov fit is for all free parameters to converge to their nominal values.

7.2.1 Asimov Fit with Parameters Starting at Nominal

The first fit performed was an Asimov fit with the starting value for each parameter at the nominal value.As seen in figure 7.1, the parameters have all been fit to their nominal values and their errors reduced somewhat. These results show the expected behaviour for the fit. The post-fit error on the template parameters provides some measure of the statistical uncertainty in the data. In general, parameters that have their prior error reduced significantly have a significant effect on the agreement between the data and model.

7.2.2 Asimov Fit with Random Starting Parameters

The previous fit has a potential issue for the fitter; as all parameters start at their best fit value, the fitter does not have to explore the parameter space significantly in

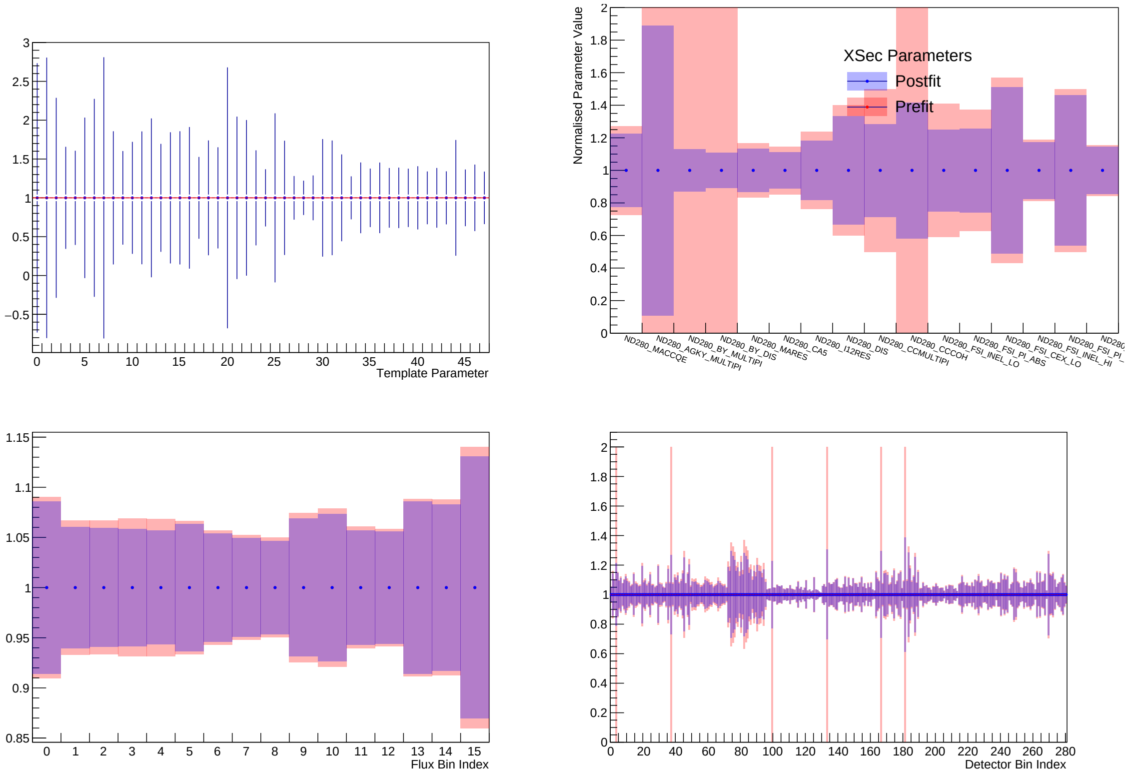

Figure 7.1: Fit results for the Asimov fit with all parameters starting at nominal.

order to converge, which can result in insufficient error coverage. There is also the potential for the fit to have convergence issues if the starting values of the parameters are sufficiently far from the best fit values. A sample of 100 fits were run with the starting values of each fit parameter randomly thrown, and all converged to the expected best fit values. Figure 7.2 shows the fit results from one of these fits, with all parameters at nominal and errors reduced as expected from the first Asimov fit.

7.3 Statistical Fluctuations and Error Coverage

The next obvious test of the fit machinery is to see how statistical fluctuations on the input data affect the fit results. If the model being fitted has insufficient freedom, certain data permutations can cause issues with convergence. In general Poissonian statistical effects are expected in the real data, and so the statistical fluctuations for this analysis are Poissonian. The best fit parameters for such fits are unknown a priori, and should be distributed around the nominal values.

The reconstructed event distributions for an example fit with Poissonian fluctuations are shown in figure 7.3, showing the nominal MC, fluctuated data, and the

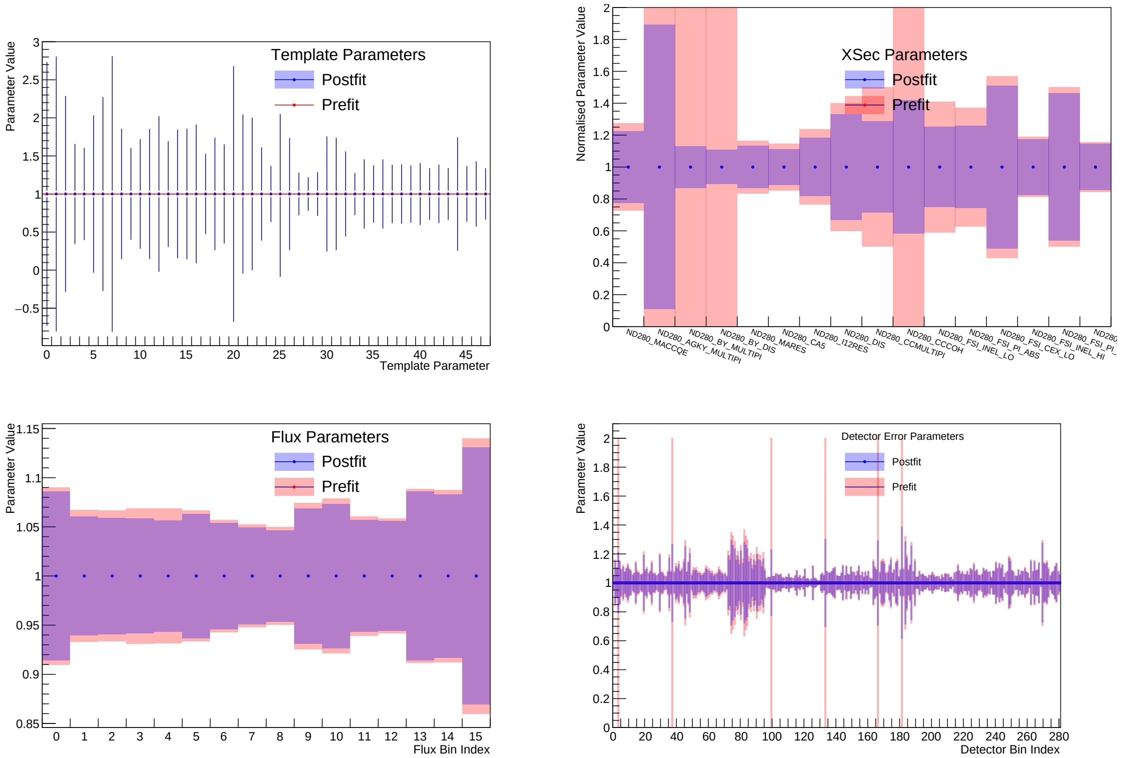

Figure 7.2: Fit results for the Asimov fit with all parameters starting at random values.

best fit value. The fit parameter results for the same fit are shown in figure 7.4. A general downward trend in the FGD2x samples can be seen to propagate to the water template parameters, resulting in a lower cross-section.

The resulting cross-sections computed from this fit are shown in figures 7.5 and 7.6. In figure 7.5 the cross-section for each true (2D) bin is shown in the top left panel, together with the correlation matrix in the top right. The bottom two cross-section plots show the integrated cross-sections in both pion momentum and angle respectively, together with the correlation matrix between each of these bins. Figure 7.6 shows the integrated cross-sections for each target as a function of the pion variables, showing both the nominal MC cross-section together with that calculated from the fake data. The correlations between these bins are the same as those in the bottom right panel of figure 7.5.

7.3.1 Error Validation

Propagation of errors is an important part of this analysis, and validating that these errors are propagated correctly and of adequate size is necessary. At various points in

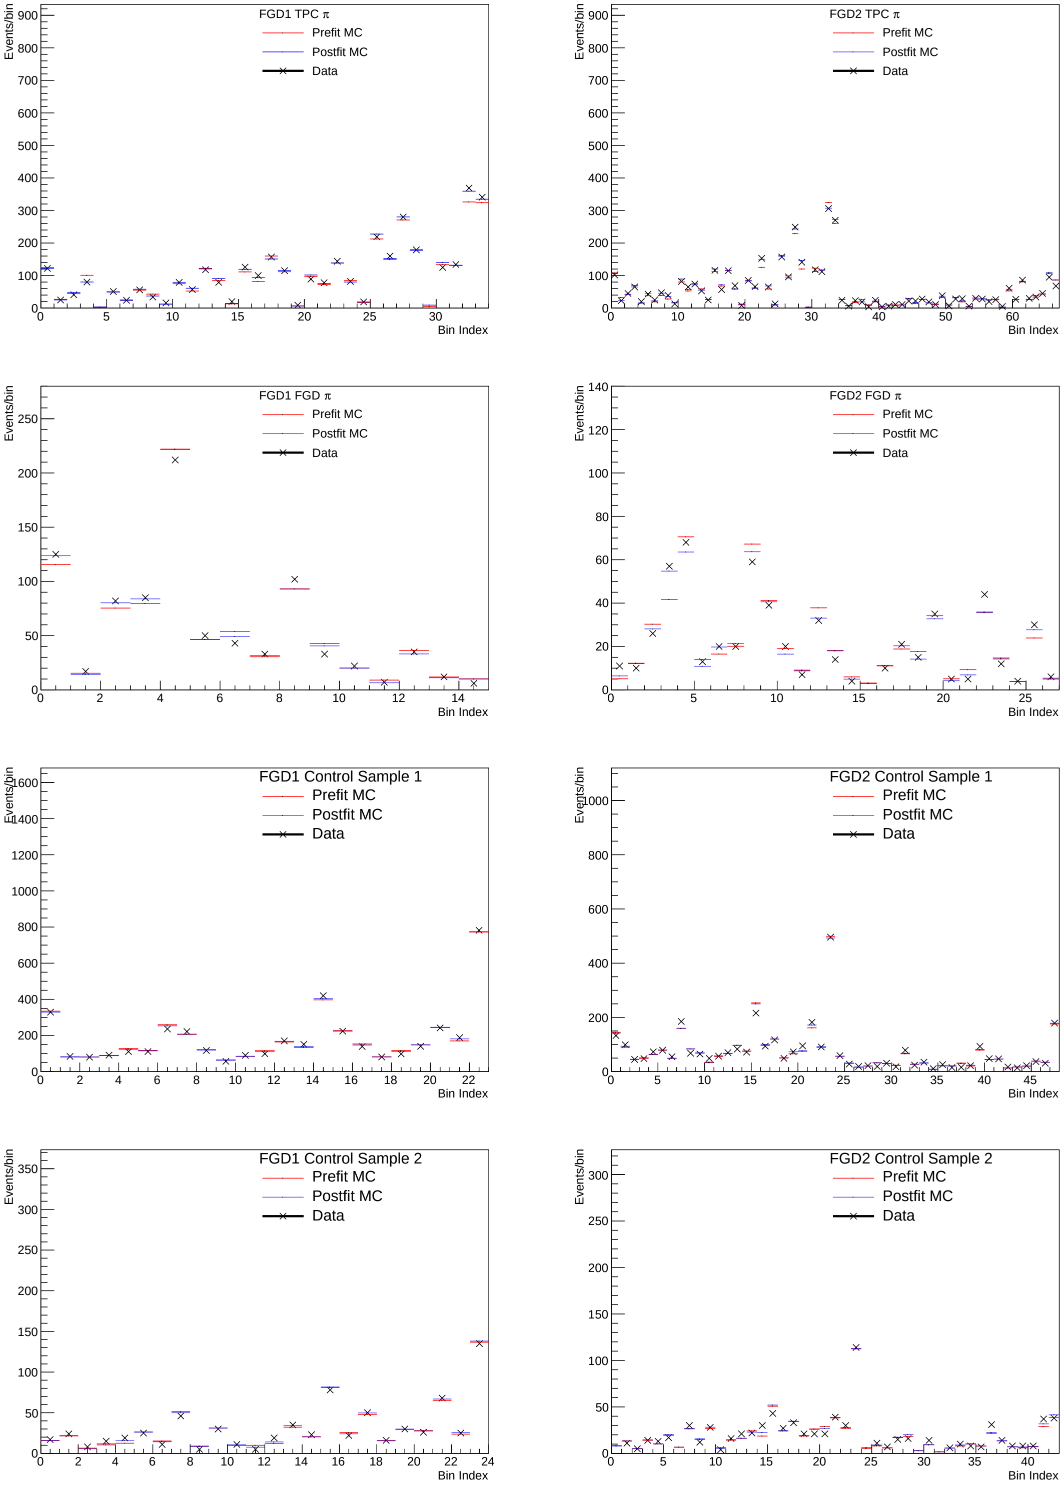

Figure 7.3: Reconstructed level comparison for the input model with statistical fluctuations applied for all 8 samples; from top to bottom, the samples are TPC π^-, FGD π^-, and control regions 1 and 2 for FGD1 (*left*) and FGD2 (*right*).

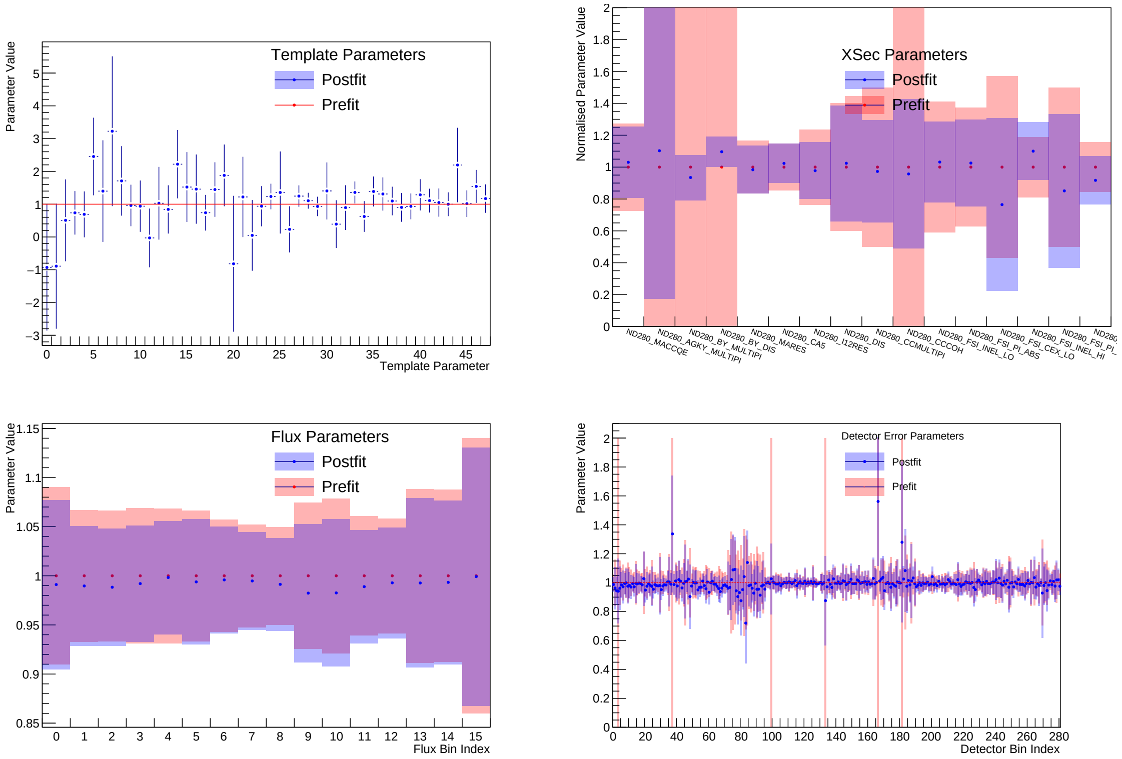

Figure 7.4: Fit results for the model fit to the data shown in figure 7.3.

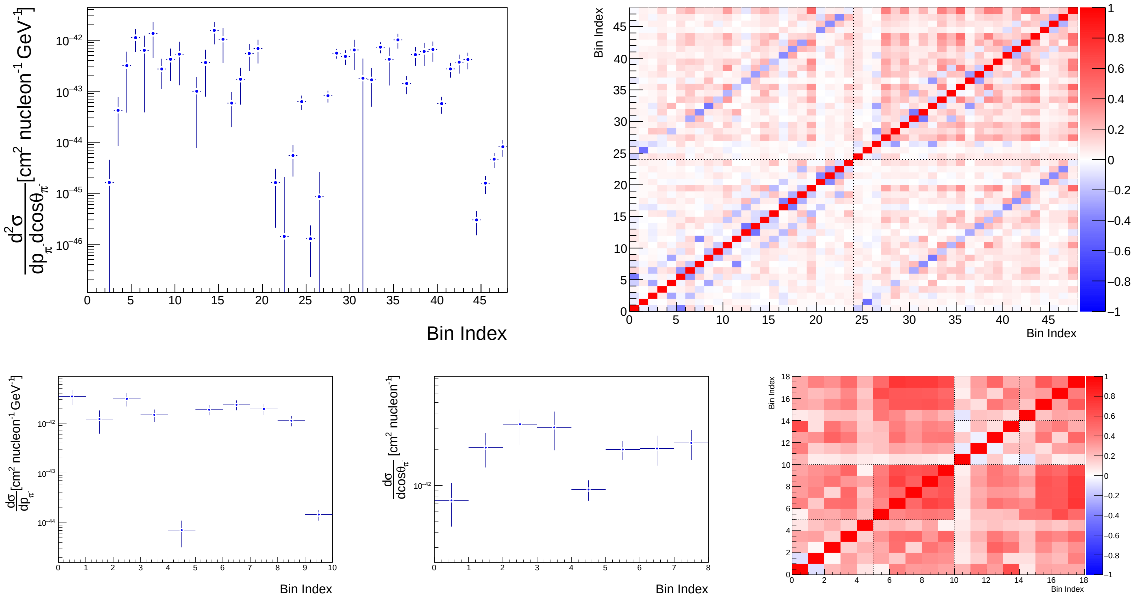

Figure 7.5: Cross-section results using the fit result shown in figure 7.4.

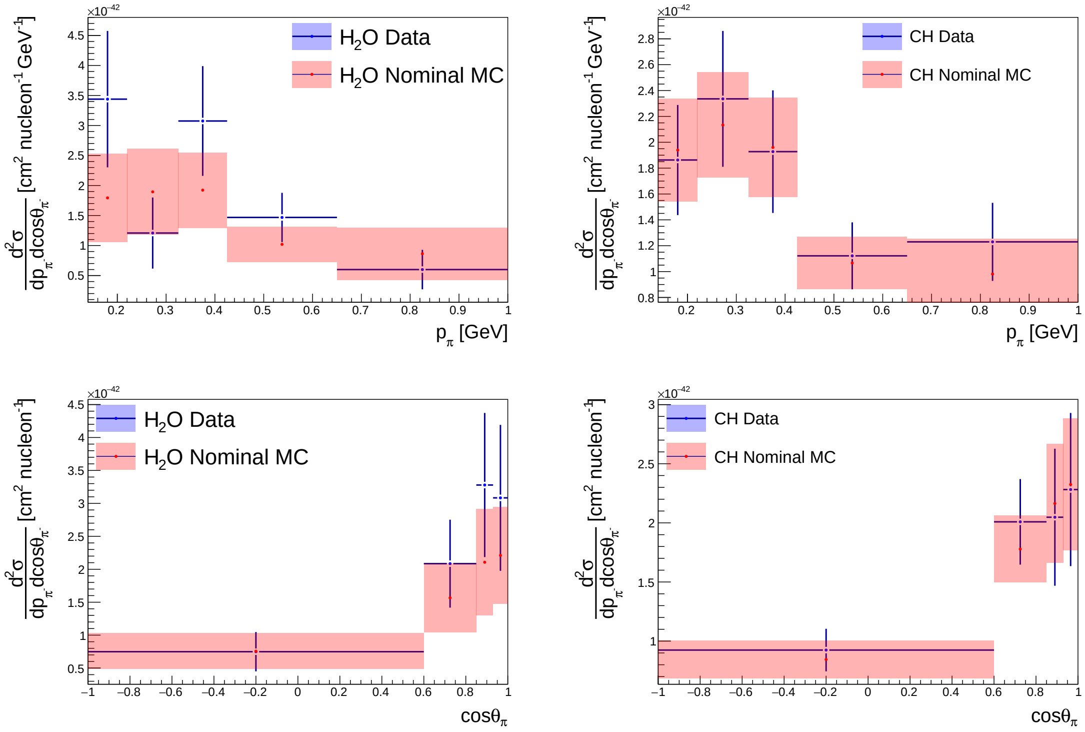

Figure 7.6: Cross section computed with statistical fluctuations applied to the input simulation, with the Asimov-fit result in red.

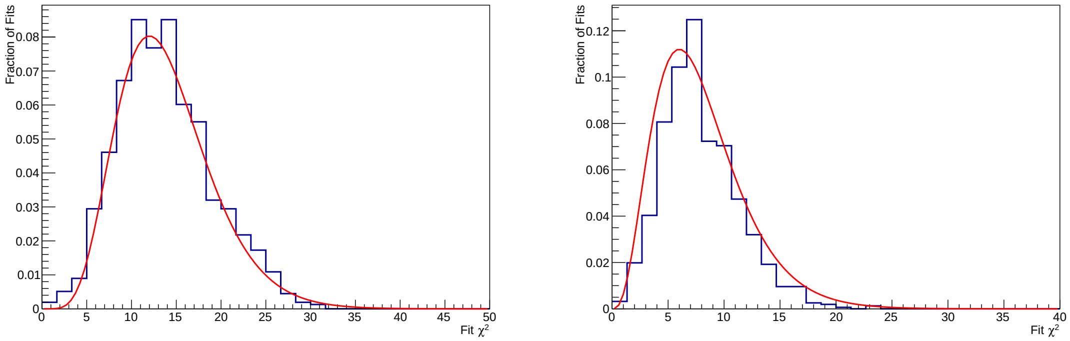

Figure 7.7: Cross-section χ^2 distributions for 1000 statistically fluctuated fits. In the left plot χ^2 is computed in the fine 2D binning, while in the right plot it is computed in the integrated single differential binning.

this analysis, errors are assumed to be Gaussian, but as touched on in the previous chapter, there are bins in which the number of events is sufficiently low that the uncertainties will be decidedly non-Gaussian. To determine whether these bins can cause issues, a sample of 1000 independent fits are performed with the data fluctuated as in the previous example fit. In these fits the detector nuisance parameters — which are Gaussian by construction — were fixed to speed up the fits.

