Constraining the T2K neutrino oscillation parameter results using data from the off-axis near detector, ND280

Implementation of a nucleon removal energy systematic uncertainty treatment in the BANFF fit

Joseph G. Walsh

Supervised by

Dr Laura L. Kormos

This thesis is submitted for the degree of Doctor of Philosophy

Physics Department

Lancaster University

April 19, 2022

Abstract

Presented in this thesis are the results of the BANFF near-detector fit as part of the T2K 2020 neutrino oscillation parameter constraint, with a focus on the implementation of four nucleon-removal energy parameters, ΔE_{rmv}. These parameters correspond to the systematic uncertainty associated with the energetic cost, E_{rmv}, of liberating a bound nucleon from the ground state of a nucleus in quasielastic neutrino scattering. Previously the dominant source of systematic uncertainty on the extraction of the neutrino mass splitting term Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO), an update of the nuclear model used for CCQE interactions at T2K from a relativistic Fermi gas model to a spectral function model and a new treatment of the systematic uncertainty on E_{rmv} has allowed the total bias on Δm^2_{32} to be reduced by a factor of 2.8, and does not impact T2K's ability to exclude leptonic CP-conservation.

The fit to the ND280 data is an essential stage of the extraction of the PMNS mixing parameters from T2K's data in which the beam and interaction cross-section models common to ND280 and Super-Kamiokande are constrained by sampling the unoscillated beam. The ND280 data are shown to be consistent with the T2K model, reporting a p-value of $p = 0.74$, an improvement on the previous ND280 fit $p = 0.5$. A study of the postfit model shows an improved consistency with a p-value of $p = 0.82$.

The impact of propagating biases in the fits to the ND280 data to the fits to Super-Kamiokande data on the constraints on Δm^2_{32} and δ_{CP} are investigated, shown to be small, and covered by an additional uncertainty term in the likelihood driven by fits to alternative models. The overall contribution of the E_{rmv} systematic uncertainty to the total variance on Δm^2_{32} using the 2020 (2018) T2K implementation was estimated to be $\sigma^2_{E_{\text{rmv}}}/\sigma^2_{\Delta m^2} = 1(6)\%$.

Contents

List of Tables

List of Figures

Acknowledgements

It would not be possible, within any reasonable number of pages, to pay proper thanks to all of the people I wish to thank for making my last four years as memorable and enjoyable as they have been, even in spite of a global pandemic; however, I must take the time to say a special thank you to the people who have been most supportive, both professionally and personally, throughout my PhD.

First and foremost, I would like to thank my supervisor, Laura, without whom I would not have even begun this PhD never mind finished this thesis! Secondly, I must thank my masters supervisor, Helen, for putting me on the neutrino physics track, and the rest of The T2K Lancaster group who have all been a pleasure to work with.

The T2K collaboration is full of wonderful people, but I would especially like to thank those in the BANFF and OA groups, Mark, Clarence, Will and especially Laura, who I shared the delights of running the near detector fits. I wish Ciro, Tristan, Kamil and Nauman, and all of the other new analysers the best on the current analysis. The T2K analysis would be nothing without the fits to the far detector too, and so I would like to thank Lukas, Ed, Kevin, Asher, Patrick, and Cristophe. Much of the work in this thesis is built of the back of the tireless effort of the NIWG, and would not be possible without Luke and Stephen, who sometimes seem to keep the whole of T2K going.

I would like to express my gratitude to Kendall, for offering me a job when I finish my PhD, which will allow me to continue to work with all of the wonderful people on T2K whilst hopefully meeting a whole bunch of new fantastic friends and colleagues as start to work on DUNE.

Throughout the past four years I have always been glad to see the familiar faces of my friends from my undergraduate course who embarked on PhD programs in neutrino physics at the same time as me, Alex, Sam and Billy. I hope we can all meet in Japan again soon. Lancaster University has been my home for more than eight years now, and as I come to leave, I know Lancaster will always feel like home to me. I would not be the person I am today without all of the wonderful people I have met here, especially in the campus bars.

I would like to thank my parents, Gerard and Jill, who always encouraged me to pursue my interests, be that physics or music or anything else, and have always made sure I had the freedom and support to do so.

Finally, I would like to thank my friends, though there seem too many to count, but especially Finn and Reuben, without whom Lancaster would have been a very different place. I wish them the best as we all leave Lancaster this year for new chapters of our lives.

Author's declaration

The work presented in this thesis is either my own, or the result of collaborative research on the T2K experiment. The focus of this work, the implementation of, fitting of, and validation of the nucleon removal energy systematic uncertainties, was conducted by myself with the exception of model development, the fake data studies, and propagation of the ND280-fit results to the Super-Kamiokande fits, which were conducted by my co-analysers. This work has not been submitted in substantially the same form or for the award of a higher degree. The work presented is currently under collaboration review, prior to submission to a journal for peer review, but has not yet been published. The results of this work have been presented publicly at conferences by myself and others.

As part of my work as a T2K collaborator, I have been involved in service work that accompanies my research. As an ND280 ECal detector expert, I spent several months maintaining and calibrating the electronics and cooling systems of the ECal, PØD and SMRD subdetectors of the ND280 which are discussed in chapter 4. During the 2018 refurbishment of the Super-Kamiokande detector, I spent two weeks working primarily on cleaning, rust-removal and Tyvek replacement in the Super-Kamiokande outer detector. The details of this are discussed in subsection 3.5.4. For the four years of my PhD, I have been jointly responsible for ensuring data from ND280 is transferred properly to the CERN LHC computing GRID for use by the whole collaboration, as discussed in subsection 4.9. In my first year, I validated the more sophisticated magnetic field simulation for particle tracks which pass through the solid iron flux-return yoke of the ND280 described in subsection 4.0.1 As a BANFF analyser, I was responsible for running the Asimov and data fits for the T2K 2020 oscillation analysis in the frequentist sequential fitting stream. This also involved running fits to large ensembles of random simulated data sets for both the bias studies and the p-value calculations. For the nucleon removal energy, I helped develop the underlying kinematic varying parameter used by all of the T2K fitters at both ND280 and Super-Kamiokande, and generated the simulated data from which it calculated the momentum shifts. I then implemented a new style of uncertainty treatment in the BANFF fitter to allow for the kinematic variations produced by the underlying parameter to be approximated by interpolated response splines for each bin in the ND280 samples. I then implemented the fake data set for the robustness studies conducted by my co-analyser. The impact of these fit results on the far detector fits were prepared by co-analysers in the relevant fitting groups. Additional description of as yet unfinished work that I have begun on the adaption of this method to allow for multi-parameter dependence is presented in chapter 10, as are initial sensitivity studies of the impact of pion kinematic information in the ND280 fit.

Chapter 1

Introduction

Over the past half century, the predictive success of the standard model (SM) [1][2][3] has led to the discovery of new particles, from the W and Z bosons in 1983 to more recent discoveries such as the Higgs boson in 2012 [4][5][6]. Using the groups and symmetries in quantum field theories to predict the existence and properties of particles from those which are already known, the SM is powerful, though to some, unsatisfyingly incomplete. The SM, which is described in section 2.1, is far from a theory of everything and fails to unify the three fundamental forces which govern the interactions it describes. There is currently no explanation for the vast differences in strengths of these forces or the range of mass scales of the SM particles. Attempts to find more fundamental frameworks which would make testable predictions of Beyond Standard Model (BSM) physics have not been fruitful.

Experimental searches have begun to find evidence for BSM physics, such as the tests of lepton universality in b quark decays at LHCb (3.1 σ) [7], or recent measurements of the muon magnetic moment by g-2 (4.2σ) [8]. Arguably the most successful field of particle physics to return BSM results is that of neutrino oscillation experiments. As neutrinos were assumed to be massless, the discovery of mass-squared-difference dependent oscillations in 1998 [9] surprised physicists, as the SM does not have a mechanism by which neutrinos can gain mass as discussed in chapter 2. Originally proposed as a simple energy-conservation book-keeping trick in radioactive decays [10], neutrinos may provide an avenue to understand the fundamental differences between matter and antimatter, as well as clues to extensions of the SM from which a more fundamental theory can be inferred. Neutrino mixing is discussed in detail in section 2.2.

In the SM there are conservation laws relating to the symmetries in the model; each of these conservation laws has an associated "good" quantum number which is conserved in interactions which obey these symmetries. The charge conjugation number, C, and the parity number, P, were assumed to be conserved until P was found to be maximally violated in weak interactions

in 1956 [11]. It was then assumed that their product, CP, was still conserved until it was found to be violated in the decays of neutral mesons indicating a difference between the behaviour of the meson and antimeson components of these systems in 1964 [12]. It is currently assumed that this CP violation corresponds to a breaking of the time, T, symmetry in particle interactions, and that the product of all three numbers, CPT, is still conserved, although experimental searches for CPT-violation are ongoing. The mechanism for neutrino oscillations set forth by Pontecorvo, Maki, Nakagawa and Sakata (PMNS), discussed in section 2.2, allows for CP-violation. This is the only place in which this can occur for fundamental particles, as opposed to the composite quark systems of mesons. A discussion of the field as a whole, and the global best fit values of the PMNS parameters is given in section 2.5.

The Tokai to Kamioka (T2K) experiment is a long-baseline neutrino-oscillation experiment which measures the rates of neutrino oscillations in a muon neutrino or muon antineutrino beam. T2K uses a near detector, "ND280", 280 m downstream of the beam production point to sample and characterise the beam before the neutrinos travel 295 km to the Super-Kamiokande (SK) detector where the neutrino components of the beam are measured again. T2K aims to place constraints on the accelerator terms in the mixing matrix of the PMNS mechanism [13][14], with sensitivity to the mixing angles between first and third, θ_{13}, and second and third, θ_{23}, mass states, the largest difference between the squares of their masses, Δm^2, and the extent of leptonic-CP violation, encapsulated in a CP-violating phase, δ_{CP} [15]. The sign of the largest Δm^2 is currently unknown, but will tell us the ordering of the neutrino masses which have yet to be measured directly. An overview of the T2K experiment is presented in chapter 3.

Presented in this thesis are contributions to the current world-leading constraints on leptonic CP-violation, and the accelerator terms of the PMNS matrix [16]. The overall analysis framework is presented in chapter 5. Specifically, this thesis describes the procedure, validation, and results of the fit to T2K's off-axis near detector ND280, in chapters 6, 7, and 8, respectively. The ND280 itself is discussed in detail in chapter 4, and the impact of this near detector fit on the oscillation parameter extraction is explored in chapter 9.

T2K builds on the work of experiments such as K2K [17] and, complements the Nobel prize-winning measurements of solar and atmospheric neutrinos by Super-K [9]. A collaboration of around 500 members from 70 institutions across 12 countries [18], the T2K experiment is the result of years of work from its conception in 1999 to the beginning of data collection in 2009 to the decade of data-taking so far. Significant ongoing work at T2K, as outlined in chapter 10, will allow for improved analysis of existing data as well as measuring new physics as detectors are upgraded or added to the experiment. This will allow T2K to take full advantage of its own data,

whilst paving the way for future experiments such as Hyper-Kamiokande and DUNE which will be able to make more precise measurements of the oscillation experiments.

1.1 Implications of CP-violation in cosmology

All of the material seen in the universe today which makes up the stars, planets and galaxies is made of matter, with very little antimatter to be found in the universe. Given that CP-symmetry is unbroken in most interactions, it is not unreasonable to assume that the universe may have begun in a symmetric state with equal quantities of matter and antimatter. If this is the case, why is there more matter in the universe than antimatter today? In 1966, Andrei Sakharov put forward a set of three conditions under which baryogenesis could occur, breaking any matter-antimatter symmetry [19]:

1. Matter must exit thermal equilibrium with the hot plasma which existed at the beginning of the universe
2. Both C and CP conservation must be violated
3. Baryon number must not be conserved

The CP-violation measured in the quark sector is not sufficient to satisfy the second condition on its own; however, if neutrinos can violate CP-conservation, and the violation is sufficiently large, this may go some way to satisfying the condition. Minimal extensions of the SM which require massive neutrinos with CP-violating oscillations are prerequisites for models such as leptogenesis, which may explain the mater antimatter asymmetry seen today.

1.2 Neutrino interactions and nuclear models

Much of the study of neutrino interactions in matter relies on an understanding of the material it interacts with at a nuclear level. Comprised of many particles and fields, the nucleus is a challenging environment in which to predict interactions of neutrinos as well as the behaviour of the products of these interactions. Therefore, the ability to accurately reconstruct information about the neutrino from the products of these interactions requires an accurate description of the nuclear potential. Until experimental methods improve enough to be able to statistically separate out events on the single proton nucleus of a hydrogen atom, this will be a sizeable contribution to the systematic uncertainty of any neutrino experiment. In the case of T2K, the choice of nuclear model has a large impact on the rate of interactions and the kinematics of the outgoing particles. A large

source of systematic error in previous T2K analyses [20], the nucleon removal energy uncertainty impacts the reconstruction of the neutrino energy which is dependent on an analogous quantity, and so if inaccurately determined, can bias the oscillation parameters. The largest impact of this bias in particular is on the value of the mass splitting term as this is extracted from the position of the oscillation maxima in the neutrino energy spectrum

Chapter 2

Theory

Neutrinos are the most abundant known massive particle in the universe, and yet their extraordinarily small mass and lack of electric charge make them incredibly difficult to study. As neutrinos couple only to the weak force and with very small interaction cross sections, neutrino detectors must either have a large mass or run for a very long time to provide sufficient opportunity to collect a statistically large enough data set to study neutrino properties. This requirement for very high-mass detectors complicates the study of neutrinos further by obscuring the interactions of neutrinos with the complicated effects of the nuclear potential. The ability to accurately predict the behaviour of neutrinos is not just dependent on an understanding of the standard quantum field theories, which are commonplace in the field of particle physics, but also on an understanding of nuclear theory which comes with its own challenges.