From each fit, a value for the χ^2 for the cross-section can be calculated using

$$\chi^2 = \left(\vec{\sigma}_{\text{toy}} - \hat{\vec{\sigma}}\right) \left(\mathbf{V}^{\text{ cov.}}\right)^{-1} \left(\vec{\sigma}_{\text{toy}} - \hat{\vec{\sigma}}\right)^T . \tag{7.1}$$

The shape of the χ^2 distribution is used here to understand whether these low statistics bins have a significant effect, as if the error coverage is skewed the distribution in χ^2 should also be skewed. Where no significant skew, shift, or secondary peaks are observed, and the distribution is in agreement with a standard χ^2 distribution, this is taken as an indication that the low event rates of some bins does not have a significant effect on the overall likelihood function, and thus on the fit.

Figure 7.7 shows the χ^2 distributions for both the finely binned 2D and 1D integrated cross sections, including correlations. The general shape of these distributions matches the expected χ^2 distribution, indicating that the Gaussian errors assumption should not have a significant impact on the error treatment for this analysis.

7.3.2 Error Coverage

To ensure that the post-fit errors correctly cover the true uncertainties of the input model, a number of statistically fluctuated data fits were performed as in the previous

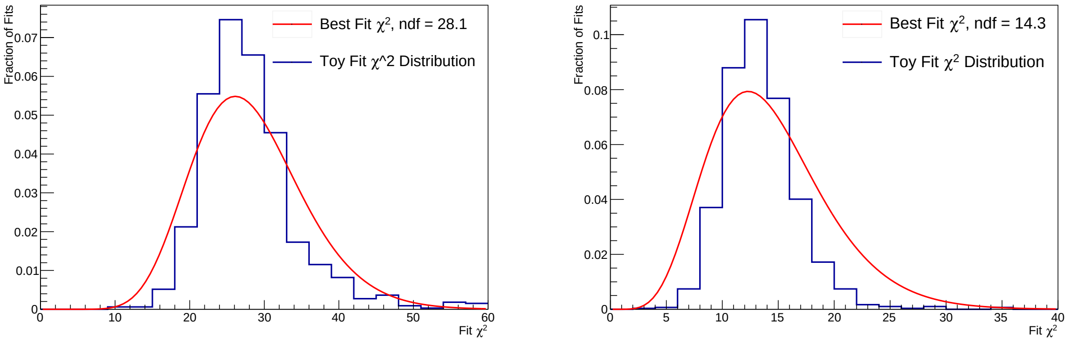

Figure 7.8: Cross-section χ^2 distributions for 1200 statistically fluctuated fits with all model parameters free. In the left plot χ^2 is computed in the fine 2D binning, while in the right plot it is computed in the integrated single differential binning. Best fit values for the number of degrees of freedom are computed for both of these distributions.

section, but with the detector nuisance parameters included. Fitting the resulting χ^2 distribution allows computing the effective number of degrees of freedom, which will be useful later in understanding how good a given fit χ^2 is to the model.

Figure 7.8 shows χ^2 distributions calculated both in the double differential fine binning and the integrated single differential binning. Best fit values for the number of degrees of freedom in these respective binnings were determined to be ndf $= 28.1$ and ndf $= 14.3$ respectively.

7.4 Signal and Model Variations

It is important to demonstrate that the analysis can accurately respond to changes in the input signal, as accurate measurement of the signal is the goal of the analysis. In a similar manner, it is also useful to show that the fit can resolve individual systematic variations without adversely affecting the overall fit. This motivates a number of fake data fits which will be performed here.

7.4.1 Signal Enhanced Fake Data

In this fit a data sample is created from the input MC where all $\bar{\nu}_\mu$CC1π^- events are weighted by 1.2 if they occur on H_2O and 1.1 if they occur on hydrocarbon. The expected result of this fit is for the template parameters to be fit to 1.2 and 1.1, and for this to propagate directly into 20% and 10% increases in the final cross-section.

As can be seen in figures 7.9 and 7.10, the fit has returned the expected result,

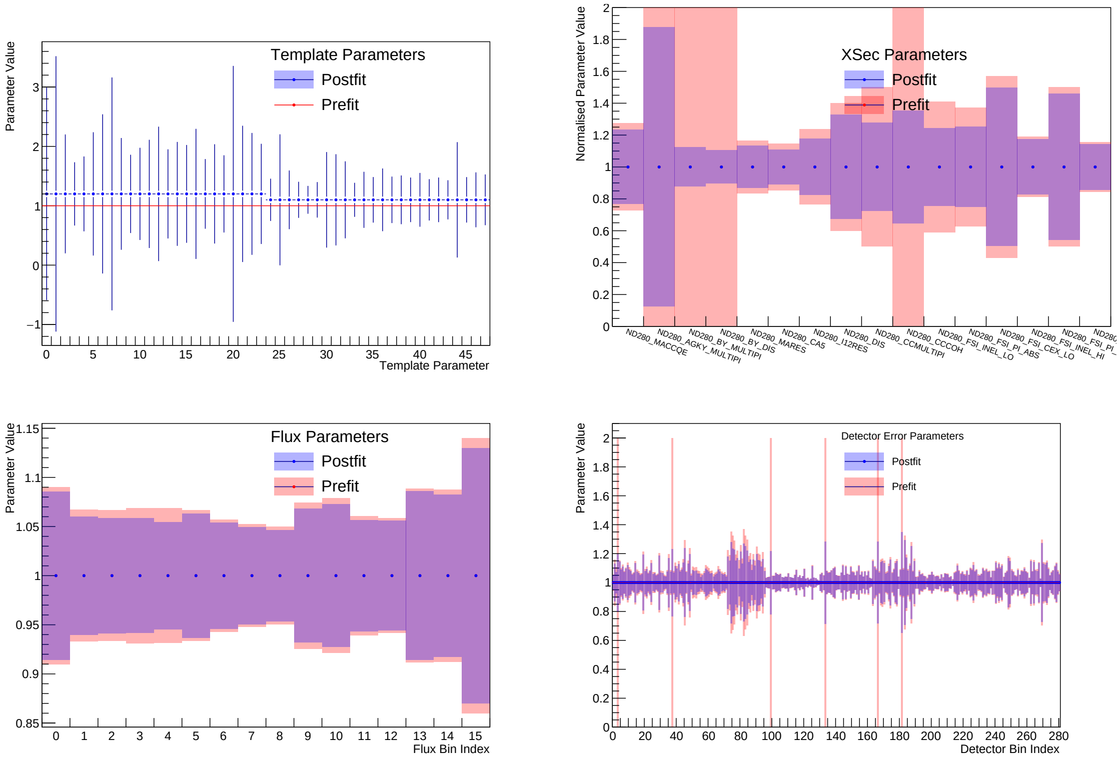

Figure 7.9: Fit results for the enhanced H_2O and CH signal fake-data sample, showing perfect reconstruction of this effect in the template parameters.

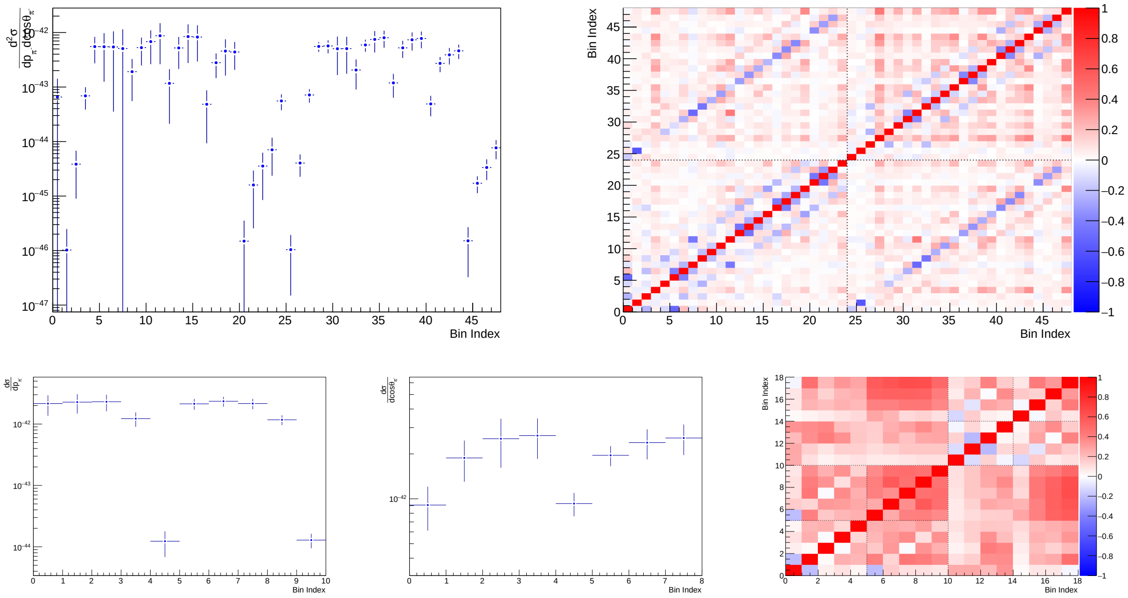

Figure 7.10: Cross-section results using the fit result shown in figure 7.9.

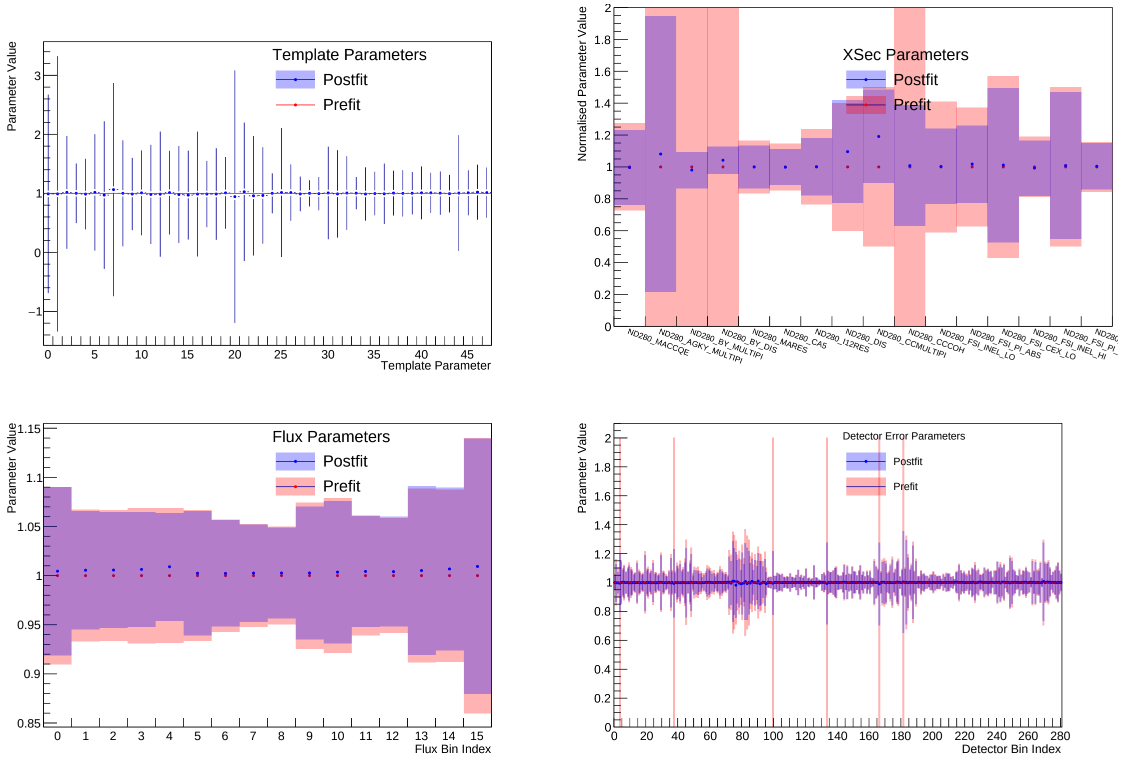

Figure 7.11: Fit results from the 20% enhanced DIS fake-data fit, with all DIS-related model parameters moving, and minimal noise seen in the other parameters.

finding the cross-section values to be enhanced by 20% and 10% for H_2O and CH respectively. This demonstrates that the fit is sensitive to the signal, and that a variation to the signal only does not impact the fitting of the background parameters.

7.4.2 DIS Enhanced Fake Data

In this context, DIS refers to all $\bar{\nu}_\mu/\nu_\mu$ CC events with an invariant hadronic mass of 2 GeV or more. All such events are weighted by 1.2 to increase this specific background, with the expected fit behaviour being the absorption of this increase into the cross-section model uncertainties with a minor χ^2 penalty. As none of the reweight dials correspond directly to a normalisation on all of these events, the overall effects on the fit are similar to a more generalised shape uncertainty.

Figures 7.11 and 7.12 show the fit parameters and the resulting cross section for this fit. The DIS and Multi-Pion cross-section dials can be seen to all have moved to varying degrees, with the FSI dials showing minor movement. Minor movements about the nominal values can also be seen in the detector and signal parameters, suggesting that the 20% normalisation increase cannot be perfectly fit by the cross

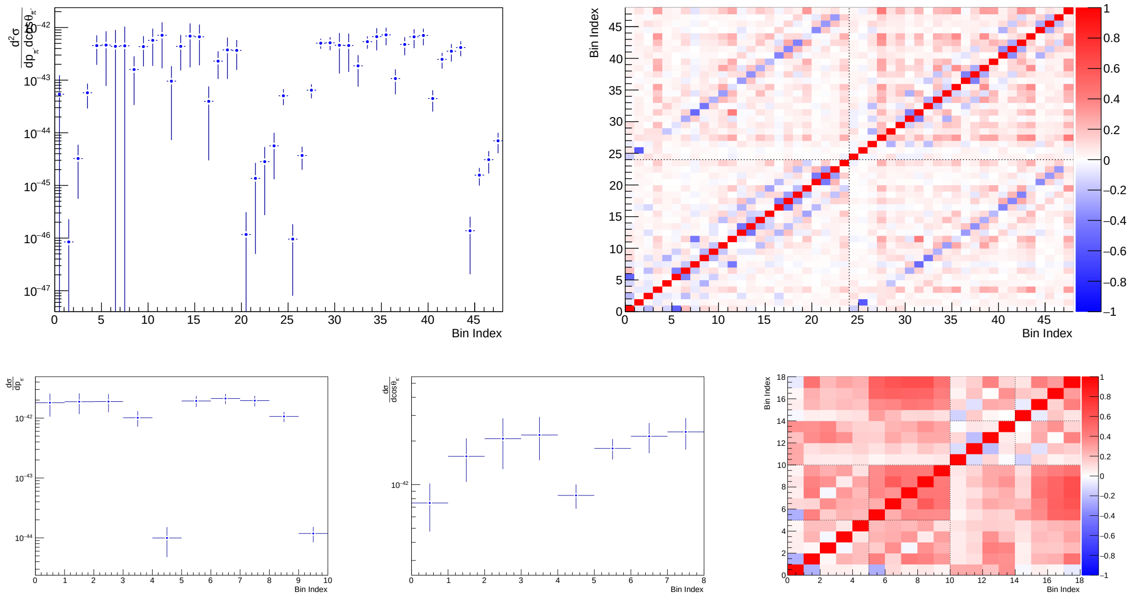

Figure 7.12: Cross-section results using the fit result shown in figure 7.9.

section dials alone, with very minor discrepancies remaining.

7.4.3 Previous Production Monte-Carlo as Fake Data

The Monte-Carlo simulation used in chapter 5 to understand the event selection is the current Monte-Carlo simulation — known as production 6T — used by T2K in the oscillation analysis. The previous mainline simulation used was production 6B, based on the same detector-level reconstruction and differing primarily in the physics models used to predict the event distributions. Of interest in this thesis was the use of the Rein-Sehgal model for coherent pion production. The Rein-Seghal prediction used in NEUT is known to heavily over-predict the coherent pion production cross section, and for $\bar{\nu}_\mu$ CC events where resonant production is suppressed relative to the ν_μ rate, the excess coherent events become a significant proportion of the $\bar{\nu}_\mu$CC1π^- signal. The expected fit result is not known exactly, as there are a number of small differences in the inputs used, but the general result should show the excess coherent pion production events being absorbed into the CC COH dial or the template parameters, as these are also signal events. Some movement of DIS and Multi-Pi dials is also to be expected, as the exact treatment of these events varies slightly between the NEUT versions used. Despite the lack of a clear expected fit result, this Monte-Carlo sample provides an opportunity to fit a well validated sample that is independent on an event-by-event basis with known differences.

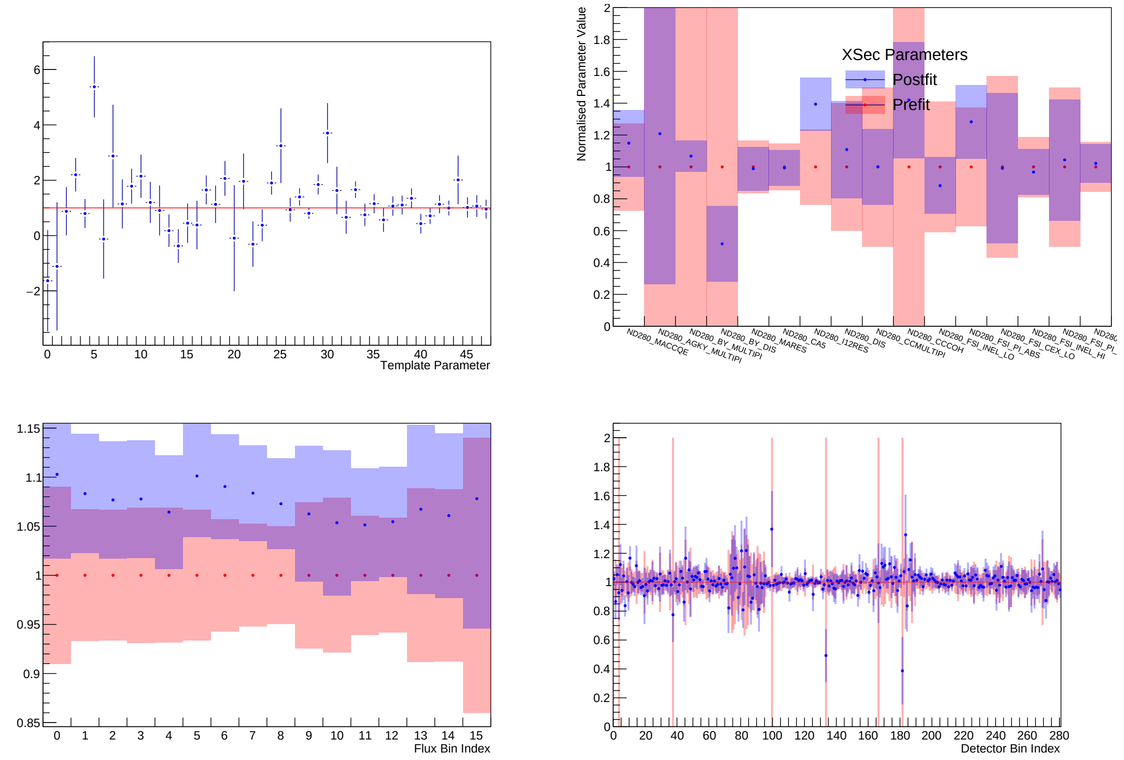

Figure 7.13: Fit result from the fit to the Production 6B Monte-Carlo, showing an increase in coherent and some effects on other related dials.

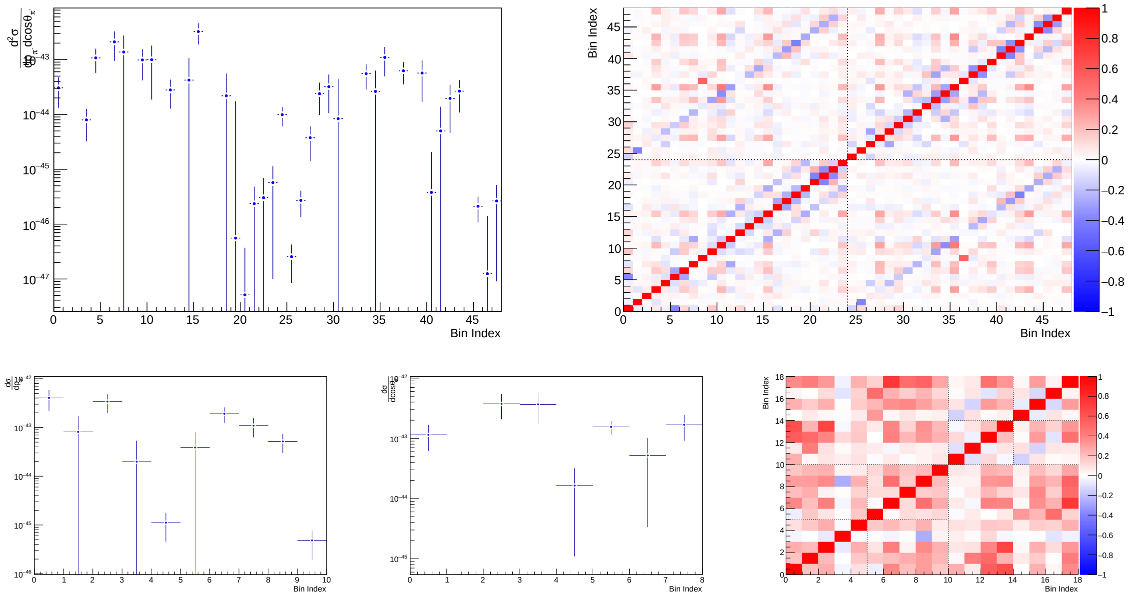

Figure 7.14: Cross-section results using the fit result shown in figure 7.13.

Large pulls in a number of cross-section dials can be seen in the post-fit parameters shown in figure 7.13, particularly in the coherent, BY DIS, and I$\frac{1}{2}$ dials, but to some extend in a number of other dials. A large systematic offset is seen in the flux parameters in this fit, which is the result of a minor bug in POT calculations for the production 6B files used. This discrepancy in POT results in a constant event weighting discrepancy for all events, which is naturally fit by scaling all flux parameters. While the correct post-fit results for the cross-section dials are not well known a-priori, these results appear to be generally as expected.

7.4.4 Real Data Control Region Un-blinding

For this fit, the real ND280 data is used, but only for the control regions. The signal region 'data' is filled with the nominal input prediction scaled to the real data POT. As the signal regions are not unblinded, the naïve expectation of this fit is that the signal parameters should not change significantly. As the data POT is significantly lower than the POT of the MC sample, the signal parameters can be expected to vary in a manner consistent with statistical fluctuations about the mean. This fit result also provides context for the full data un-blinding, showing how the control regions move the fit parameters prior to any effects from the signal region.

Figure 7.15 shows the post-fit parameter values for the control region fit. Modest movement is seen in the DIS dials, resonant and coherent pion production dials, and the FSI dials. The ν_μ background flux parameters are also seen to be pulled $3-4\%$ below nominal, with a mild shape effect seen in the $\bar{\nu}_\mu$ flux. Significant fluctuations about the nominal can be seen in the template parameters. These effects are expected of this fit; as the second control region is intended to constrain DIS and multi-pion models, it is expected that the dials most relevant to these should be constrained. As the first control region was constructed to constrain the ν_μ background in the selection, it is also expected that the ν_μ flux parameters should see more movement than their $\bar{\nu}_\mu$ counterparts, as the $\bar{\nu}_\mu$ parameters are held close to nominal by the signal region events which are not varied. Fluctuations in the template parameters are expected in order to balance the changes induced in the signal region modelling by the control regions.

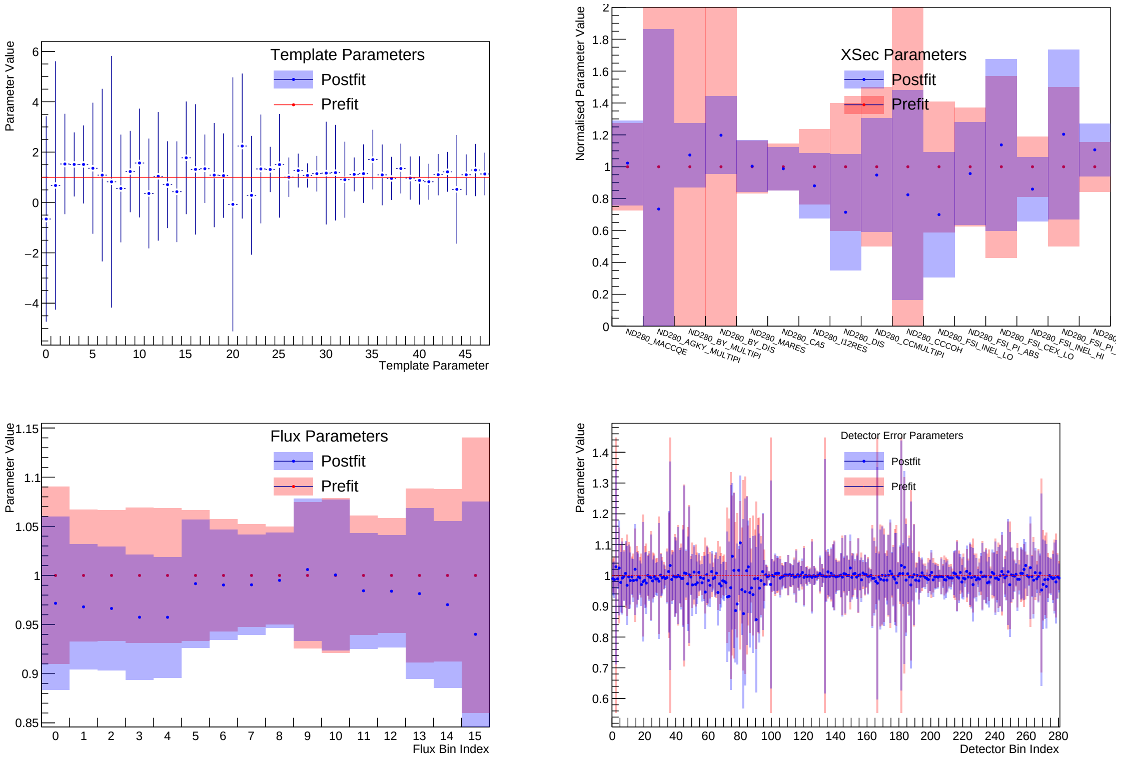

Figure 7.15: Fit result from the fit to the unblinded control regions with nominal signal samples, showing an increase in coherent and some effects on other related dials.

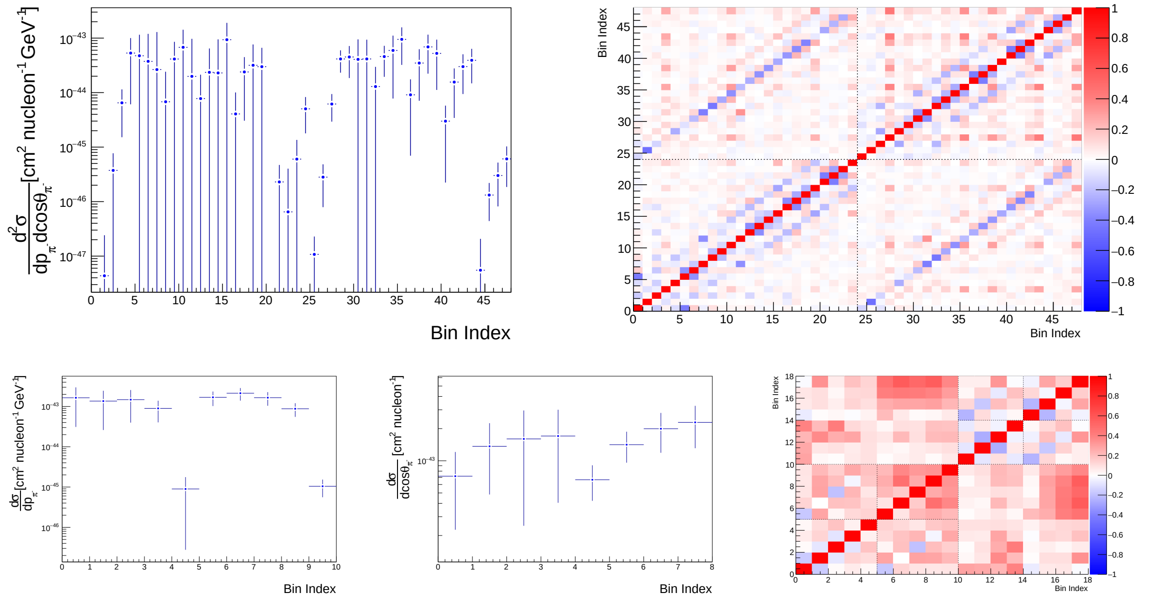

Figure 7.16: Cross-section results using the fit result shown in figure 7.15 with both control regions unblinded.

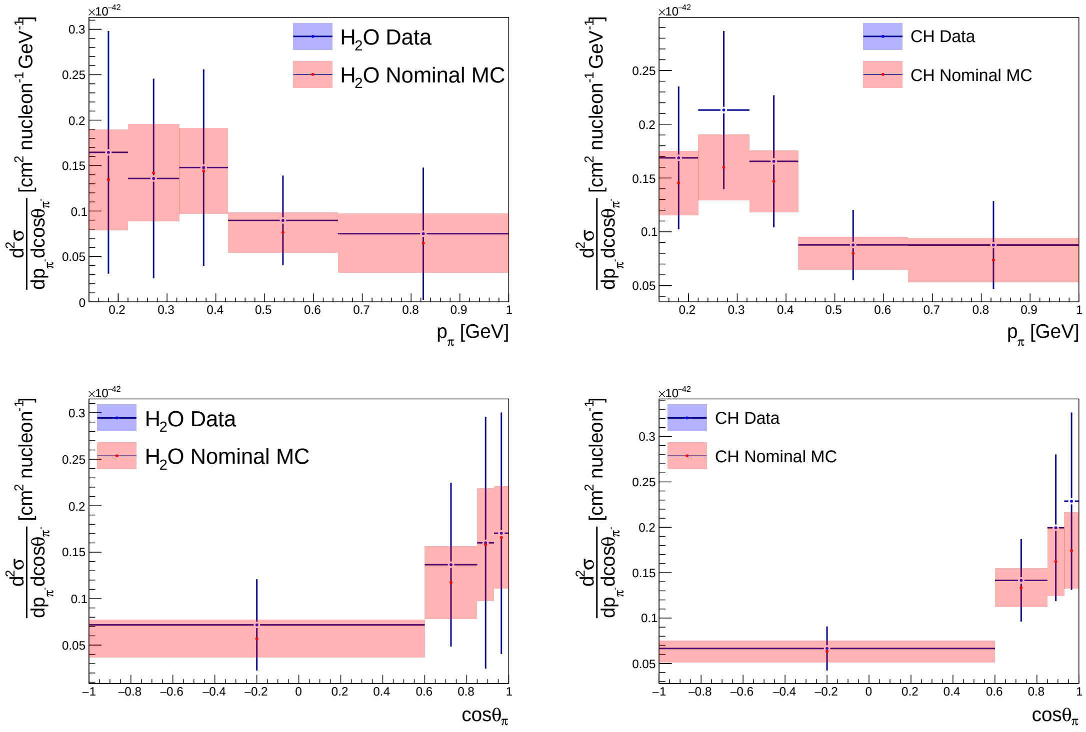

Figure 7.17: Cross section results with both control regions unblinded in blue, with the Asimov-fit result in red.

Chapter 8

Results

This chapter presents the final cross-section result, as determined by fitting to the available ND280 data.

8.1 Reconstructed Event Rates

Figure 8.1 shows the measured event rate in each reconstructed bin, together with the pre-fit MC prediction and the post-fit MC prediction from the fit in the following section. From these plots both the nominal data/MC agreement can be seen, together with the improvement in the agreement from the likelihood fit.

8.2 Data Fit Result

Figure 8.2 shows the result of the fit on the model parameters and errors, together with the correlation matrix between these parameters. Notable in this fit is the wide range of values and uncertainties associated with the template parameters. In real data this is expected, as some of the parameters are constrained by a relatively small fraction of the observed events (e.g. high angle regions). Despite appearing unwieldy, when combined in the manner detailed in chapter 4, these should combine to form a much more regular cross-section with reasonable looking errors. Moreover, the cross-section model used for signal events in this analysis is understood to be lacking in freedom, as contemporary neutrino-induced pion production models are under-developed relative to many CCQE models, and any behaviour which cannot be fit by the underlying model is likely to be absorbed primarily by the template parameters.

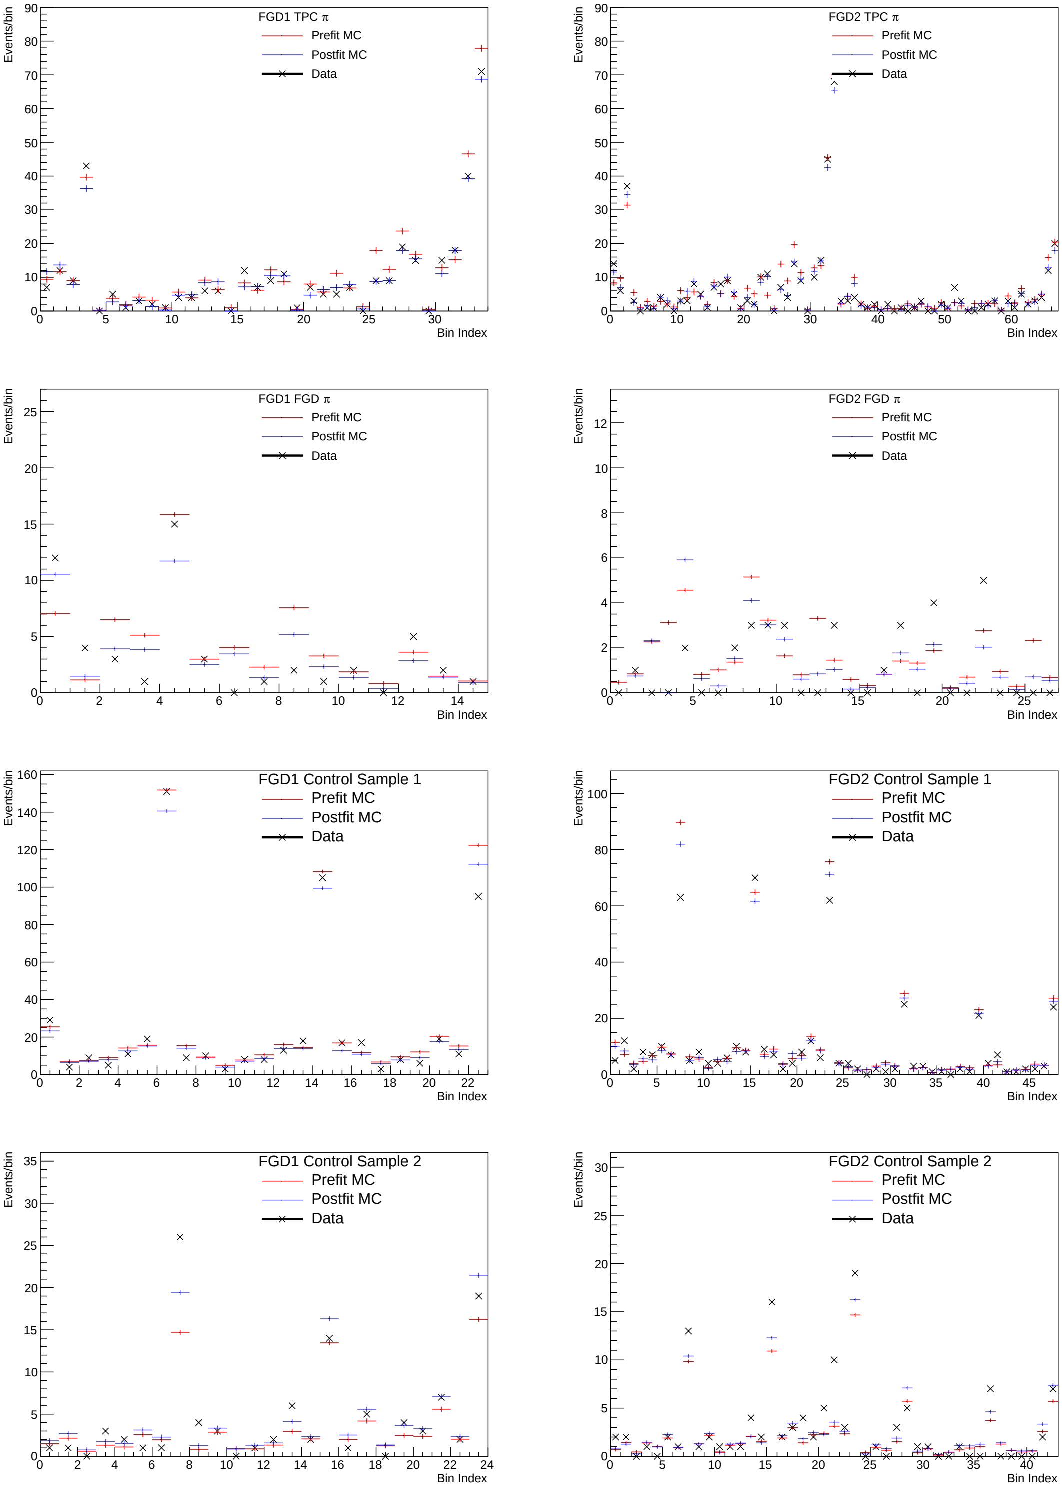

Figure 8.1: Reconstructed level samples for FGD1 (*left*) and FGD2 (*right*), samples are TPC π^-, FGD π^-, control region 1 and control region 2 from top to bottom.

χ^2 source	χ^2	ndf	χ^2/ndf
FGD1 TPC π^-	18.26	34	0.537
FGD1 FGD π^-	20.34	15	1.36
FGD1 Control Region 1	20.53	23	0.893
FGD1 Control Region 2	20.25	24	0.844
FGD2 TPC π^-	35.61	67	0.531
FGD2 FGD π^-	29.27	27	1.08
FGD2 Control Region 1	37.66	48	0.784
FGD2 Control Region 2	38.76	43	0.901
Detector Parameters	7.24	281	0.0257
XSec Model Parameters	1.99	16	0.124
Flux Parameters	0.457	16	0.0285
Statistical	220.7	–	–
Systematic	9.69	–	–
Total	230.4	–	–

Table 8.1: Post-fit χ^2 contributions in the fit broken down by sample and parameter type.

Other effects are also seen in the post-fit parameters. The flux parameters show slight movement, with the $\bar{\nu}_\mu$ component showing a very similar post-fit shape to that seen in the fake-data fit with control regions unblinded in section 7.4.4. The ν_μ component also shows similar behaviour with a few-% decrease and minor shape effects seen. Most cross-section model dials have not moved too far from their priors, although dials relating to pion production and DIS see some movement and significant error constraints. The detector parameters show a level of noise that was not seen in the fits of the previous chapter, but this is to be expected as only two of the fits from the previous chapter (statistical fluctuations and control region unblind) vary the input in a manner that adds noise relative to the underlying MC prediction. A secondary effect from this noise is the relative lack of constraint seen in the post-fit detector errorbands, which are not significantly reduced from the prior errors in general.

The final χ^2 contributions to the fit from the data samples and parameter types are listed in table 8.1, where the majority of the contribution can be seen to be statistical in nature.

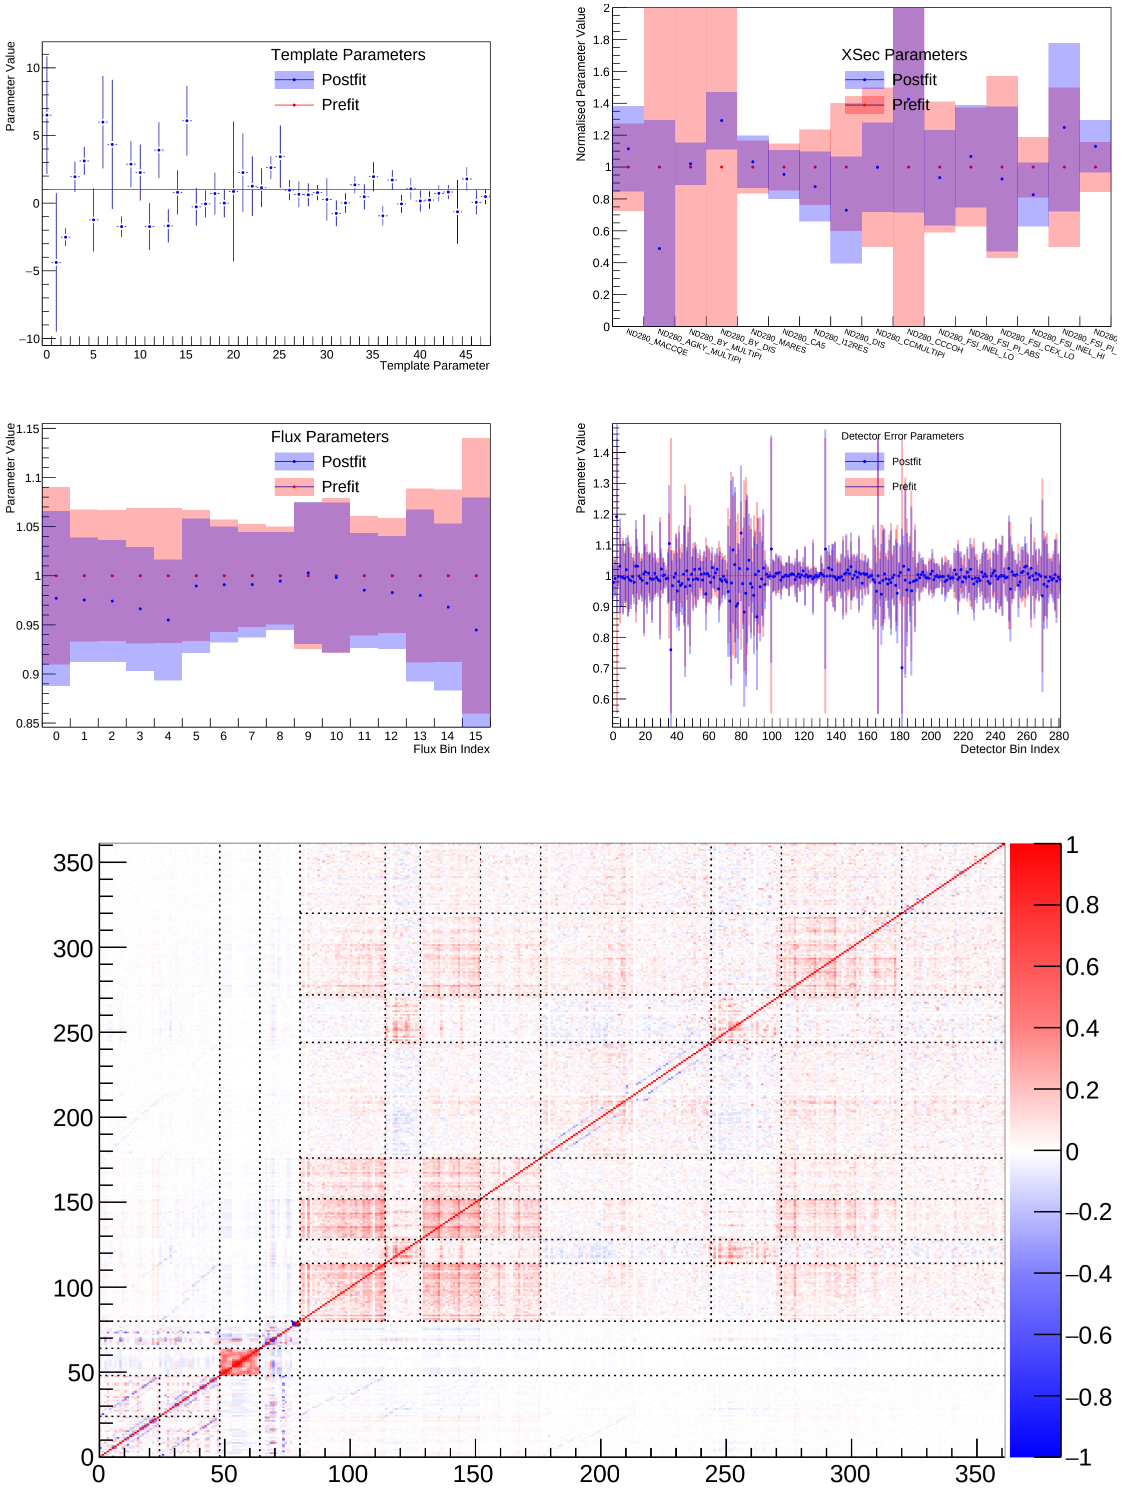

Figure 8.2: Fit results for the model fit to the data shown in figure 7.3.