2.1 The standard model of particle physics

The current accepted theoretical framework used to describe particle physics is the "Standard Model" (SM) [1][2][3][21], that aims to describe all of the fundamental building blocks of the universe as well as the interactions between them. Although the SM is considered to be an incomplete theory, it has made many successful predictions such as the existence of the weak W and Z bosons which were discovered by the UA1 and UA2 experiments at the Super Proton Synchrotron in 1983 [4][5][22][23], and the Higgs boson which was confirmed by the ATLAS and CMS experiments in 2012 [6][24].

Particles in the SM are categorised into two groups based on their spin type and corresponding statistics as shown in figure 2.1. Integer-spin particles that obey Bose-Einstein statistics are bosons, and half-integer-spin particles that obey Fermi-Dirac statistics are fermions. Bosons are the exchange particles which mediate the fundamental forces or interactions between particles,

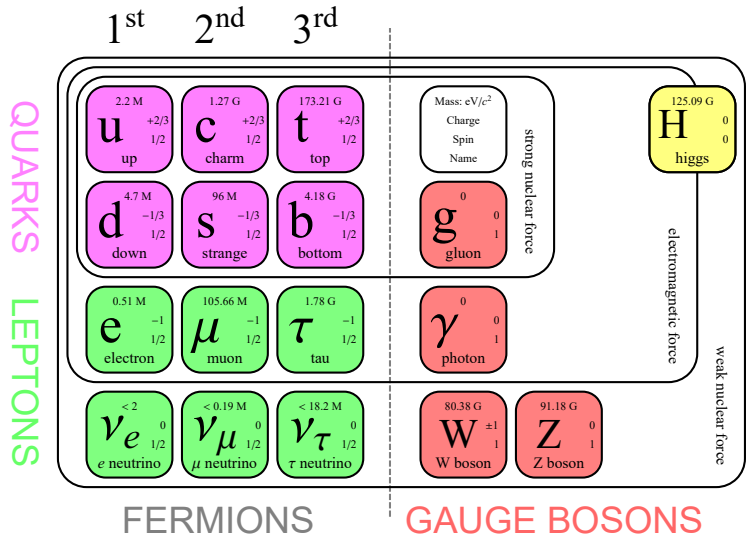

Figure 2.1: Particles in the SM and their properties, categorised by their interactions, flavour, charge and spin. The spin-1/2 fermions come in three generations which are ordered by mass (except for the neutrinos). The spin 1 gauge bosons mediate the electromagnetic, weak and strong forces, and the spin 0 scalar Higgs boson mediates the coupling of the massive particles with the Higgs field [25].

and fermions are the fundamental building blocks of matter that, when bound by bosons, form nuclei, atoms, and molecules. Fermions can be split into two categories, quarks and leptons, which are further subdivided into three generations or families each. Each generation of quarks has a quark with a +2/3 electric charge[1] and a quark with a -1/3 electric charge, and each family of leptons contains a negative (-1) charged lepton and a neutral lepton, or neutrino. The positively-charged quarks are, in the generations ordered by increasing mass: up, charm and top (u, c, t); the negatively-charged quarks are: down, strange and bottom[2] (d, s, b). The charged leptons are the electron, muon and tau, (e,μ,τ) and the neutral leptons are their associated neutrinos, the electron neutrino, muon neutrino and tau neutrino (ν_e,ν_μ, ν_τ) respectively. Each of the particles in the SM has a corresponding antiparticle, such as the antiup quark $\bar{u}$, with opposite charge and parity, though for neutral particles such as the photon this may be itself.

There are three known forces (other than gravity which is not described in the SM) which act on fundamental particles as described by the SM. Each of these forces corresponds to an interaction which conserves specific properties, or quantum numbers, between initial and final states. These forces are discussed in descending order of strength here. The strong force, mediated by gluons, couples to colour charge. Colour is a property of the six quarks and the gluons and

[1] One unit of electric charge is defined as equal to the charge of a single positron.

[2] The top and bottom quarks are sometimes referred to as the truth and beauty quarks.

comes in three types called red, blue and green. Colour can also come in the anti-colour forms anti-red, anti-blue and anti-green. The strong force requires all observable particles to be in colour-neutral states, either in configurations of trios (r,g,b) or pairs ((r, $\bar{r}$), (g, $\bar{g}$) or (b, $\bar{b}$)) of quarks. The strong interaction has a short range on the order of femtometers and increases in strength with distance. As such, if quarks are separated from each other, the energy stored in the strong potential increases; this happens until there is enough energy stored in the potential to pair-produce new quarks resulting in two (or more) colour-neutral systems and so bare quarks are not observed. These colour-neutral composite particles, or hadrons, can be classified into baryons (three valence quarks) and mesons (two valence quarks), and must have an integer sum electric charge.

Though the strong force within hadrons such as protons and neutrons is mediated by gluons, mesons, as integer-spin particles, can mediate the force in larger systems such as nuclei. Though baryons and mesons are the most simple and common possible configurations of quarks, more complex systems such as tetraquarks and pentaquarks are not forbidden by the SM. Several pentaquark systems were recently discovered by the LHCb experiment in 2015 and 2019 [26][27].

The electromagnetic force, mediated by the photon, couples to electric charge and has the furthest range of any of the forces in the SM. The electromagnetic force binds electrons to nuclei to form atoms and the complex electronic configurations which give rise to chemistry and macroscopic structure. It is the electromagnetic interaction which is responsible for most of the typical forces people think of in their day-to-day lives as well as the electromagnetic spectrum from radio waves, to visible light, to gamma rays.

The weak force, mediated by the $W^{\pm}$ and Z^0 bosons, couples to particles with weak isospin, T, or more precisely, the third component of the weak isospin, T_3. The weakest force aside from gravity, the weak force is responsible for many of the radioactive decays of unstable nuclei. Able to couple between quarks and leptons, the weak interaction allows for unstable nuclear states to decay semileptonically into more energetically favourable states by converting a proton into a neutron or vice versa, emitting a charged lepton and neutral neutrino. Each of the fundamental fermions carries a flavour, though the flavour states of the particles are not always the same as the mass states of the particles. The relationship between the mass and weak eigenstates of the quarks and leptons is described by the CKM and PMNS matrices, respectively. Two types of weak interaction can occur, charged-current interactions in which the $W^{\pm}$ boson exchanges electric charge between two fermions, or neutral-current interactions in which only energy and momentum are transferred via the Z^0 boson. Since the electric charge is a quantum number and a property of a particle, the charged-current interaction mediates a flavour change. This may be from up to down, or from electron to electron neutrino as shown in figure 2.2. The weak interaction does

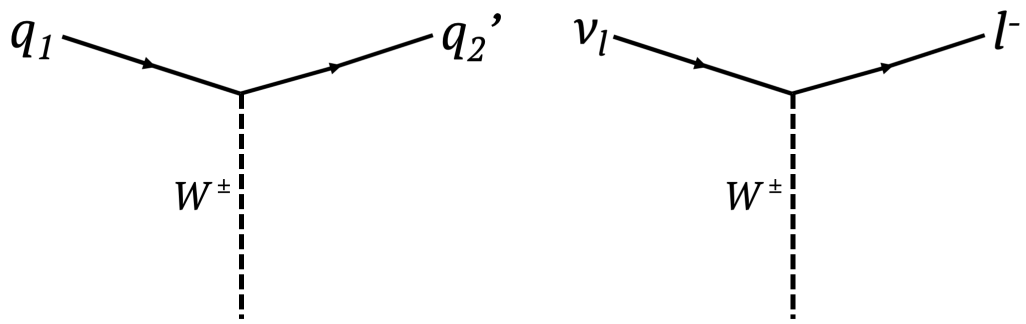

(a) Weak W quark coupling (b) Weak W lepton coupling

Figure 2.2: Interaction vertices of the $W^\pm$ boson in CC weak interaction with quarks and leptons. Here, q_1 and q_2', represent the two weak eigenstates of a quark family e.g. (u,d'), and the primed quark is a linear admixture of the three $-1/3$rd charged quarks. For leptonic weak interactions, the weak boson always interacts with a charged lepton and a neutrino of the same lepton flavour at the same vertex e.g. (e^-,ν_e).

conserve lepton flavour (electron, muon, tau), and so an initial state with a muon which exchanges its charge via a $W^\pm$ boson will result in a muon neutrino in the final state, thus producing a muon flavour lepton. This property of the weak interaction is key to the study of neutrinos, as it allows the flavour of the neutrino in an interaction to be inferred from the measured flavour of the charged lepton. At high enough energies, the electromagnetic and weak forces can be shown to be unified into one electroweak theory. The electroweak force couples to the weak hypercharge, Y_W, which is a linear combination of the electric charge and weak isospin. This unification has driven searches for predicted outcomes of theories which attempt to further unify the electroweak and the strong interaction in one Grand Unified Theory (GUT) which may explain currently unexplained phenomena in the SM.

Finally, though not a mediator of a force in the same way as the other bosons, the Higgs boson H^0 couples via the Higgs mechanism to any fundamental particle in the SM with mass. As the Higgs is a spin 0 particle, interactions between the Higgs and a massive particle must flip the helicity, or handedness, of a particle to conserve angular momentum. Helicity is the alignment of the particle's spin and its momentum, and mass can be considered to be the rate of helicity flipping. The mass acquired through the Higgs mechanism is known as their Dirac mass. If the spin and momentum of a particle are aligned then it is said to have left-handed helicity; if the spin and momentum are antialigned then it is said to have right-handed helicity. As the helicity is constructed from the momentum of the particle, a boost to a reference frame in which the particle's direction has flipped will also flip the helicity. This is not possible for massless particles.

In the massless limit, in which the helicity cannot be flipped, this is equivalent to the chirality of the particle. Chirality is an intrinsic property which determines how the wave functions describing the particles in behave under Lorentz transformations. Only left handed neutrinos and right handed antineutrinos have been observed in nature, and so in the SM the neutrino is considered to be massless [28]. However, as shall be discussed in section 2.2, recent experiments such as Super-Kamiokande and SNO have shown that neutrinos cannot be massless [9][29] and so this is a key failing of the SM, which does not explain if neutrinos have non-Dirac masses, i.e. Majorana masses, or if they do have Dirac masses, why they are so small that we have not observed the right-handed helicity neutrinos. Whilst the Higgs mechanism explains how each of the fundamental particles acquires mass, in addition to the Dirac mass of valence quarks, composite systems such as baryons have mass associated with the bound system and spontaneously-produced virtual quark-antiquark pairs known as "sea quarks". The majority of the mass of the proton, and other baryons, is associated with these and does not come from the three valence quarks and their interaction with the Higgs.

2.1.1 Parity symmetry, and neutrino handedness

Parity, P, in quantum physics is a quantum number which describes the behaviour of an object undergoing a parity transformation in which its spacial coordinates are flipped such that

$$x \to -x, \qquad y \to -y \qquad \text{and} \qquad z \to -z. \tag{2.1}$$

A parity transformation is essentially the reflection of an object, and for asymmetric objects the parity-transformed object can not be remapped by translation and rotation alone onto the original object, and an object cannot be mapped by translation and rotation onto its mirrored object as shown in figure 2.3. For example, a left hand cannot be translated into the position of a right hand so that they overlap completely. Importantly, axial vector or pseudovector properties such as angular momentum do not change under parity transformation.

For fundamental point-like particles, there is no spacial extent for the parity transformation to impact; however, the parity transformation does impact the particles' momentum. A parity-conserving interaction will therefore change the direction of travel of a particle when the coordinate basis of the interaction is reversed. The spin of a particle, much like angular momentum, is an axial vector and so does not change under a parity transformation, and so the spin and the momentum of a particle should be uncorrelated in any interaction which obeys parity symmetry. This is called P-conservation.

In 1956, Chien-Shiung Wu [11] tested P-conservation in weak interactions by using a magnetic field to align cobalt-60 atoms and polarise the decays of the atoms as shown in figure 2.4. Cobalt-

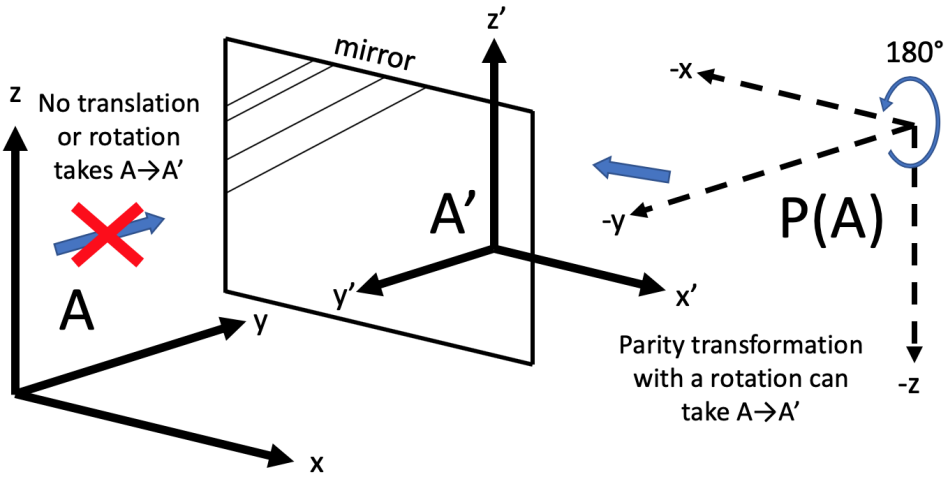

Figure 2.3: No combination of translations and rotations can map an object onto its mirrored object, but a parity transformed object can be. If the each of the coordinates x, y, z, of the basis A are inverted under parity transformation to give $P(A)$ with coordinates $-x, -y, -z$, the basis can then be mapped onto the mirror A' with $x' = x,\ y = -y',\ z' = z$ by a rotation of $180°$.

60 is a beta emitter and undergoes the decay

$$^{60}\text{Co} \rightarrow {}^{60}\text{Ni} + e^- + \bar{\nu}_e + 2\gamma \tag{2.2}$$

to Nickel-60, releasing an electron and an antineutrino. If parity is conserved the relative numbers of electrons and photons emitted in a given direction should remain roughly constant under a reversal of the direction of the magnetic field. This ratio is used because the emission of photons should be isotropic and so deviations in this isotropy from misalignment of the cobalt nuclei can be accounted for.