8.3 Data Cross-Section Result

The final differential cross-section results in p_{π^-} and $\cos\theta_{\pi^-}$ are shown in figure 8.3 together with the correlation matrix for all bins. Upon first inspection the agreement — particularly the shape — between the data and MC appears poor. Inspection of the cross-section correlation matrix shows a number of anti-correlated bins, with these generally corresponding to the bins appearing to challenge the model. A value for χ^2 can be computed between the nominal and final cross-sections as

$$\chi^2 = (\sigma_{\text{nom.}} - \sigma_{\text{data}})\, \mathcal{V}^{-1}_{\text{nom.}}\, (\sigma_{\text{nom.}} - \sigma_{\text{data}})^T\,. \tag{8.1}$$

For the cross-section presented here, a value of $\chi^2 = 87.31$ is obtained. The number of degrees of freedom was determined to be 14.3 in section 7.3.2, giving a $\chi^2/\text{ndf} = 6.1$. This reduced χ^2 values indicates a reasonable level of tension between the data and the nominal prediction. A level of tension between the data and simulation is expected, as a fundamental limitation of the input model is the inability to vary pion kinematics on an event-by-event basis, a limitation of the theoretical models and their implementations in current neutrino interaction generators.

8.4 Future Work and Potential Improvements

The method of template parameters in a log-likelihood fit as a cross-section extraction technique is generally rather robust and powerful, but various limitations exist in the analysis in this thesis. A major consideration in this analysis is the evaluation of uncertainties in the fit parameters from the likelihood surface, which is done by assuming that the surface is a multivariate Gaussian about the best fit point. For errors that are non-Gaussian, some bias related to the difference between the actual error surface and a Gaussian can be expected, although accounting for such a bias cannot be done in a straightforward manner. While attempts have been made to mitigate any such biases, many of the parameters used in the fit have uncertainties which are known to be non-Gaussian to varying degrees, in particular some cross-section model dials which are not defined beyond some boundary, and the statistical errors which are Poissonian in nature. A potential improvement is to use a Markov chain Monte-Carlo method to sample the likelihood distribution and better parameterise the likelihood surface.

A goal of this analysis was to perform a finely-binned cross-section measurement

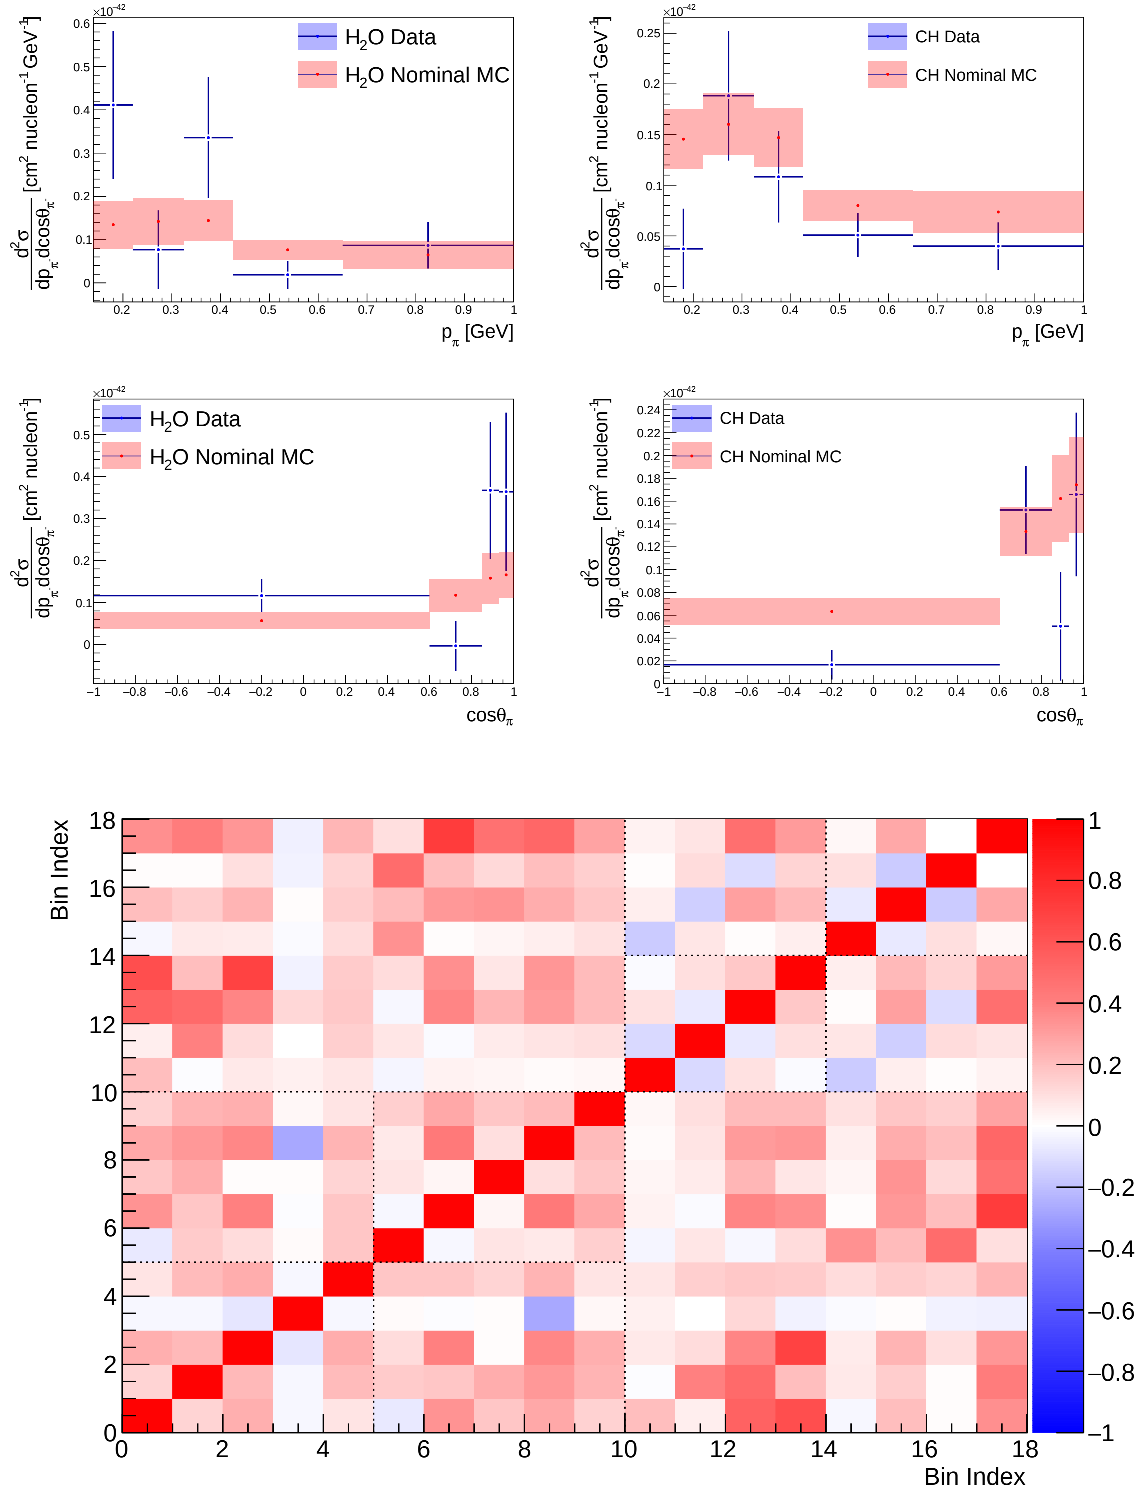

Figure 8.3: Cross section computed with statistical fluctuations applied to the input simulation, with the Asimov-fit result in red.

to be collapsed into coarser bins. Due to the highly non-uniform event rates in various parts of the kinematic phase space, some bins will inevitably have a very low proportion of events (both measured and simulated), which results in large statistical uncertainties. Low simulated event numbers in a bin limits the ability of the fitter to accurately characterise the reco.-truth mapping, and will usually correspond to a bin with a very low data rate, which will also result in a non-Gaussian error profile for that bin. For a given analysis of this form, these problems can be significantly mitigated by the use of increased Monte-Carlo statistics, and future analyses can benefit from both future data and Monte-Carlo runs.

Purely as a $\bar{\nu}_\mu$CC$1\pi^-$ measurement, various future improvements in both the selection and reconstruction, as well as the detector simulation are foreseen. The phase space of the measured cross-section contains a ‘hole’ in the $\cos\theta_{\pi^-}$ distribution in the region $(-0.35, +0.35)$ as understanding of the particle Pull distributions for high angle events is lacking. Work is presently underway to constrain events with pions exiting in this region by using cosmic muon control samples, and future pion production measurements, together with proton kinematic measurements are expected to benefit from an expanded angular phase space.

Finally, this cross-section measurement could be performed as part of a joint analysis in future, where multiple cross-sections (e.g. CC0π and CC1π) can be measured simultaneously. Such an analysis has the benefit of correlating various parameters — most importantly the flux — between these measurements, providing stronger constraints on the measurement.

8.5 Conclusions

This thesis has presented a joint measurement of the $\bar{\nu}_\mu$CC$1\pi^-$ cross-sections on water and hydrocarbon performed with the ND280 detector. This is the first such measurement at T2K, and makes use of all available ND280 $\bar{\nu}_\mu$-beam data, together with up-to-date calibrations, detector simulation, and neutrino interaction models together with a simultaneous log-likelihood fit to determine the cross section. A variety of mock datasets are used to validate various aspects of the analysis, as well as to understand the level of agreement between the result and the input model.

A method of collapsing a double differential cross-section into two correlated single differential cross-sections has been developed for this analysis and is presented here, although this method was retroactively applied to the $\nu_e, \bar{\nu}_e$ cross-section that has since been presented in [43]. Presenting the cross-section as a function of the outgoing π^- kinematic variables will allow model comparisons in variables which are more sensitive to both nuclear effects and the underlying neutrino-induced pion production physics. As pion production models are developed further in the coming years, it is hoped that this result will constrain theoretical models for pion production in a manner helpful to their improvement. Precise, modern measurements of pion production will be important in future neutrino oscillation experiments such as DUNE which will have a much higher event rate contribution from single and multiple pion production events due to the much higher neutrino energies.

Appendix A

Appendices

A.1 Detector Systematics

For each individual source of detector systematic uncertainty in the selection presented in chapter 5, the relative error distributions in reconstructed muon candidate momentum are presented here. Descriptions of each systematic can be found in the corresponding reference found in table 5.2. Table A.1 provides an index for this appendix, indicating the corresponding figures for each systematic.

Systematic	**Figure**
B-Field Distortions	Fig. A.1
TPC Charge ID Eff.	Fig. A.2
FGD Hybrid Tracking Eff.	Fig. A.3
FGD PID	Fig. A.4
Michel Electron Eff.	Fig. A.5
TPC Momentum Resolution	Fig. A.6
TPC Momentum Scale	Fig. A.7
OOFV Background	Fig. A.8
Pile-up	Fig. A.9
Pion Secondary Interactions	Fig. A.10
Proton Secondary Interactions	Fig. A.11
TPC Cluster Eff.	Fig. A.12
TPC-FGD Matching	Fig. A.14
TPC PID	Fig. A.15
TPC Tracking Eff.	Fig. A.16

Table A.1: Index of detector systematic figures found in this appendix.

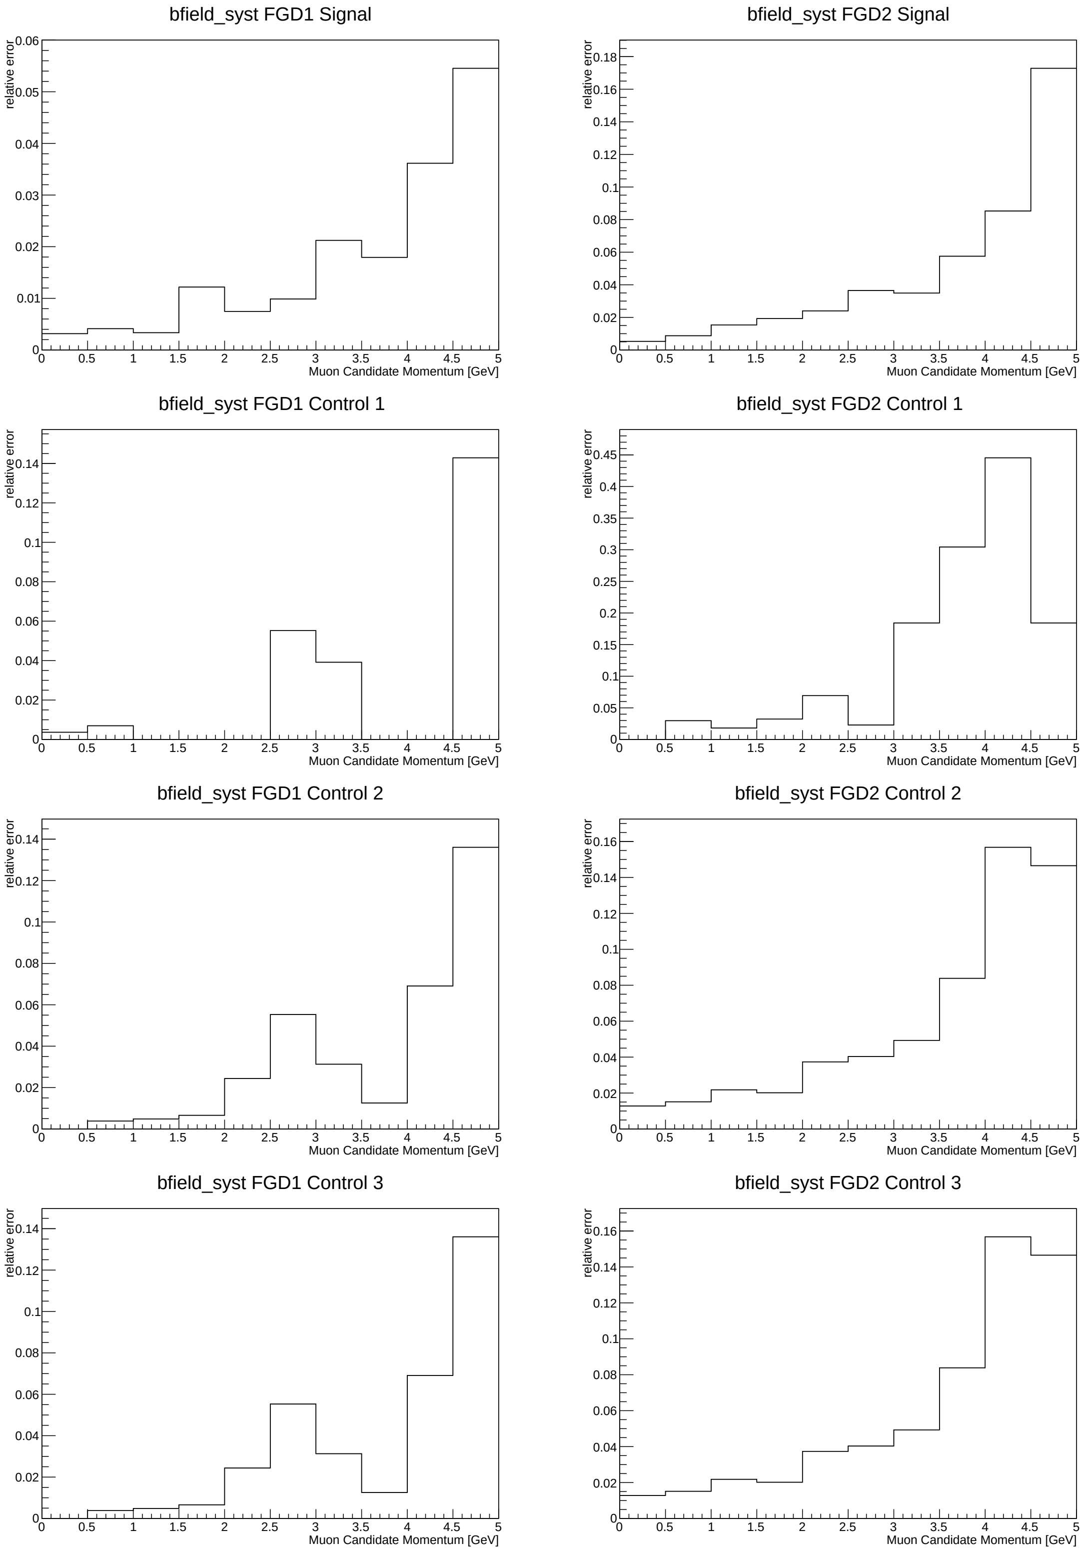

Figure A.1: Relative error from magnetic field uncertainties as a function of reconstructed μ^+ candidate momentum.

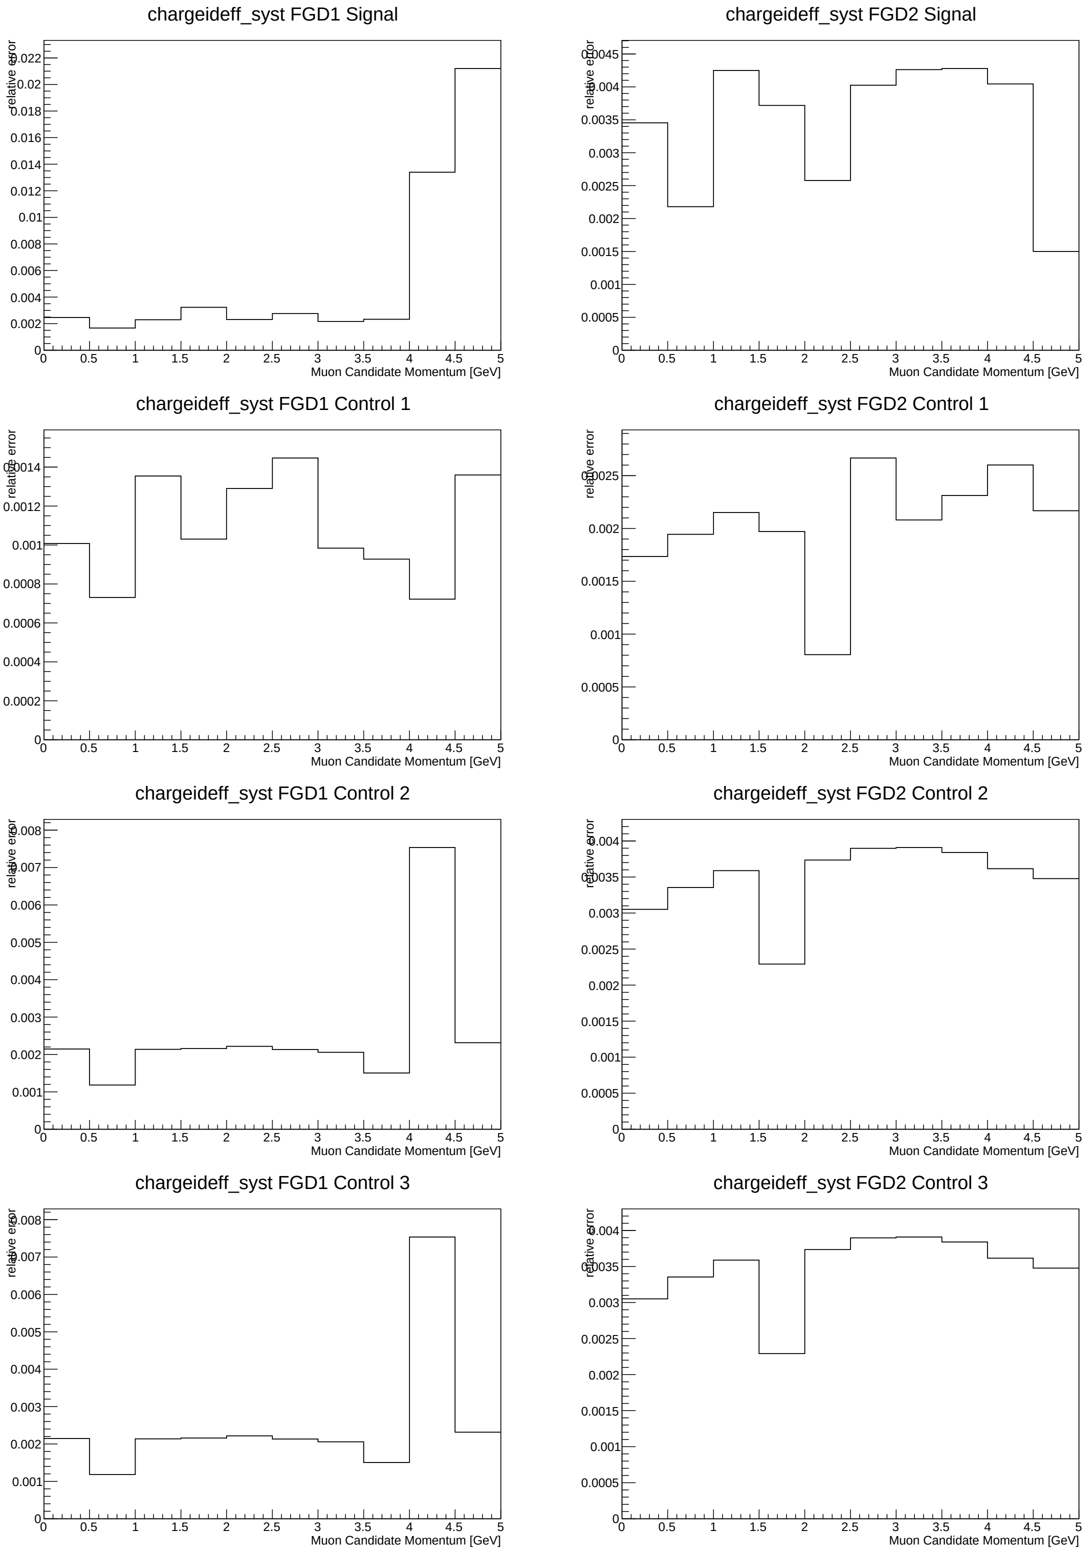

Figure A.2: Relative error from charge identification uncertainties as a function of reconstructed μ^+ candidate momentum.

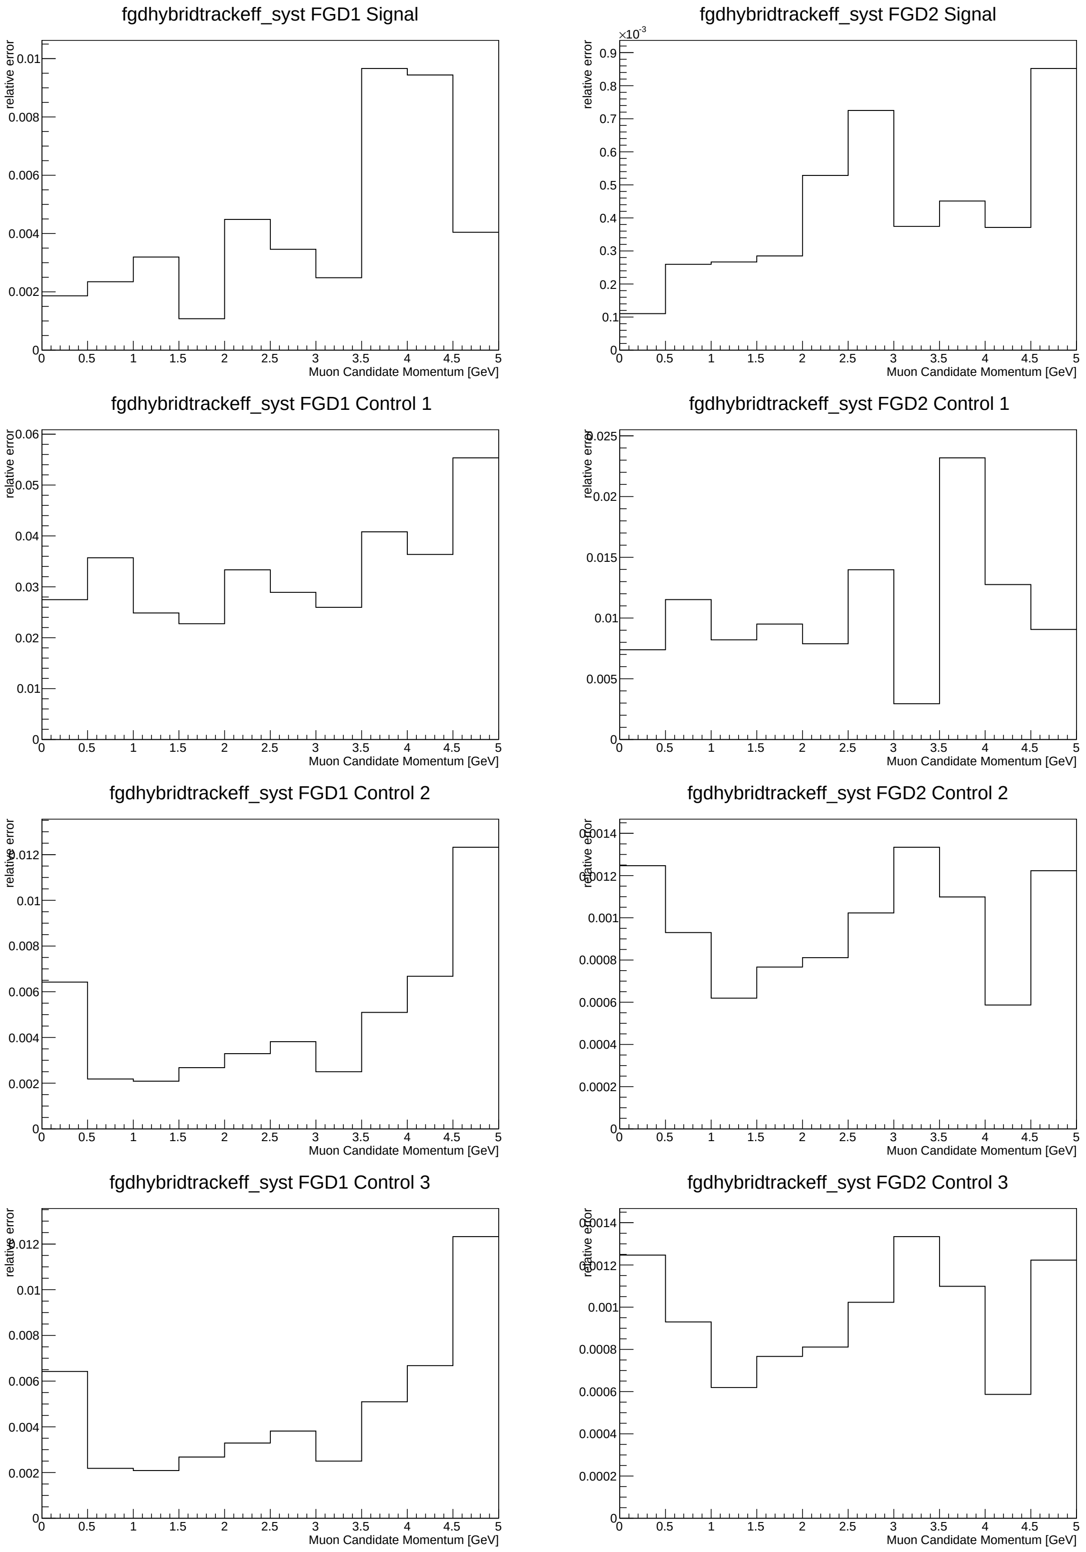

Figure A.3: Relative error from FGD hybrid track uncertainties as a function of reconstructed μ^+ candidate momentum.

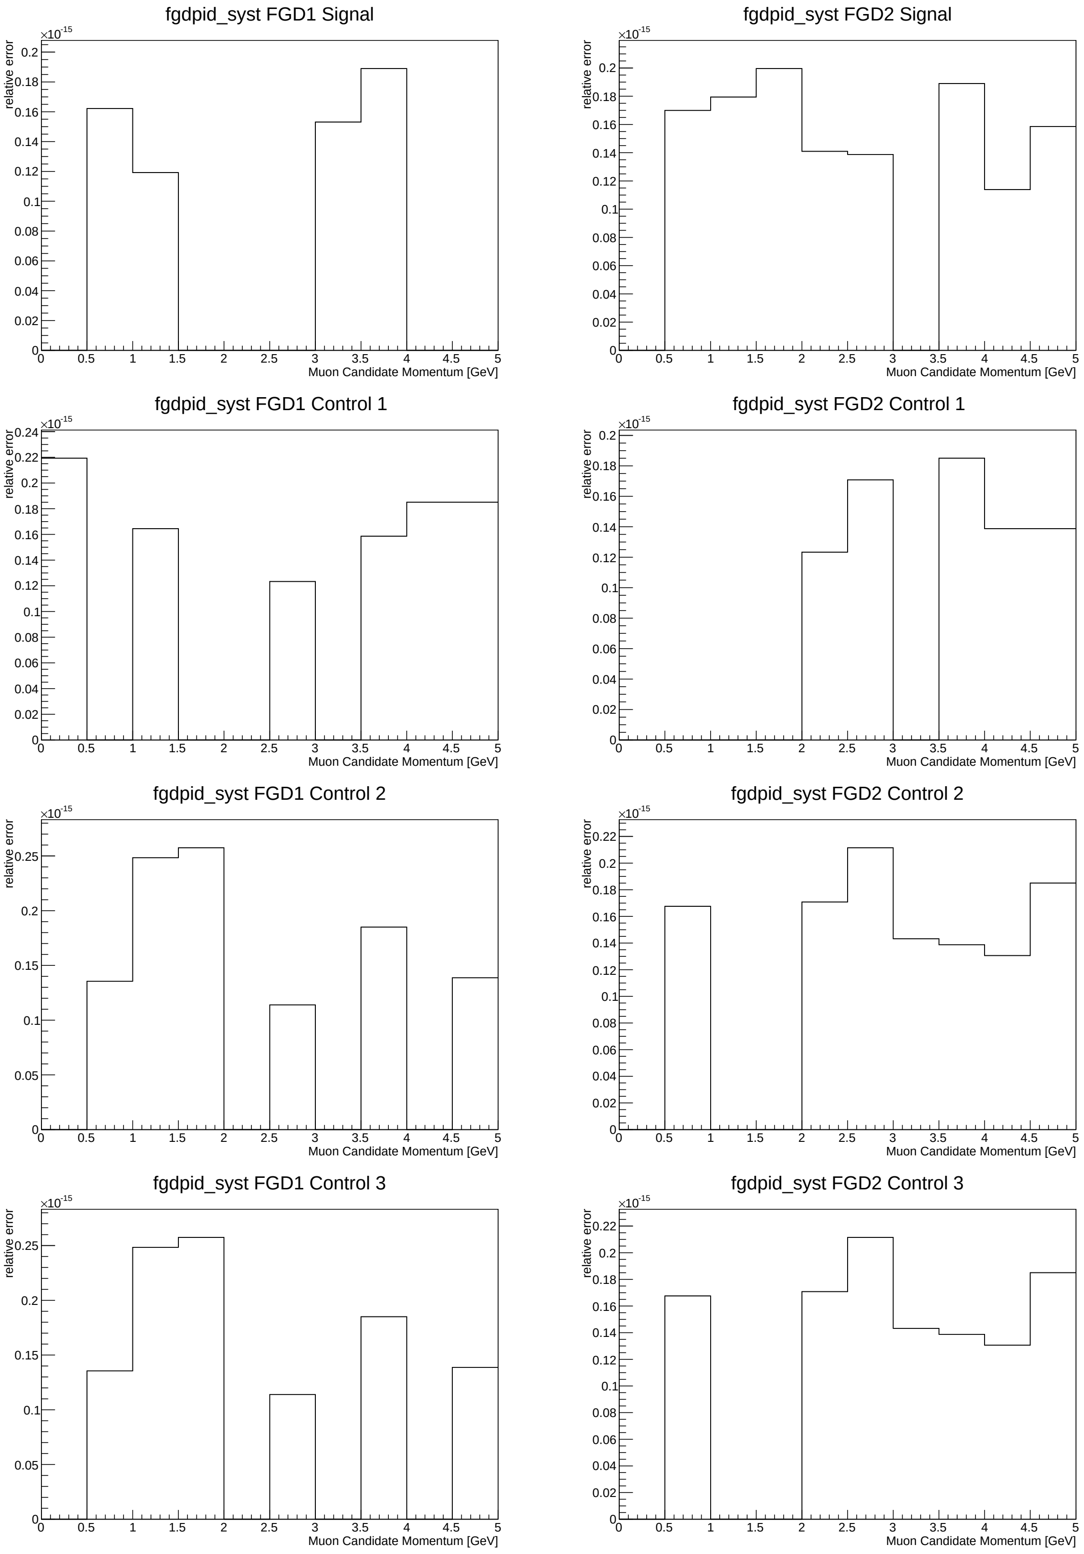

Figure A.4: Relative error from FGD PID uncertainties as a function of reconstructed μ^+ candidate momentum.

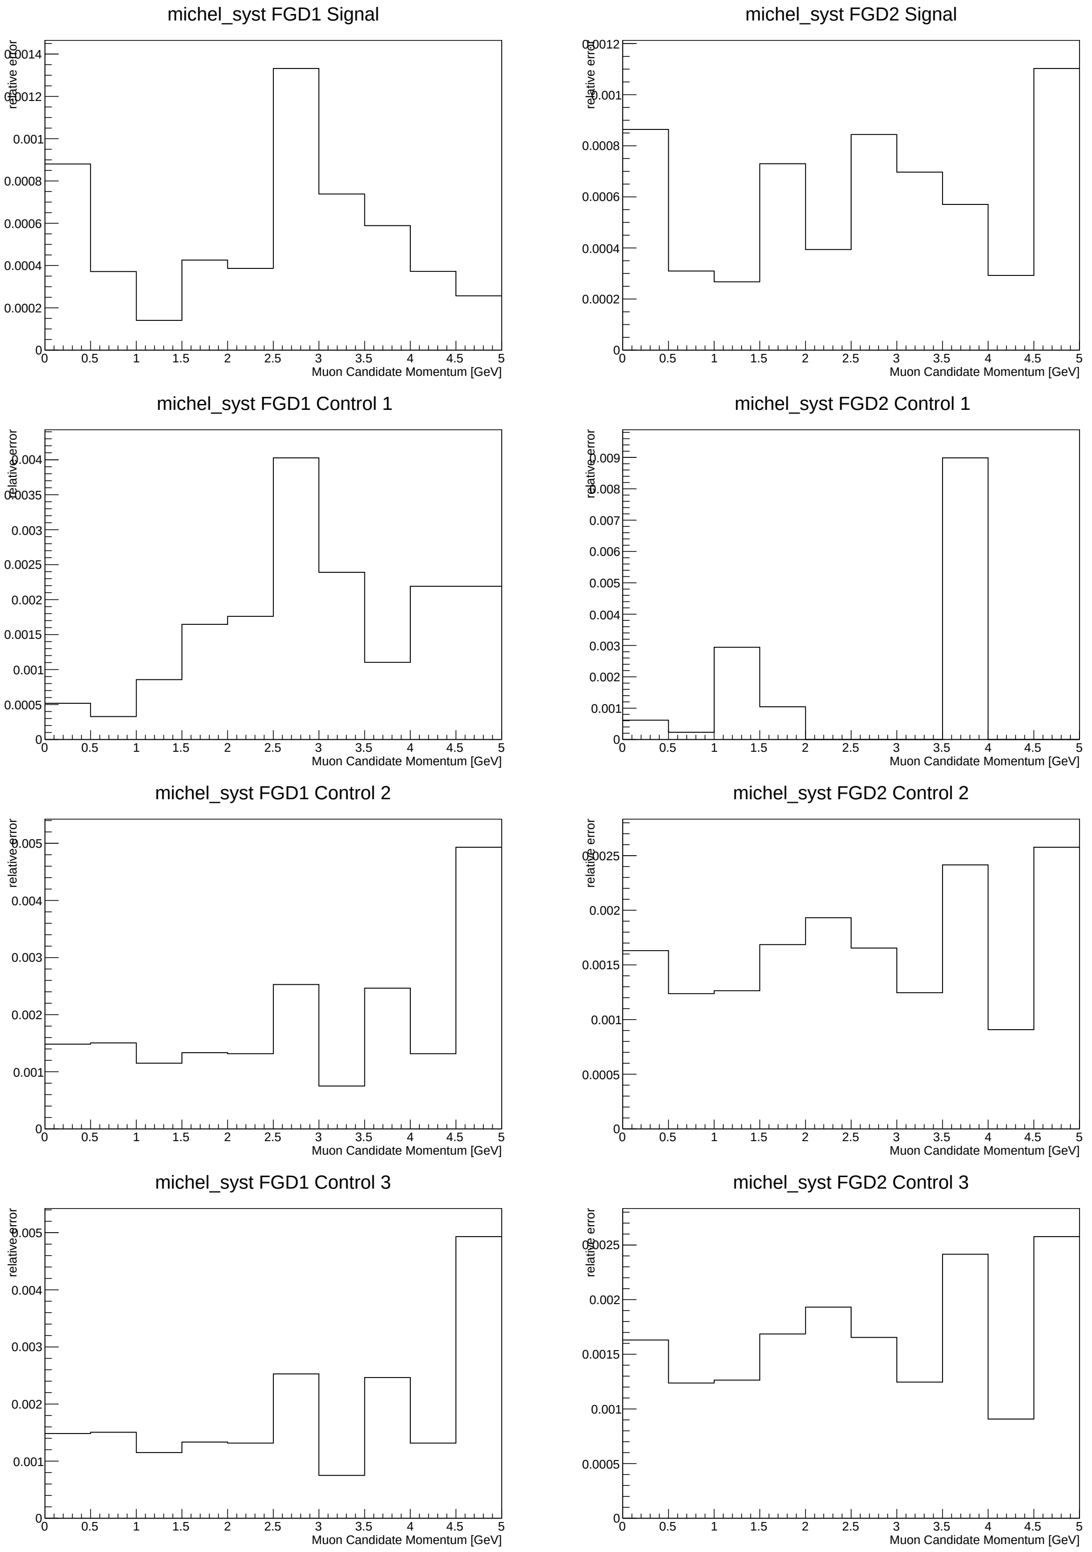

Figure A.5: Relative error from Michel electron modelling uncertainties as a function of reconstructed μ^+ candidate momentum.

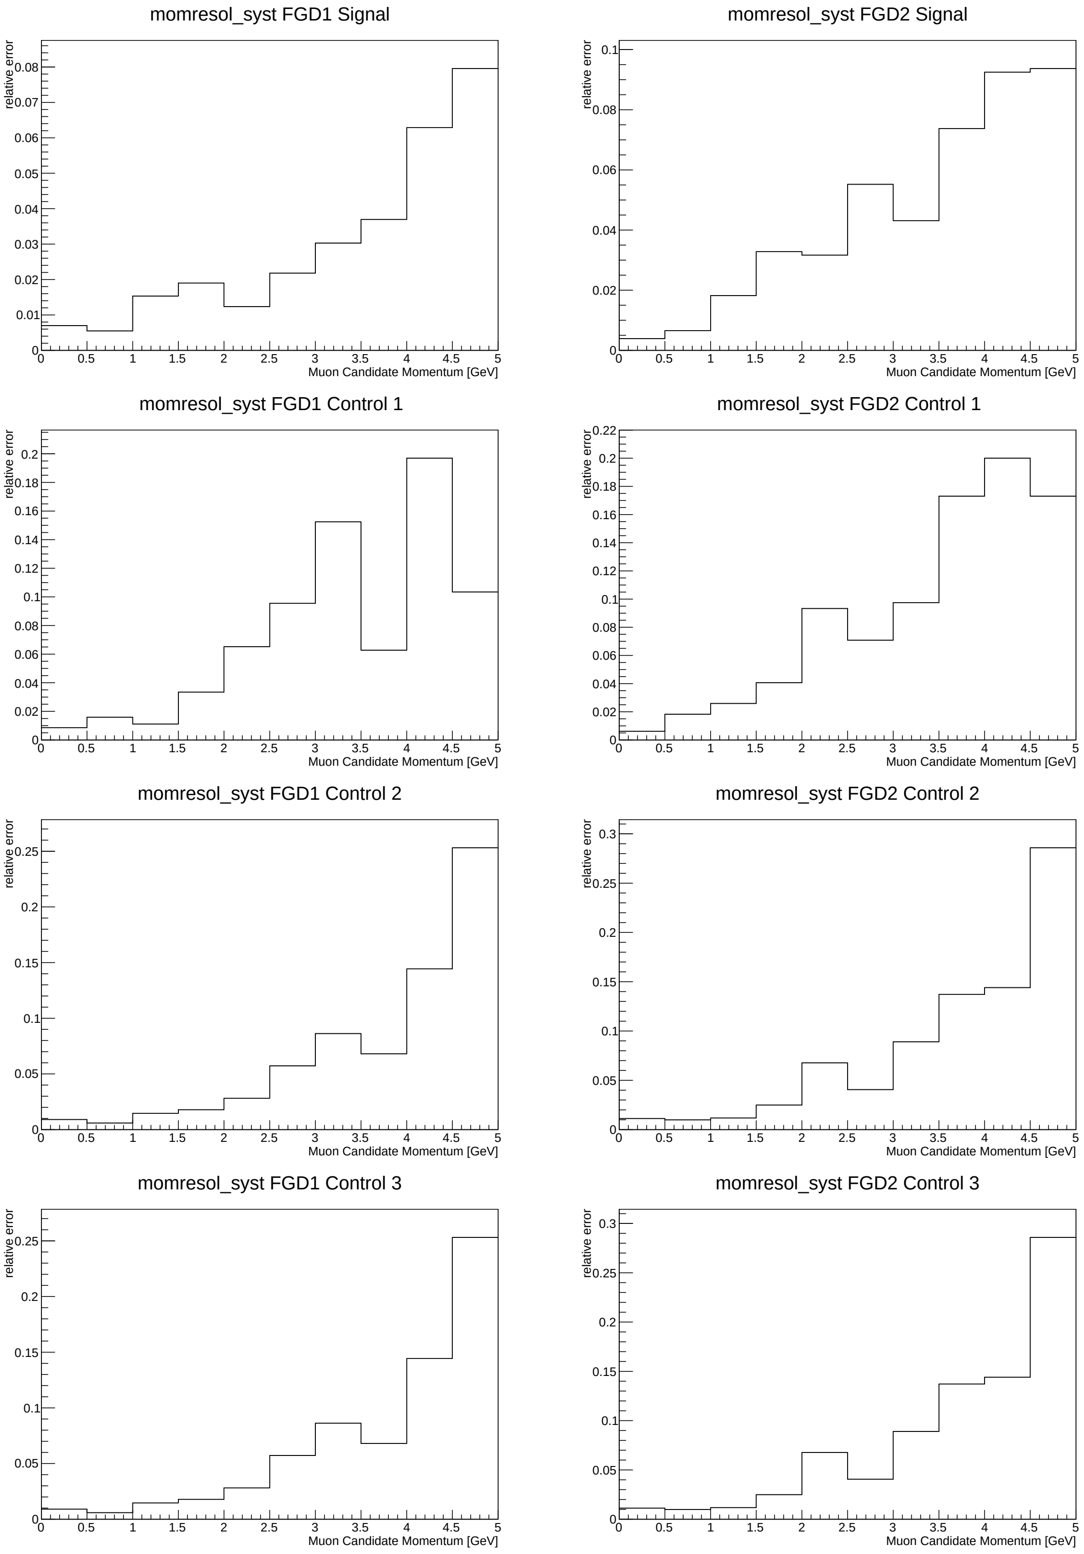
momresol_syst FGD1 Signal

momresol_syst FGD2 Signal

momresol_syst FGD1 Control 1

momresol_syst FGD2 Control 1

momresol_syst FGD1 Control 2

momresol_syst FGD2 Control 2

momresol_syst FGD1 Control 3

momresol_syst FGD2 Control 3

Figure A.6: Relative error from TPC momentum resolution uncertainties as a function of reconstructed μ^+ candidate momentum.