Wu observed that the direction of electrons emitted in the decays were far more anisotropic than that of the photons and highly favoured decays in the opposite direction to that of the nuclear spin. This result suggests that not only does the weak force prefer a handedness, it does so maximally, only coupling to left-handed particles and right-handed antiparticles.

2.1.2 Antimatter and CP symmetry

Each of the fundamental particles in the SM has an antimatter counterpart. The properties of antiparticles are the same as for the particle except with all of their charges flipped, i.e. positive to negative and vice versa. As such they have the opposite sign charge and have their physical orientation mirrored. Neutral particles such as the Z^0 or neutral femions may be their own antiparticle, making them a "Majorana" particles. There are known Majorana fermions which are not funda-

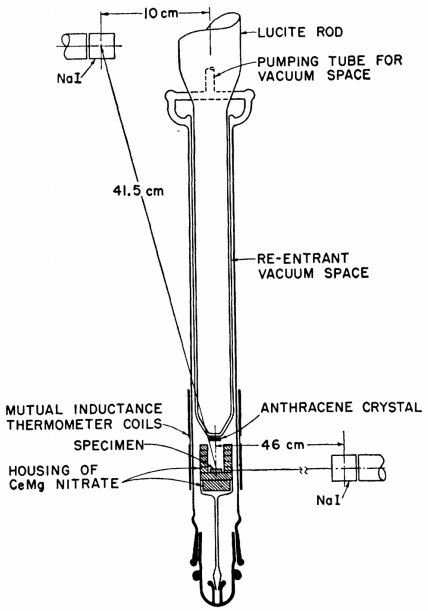

Figure 2.4: Experimental apparatus used in by C.S. Wu to demonstrate parity violation in the decay of ^{60}Co nuclei. A radioactive ^{60}Co source was cooled in a vacuum chamber to near absolute zero within a solenoid magnet coil. The uniform magnetic field of the coil was used to align the spins of the ^{60}Co atoms, polarising their decays. During the decay, neutron emits a beta particle producing a proton in what is now a metastable ^{60}Ni nucleus. This metastable nucleus then decays to the ground state by emitting a pair of back-to-back photons. An anthracene crystal was used to detect the beta particles emitted in the direction of the magnetic field alignment, the light from which is directed into a photomultiplier tube by a lucite rod. Two NaI scintillator detectors were also used to measure the anisotropy of the photons emitted for comparison [11].

mental, such as neutral mesons made of quark anti-quark pairs such as π^0 and $\bar{\pi}^0$. It is unknown if there are any fundamental Majorana fermions. As the only neutral fundamental fermions, neutrinos may be Majorana particles and this is currently being investigated by multiple experiments. It was assumed that matter and antimatter symmetry, or CP-symmetry, was a fundamental symmetry of nature and the quantum number CP was conserved in all interactions; however, in 1964 it was demonstrated that CP-symmetry was violated in the decays of the neutral kaon system.

The neutral kaon K^0 is a weak eigenstate in which the kaon is produced and interacts in weak interactions. K^0, however, is not a mass eigenstate, and so with no well-defined mass the kaon does not propagate through space in this state. Instead, the weak kaon state is a quantum superposition of two mass eigenstates with different mean lifetimes, K-short K_S^0 and K-long K_L^0. The different decay lengths of these mass eigenstates allowed for a statistical separation of the different components of the K^0 system. In 1955 Murray Gell-Mann and Abraham Pais proposed a mechanism for neutral meson oscillations [30] later confirmed in the K^0 system by Christenson, Cronin, Fitch and Turlay [12]. Gell-mann and Pais hypothesised that a θ^0 massive boson and its distinct (non-Majorana) $\bar{\theta}^0$ antiparticle, are linear admixtures of true "particles" θ_1 and θ_2 with well-defined masses and lifetimes. In the context of the neutral kaon system this is understood to mean that we can express these massive particle states as

$$|K_1^0\rangle = \frac{|K^0\rangle + |\bar{K}^0\rangle}{\sqrt{2}} \quad \text{and} \quad |K_2^0\rangle = \frac{|K^0\rangle - |\bar{K}^0\rangle}{\sqrt{2}}. \tag{2.3}$$

The K_L^0 component of the K^0 state, which had $CP = -1$, was observed decaying to two pions giving a final state with $CP = +1$, violating CP-symmetry. This can be explained by the GIM mechanism, put forward by Glashow, Iliopoulos and Maiani [31], if K_L^0 and K_S^0 are weak eigenstates but not CP-eigenstates, and are related to the mass eigenstates by the relations

$$|K_L^0\rangle = |K_2^0\rangle + \varepsilon\,|K_1^0\rangle \quad \text{and} \quad |K_S^0\rangle = |K_1^0\rangle - \varepsilon\,|K_2^0\rangle \tag{2.4}$$

where ε is the degree of CP-violation in the system.

2.2 Neutrino oscillations

In 1957 Bruno Pontecorvo proposed that, similarly to the neutral kaon system, as neutral particles with distinct antiparticles, neutrinos would be able to undergo mixing [32][13] and we could express the ν and $\bar{\nu}$ states of the neutrino as a linear admixture of defined neutrino mass states ν_1 and ν_2

$$|\nu\rangle = \frac{|\nu_1\rangle + |\nu_2\rangle}{\sqrt{2}} \quad \text{and} \quad |\bar{\nu}\rangle = \frac{|\nu_1\rangle - |\nu_2\rangle}{\sqrt{2}}. \tag{2.5}$$

Upon the discovery of multiple neutrino types and with the assumption of the conservation of neutrino handedness forbidding $\nu \leftrightharpoons \bar{\nu}$, Ziro Maki, Masami Nakagawa, and Shoichi Sakata [14] formulated a description for two-flavour neutrino mixing with one mixing angle θ similar to the Cabibbo angle θ_c in the quark sector [33]. The neutrino flavour eigenstates are expressed as mixtures of the neutrino mass eigenstates as

$$\nu_e = \nu_1 \cos\theta - \nu_2 \cos\theta \tag{2.6}$$

$$\nu_\mu = \nu_1 \cos\theta + \nu_2 \cos\theta \tag{2.7}$$

which can also be expressed in matrix form

$$|\nu_\alpha\rangle = U\,|\nu_i\rangle = \begin{bmatrix} U_{\alpha i} & U_{\alpha j} \\ U_{\beta i} & U_{\beta j} \end{bmatrix} |\nu_i\rangle \tag{2.8}$$

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}\,|\nu_i\rangle \tag{2.9}$$

where $|\nu_\alpha\rangle$ represents the flavour eigenstates, $\alpha = e, \mu$ and $|\nu_i\rangle$ represents the mass eigenstates $i = 1, 2$. As with kaons, the time evolution of this system (in a vacuum) can be solved to calculate the probability of a neutrino interacting as a given flavour state after travelling for some distance. For vacuum oscillations we can solve the time-dependent Schödinger equation for a given mass eigenstate by setting the potential term to zero

$$i\frac{\partial}{\partial t} E\,|\nu_i(x,t)\rangle = -\frac{1}{2m_i}\frac{\partial^2}{\partial x^2}\,|\nu_i(x,t)\rangle \tag{2.10}$$

which has a plane-wave solution of the form

$$|\nu_i(x,t)\rangle = \exp(-i(Et - \vec{p}\cdot\vec{x}))\,|\nu_i(0,0)\rangle\,. \tag{2.11}$$

In the ultra-relativistic limit $p >> m$ where $\vec{p}$ and $\vec{x}$ are colinear, and setting $\hbar = c = 1$ yields,

$$E = \sqrt{p^2 + m^2} \to E \approx p + \frac{m^2}{2p} \approx E + \frac{m^2}{2E} \quad \text{and} \quad t \to L$$

as the time dependence is effectively a distance dependence with $L = x(t) - x(0)$ expressed in units of inverse eV. The distance evolution of the neutrino mass state is given by

$$|\nu_i(L)\rangle = \exp\left(-i\frac{m_i^2}{2E}L\right)|\nu_i(0)\rangle \tag{2.12}$$

and so using equation 2.9 we can find the dependence of the flavour state, $|\nu_\alpha\rangle$, on L:

$$|\nu_\alpha(L)\rangle = \sum_i U_{\alpha i} \exp\left(-i\frac{m_i^2}{2E}L\right)|\nu_i(0)\rangle\,. \tag{2.13}$$

For the initial state the exponential component goes to identity and so by inverting U we find the components of the initial flavour state that is in a given mass state to be

$$|\nu_i(0)\rangle = \sum_\alpha U^*_{\alpha i} |\nu_\alpha(0)\rangle \tag{2.14}$$

which makes the general final state $|\nu_\beta(L)\rangle$ be dependent on the initial state $|\nu_\alpha(0)\rangle$. By substituting expression 2.14 for the initial state in equation 2.13, and allowing the final state flavour β to differ from the initial state flavour α, yields

$$|\nu_\beta(L)\rangle = \sum_i U_{\beta i} \exp\left(-i\frac{m_i^2}{2E}L\right) \sum_\alpha U^*_{\alpha i} |\nu_\alpha(0)\rangle \tag{2.15}$$

and so taking the inner product of a final state $|\nu_\beta(L)\rangle$ with the initial state $|\nu_\alpha(0)\rangle$

$$\langle\nu_\beta(L)|\nu_\alpha(0)\rangle = \sum_\gamma \sum_i U_{\gamma i} \exp\left(i\frac{m_i^2}{2E}L\right) U^*_{\beta i} \langle\nu_\gamma(0)|\nu_\alpha(0)\rangle \tag{2.16}$$

where the other initial flavour components, γ, of $|\nu_\beta(L)\rangle$ can be summed over leaving only a summation over the mass states

$$\langle\nu_\beta(L)|\nu_\alpha(0)\rangle = \sum_i U_{\alpha i} \exp\left(i\frac{m_i^2}{2E}L\right) U^*_{\beta i} \tag{2.17}$$

to arrive at the probability of a neutrino produced in a weak interaction with flavour α being subsequently detected in another weak interaction as flavour β

$$P(\nu_\alpha \to \nu_\beta) = \left| \sum_i U_{\alpha i} U^*_{\beta i} \exp\left(i\frac{m_i^2}{2E}L\right) \right|^2 \tag{2.18}$$

which when expanded out is

$$P(\nu_\alpha \to \nu_\beta) = \sum_i U^*_{\alpha i} U_{\beta i} \sum_j U_{\alpha j} U^*_{\beta j} \exp\left(+i\frac{m_i^2}{2E}L\right) \exp\left(-i\frac{m_j^2}{2E}L\right). \tag{2.19}$$

This allows us to see the dependence of oscillations on the mass of the neutrinos

$$P(\nu_\alpha \to \nu_\beta) = \sum_i \sum_j U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j} \exp\left(+i\left(\frac{\Delta m_{ij}^2}{2E}\right)L\right) \tag{2.20}$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ is the difference between the square of the masses of the mass states. Expressing this in real and imaginary components we can see the oscillatory nature of the probability

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) =& \delta_{\alpha\beta} - 4\sum_{i>j} \mathrm{Re}\left(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}\right) \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \\ &+ 2\sum_{i<j} \mathrm{Im}\left(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}\right) \sin\left(\frac{\Delta m_{ij}^2 L}{2E}\right). \end{aligned} \tag{2.21}$$

This yields the two-flavour oscillation probability of result of

$$P(\nu_\alpha \to \nu_\beta) = \sin^2 2\theta \sin^2\left(\frac{1.27\Delta m_{ij}^2 L}{2E}\right) \tag{2.22}$$

where θ is the mixing angle between the two mass states and $\alpha \neq \beta$. A constant of 1.27 is used to convert from natural units, where $\hbar = c = 1$, to SI units so L can be given in km which is more convenient for the construction of experiments. The disappearance probability can be found from subtracting the appearance probability from unity

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2 2\theta \sin^2 \left(\frac{1.27\Delta m^2_{ij} L}{2E} \right) \tag{2.23}$$

The two flavour formalism was later extended to the three flavour formalism by Pontecorvo [13] to include the tau neutrino. In the three-flavour case the PMNS mixing matrix U can be expressed as the product of three rotation matrices

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.24}$$

with three rotation angles θ_{ij} describing the mixing between the three mass states $i, j = 1, 2, 3$, and a CP-violating phase, δ_{CP}. For a beam of pure muon neutrinos in the three-flavour case, the survival probability of a muon neutrino is then

$$P(\nu_\mu \to \nu_\mu) \simeq 1 - \cos^4(\theta_{13}) \sin^2(2\theta_{23}) \sin^2 \left(1.27\Delta m^2_{32} \frac{L}{E_\nu} \right) \tag{2.25}$$

and the appearance probability of an electron neutrino is

$$\begin{aligned} P(\nu_\mu \to \nu_e) \simeq & \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2 \left(1.27\Delta m^2_{32} \frac{L}{E_\nu} \right) \\ & \mp 1.27\Delta m^2_{32} \frac{L}{E_\nu} 8 J_{CP} \sin^2 \left(1.27\Delta m^2_{32} \frac{L}{E_\nu} \right) \end{aligned} \tag{2.26}$$

where J_{CP} is the Jarlskogg invariant which encapsulates CP violation as

$$J_{CP} = \cos\theta_{12} \cos^2\theta_{13} \cos\theta_{23} \sin\theta_{12} \sin^2\theta_{13} \sin\theta_{23} \sin\delta_{CP} \tag{2.27}$$

which is proportional to the degree of the CP-violation in the mixing probability via the term δ_{CP}. Equations 2.25 and 2.26 are valid in a vacuum, but matter effects relating to the electron density of matter must be taken into account for neutrinos travelling long distances through matter such as in measurements of atmospheric neutrinos which have passed through the earth. These formulae also assume that the mass eigenstates have the same momenta and energy, allowing the mass states to be modelled as plane waves.