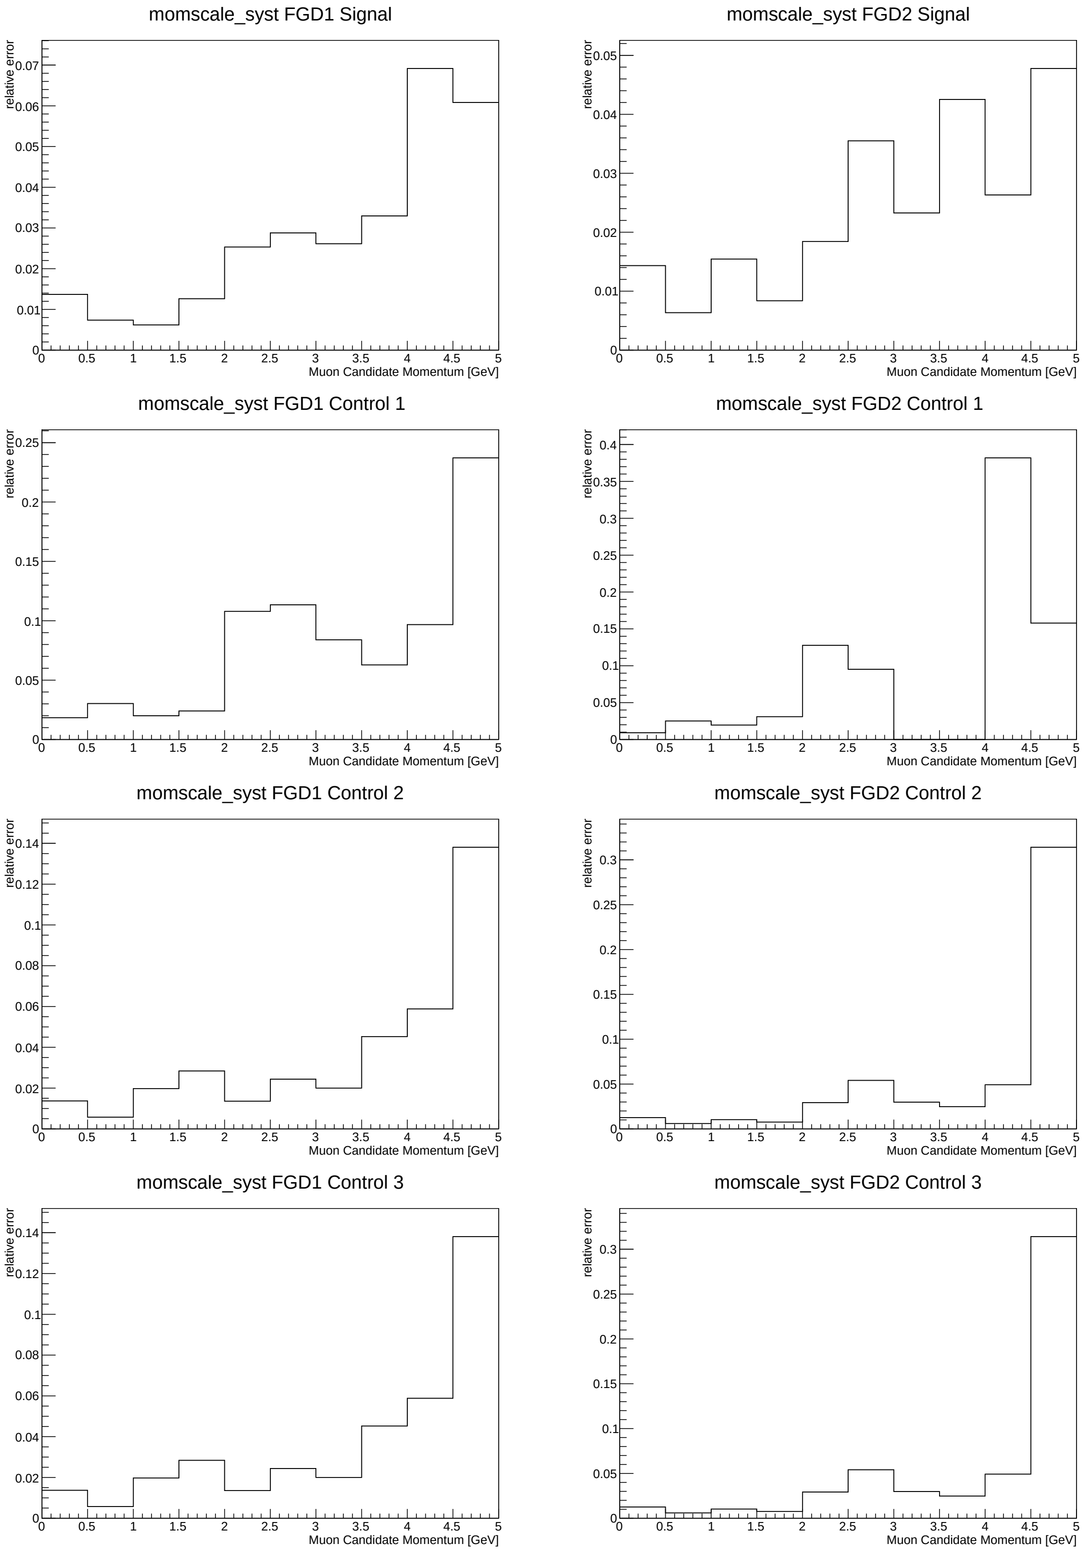

Figure A.7: Relative error from TPC momentum scale uncertainties as a function of reconstructed μ^+ candidate momentum.

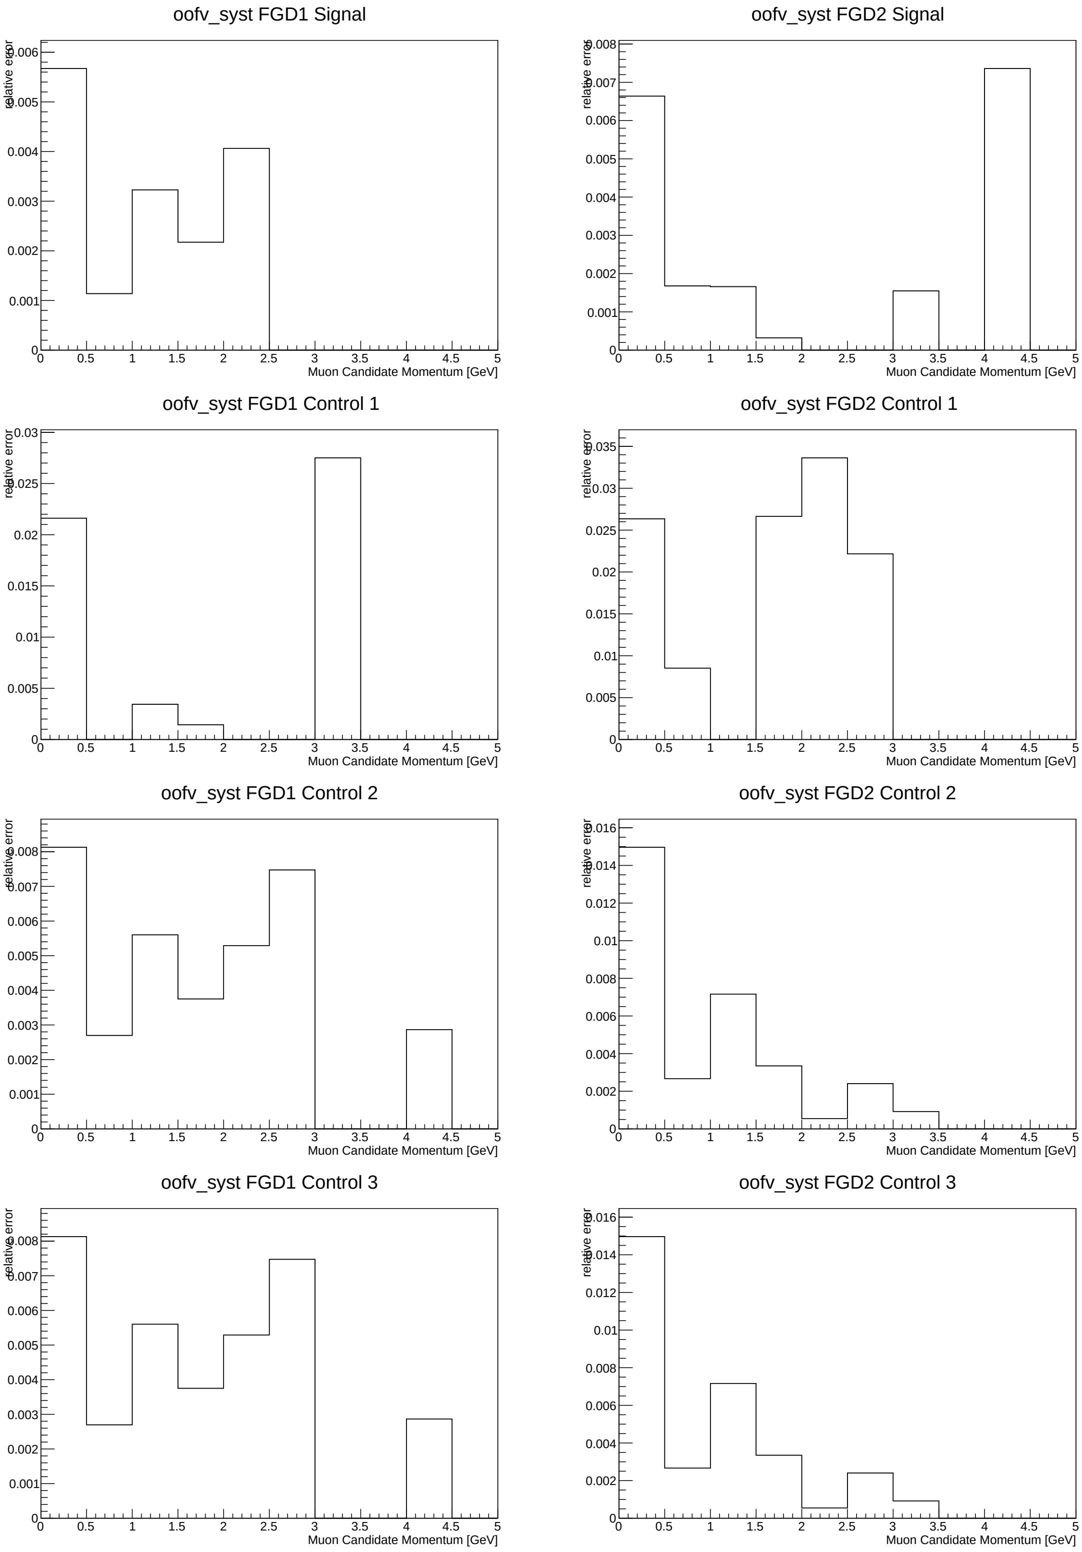

Figure A.8: Relative error from fiducial volume related uncertainties as a function of reconstructed μ^+ candidate momentum.

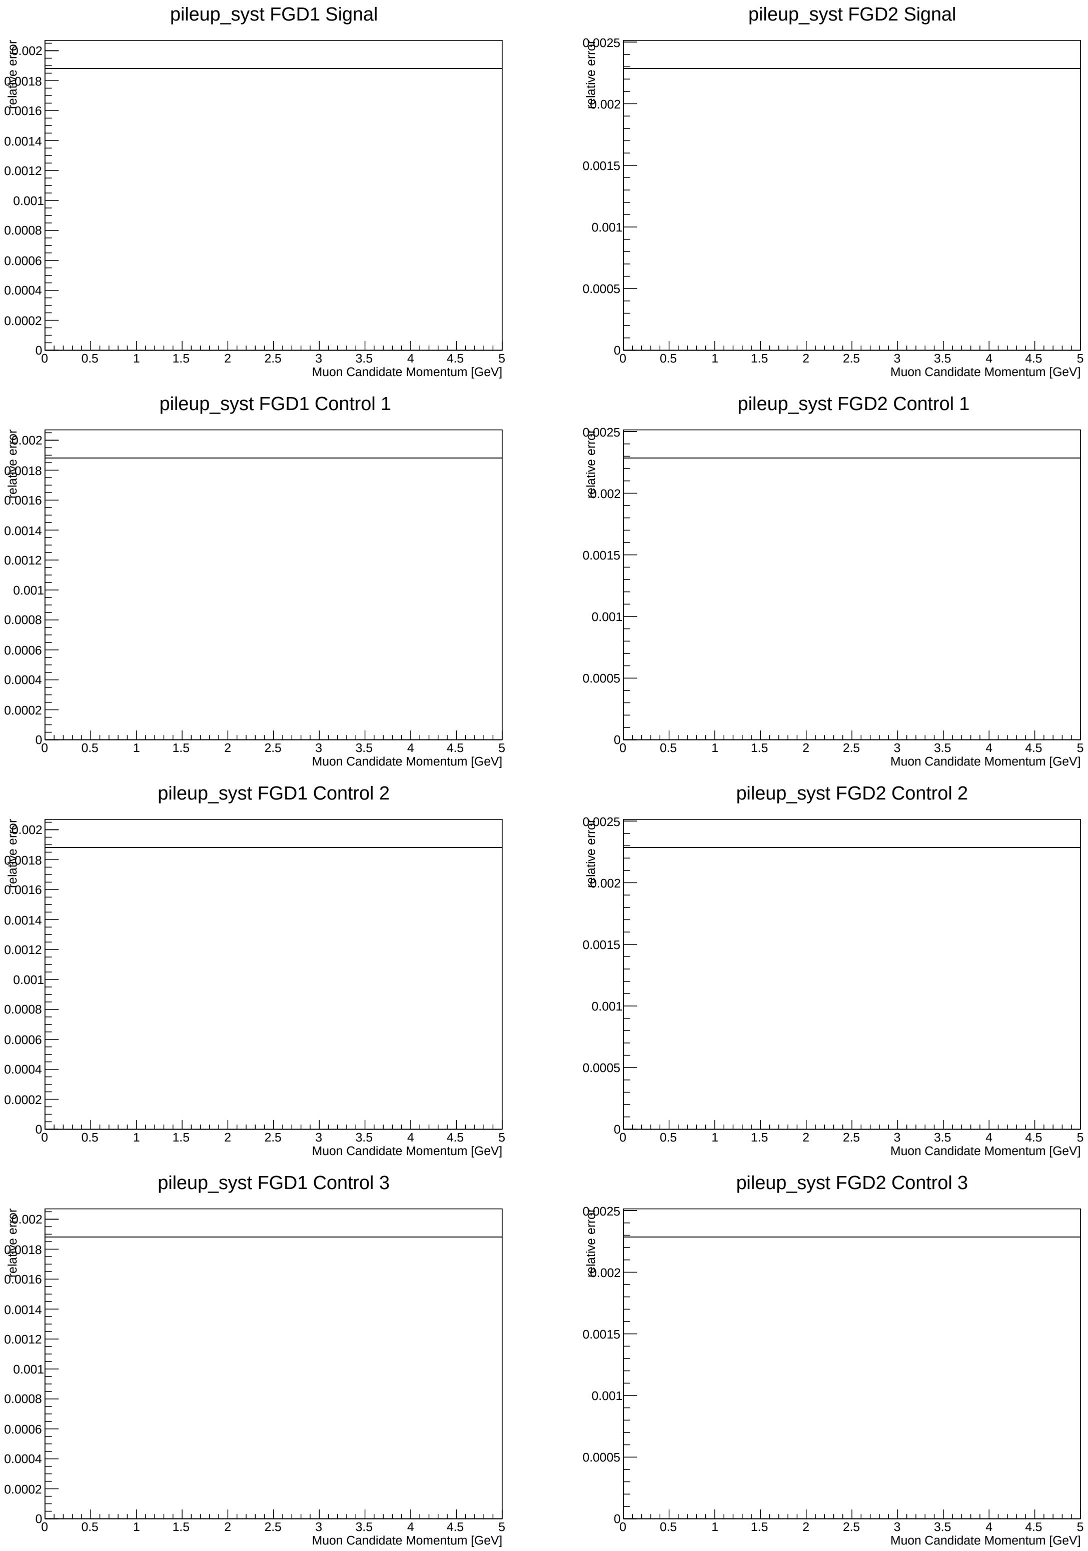

Figure A.9: Relative error from detector pileup uncertainties as a function of reconstructed μ^+ candidate momentum.

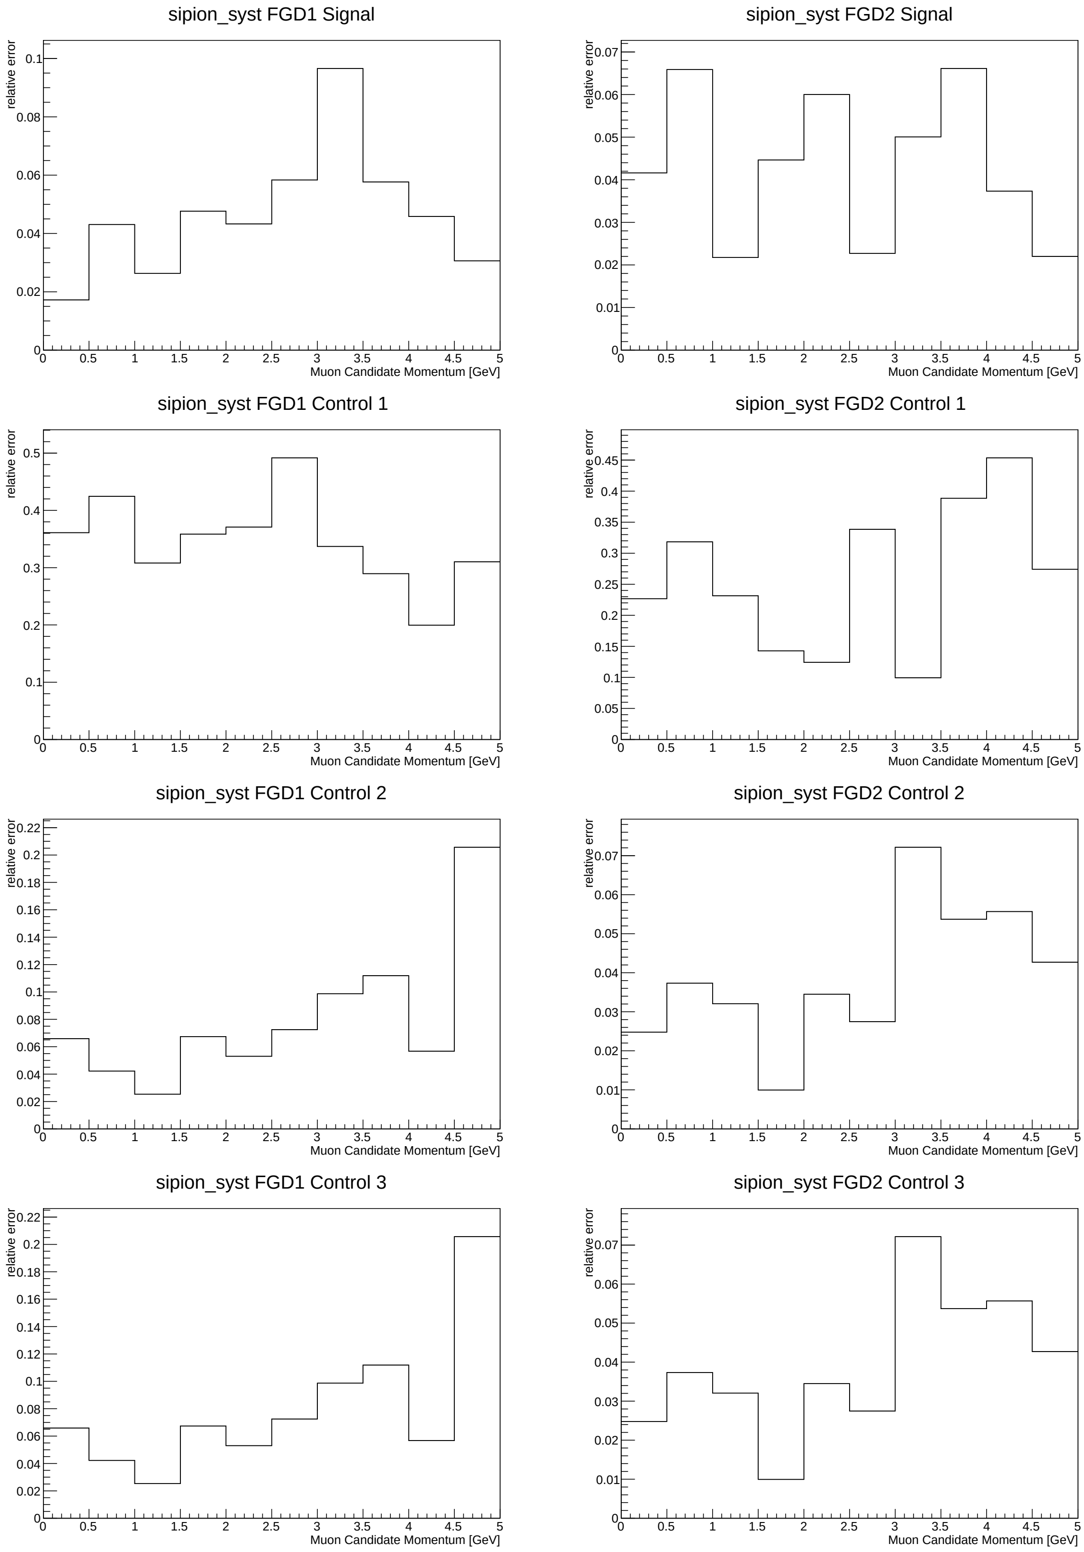
sipion_syst FGD1 Signal

sipion_syst FGD2 Signal

sipion_syst FGD1 Control 1

sipion_syst FGD2 Control 1

sipion_syst FGD1 Control 2

sipion_syst FGD2 Control 2

sipion_syst FGD1 Control 3

sipion_syst FGD2 Control 3

Figure A.10: Relative error from pion secondary interaction uncertainties as a function of reconstructed μ^+ candidate momentum.