The observation of neutrino mixing therefore proves that at least two of the neutrino mass states have non-degenerate masses, and provides opportunity to test if neutrino mixing violates CP-symmetry. The mechanism by which neutrinos acquire mass is as yet unknown, but may provide

insight to theorists as to how to extend the SM to a more complete picture of the fundamental building blocks of the universe.

The sign of the second term in the appearance probability in equation 2.26 is denoted by a $\mp$, which indicates the how the sign of this term differs for neutrinos and antineutrinos, and relates to the conjugation of the terms in the matrix when the formula is derived for the antineutrino states. Positive values of δ_{CP} will therefore result in a suppression of electron neutrino appearance and an enhancement of electron antineutrino appearance in a muon neutrino or antineutrino beam respectively, when compared to the CP-conserving case. These may be difficult to directly compare, however, if the differences in the neutrino and antineutrino interaction cross-sections in matter are not taken into account.

2.2.1 Neutrino mass ordering

Observation of neutrino oscillations confirm that at least two of the neutrino masses are nonzero and allow us to measure the difference of the squares of the masses, but they do not tell us the absolute masses of the neutrinos. They can tell us the ordering of the neutrino masses; however, to first order, the mass squared splitting term in the oscillation probability formula is encapsulated in the square of a sine function, for which there are multiple solutions.

The sign and value of the smaller mass squared splitting term Δm^2_{21} (sometimes called Δm^2_{sol}) has been determined to be positive by measuring solar-neutrino mixing and long-baseline reactor-neutrino mixing [29][34]. Matter effects in solar neutrino mixing enhance the bare mass squared splitting terms and allow for the sign of Δm^2_{21} to be determined more easily. From this the order of the first and second neutrino masses can be inferred to be $m_1 < m_2$.

The size of the splitting between the third mass state and the other mass states is sufficiently large that $|\Delta m^2_{32}| \approx |\Delta m^2_{31}|$ (sometimes called Δm^2_{atm}) and can be determined from atmospheric and accelerator experiments in which the unoscillated neutrinos are predominantly muon type. There are two solutions for the measurement of the largest mass squared splitting term to first order giving two possible mass orderings[3] as shown in figure 2.5: the "normal ordering" (NO) is defined as $m_1 < m_2 < m_3$; and the "inverted ordering" (IO) is defined as $m_3 < m_1 < m_2$. In this thesis the convention $\Delta m^2_{\text{atm}} = +\Delta m^2_{32} \approx -\Delta m^2_{31}$ will be used. The sign of this splitting term is also degenerate with the octant of θ_{23}, and the CP-violating phase δ_{CP}. Resolving the mass ordering will allow for a more precise measurement of θ_{23} and δ_{CP}.

[3] Mass ordering may sometimes be referred to as the "mass hierarchy", NO=NH and IO=IH

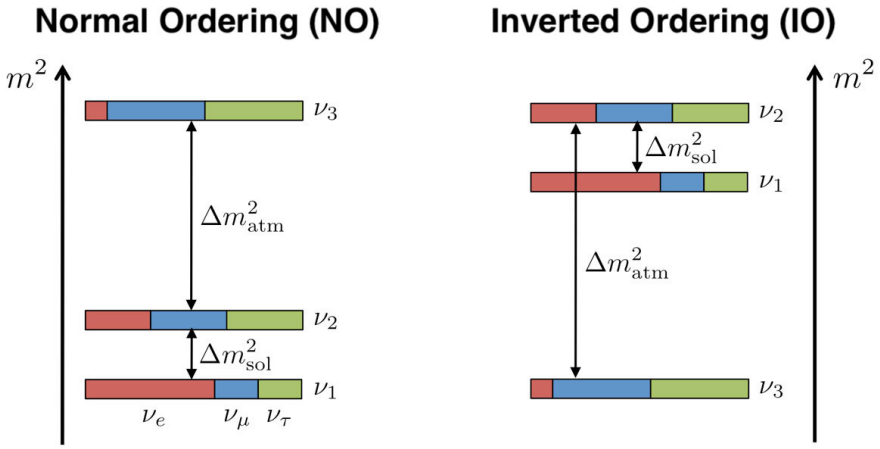

Figure 2.5: Possible neutrino mass orderings depending on the sign of the atmospheric mass squared splitting, Δm^2_{atm}. In the normal ordering ($m_1 < m_2 < m_3$) Δm^2_{atm} is referred to as Δm^2_{32} and in the inverted ordering ($m_3 < m_1 < m_2$) it is referred to as $-\Delta m^2_{32}$ which is approximately equal to Δm^2_{31} due to the small size of the solar mass squared splitting, Δm^2_{sol}. Δm^2_{sol} is equivalent to Δm^2_{21}, as we know the ordering of the first and second neutrino masses to be $m_1 < m_2$) [35].

2.3 Neutrino-nuclear interactions

Different neutrino interactions dominate at different neutrino energies and momentum transfers, from elastic or quasielastic scattering at lower neutrino energies, to Deep Inelastic Scattering (DIS) at higher neutrino energies. In elastic scattering only momentum is transferred to the nucleon or nucleus, and in quasielastic scattering a single charged lepton of the same flavour as the incident neutrino is produced. In DIS, the target nucleon is destroyed and the constituent quarks undergo hadronisation which may produce a large multiplicity of bound hadronic states. In addition to these, there are resonant processes in which the energy of the target nucleon is raised to an excited baryonic state, and coherent and diffractive processes, all of which can produce hadronic products. Single resonant pion production dominates at neutrino energies between the (quasi)elastic and DIS regimes. The initial interaction of the neutrino with a nucleus is known as the “hard scatter”. For (quasi)elastic and resonant interactions the hard scatter can be considered to take place on an unbound nucleon in the impulse approximation; however, as these processes take place within a nucleus, their products may undergo subsequent interactions within the nuclear potential, such as absorption, charge exchange, or further hadronic production. These additional processes are called Final-State Interactions (FSI). In the impulse approximation, there must be sufficient energy transferred to overcome the nucleon removal energy and free the nucleon from its nuclear ground state. The nuclear remnant is then modelled separately to the nucleon involved in the hard scatter.

This nuclear remnant may also be left in a metastable state which may decay by emitting photons.

2.3.1 Charged-and neutral-current interactions

Neutrino interactions are mediated by the charged and neutral weak bosons and so can be classified as charged- or neutral-current interactions. For charged-current (CC) interactions the neutrino must be energetic enough to produce the mass of the charged lepton in the final state, and so the cross-section is dependent on the flavour of the neutrino. For neutrino experiments such as T2K, in which knowledge of the flavour of the observed neutrino is essential, understanding these cross-section differences is of vital importance. Neutral-current (NC) interactions do not have the same final-state phase-space considerations as CC interactions but whilst they do not produce a flavour-taggable charged lepton, they do contribute to background processes. The different interaction channels are presented here roughly in order of their momentum transfer. Since T2K is primarily concerned with CC interactions, the focus of this section will be on CC interaction models, although there are NC equivalents to most of these interactions.

2.3.2 Quasi-elastic interactions

At neutrino energies below approximately $E_\nu = 1$ GeV, quasielastic (QE or CCQE) interactions dominate the total CC interaction cross section. CCQE events produce a single charged lepton which typically carries most of the energy of the incident neutrino, but produces no mesons in the final state. In the LLewelyn Smith formalism presented in [36], the QE interaction cross section is factorised into neutrino-nucleon interaction and the nuclear model according to the impulse approximation shown in figure 2.6. The neutrino interacts with a single nucleon which is ejected leaving one hole in the nuclear remnant, and so CCQE interactions described by this model are sometimes referred to as "one-proton one-hole" (1p1h) interactions. The strength and shape of the neutrino nucleon interaction is largely determined by vector and axial-vector nucleon form factors [37].

Whilst statistical correlations between nucleons can be largely ignored, a non-negligible number of multi-nucleon QE interactions with a mesonless final state are from neutrino interactions with bound pairs of nucleons, in "two-proton two-hole" (2p2h) interactions in which two nucleons are ejected from the nucleus. T2K models these interactions using the "Valencia" model described by Nieves et al. [38]. These interactions are topologically similar to 1p1h events but have different final-state kinematics and so will bias any QE estimation of the neutrino energy for the 1p1h model.

The equivalent NC process to CCQE is neutral current elastic interaction (NCE). This can be

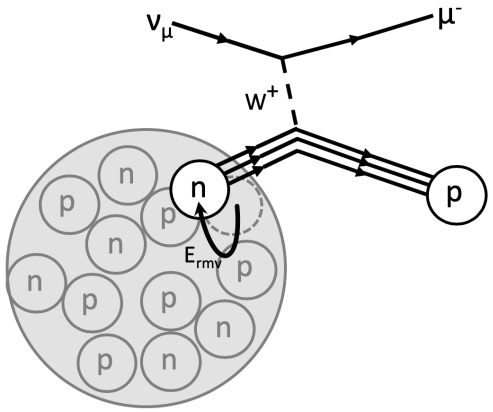

Figure 2.6: Diagram showing 1p1h quasielastic scattering in the impulse approximation the interaction of the incident neutrino on an individual nucleus, or the hard scatter, is treated as separate from the rest of the nucleus. An energy cost, E_{rmv}, is associate with liberating the bound nucleon from the nuclear ground state and its value effects the expected kinematics of the final state charged lepton. This is related to missing energy observed in electron scattering data. The free nucleon is ejected leaving a hole in the nuclear remnant.

modelled using the same formalism, but has different phase space considerations. The final-state lepton is a neutrino and so can be treated as massless, although momentum transfer can still eject a nucleon leaving a remnant.

2.3.3 Single meson production

Single pion production (SPP) is is the dominant interaction topology in the 1 GeV $< E_\nu <$ 10 GeV range and second dominant to CCQE at energies just above 0.5 GeV, and to DIS at energies below 100 GeV. The neutrino energy in SPP processes is less trivial to reconstruct than in the QE case due to the higher momentum transfer to the nucleon. The largest SPP contribution is resonant pion production, in which the neutrino excites the incident nucleon to a delta baryon state, which subsequently decays to a meson or a photon and a baryon as described by Rein and Seghal [39]. Low-momentum positively-charged pions decay via a muon and muon neutrino to a Michel electron (positron) and electron neutrino. This same process can produce photons, or heavier mesons such as kaons and etas if the deexcitation energy is sufficient.

In addition to the resonant production mode, coherent production in which the neutrino interacts with the nucleus as a whole as described by Berger and Seghal [40], and diffractive production described by Rein [41], can produce pions without exciting the nucleus. In these cases, the nucleon is not treated independently of the nucleus.

NC neutral pion production is a background to measuring the rate of CC electron neutrino interactions, as asymmetric decay or reconstruction of the decay photons of π^0 are indistinguishable from electrons in water-Cherenkov based detectors. As in the QE case, the NC pion production processes can be described by the same models put forward by Rein, Seghal and Berger.

2.3.4 Multiple pion production and Deep Inelastic Scattering

At high neutrino energies and momentum transfers, the nucleon may be excited to much higher energies in which the deexcitation can produce multiple hadrons if there is sufficient energy produced in the deexcitation. If the energy is high enough the neutrino can probe the structure of the nucleon in deep inelastic scattering (DIS), causing the quarks to become unbound, destroying the nucleus and leaving the constituent quarks and gluon to undergo hadronisation. As the momentum transfer moves into the inelastic regime, the interaction cross section becomes dependent on the underlying structure of the nucleon which is described by Parton Distribution Functions (PDFs) [42]. The hadronisation of the constituent quarks then typically produces a multi-pion or multi-mesonic final state.

2.4 Nuclear models

Due to the small interaction cross sections of neutrinos, typically on the order of 10^{-42} cm^2, there must be targets with sufficient mass available for neutrinos to interact with inside detectors. Practically, this means that most interactions take place on nuclear targets with complex potentials. At T2K beam energies neutrinos penetrate the nuclear structure but do not typically probe bare quarks. Nuclear effects are a non-negligible source of systematic error on the prediction of the final-state lepton kinematics and so appropriate uncertainties must be used to avoid biasing the reconstructed neutrino energy.

Typical target materials in neutrino experiments are plastic scintillator such as polystyrene (hydrocarbon), water, or argon. Although experiments may also use much heavier elements such as lead or iron to cause more interactions or trigger more energy conversion within a smaller volume detector, target choice is often also dependent on instrumentability; more interactions will take place in a detector which is largely made of lead, but particle tracking will be more difficult if the tracking components of the detector cover less of the volume.

Full calculations of interactions within nuclear potentials are computationally expensive and so neutrino experiments typically have used simple models whilst their uncertainties have been statistically dominated. As T2K and other modern experiments are still statistically limited, the use of these simple models is not currently a limiting factor on the extraction of the oscillation

parameters; however, this is not to say that there is no effect, and as the field moves to larger experiments such as Hyper-K [43] or DUNE [44] which can collect more data, a well-constrained nuclear model will be necessary for precision measurements of the PMNS parameters.

2.4.1 Shell model

The shell description of the nucleus is analogous to the electronic shell description of the atom which gives rise to the periodicity of the elements [45]. The shell model describes the possible bound states of the nuclear potential, taking into account the shape of the overall nuclear potential, but also interactions between different nuclei. Nuclei are arranged in orbital pairs which can explain phenomena relating to odd or even numbers of protons and neutrons, and these orbitals build up stable configurations, or shells, when they correspond to "magic numbers" of nucleons, which deviate from the predicted behaviour of simpler models such as the liquid-drop or Fermi gas models. A nucleus is said to be magic if it has a magic number of protons or neutrons, and can be doubly magic if it has a magic number of each.

Unlike in the electronic shell structure of the atom, in which the most stable energy configurations of the atom correspond to filling shells which generally correspond directly to the principal quantum number, the most stable nuclear configurations do not typically correspond to these due to more overlapping subshell structure. The principal quantum numbers, which are solutions to the radial component of the Schrodinger equation, often directly correspond to the shell structure in the atom as the nucleus is point like in comparison to the shells of the electrons. For a nucleon, however, the nucleus is no longer point like, and effects such as spin-orbit coupling are significantly larger, leading to more overlap between the energy levels corresponding to different principal quantum numbers. This much larger overlap leads to groups of subshells corresponding to different principal quantum numbers, somewhat arbitrarily, forming the most stable energy configurations, or shells. For example, the $1d$ and $2s$ energy levels overlap forming one shell which can hold up to 12 electrons. The next highest shell consists of a single $1f7/2$ orbital, and the next highest contains orbitals from the 1f, 2p and 1g levels. The number of nuclei in the ground state of a nucleus which these stable shells correspond to are referred to as the magic numbers. The energy levels resulting from the shell structure are shown in figure 2.7. Whilst a full shell model is not implemented at T2K, the known resulting effects of a shell model can guide the development of corrections and uncertainties applied to models such as the spectral function, which aim to give enough freedom to the T2K interaction model to cover these possible effects on T2K data.

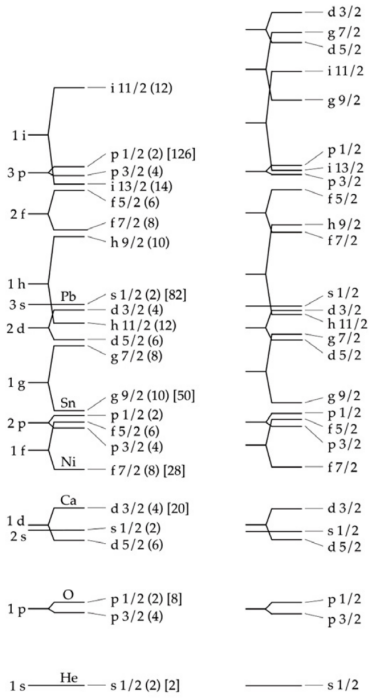

Figure 2.7: Energy levels ordering for protons (left) and neutrons (right) filling according to the shell model [45]. Differences between protons and neutrons are due to the inclusion of the extra Coulomb potential for protons. Each level can be occupied by two protons (neutrons) in opposing spin configurations. Energy level structure arises from solutions of the Schrodinger equation for the nuclear potential, which results in splitting due to the coupling between the spin and angular momentum of the nucleon [45].

2.4.2 Relativistic Fermi Gas model

One of the simplest nuclear models which predicts the distributions of the nucleon momentum is the (Global) Relativistic Fermi Gas (GRFG) model [46]. Nucleons are spin 1/2 particles and obey Fermi-Dirac statistics, so this model treats the nucleus as a Fermi gas within a boundary with the corresponding density of states. From this picture the average nucleon momentum and energy can be found.

2.4.3 Spectral function model

The spectral function (SF) model is a semi-empirical model that aims to bridge the gap between the one-dimensional RFG description, and more sophisticated shell-like descriptions such as mean field models [47], using kinematic observable information from electron-nuclear ($ee'p$) scattering data [48] and taking theoretical considerations into account [49].

In elastic $ee'p$ scattering experiments, the missing energy in an interaction, E_{miss}, can be measured by comparing the sum of the energy of the initial and final state particles. The initial state energy of the system is

$$E_i = E_e + M_N \tag{2.28}$$

where E_e is the initial state lepton energy and M_N is the mass of the nucleus which is assumed to be at rest. The final state energy is given by the energies of the outgoing electron, $E_{e'}$, the nuclear remnant E_{rem}, and the ejected nucleon E_n and can be expressed as

$$E_f = E_{e'} + E_{\text{rem}} + E_n, \tag{2.29}$$

where the energy of the remnant

$$E_{\text{rem}} = M_{\text{rem}} + T_{\text{rem}} = M_{\text{rem}} + \sqrt{p_{\text{miss}}^2 + M_{\text{rem}}^2} - M_{\text{rem}} \tag{2.30}$$

is the sum of it's mass, M_{rem}, and kinetic energy, T_{rem}. Here, T_{rem} is calculated using the missing momentum in the system

$$p_{\text{miss}} = |\vec{p}_e - \vec{p}_{e'} - \vec{p}_n| \tag{2.31}$$

where $\vec{p}_e$, $\vec{p}_{e'}$ and $\vec{p}_n$ are the momentum of the incident electron, final-state electron, and ejected nucleon respectively. If the missing energy E_{miss} is then defined as the difference between the sums of the initial and final state energy, then

$$E_{\text{miss}} = E_e + M_n - E_{e'} - E_n - T_{\text{rem}}. \tag{2.32}$$

The difference in the mass terms of the nucleus and the nuclear remnant is simply the mass of the ejected nucleon, $M_n = M_N - M_{\text{rem}}$, as the remnant has not yet decayed to a ground state. If the

energy of the incident electron is fixed in a monoenergetic beam, then the the values of E_{miss} are only dependent on $E_{e'}$, E_n and p_{miss}, the last of which can be measured.

The SF is a two dimensional function which describes the probability of an interaction taking place with a given E_{miss} and p_{miss}. The function is built from the theoretical predictions of a Mean Field (MF) model constrained by electron scattering data, and short-range correlated (SRC) nucleons. SRC nucleons are strongly repulsive pairs with typically larger momenta than that of the Fermi momentum and so they populate much of the tail of the E_{miss} and p_{miss} distributions. The SF probability. $P(\vec{p}_{\text{miss}}, E_{\text{miss}})$, is then given by the linear addition of the two components

$$P(\vec{p}_{\text{miss}}, E_{\text{miss}}) = P_{\text{MF}}(\vec{p}_{\text{miss}}, E_{\text{miss}}) + P_{\text{corr.}}(\vec{p}_{\text{miss}}, E_{\text{miss}}) \tag{2.33}$$

where $P_{\text{MF}}(\vec{p}_{\text{miss}}, E_{\text{miss}})$ is the MF component and $P_{\text{corr.}}(\vec{p}_{\text{miss}}, E_{\text{miss}})$ is the SRC component. The resultant distribution is shown in figure 2.8. A comparison of the SF to the E_{miss} predicted by MF models [50][47] is shown in figure 2.9

The removal energy, E_{rmv}, or the energy required to remove a nucleon from its ground state potential in $\nu - A$ scattering, is analogous to the E_{miss} measured in $ee'p$ scattering. In this case, the p_{miss} is assumed to associated with the unaccounted momentum of the initial state nucleon, k. Once the differences in masses and phase space are accounted for, the SF measured in $ee'p$ scattering can be used to better predict probability of neutrino interactions taking place with a given set of final state momenta. If modelling QE processes according to the SF nuclear model, this dependence on the removal energy has information about the nuclear shell structure "baked-in" giving a more accurate prediction of the energy available for the products of the interaction. The SF used in the T2K model for CCQE interactions is outlined by Benhar et al. in [49]. Modelling the kinematics of the final-state lepton for a neutrino interaction of a given energy E_ν and nuclear target is key to modelling the correct reconstructed energy spectrum in a neutrino experiment. A bias in the energy estimation will lead to a bias in the Δm^2_{32} constraint which depends on accurately determining the position of the oscillation maximum in true E_ν.

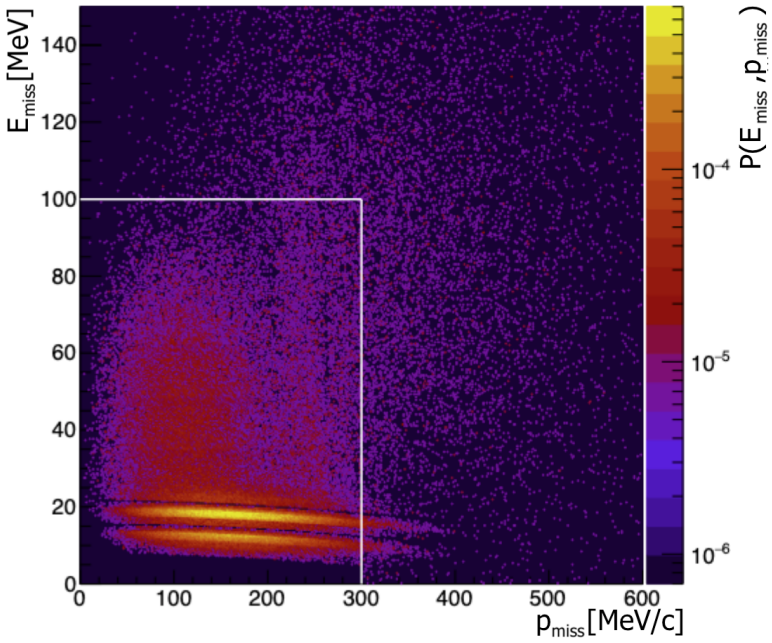

Figure 2.8: Two dimensional SF for oxygen built from theoretical predictions of mean field (MF) models highly constrained by electron scattering experiment data [49] describing the probability of an interaction with a nucleon resulting in a missing energy E_{miss} and missing momentum p_{miss}. In CCQE interactions these can each be considered approximately equal to the nucleon removal energy, E_{rmv}, and the initial nucleon momentum, k.The region inside the white box corresponds to the shell structure prediction of the MF models, and the shape of the region outside of this describes short-range correlated (SRC) nucleons. In the MF region, the two sharp $1p$ orbitals can be seen as peaks below 20 MeV in E_{miss}, and the diffuse $1s$ shell is spread out up to around 80 MeV. The SRC contributions are more spread across the function.

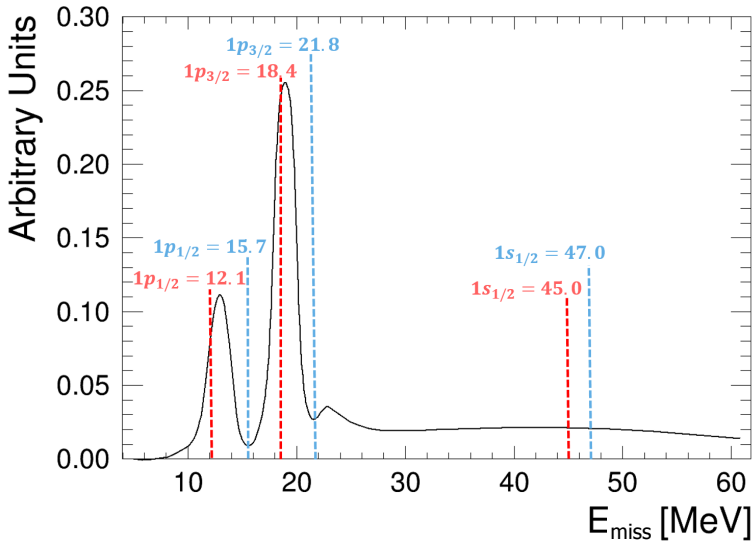

Figure 2.9: Comparison of the SF E_{miss} distribution (black line) for oxygen, to the E_{miss} predictions for different nuclear shells from MF models [50][47] (dashed lines) for protons (red) and neutrons (blue) with values given in MeV. The sharp peaks correspond to the more loosely bound outer p shells and the more diffuse tail corresponds to the higher-energy, diffuse inner s shell.

2.5 Neutrino experiments past and present

There are many ways in which neutrinos are produced in natural and artificial processes which allow for different avenues to investigate their properties. Each of these avenues, however, has constraints. Most significantly, their interaction cross section generally increases with energy [51], meaning the more abundant, lower energy neutrinos are difficult to study due to their low interaction rates. The mechanism by which they are generated may also favour one neutrino flavour over another, or the distance they travel before reaching detectors on earth may lead to mixing obscuring the processes that generated them.

2.5.1 The cosmic neutrino background

The most abundant neutrinos in the universe are those in the Cosmic Neutrino Background (CNB) which were produced shortly after the big bang when neutrinos froze out of thermal equilibrium with the hot plasma of quarks and gluons [52]. At this point the universe went from opaque to transparent to neutrinos and the neutrinos in the CNB today have travelled the farthest, and are the oldest in the universe. As the universe has expanded, their energy has decreased due to adiabatic cooling and are now below 10^{-3} eV in energy. This is far below the threshold for CC interactions as there is not enough energy to produce the mass of any of the charged leptons, and so CNB neutrinos can only interact through the NC.

2.5.2 Reactor neutrinos

Reactor neutrinos are produced through fission processes in nuclear-power reactors. These are typically associated with β-decay and so are $\nu_e(\bar{\nu}_e)$. Reactor neutrinos are slightly higher energy than those produced in natural geological processes, but are still relatively low in the MeV range. Reactors do provide high intensity sources of neutrinos where the process through which they are produced can be monitored directly through the power output of a reactor.

The first direct detection of a neutrino in 1956 by Cowan and Reines took advantage of the high flux of neutrinos from a nuclear reactor and used tanks of cadmium-chloride solution sandwiched between tanks of liquid scintillator instrumented by photomulitplier tubes (PMTs) [53][54]. The inverse beta decay process

$$\bar{\nu}_e + p \rightarrow n + e^+ \tag{2.34}$$

results in the production of two 0.51 MeV photons when the positron annihilates with an electron in the detector which can be detected in coincidence by the PMTs in two adjacent tanks of scintillator. In addition to these initial, or 'prompt' photons, the capture of a neutron on a cadmium-108 nucleus produces an additional photons from the decay of the metastable cadmium-109,

$$n + {}^{108}\mathrm{Cd} \rightarrow {}^{109}\mathrm{Cd}^* \rightarrow {}^{109}\mathrm{Cd} + N\gamma \tag{2.35}$$

with the total energy of these photons summing to 9 MeV. The neutron capture takes place within a few microseconds of the initial coincident photon detection, and so the detection of the positron annihilation followed immediately by the neutron capture indicates the interaction of an antineutrino in the detector, indicating the detection of the products of equation 2.34. The approach taken by Cowan and Reines is similar to that taken by modern experiments, such as alternating water and scintillating materials as in ND280, and ongoing work at SK using gadolinium doped water for neutron capture.

Reactors are also useful for studying short baseline neutrino oscillations due to their high intensity and low energy. It is also typical for multiple nuclear reactors to be found in relatively close proximity, offering a range of known baselines. The Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) experiment is a liquid scintillator detector built in the cavity previously used by the Kamioka Nucleon Decay Experiment (KamioKANDE) in the Kamioka mountains on the west coast of Japan [55]. KamiokaNDE studies the neutrinos from 55 reactors operated in Japan's large commercial nuclear power sector. This gives a range of baselines from which oscillations can be studied. KamLAND has made a precise measurement of the Δm^2_{21} and measurements of $\sin^2\theta_{12}$ consistent with solar experiments.