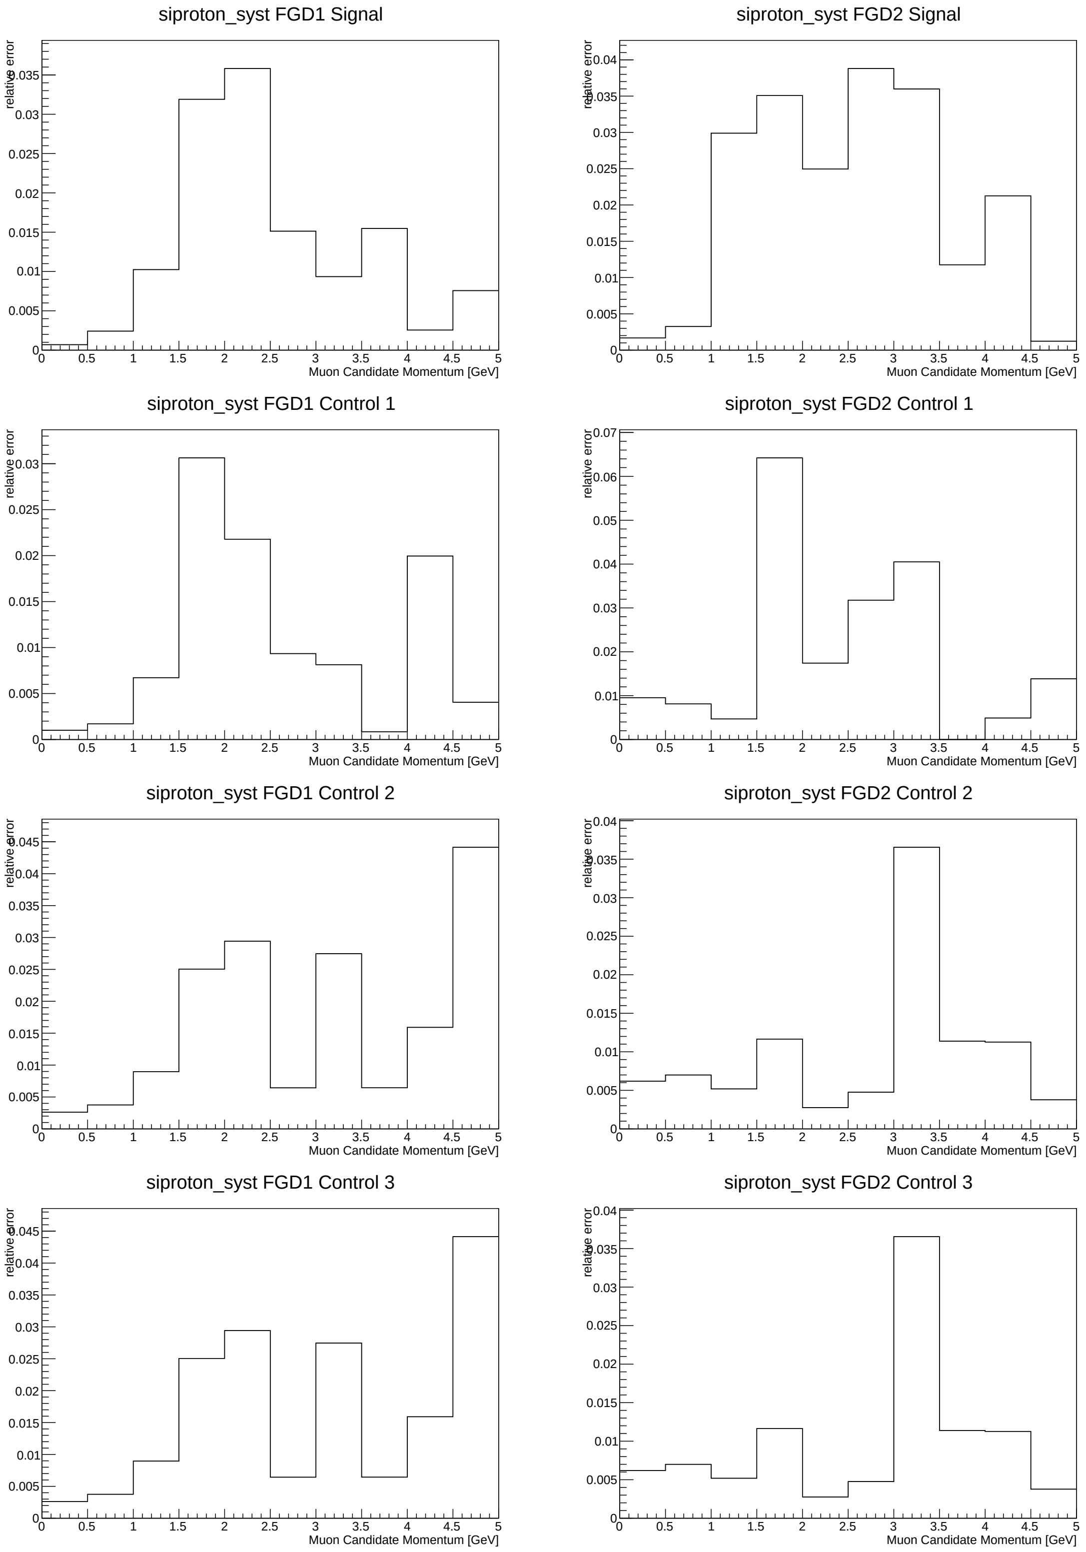
siproton_syst FGD1 Signal

siproton_syst FGD2 Signal

siproton_syst FGD1 Control 1

siproton_syst FGD2 Control 1

siproton_syst FGD1 Control 2

siproton_syst FGD2 Control 2

siproton_syst FGD1 Control 3

siproton_syst FGD2 Control 3

Figure A.11: Relative error from proton secondary interaction uncertainties as a function of reconstructed μ^+ candidate momentum.

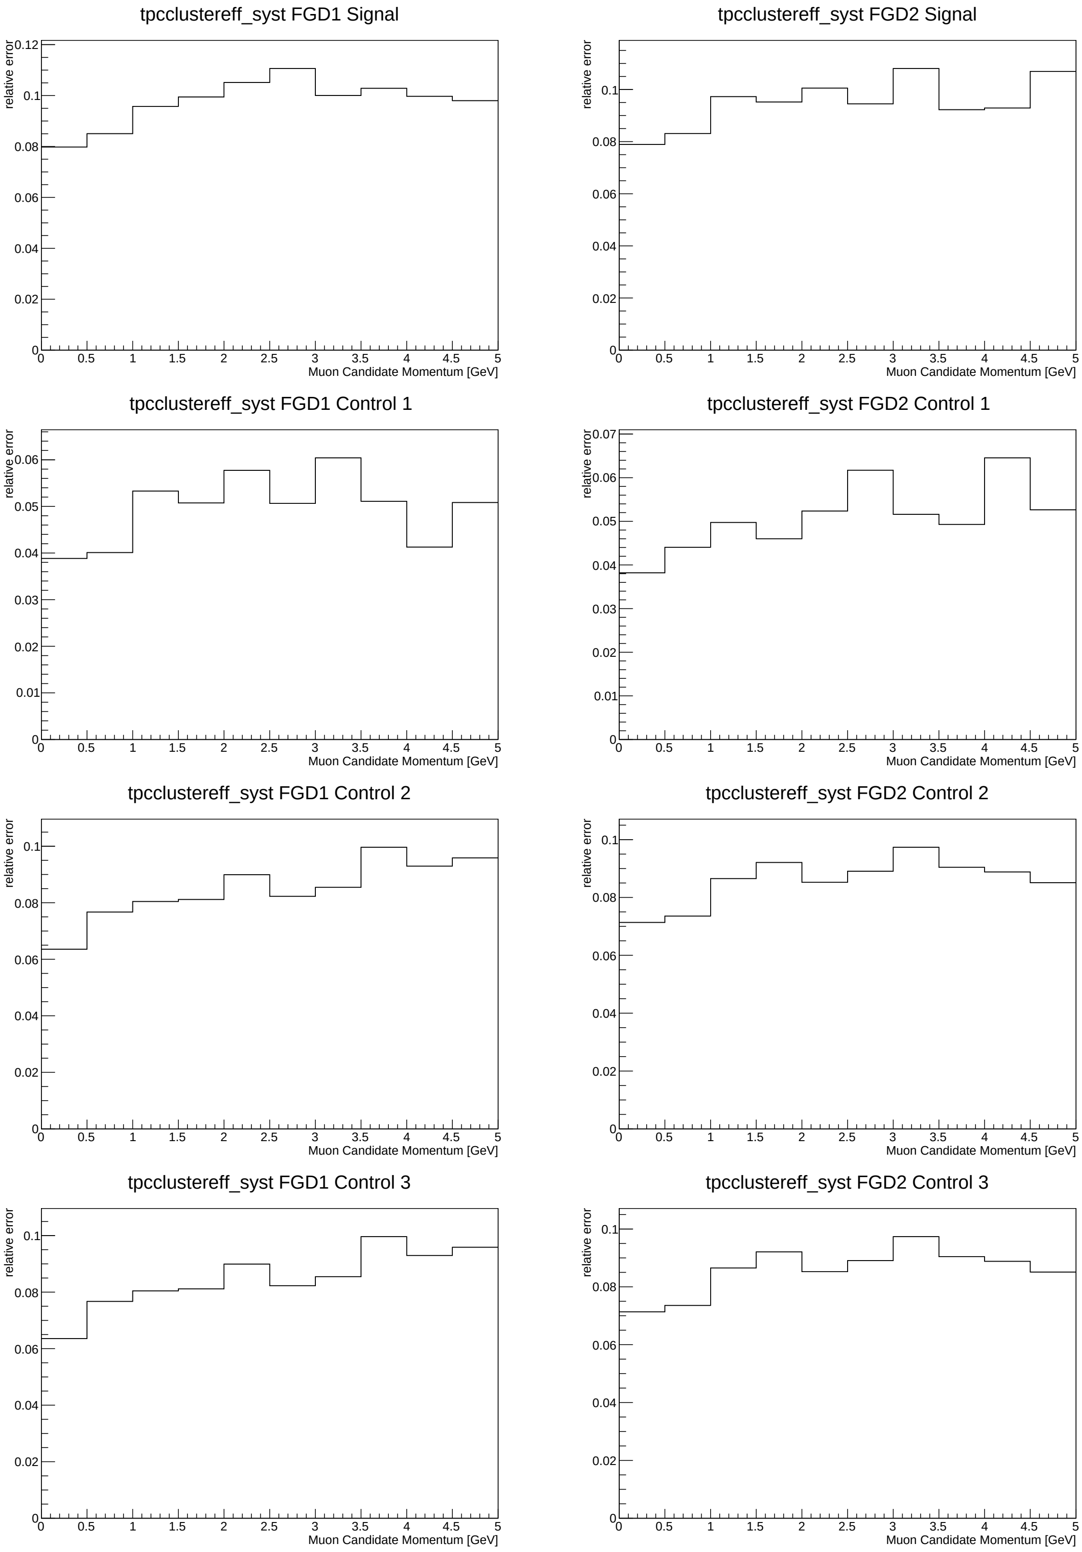

Figure A.12: Relative error from TPC hit-clustering uncertainties as a function of reconstructed μ^+ candidate momentum.

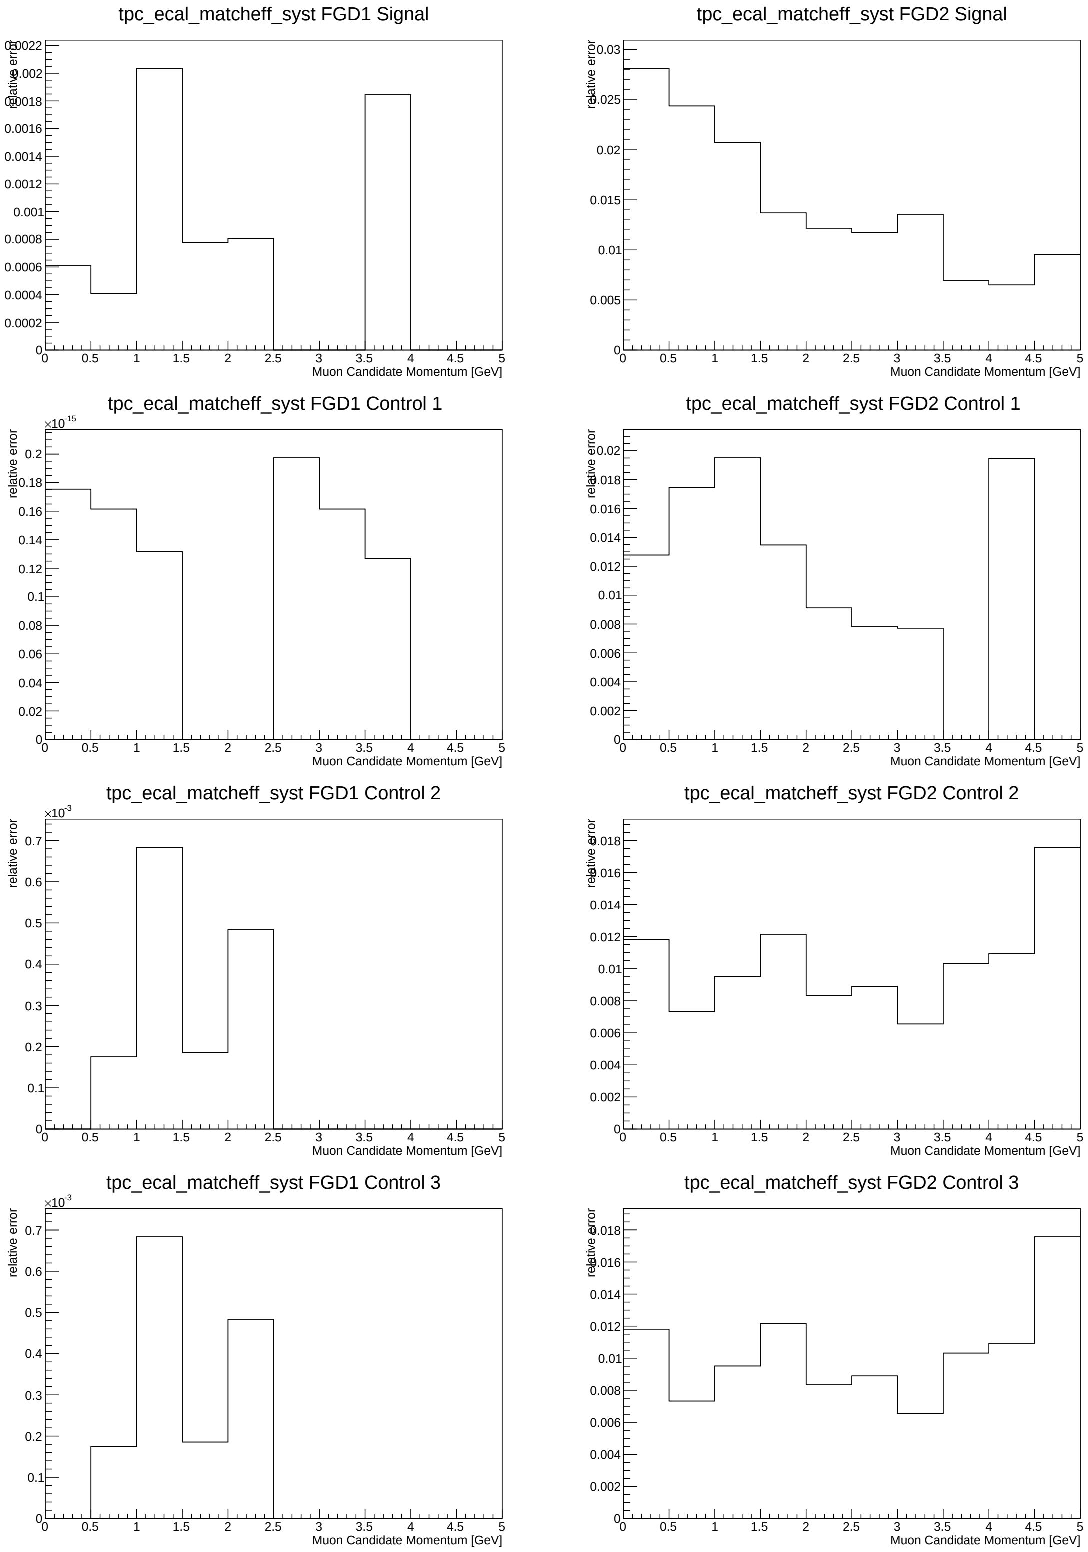

Figure A.13: Relative error from TPC-ECal track matching uncertainties as a function of reconstructed μ^+ candidate momentum.

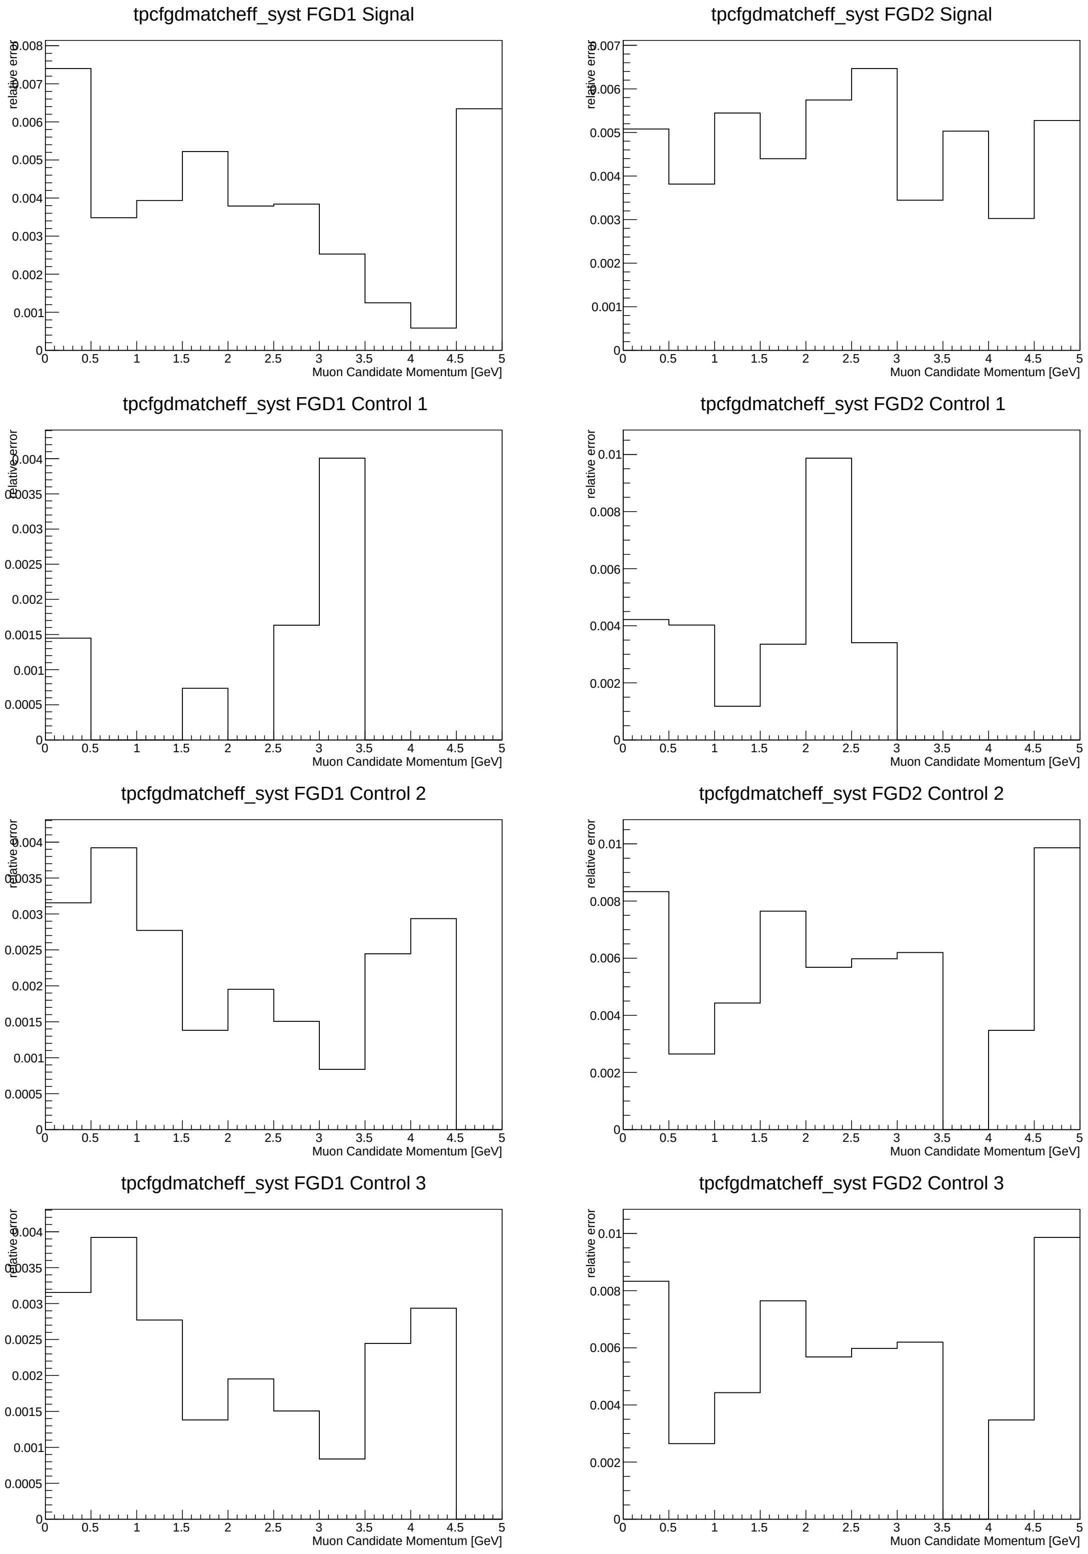

Figure A.14: Relative error from TPC-FGD track matching uncertainties as a function of reconstructed μ^+ candidate momentum.