A gadolinium doped liquid scintillator-based reactor experiment, the Chooz experiment [56],

measured the $\bar{\nu}_e$ spectrum from the Chooz-B nuclear power plant in near the Belgian border of France at a baseline of 1.05 km from two reactors. This allowed for the use of direct thermal power measurements of the reactors of the two reactors to constrain the neutrino flux modelling of the experiment. The Chooz experiment demonstrated that the mixing between $\bar{\nu}_e$ and $\bar{\nu}_\mu$ was not large enough to explain the deficit in upward going atmospheric neutrinos, by excluding large values of the mass splitting for large values of the mixing angle (and vice versa) between the first neutrino mass states and any other. In the three flavour framework this leaves only the possibility of mixing between the ν_μ and ν_τ.

The South Korean Reactor Experiment for Neutrino Oscillations (RENO) uses two identical gadolinium doped liquid scintillator detectors placed 290 m and 1.4 km from the beam to measure oscillations from a known source [57][58]. This near-to-far extrapolation allows for a reduction in the systematic uncertainty by comparing the ratios of the flux for the same neutrino source at different baselines directly.

Double Chooz took repurposed the original site of the Chooz reactor to perform a similar near-to-far extrapolation based analysis as RENO [59]. The 1.05 km baseline site became the far detector of the Double Chooz analysis, with a nearly identical detector placed just 250-300 m away from the two detector cores. The precise measurement of the $\bar{\nu}_e$ flux using the near detector allowed for a much more precise measurement of θ_{13} than the exclusions set by Chooz.

The Daya Bay experiment in China uses a series of eight identical detectors placed in groupings that allow for measurements range of short baselines from 365 m to 1663 m from the reactors at three separate nuclear power plants in the Daya Bay nuclear power plant complex [60]. This multi-detector and multi-baseline approach has allowed Daya Bay to make the most precise measurement of $\sin^2 2\theta_{13}$. The successor to Daya Bay, JUNO, is expected to begin taking data in 2023 and make sub-percent precision measurements of the reactor mixing parameters, and aims to determine the neutrino mass ordering [61]. Currently, T2K and NOvA use the PDG global best-fit value of θ_{13} taken from reactor experiments [34] to better constrain δ_{CP} as discussed in section 5.3. In future, T2K plans to use the $\sin^2 2\theta_{13}$-Δm^2_{32} likelihood surfaces published by the reactor experiments directly.

2.5.3 Solar neutrinos

Solar neutrinos are produced by various fusion chain processes in the sun, and the energy and flux of these differs. There are two main chains of reactions which produce neutrinos in the sun, the *pp* chain and the CNO chain. The most common solar neutrinos come from the proton-proton or *pp* interaction in the *pp* chain with energies below 400 keV, which is still not energetic

enough to produce an electron in a charged current interaction. The highest energy neutrinos are from helium-3-proton or *hep* interactions from the same chain and can reach energies of up to 19 MeV, though these have the lowest flux. More common, with slightly smaller maximum energies are from the decay of boron-8 (^{8}B) nuclei. Neutrinos in the Sun undergo matter effects during oscillations as they travel from the core to the surface in a highly electron-dense material as demonstrated by Mikheyev, Smirnov [62] and Wolfenstein [63]. These present themselves as energy-dependent modifications to the vacuum oscillation probability which can resonantly enhance oscillations in the Mikheyev-Smirnov-Wolfenstein (MSW) effect.

In 1965 Ray Davis Jr. looked for solar neutrinos using a large tank of 380 m^3 perchloroethylene (a common cleaning fluid) roughly 1.5 km underground at the Homestake mine to reduce backgrounds from cosmic rays and other sources [64]. ν_es would interact with the chlorine-37 isotopes in the tank and produce argon-37 as a neutron was changed into a proton in a CC interaction

$$\nu_e\,^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Cl}^* \rightarrow {}^{37}\mathrm{Ar} + e^-. \tag{2.36}$$

The argon could be filtered out of the solution and slowly collected. The decay of this argon-37 isotope could then be monitored to count the number of argon atoms collected. Davis did indeed confirm the existence of solar neutrinos provide evidence that the Sun was powered by fusion in its core. However, the number of (electron) neutrinos detected was found to be roughly a third of that calculated by John N. Bahcall using the standard solar model (SSM). This was referred to as the Solar neutrino problem, and was confirmed by other experiments using a similar radiochemical detection method using gallium, SAGE [65] (1989-Present) and GALLEX [66] (1991-1997). In 1989, the Kamiokande-II experiment [67] (an upgrade to the Kamioka Nucleon Decay Experiment) a 3000 tonne water Cherenkov detector also observed a deficit of solar neutrinos, but in this case the observed number of neutrinos was approximately half that of the prediction, not a third. In 2001 this was proven to be due to ν_es oscillating to other flavours by the SNO experiment, which was also able to measure the total three-flavour flux through NC interactions in a heavy-water detector [29]. The discrepancy between the radiochemical and water Cherenkov detectors was determined to be due to the MSW effect strength at different neutrino energies.

2.5.4 Atmospheric neutrinos

Atmospheric neutrinos, produced by cosmic ray bombardment of the atmosphere, cover a wide range of energies. Though they have a peak energy in the 10-100 MeV range, they can also be found up to the TeV scale. These are a mixture of ν_μ and ν_e. As the production rate of atmospheric neutrinos should be isotropic, and neutrinos can pass through the earth, the zenith

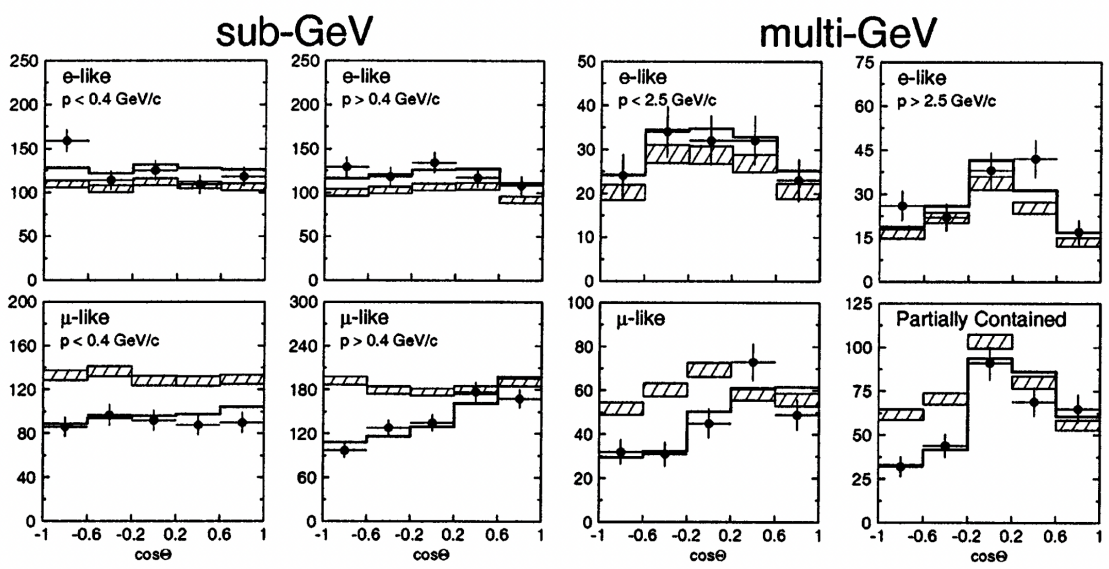

Figure 2.10: Event rate distributions binned in zenith angle for the SK atmospheric samples which demonstrated ν_μ disappearance through the event rate dependency on the angle, and therefor the distance travelled through the earth. A clear deficit in upward going ν_μ candidates can be seen. The ν_e appearance channels show no strong dependency on the angle, indicating the ν_μ are oscillating to another flavour [9].

angle of neutrinos at any point near or on the Earth's surface will correspond to a varying baseline of neutrino oscillations.

Kamiokande-II was also able to make an initial measurement of a discrepancy between upward-going and downward-going flux from atmospheric neutrinos, but was not able to make a precise measurement of the disappearance. The Super-Kamiokande experiment, with a scaled up and more sophisticated detector (described in detail in section 3.5.) definitively measured the ν_μ disappearance [9]. It was found that although the ν_μ flux had a clear deficit (as shown in figure 2.10, there was no strong indication of oscillation to ν_e, suggesting that most of the neutrinos were oscillating to ν_τ (in a three-flavour formalism).

2.5.5 Accelerator neutrinos

Neutrinos can be produced artificially by accelerators. These are typically in the 100-MeV to 10-GeV range and are mostly ν_μ. Producing neutrinos in this way provides a controllable and high intensity source of neutrinos to study. This makes accelerators useful for studying neutrino interactions as well as oscillations.

In 1962 the Alternating Gradient Synchrotron (AGS) experiment at Brookhaven found that neutrinos produced by the decay of pions produced in proton collisions produced only muons [68],

in contrast to those detected in early reactor experiments which produced electrons. This demonstrated that there was more than one type of neutrino, and that these were associated with the different charged leptons.

In 1960, Mel Schwartz proposed a method by which high energy proton collisions would produce charged hadrons which would semi-leptonically decay to μ^+ and ν_μ (as well as μ^-, $e^\pm$ and $\bar{\nu}$) [69]. A series of improvements allowed over the years at different accelerators led to the modern neutrino beam. Hadrons could be focused and collimated not only to select for a narrower range of neutrino energies, but also would be able to select for the charge of the hadron and produce either neutrinos or antineutrinos [70]. Targets are typically made from graphite, aluminium or beryllium. A more detailed description of a modern neutrino beam, as used in T2K, is given in section 3.1.

Accelerator-based neutrino experiments have been used for a range of neutrino-interaction measurements, such as those made by the bubble chamber experiment Gargamelle which discovered the NC interaction in 1973 [71], and more modern experiments like MINERvA [72] which made detailed cross-section measurements in a range of materials.

Although the most common neutrino production method at accelerators uses the decays of pions to ν_μ, neutrinos can be produced though the decays of other mesons. The Direct observation of the nu tau (DONuT) experiment at Fermilab used the Tevatron proton beam to produce charmed mesons (D_s) which can undergo decay to ν_τ [73]. DONUT used an emulsion-hybrid detector with very high spacial resolution to identify the short tracks left by the τs before they could decay.

Accelerators also opened up the opportunity to do more precise measurements of the oscillations of ν_μs to other flavours by achieving a high intensity flux and being able to control the initial neutrino energy well. The ability to control the energy allowed for experiments to tune their L/E_ν ratio far more, and target the oscillation probability maxima. The KEK to Kamioka (K2K) experiment was the first long-baseline oscillation experiment which used a neutrino beam produced at KEK in Tsukuba which was pointed toward the SK detector with a peak energy of 1.3 GeV and a baseline of 250 km [17]. A near detector using the same water Cherenkov detector design as SK but scaled down to 1 kT was used to constrain the beam and interaction uncertainties. K2K set limits on the allowed ranges of the $\sin^2\theta_{23}$ and Δm^2_{32} with a best-fit value of $2.8^{+0.7}_{-0.9} \times 10^{-}3$ eV^2 consistent with maximal mixing. This was consistent with the SK atmospheric measurement.

Another long-baseline experiment, MINOS, with a higher energy and longer baseline was conducted between FermiLab and the Soudan mine [74]. MINOS had two magnetised steel-scintillator calorimeter detectors and used the FermiLab NuMI beam. The baseline of the experiment was 735 km and had peak energy of 3 GeV. The NuMI beam was also able to run at higher

energies, but as it became clear from other measurements that Δm^2_{32} was small, MINOS only ran in the "low" energy mode. An upgrade to MINOS, MINOS+, ran with a higher energy beam with a peak of 7 GeV, making more precise measurements of the standard three flavour mixing parameters [75] and searched for larger mass squared splittings that may be associated with heavier (sterile) neutrino mass state [76].

After the largest mass-squared splitting was found to be small, and the mixing angle to be non-zero, experiments could be designed to more precisely measure the reactor PMNS parameters. Significantly, the large size of θ_{13} measured by reactor experiments suggested that long-baseline accelerator experiments would be sensitive to the CP-violating phase, δ_{CP}, through $\nu_e(\bar{\nu}_e)$ appearance in $\nu_\mu(\bar{\nu}_\mu)$ beams. Two currently running long-baseline accelerator experiments are T2K [15], in the J-PARC beam, and NOvA [77], in the FermiLab NuMI beam with its near detector in the same detector hall as MINERvA and MINOS+. These experiments have different baselines, neutrino energies, and detector designs and provide complimentary constraints on the values of θ_{13}, θ_{23}, δ_{CP}, and $\Delta m^2_{32}(\text{NO})/|\Delta m^2_{31}|(\text{IO})$. Both experiments use narrow-band off-axis beams, but NOvA's beam energy is focused at 2 GeV with a baseline of 810 km, and T2K's is focused at 0.6 GeV with a baseline of 395 km. Both experiments are tuned to probe oscillations at the first oscillation maximum between ν_μ and ν_τ, but these differences both drive, and are driven by, decisions about the detector designs. NOvA has more high-energy interactions with DIS and multi-π production and the detectors are designed to capture as much of the energy as possible, using large liquid scintillator-tracking detectors for both near and far detectors. T2K uses the pre-existing SK detector (see section 3.5), which is most efficient at reconstructing quasielastic interactions with a single charged particle in the final state, and uses a more sophisticated tracking detector, ND280 (see chapter 4), to sample the unoscillated beam and constrain model uncertainties as is discussed in the analysis presented in this thesis. This leads to different sensitivity to the PMNS model parameters for each detector due to both the interaction energies and different matter effects. These sensitivities are themselves dependent on the true values of the oscillation parameters. A comparison of the current T2K and NOvA results is given in section 10.7.2 where the benefits of a joint fit are explored.

Two new experiments, Hyper-Kamiokande [43] (HK) and the Deep Underground Neutrino Experiment (DUNE) [44] are under construction which aim to make precise measurements of the remaining oscillation parameters and determine the neutrino mass ordering. HK will follow T2K using the upgraded 1.5 MW J-PARC beam and near detector, as well as an Intermediate Water-Cherenkov Detector (IWCD) with a variable off-axis angle at 1 km from the target. The new far detector for HK will use the same baseline and off-axis angle as T2K, but on the other side of

the beam, and have a water mass 10× that of SK. The quantum efficiency of the HK PMTs is expected to be double that of SK. All of these factors combine to a much larger collected data set with more well-understood systematic uncertainties. In addition to the beam measurement, HK will be able to make a more precise atmospheric-mixing measurement than SK. DUNE will use a 40 kilotonne Liquid Argon Time Projection Chamber (LArTPC) detector and be the largest of its kind. LArTPCs have incredibly precise resolution and very detailed information about neutrino interactions can be extracted. DUNE will use an on-axis 1.3 MW beam with a peak energy of 3 GeV for a baseline of 1300 km. A series of near detectors will measure the beam intensity and also feature variable-off-axis positions. Like HK, DUNE will be able to make other measurements such as atmospheric-mixing measurements and sterile searches.

2.5.6 Supernova neutrinos

Supernovae produce a high intensity flux of neutrinos during gravitational core collapse, prior to the emission of visible light [78]. This makes supernova neutrinos a useful tool for astronomers to give them an early warning. In 1987, Kamiokande-II collected data on the only supernova which has been observed in the Milky Way since the discovery of the neutrino, Supernova 1987A, observing 11 neutrino events [79]. Many neutrino detectors are on standby for supernovae neutrinos, which if close enough will result in many neutrinos events within a short time period due to the high flux. Neutrinos are produced by supernova before the visible burst and so this may act as an early warning for astronomers. The SuperNova Early Warning System (SNEWS) takes triggers from a range of experiments including SK, IceCube, KamLAND, and Daya Bay [80].

2.5.7 Astrophysical and cosmogenic neutrinos

At high energies, neutrinos come from Active Galactic Nuclei (AGN) and cosmogenic processes. AGN are point-like sources in the sky and produce neutrinos ranging from the GeV to the PeV range [81]. Cosmogenic neutrinos are produced by ultra-high energy cosmic-ray protons interacting with photons in the cosmic microwave background and are above a PeV in energy [82]. The flux of these high energy neutrinos are extremely low, but their cross-section is much higher than most neutrinos produced in other mechanisms.

Two 1-km-cubed scale experiments search for the high-energy and ultra-high-energy neutrinos from astrophysical and cosmogenic sources. IceCube [83] in the antarctic and ANTARES [84] in the Mediterranean instrument large volumes of ice and water respectively with optical modules containing PMTs. The two experiments are complimentary as each detector has better sensitivity to point-like sources in the sky on the other side of the earth where atmospheric and cosmic-ray

backgrounds are reduced.

2.5.8 The number of neutrinos and sterile neutrinos

Collider experiments were able to determine on the number of weakly interacting light neutrino generations through the decay width of the Z^0. Since the Z^0 is neutral it should be able to decay to all weakly interacting $\nu\bar{\nu}$ pairs with a combined mass less than that of the Z^0, m_{Z^0}. The LEP experiments, ALEPH, DELPHI, L3 and OPAL determined to high precision that there are only three generations of neutrinos that the Z^0 could decay to [85]. This is consistent with a doublet lepton sector and the known number of charged lepton flavours. There may, however, be neutrinos that are either not weakly interacting (sterile) or weakly interacting and incredibly massive ($m > m_{Z^0}/2$).

A larger mass would lead to much larger mass-squared differences and so searches for short-baseline oscillations may find neutrinos oscillating to as yet unknown heavy mass states. The Liquid Scintillator Neutrino Detector (LSND) was a short baseline (29.7 m) accelerator experiment which reported a 3σ excess over the expected backgrounds [86]. The MiniBooNE experiment intended to independently verify this excess, with a different detector principle (mineral oil scintillation and Cherenkov light), MiniBooNE also reported an excess consistent with a oscillations to a sterile neutrino with a significance of 4.8σ [87]; however, the LSND and MiniBooNE results were not consistent with each other [88]. MicroBooNE, a precision LArTPC detector has so far been unable to find an excess consistent with either LSND or MiniBooNE [89], but has also ruled out possible background processes proposed to explain the MiniBooNE excess [90].

A series of short-baseline reactor experiments [88] found significant disagreement with the predicted flux by the leading theoretical models by Mueller et al.[91] and Huber [91]. This is referred to as the "reactor anomaly", and may be consistent with a 3+1 (three weakly interacting and one sterile neutrino) sterile neutrino model. However, neither the non-mixing or mixing hypotheses can explain all of the features of the observed reactor flux spectrum. The 3+1 model is also in tension with cosmological limits which are sensitive to the number of degrees of freedom associated with the number of particle species in the early universe.

2.5.9 Neutrinoless double beta decay

Several experiments are searching for neutrinoless double beta decay, in which two beta particles are observed emitted back to back from a radioactive decay indicating that there is no missing momentum carried away by the neutrinos as there is in standard double beta decay. This would demonstrate that neutrinos are their own antiparticle, or "Majorana particles" [92], and are carrying

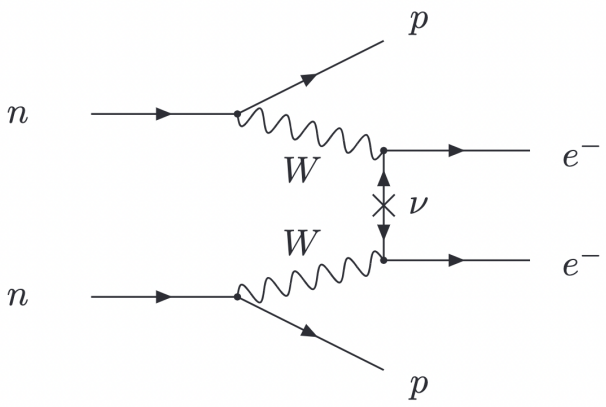

Figure 2.11: Feynman diagram for the neutrinoless double beta decay process in which the neutrino, being Majorana, can mediate the interaction. The two final-state beta particles must be emitted back to back to according to conservation of momentum [93].

the lepton flavour current between the two beta particles as shown in the Feynman diagram in figure 2.11. This would provide a non-Higgs mechanism for neutrinos to acquire mass. If neutrinos are Majorana, this would also introduce two additional complex CP-violating phases to the PMNS matrix, but standard long-baseline neutrino oscillation experiments are not sensitive to these. A series of experiments (GERDA, Majorana, LEGEND, KamLAND-zen, NEMO-3) [93] have set limits on the rate of neutrinoless double beta decay for different isotopes but it has not yet been observed. Typical half-lives are of the order of 10^{26} years or longer [94]. The upgrade to the SNO detector, SNO+ will search for neutrinoless double beta decay from Tellurium-130 [95].

2.5.10 Absolute masses of the neutrinos

Whilst neutrino oscillations are able to determine the difference between the squares of the masses, and therefore set a lower bound on the mass of two of the three known neutrino mass states, they cannot make a statement on the absolute mass of the neutrino nor exclude a zero-mass lightest neutrino mass state. Cosmological measurements are able to place upper limits on the combined masses of the neutrinos through measurements of the cosmic microwave background and baryonic acoustic oscillations [96]. These do presume a specific cosmological model, typically with a dark energy-like cosmological constant Λ and Cold Dark Matter (ΛCDM).

Some experiments aim to directly constrain the masses of the neutrino with highly precise spectrometers. The latest of these, the KArlsruhe TRItium Neutrino (KATRIN) experiment, uses a precise measurement of the tail of the β-decay spectrum of tritium to set an upper limit on the effective mass of the ν_e of 0.8 eV/c^2 at a confidence level of 90% [97][98]. If the neutrino is

not massless, then the $\beta - decay$ spectrum must fall off and truncate compared to the massless neutrino prediction, as the neutrino must carry some of the energy from the decay away in its mass.

2.5.11 Global best fit

Several groups publish global best-fit values of the three-flavour PMNS values using the published results and data releases of different neutrino experiments [93]. Results from the NuFit group [99][100] (including the SK-Atmospheric $\Delta\chi^2$ surface) are presented in table 2.1. The mixing parameters can be broadly split into the solar, atmospheric, and reactor/accelerator terms though there is overlap.

The solar mixing angle, θ_{12} is best constrained by the SNO and SK experiments, whilst the size of the corresponding mass-squared splitting, Δm^2_{21}, comes from the long-baseline reactor experiment, KamLAND, which is not sensitive to as strong matter effects. However, KamLAND is less sensitive to θ_{12} than the solar measurements. θ_{12} is large at 33.44° and Δm^2_{21} is very small, on the order of 10^{-5} eV^2. This also sets a lower bound for the second lightest neutrino mass on the order of 10^{-3} eV.

Atmospheric mixing measurements, such as SK's, are dependent on θ_{23}, Δm^2_{32}, θ_{13} and δ_{CP}, but are not dominant constraints on these parameters an reactor experiments and long-baseline accelerator experiments have made more precise measurements.

The θ_{13}, is most precisely measured by the medium-baseline reactor experiments like Daya Bay. Reactor experiments are also somewhat sensitive to Δm^2_{32}, but this can be most tightly constrained by the ν_μ disappearance measurements at long-baseline accelerator experiments such as T2K and NOvA. Global fits of θ_{23} favour the upper octant, slightly away from maximal mixing, but is still large. Octant determination is one of the goals of the long-baseline experiments, as there is an eight-fold degeneracy between the octant, the mass ordering and δ_{CP}. T2K and NOvA also provide the best constraints on δ_{CP}, but this is aided by the inclusion of the reactor constraint on θ_{13}. The best fit point for δ_{CP} in the normal ordering is close to the CP-conserving value of 180° (π), but in the inverted ordering is close to the maximal enhancement of ν_e appearance over $\bar{\nu}_e$ in long baseline accelerator experiments at 270° ($3\pi/2,-\pi/2$). Whilst both T2K and NOvA individually show a slight preference for NO, the global best fit shows a slight preference for IO. This is most likely driven by T2K's stronger constraint on δ_{CP} at values close to 270°. This is discussed in more detail in section 10.7.2.

	Normal Ordering		Inverted Ordering	
	bfp $\pm 1\sigma$	3σ	bfp $\pm 1\sigma$	3σ
$\sin^2\theta_{12}$	$0.304^{+0.012}_{-0.012}$.	$0.269 \to 0.343$	$0.304^{+0.013}_{-0.012}$	$0.269 \to 0.343$
$\theta_{12}/^\circ$	$33.44^{+0.77}_{-0.74}$	$31.27 \to 35.86$	$33.45^{+0.78}_{-0.75}$	$31.27 \to 35.87$
$\sin^2\theta_{23}$	$0.573^{+0.016}_{-0.020}$	$0.415 \to 0.616$	$0.575^{+0.016}_{-0.019}$	$0.419 \to 0.617$
$\theta_{23}/^\circ$	$49.2^{+0.9}_{-1.2}$	$40.1 \to 51.7$	$49.3^{+0.9}_{-1.1}$	$40.3 \to 51.8$
$\sin^2\theta_{13}$	$0.02219^{+0.00062}_{-0.00063}$	$0.02032 \to 0.02410$	$0.02238^{+0.00063}_{-0.00062}$	$0.02052 \to 0.02428$
$\theta_{13}/^\circ$	$8.57^{+0.12}_{-0.12}$.	$8.20 - 8.93$	$8.60^{+0.12}_{-0.12}$	$8.24 \to 8.96$
$\delta_{\rm CP}/^\circ$	197^{+27}_{-24}.	$120 \to 369$	282^{+26}_{-30}	$193 \to 352$
$\Delta m^2_{21}/10^{-5}\text{eV}^2$	$7.42^{+0.21}_{-0.20}$	$6.82 \to 8.04$	$7.42^{+0.21}_{-0.20}$	$6.82 \to 8.04$
$\Delta m^2_{3l}/10^{-3}\text{eV}^2$	$+2.517^{+0.026}_{-0.028}$	$+2.435 \to +2.598$	$-2.498^{+0.028}_{-0.028}$	$-2.581 \to -2.414$

Table 2.1: Best fit values for the PMNS parameters from a global fit to experimental results, including the SK-Atmospheric likelihood surface, by the NuFit group [93][99][100]. Most values are consistent across mass orderings with the exception of Δm^2_{3l}, and $\delta_{\rm CP}$. $\delta_{\rm CP}$ is close to CP-conserving values in the NO, and close to maximally CP-violating values in IO.

Chapter 3

T2K

The Tokai to Kamioka (T2K) experiment is a long-baseline narrow-band neutrino oscillation experiment which probes neutrino oscillation physics through muon (anti)neutrino disappearance and electron (anti)neutrino appearance in a muon (anti)neutrino beam. Figure 3.1 shows the path of the T2K beam along its 295-km journey through the Earth's crust as it travels from the Japan Proton Accelerator Research Complex (J-PARC) on the east coast of Japan to the Super-Kamiokande detector on the west coast. T2K is able to run in neutrino or antineutrino beam modes and so can constrain the level of CP violation in the lepton sector through comparisons of electron neutrino and antineutrino appearance samples to the CP conserving prediction. Muon neutrinos or antineutrinos are produced by firing 30-GeV protons from the J-PARC main ring accelerator onto a 90-cm-long graphite target producing hadrons, primarily pions and kaons, which are selected by charge using magnetic horns and allowed to decay leptonically to antimuons and neutrinos (or muons and antineutrinos).