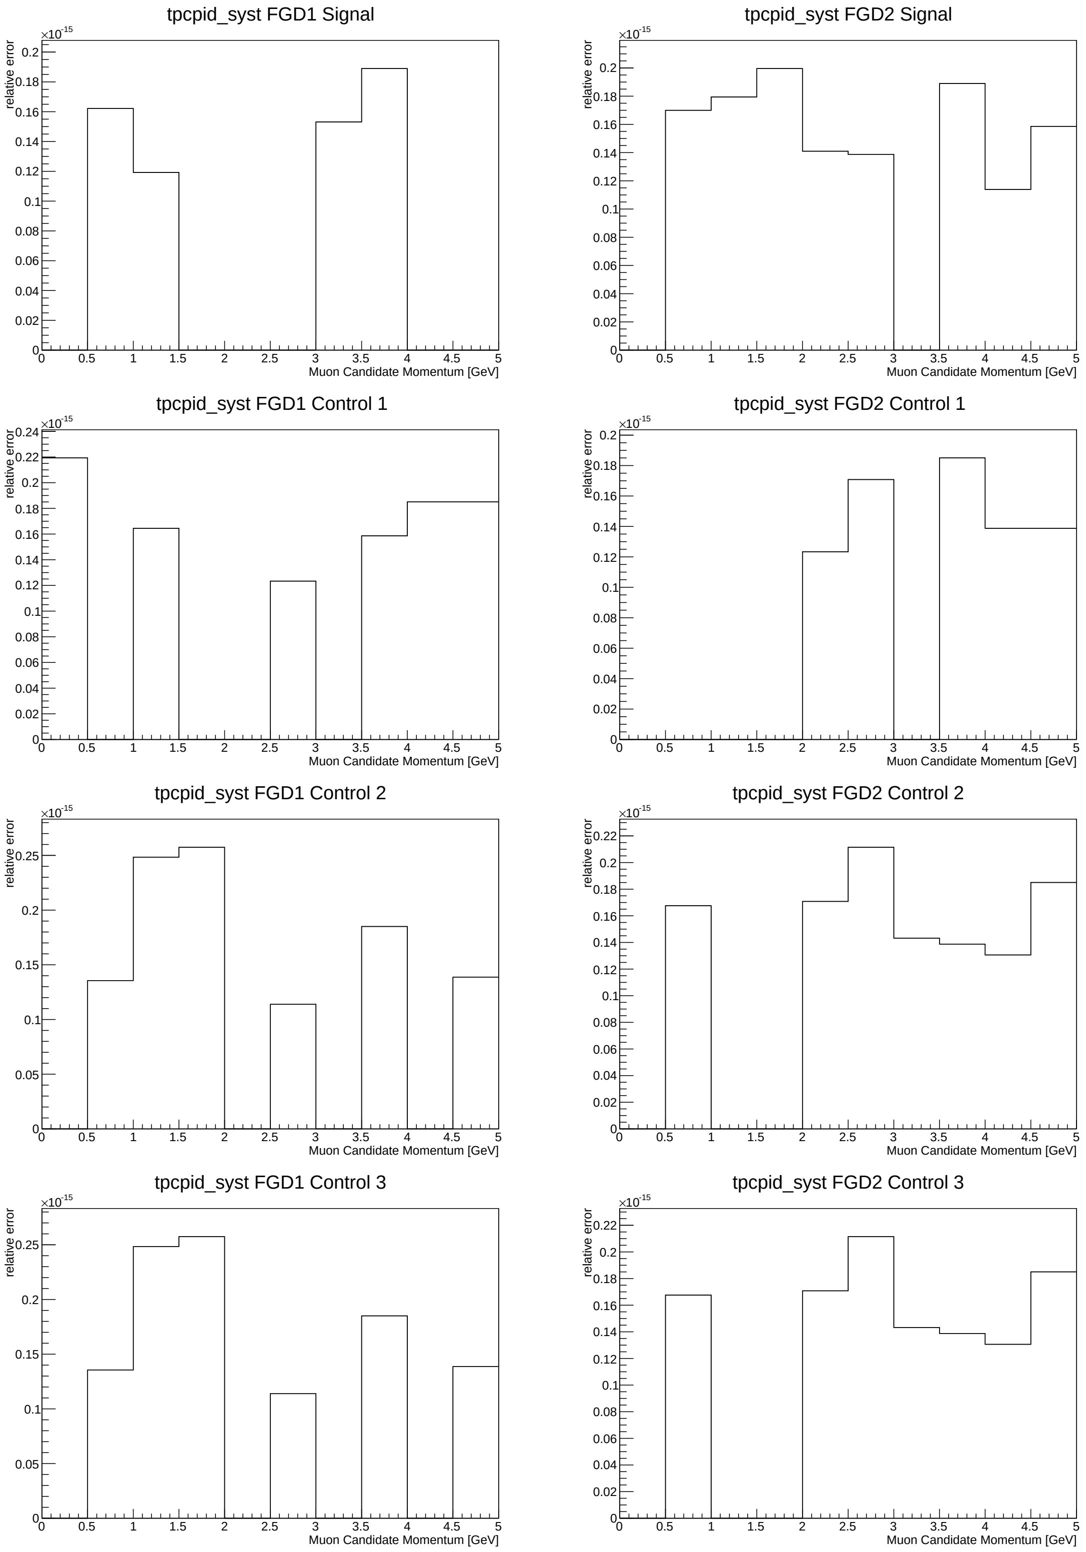

Figure A.15: Relative error from TPC PID uncertainties as a function of reconstructed μ^+ candidate momentum.

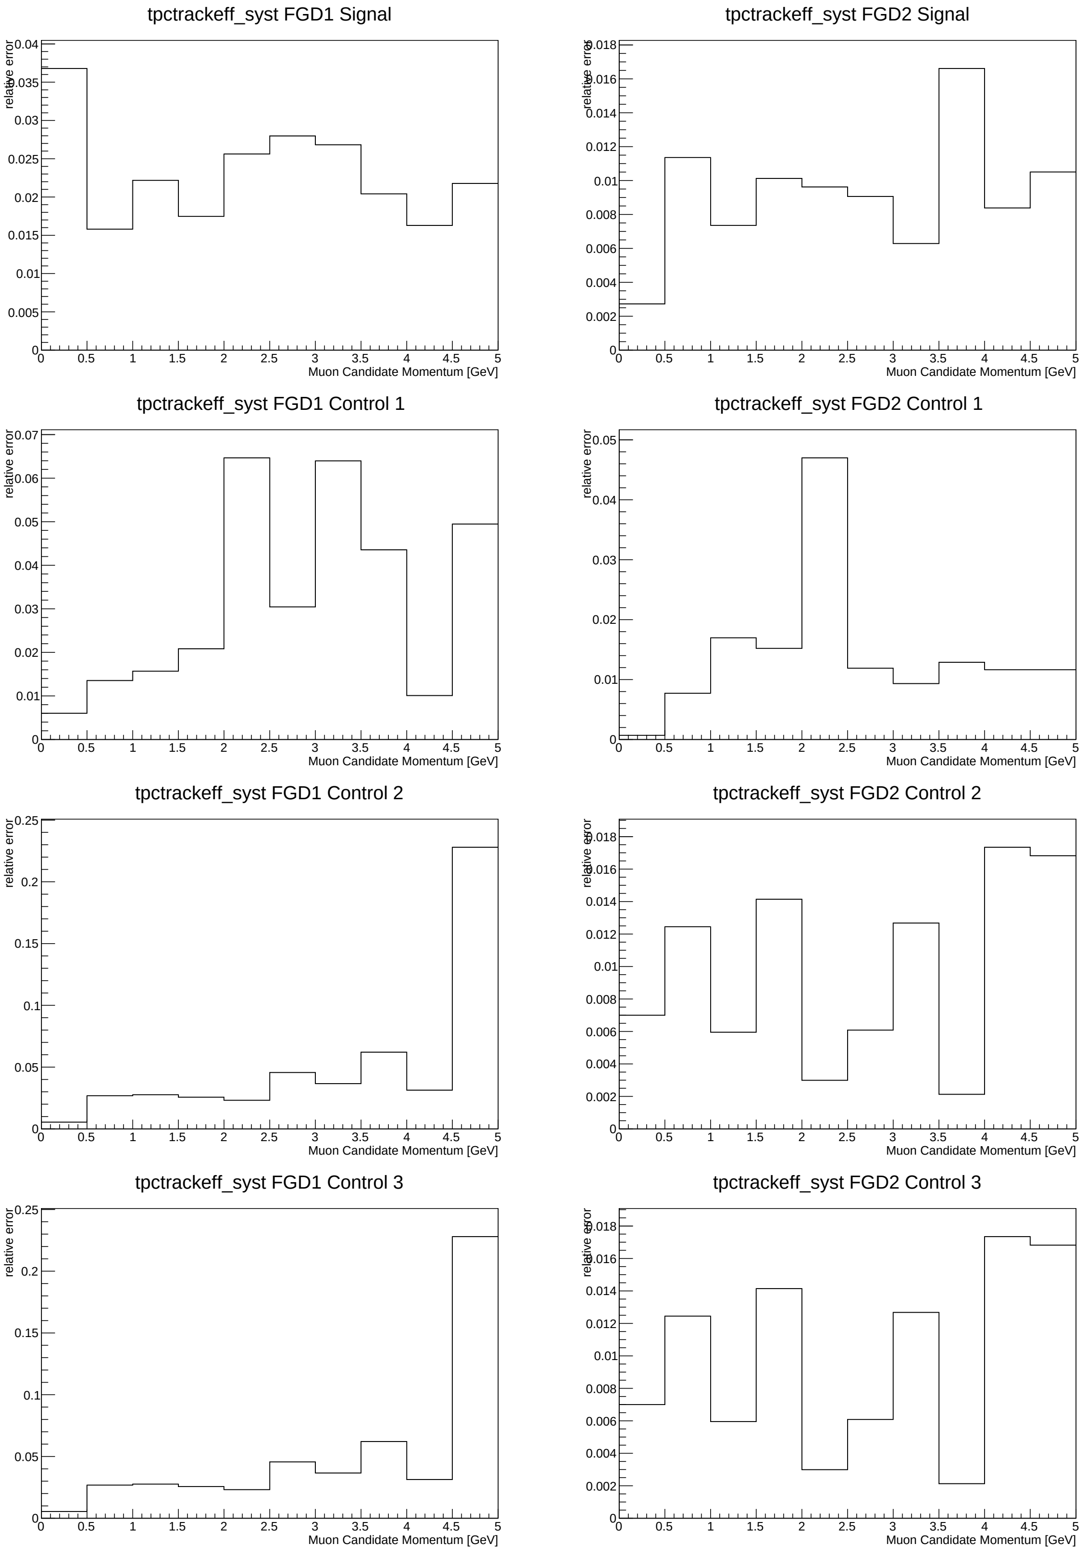

Figure A.16: Relative error from TPC reconstruction efficiency uncertainties as a function of reconstructed μ^+ candidate momentum.

Bibliography

[1] *The Reines-Cowan Experiments: Detecting the Poltergeist*, F. Reines and C. Cowan, Los Alamos Science vol. 25 (1997)

[2] *Detection of the Free Neutrino: a Confirmation*, C. Cowan et al., Science vol. 124 (3212) 103104 (1956)

[3] *The Brookhaven Alternating Gradient Synchrotron*, R. A. Beth, C. Lasky, Science 128:3336 (1958)

[4] *Observation of Tau Neutrino Interactions*, DONUT Collaboration, Phys. Lett. B504 218-224 (2001)

[5] *Precision Electroweak Measurements on the Z Resonance*, ALEPH Collaboration et al., Phys. Rept. 427 257-454 (2006)

[6] *Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector*, B. T. Cleveland et al., The Astrophysical Journal vol. 496 (1) 505 - 526 (1998)

[7] *Where do we stand with solar neutrino oscillations?*, J. N. Bahcall, P. I. Krastev, A. Y. Smirnov, Phys. Rev. D 58:096016 (1998)

[8] *Solar neutrino flux measurements by the Soviet-American gallium experiment (SAGE) for half the 22-year solar cycle*,

[9] *Results of the whole GALLEX experiment*, M. Cribier et al., Nuclear Physics B Proceedings Supplements vol. 70 (1) 284 291 (1999) J. N. Abdurashitov et al., Journal of Experimental and Theoretical Physics vol. 95 (2) 181193 (2002)

[10] *Evidence for oscillation of atmospheric neutrinos*, Y. Fukuda et al. (The Super-Kamiokande Collaboration), Phys. Rev. Lett. 81 1562-1567 (1998)

[11] *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, Q. R. Ahmad et al. (The SNO Collaboration), Phys. Rev. Lett. 89 011301 (2002)

[12] *Inverse beta processes and nonconservation of lepton charge*, Sov. Phys. Journ. Exp. Theor. Phys. 34:247 (1958)

[13] *Remarks on the Unified Model of Elementary Particles*, Z. Maki, M. Nakagawa, S. Sakata, Prog. Theor. Phys. 28:5 (1962)

[14] *The Review of Particle Physics*, P. A. Zyla et al., Prog. Theor. Exp. Phys. 2020 083C01 (2020)

[15] *The Mikheyev-Smirnov-Wolfenstein (MSW) Effect*, A. Smirnov, arXiv:1901.11473 [hep-ph] (2018)

[16] *Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects*, P. F. de Salas, S. Gariazzo, O. Mena, C. Ternes, M. Tórtola, Front. Astron. Space Sci. 5 36 (2018)

[17] *Neutrino Mass Hierarchy*, X. Qian, P. Vogel, Prog. Part. Nucl. Phys. 83 (2015)

[18] *Short Baseline Neutrino Oscillation Experiments*, T. Katori, J. Phys. Conf. Series 598 (2013)

[19] *First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA*, M. A. Acero et al., Phys. Rev. Lett. 123, 151803 (2019)

[20] *T2K Technical Note 399: Super-Kamiokande Data Quality, MC, and Systematics in Run10*, The T2K Collaboration (2020)

[21] *Constraint on the matterantimatter symmetry-violating phase in neutrino oscillations*, K. Abe et al. (The T2K Collaboration), Nature 50, 339-344 (2020)

[22] *Hyper-Kamiokande Design Report*, K. Abe et al., 1805.04163 [physics.ins-det] (2018)

[23] *The Nuance neutrino physics simulation, and the future*, D. Casper et al., Nucl. Phys. Proc. Suppl. 112 161170 (2002)

[24] *From eV to EeV: Neutrino Cross Sections Across Energy Scales*, J. A. Formaggio, G. P. Zeller, Rev. Mod. Phys. 84 (2012)

[25] *Neutrino reactions at accelerator energies*, C. H. Llewellyn Smith, Physics Reports vol. 3 (5) 261379 (1972)

[26] *Measurement of the Neutrino Neutral-Current Elastic Differential Cross Section*, The MiniBooNE Collaboration, Phys. Rev. D82 092005 (2010)

[27] *Neutrino Excitation of Baryon Resonances and Single Pion Production*, D. Rein and L.M. Sehgal, Annals Phys. 133 (1981) 79-153

[28] *Single pion production in neutrino-nucleon interactions*, M. Kabirnezhad, Phys. Rev. D 97 013002 (2018)

[29] *Weak Pion Production off the Nucleon*, E. Hernandez, J. Nieves, MV̇alverde, Phys. Rev. D76 033005 (2007)

[30] *Coherent π^0 production in neutrino reactions*, D. Rein and L.M. Sehgal, Nuc. Phys. B 223-1 (1983)

[31] *PCAC and coherent pion production by low energy neutrinos*, C. Berger and L.M. Sehgal, Phys. Rev. D 79 053003 (2009)

[32] *Nuclear theory and event generators for charge-changing neutrino reactions*, J.W. Van Orden and T.W. Donnelly, Phys. Rev. C 100 044620 (2019)

[33] *Two-nucleon spectral function in infinite nuclear matter*, O. Benhar and A. Fabrocini, Phys. Rev. C 62 034304 (2000)

[34] *NuSTEC White Paper: Status and Challenges of Neutrino-Nucleus Scattering*, L. Alvares-Ruso et al., Progress in Particle and Nuclear Physics 100 (2018)

[35] *A neutrino interaction simulation program library NEUT*, Y. Hayato, Acta Physica Polonica B 40, 2477 (2009)

[36] *The GENIE Neutrino Monte Carlo Generator*, C. Andreopolous et al., Nucl. Instrum. Meth. A614 87-104 (2010)

[37] *Effects of final-state interactions in neutrino-nucleus interactions*, T. Golan et al., Phys. Rev. C 86, 015505 (2012)

[38] *Transport-theoretical description of nuclear reactions*, O. Buss et al., Phys. Rep. 512 1–2 (2012)

[39] *Inclusive charged-current neutrino-nucleus reactions*, J. Nieves et al., Phys. Rev. C 83 045501 (2011)

[40] *Evidence for Oscillation of Atmospheric Neutrinos*, Y. Fukuda et al. (The Super-Kamiokande Collaboration), Phys. Rev. Lett. vol. 81 (1998)

[41] *Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector*, K. Abe et al. (The T2K Collaboration), Phys. Rev. D vol. 96 (2017)

[42] *Constraint on the matterantimatter symmetry-violating phase in neutrino oscillations*, K. Abe et al. (The T2K Collaboration), Nature 580, 339344 (2020)

[43] *Measurement of the charged-current electron (anti-)neutrino inclusive cross-sections at the T2K off-axis near detector ND280*, K. Abe et al. (The T2K Collaboration), Journal of High Energy Phys. 2020, 114 (2020)

[44] *First combined measurement of the muon neutrino and antineutrino charged-current cross section without pions in the final state at T2K*, K. Abe et al., Phys. Rev. D 101 112001 (2020)

[45] *RecPack, a general reconstruction toolkit*, A. Cervera-Villanueva and J.J. Gómez-Cadenas, ICATPP11 Proceedings (2011)

[46] *The T2K Neutrino Flux Prediction*, K. Abe et al. (The T2K Collaboration), Phys. Rev. D 87, 012001 (2013)

[47] *J-PARC TDR*, KEK and JAERI, https://cds.cern.ch/record/747209/files/34072617.pdf, 2003 (accessed May 2020)

[48] *Design and performance of the muon monitor for the T2K neutrino oscillation experiment*, K. Matsuoka et al., Nucl. Instrum. Meth. A624, 591 (2010).

[49] *Overview of the FLUKA code*, G. Battistoni et al., Ann. Nuc. Energy 82 (2015)

[50] *The GEANT - CALOR interface and benchmark calculations of ZEUS test calorimeters*, C. Zeitnitz and T. A. Gabriel, Nucl. Instrum. Meth. A 349 (1994)

[51] *GEANT: Detector Description and Simulation Tool*, R. Brun et al., Technical Report No. CERN-W5013 (1994), https://cds.cern.ch/record/1082634 (accessed June 2020)

[52] *Measurements of Cross Sections and Charged Pion Spectra in Proton-Carbon Interactions at 31 GeV/c* NA61/SHINE Coll., Phys.Rev. C84 (2011) 034604

[53] *Results from the HARP experiment*, M. G. Catanesi et al., Nuc. Phys. B 168:1 (2007)

[54] *The T2K ND280 Off-Axis Pi-Zero Detector*, S. Assylbekov et al., Nucl. Instrum. Meth. A 686 (2012)

[55] *The T2K Side Muon Range Detector*, S. Aoki et al., Nucl. Instrum. Meth. A 698 (2013)

[56] *First Proton Anti-proton Collisions in the CERN SPS Collider*, Staff of the CERN Proton-Antiproton Project, Phys. Lett. B 107:306-309 (1981)

[57] *The T2K Experiment*, K. Abe et al. (The T2K Collaboration), Nucl. Instrum. Meth. A 659, 106 (2011)

[58] *Time Projection Chambers for the T2K Near Detectors*, N. Abgrall et al. (The T2K ND280 TPC Collaboration), Nucl. Instrum. Meth. A 637 (2011)

[59] *The T2K fine-grained detectors*, P. A. Amaudruz et al. (T2K ND280 FGD Collaboration), Nucl. Instrum. Meth. A 696, 1 (2012)

[60] *Development of Multi-Pixel Photon Counters*, M. Yokoyama et al., arXiv:physics/0605241 [physics.ins-det]

[61] *The Electromagnetic Calorimeter for the T2K Near Detector ND280*, D. Allan et al., arXiv:1308.3445v2 [physics.ins-det]

[62] *Progress in Measurements of 0.110 GeV NeutrinoNucleus Scattering and Anticipated Results from Future Experiments*, K. Mahn, C. Marshall, and C. Wilkinson, Ann. Rev. Nucl. Part. Sci. 68.

[63] *MEASUREMENT OF THE CHARGED CURRENT MUON NEUTRINO DIFFERENTIAL CROSS SECTION ON SCINTILLATOR WITH ZERO PIONS IN THE FINAL STATE WITH THE T2K ON/OFF-AXIS NEAR DETECTORS*, A.B. Cudd, Michigan State University Doctoral Thesis (2020)

[64] *Introduction to Unfolding: A Statistician's Perspective* M. Kuusela, PhyStat-ν 2019, https://indico.cern.ch/event/735431/contributions/3275244/attachments/1784103/2904689/Ph
(accessed May 2021)

[65] *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, S. Wilks, The Annals of Mathematical Statistics 9: 6062 (1938)

[66] *Minuit 2*, F. James and M. Winkler, CERN (2018), https://root.cern.ch/guides/minuit2-manual

[67] *Study of the Tracker ECal systematics*, T2K Technical Note 279.

[68] *ECal* π^0 *veto studies*, T2K Technical Note 270.

[69] ν_μ *CC event selections in the ND280 tracker using Run 1+2+3+4+5*, T2K Technical Note 212.

[70] *HighLAND tutorial*, A. Izmaylov and A. Cervera, https://twiki.cern.ch/twiki/pub/Main/HighlandAnalysis/highland2_intro.pdf, June 2017 (accessed May 2021)

[71] *Pion-nucleus single-charge exchange above the resonance*, E. Oset and D. Strottman, Phys. Rev. C. 42, 6 p. 2454-2462 (1990)

[72] *Characterisation of nuclear effects in muon-neutrino scattering on hydrocarbon with a measurement of final-state kinematics and correlations in charged-current pionless interactions at T2K*, K. Abe et al., Phys. Rev. D 98 032003 (2018)