Whilst the on-axis neutrino beam has a wide-band energy spectrum, placing detectors at an off-axis angle can both narrow the peak of the spectrum and enhance it at the desired energies. As

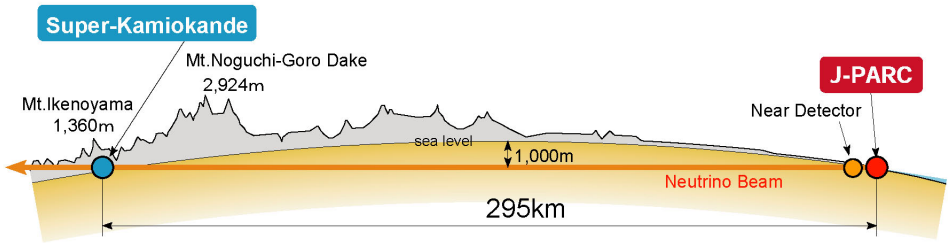

Figure 3.1: Cross-section view of Japan along the T2K baseline showing the direction of the beam through the Earth's crust from J-PARC on the east coast of Japan to Super-Kamkiokande on the west coast.

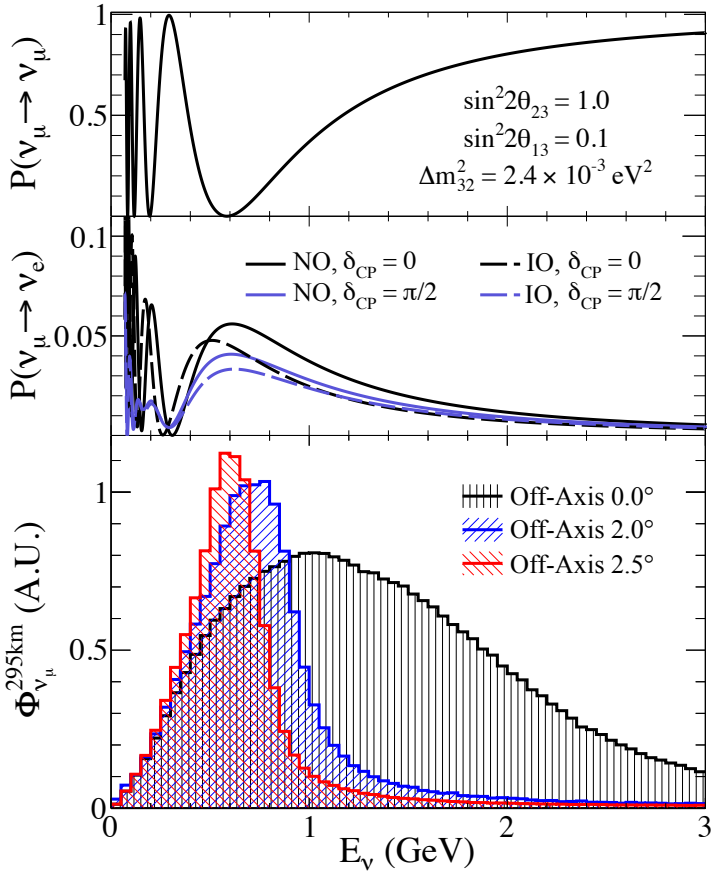

Figure 3.2: Top: Probability of muon neutrino survival over the T2K baseline versus neutrino energy for given oscillation parameter values. Middle: Probability of a muon neutrino oscillating to an electron neutrino versus neutrino energy, for normal ordering (NO), inverted ordering (IO) and diffeent values of δ_{CP}. Bottom: Neutrino energy spectrum for the T2K beam at given off-axis angles.

shown in figure 3.2, this allows the beam to be tuned to maximise the chance of oscillations for the T2K baseline. One benefit of T2K's chosen off-axis angle of 2.5° is not just that it tunes the value of L/E_ν to maximise oscillations, but it also selects lower energy neutrinos ($E_{\text{peak}} = 600$ MeV) meaning that the most likely mode of neutrino interaction is quasi-elastic. These interactions typically have a simple topology for which we can most reliably reconstruct the energy of the neutrino. The off-axis angle also suppresses the ν_e contamination of the beam due to the different hadronic parent kinematics.

T2K has a suite of near detectors 280 m downstream of the target which characterise the un-oscillated beam. One of the near detectors is the on-axis near detector INGRID which measures the beam intensity and direction, and several off-axis detectors, ND280, WAGASCI, and Baby-MIND, which carry out detailed neutrino interaction studies as part of the oscillation analysis and a wider physics program.

The Super-Kamiokande (SK) detector is a 50-kton water-Cherenkov detector in the Kamioka mountains 295 km away from J-PARC. An experiment in and of its own right, the SK collaboration make measurements of atmospheric and solar neutrino oscillations, as well as searching for proton decay and supernovae neutrinos. The T2K experiment uses the SK detector as a far detector to sample the relative rates of muon neutrino disappearance and electron neutrino appearance in the beam at the first oscillation maximum for T2K energies. SK and the ND280 near detector are both at a 2.5° off-axis angle to the beam direction, and therefore see the same beam, allowing a strong constraint to be placed on the neutrino beam flux and cross-section modelling at SK, attained through fitting the models to the ND280 data. There is some cancellation of systematic uncertainties between the two detectors due to correlations of the flux parameters. The advantages of these constraints and the methods used to obtain them are discussed in detail in chapters 6, 7 and 8.

3.1 Beam

The T2K beam is created by firing protons at a graphite target and allowing the hadronic products to decay leptonically. Negative hydrogen ions are accelerated by a linear accelerator or LINAC to 400 MeV. Then the two H^- electrons are removed by charge-stripping foils producing H^+ ions or bare protons during injection into the Rapid Cycling Synchrotron (RCS). The RCS accelerates the protons up to 3 GeV with a 25-Hz cycle frequency. From the RCS the protons are then injected into the Main Ring (MR) which accelerates the protons up to 30 GeV. Protons are grouped in eight bunches and can be extracted in either fast extraction or slow extraction which are used for neutrino production and hadron experiments, respectively. In slow extraction a ribbon is used

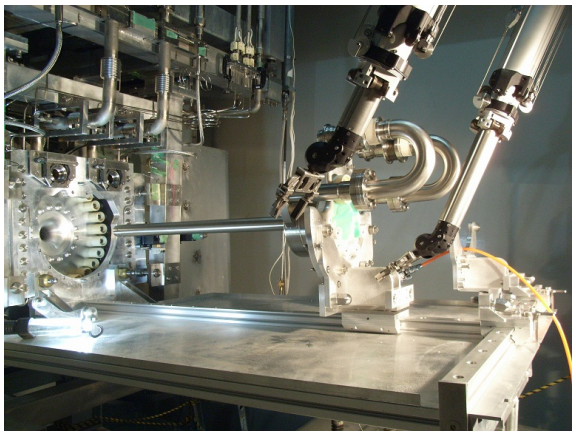

Figure 3.3: The T2K graphite target shown here outside of the first magnetic horn shown on the left of the picture. Due to the high radiation levels in the target hall, handling of the target must be done using a remote mechanical system, seen on the right, in event of target failure. During beam running, the target sits inside the first horn to facilitate the capture of all hadrons with the correct charge.

to split the bunch as a kicker magnet directs a portion of each bunch to the hadron-experiment beamline over multiple turns. In fast extraction the eight bunches are extracted in one turn by five kicker magnets and directed towards the primary proton beamline of the neutrino beamline. The beam may be aborted by using the fast extraction kicker magnets to direct the beam to a beam dump.

The primary proton beamline directs the protons towards a graphite target measuring 90 cm in length and 2 cm in diameter as can be seen in in figure 3.3. Proton collisions on carbon produce an array of hadrons, including charged pions and kaons. A series of three electromagnetic horns, as shown in figure 3.4, is used to select $\pi^{\pm}$ depending on the horn polarity. The "Forward Horn Current" (FHC) running mode selects π^{+} and the "Reverse Horn Current" (RHC) selects π^{-}. Complete separation of charged particles in this manner is not possible and so there will always be some contamination of other particles, of kaons and wrong sign pions in particular. The wrong sign contamination is greater in the RHC mode due to the fact the pion production process starts with positively-charged protons incident on carbon nuclei with a net positive electric charge. The desired pions that are produced in the target are then allowed to travel through a 100-m-long volume in which they can decay leptonically, with positive (negative) pions decaying to antimuons (muons) and muon (anti)neutrinos. The number of neutrinos produced cannot be directly measured; therefore, the metric used to assess the total quantity of neutrinos we expect to have collected in a given run period is the number of protons on target, or POT. This assumes that the rate of neutrinos produced is relatively stable when compared to the rate of protons incident

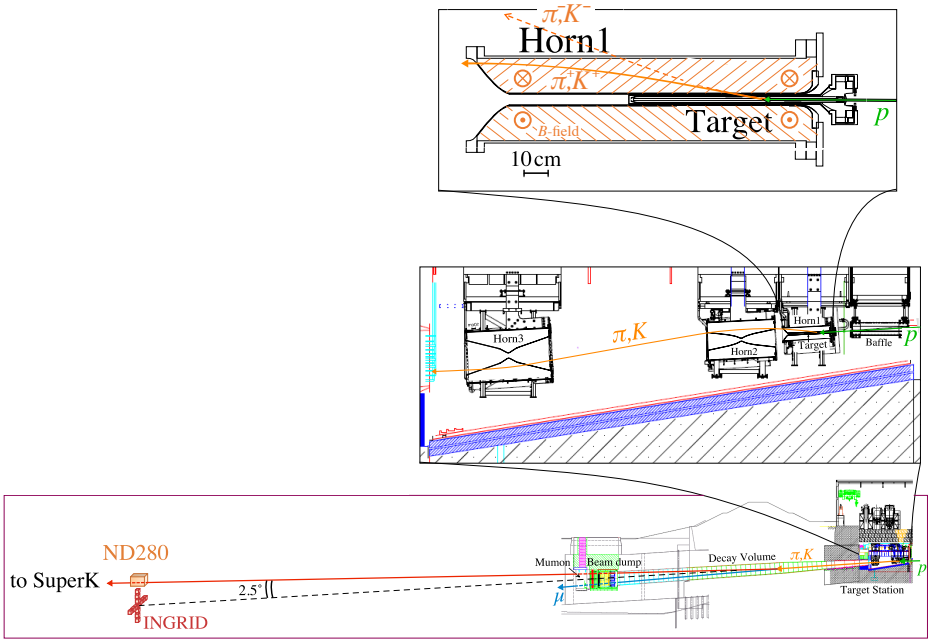

Figure 3.4: Schematic of the neutrino production target as situated within the first of three magnetic horns. This series of horns is used to focus and select charged pions of the desired sign which are directed towards a 100-m decay volume. Here, the pions decay to antimuons and muon neutrinos for a forward horn current, or muons and muon antineutrinos for a reverse horn current. Remaining particles are swept up by a beam dump to prevent further decays whilst allowing neutrinos to pass through to the various detectors. Neutrinos are then first sampled 280 m downstream of the target by INGRID and ND280.

on the target over time. The charged pions decay via the process

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{3.1}$$

to the desired muon neutrino, but also produces an anti-muon which can undergo further decays. A beam dump is positioned to sweep up as many of these muons as possible before they can decay to electrons in the process

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \tag{3.2}$$

which can generate an electron neutrino background in the beam. Similarly, the decay of kaons in this volume will also produce electron neutrino backgrounds via the process

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e \tag{3.3}$$

although kaons may also decay to muons making modelling of the muon neutrino flux more complicated at higher energies where this becomes the dominant source.

3.1.1 External hadron measurements at NA61/SHINE

To correctly predict the flux of the T2K neutrino beam the relative contributions and kinematics of the neutrino parents must be understood. Modelling of hadron production in the graphite target is a non-trivial problem in part due to a lack of data on p-C interactions in the 30 GeV range used at T2K, and so measurements of the hadron production from a graphite target are needed. Since direct measurements of hadron kinematics is not possible in the T2K target hall, external measurements from the NA61/SHINE [101] experiment are used to constrain these models. For analyses up until the 2020 result, T2K used "thin target" data and simulated secondary interactions within the target using Monte Carlo simulations and the FLUKA package [102]. For the 2020 analysis, hadron production in a T2K-replica 90-cm-long target was measured at NA61/SHINE and was used to reduce the uncertainty in the neutrino flux modelling in FLUKA.

A prediction of the FHC muon neutrino and electron neutrino flux at SK is shown in figure 3.5, broken down by the parent particle. As can be seen from the difference between the relative contributions to the ν_μ and ν_e spectra, the majority of ν_μ come from the decay of pions, though significant contributions in the higher energy tail of the flux come from the decays of kaons. The dominant source of the inherent ν_e background (i.e. ν_e not from oscillations but rather intrinsic to the beam) at the far detector come from the secondary decay of muons, but again with significant contributions from charged kaons as well as some from the K long component of the neutral kaon system.

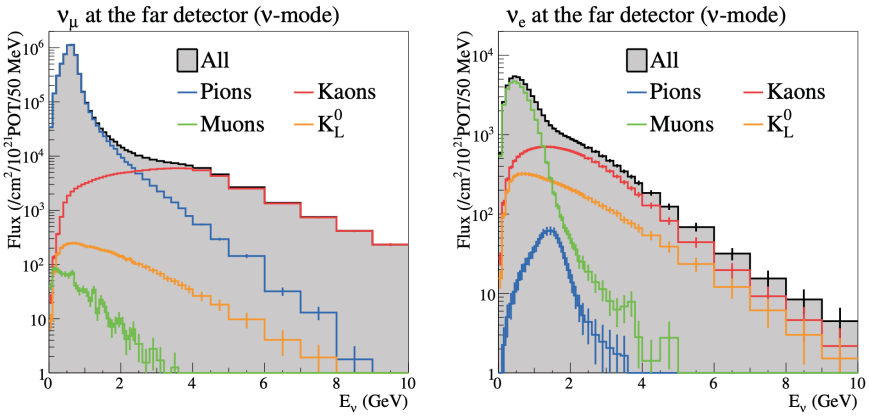

Figure 3.5: The predicted muon and electron neutrino spectra at SK broken down by neutrino parent for the FHC beam mode constrained by measurements of hadron production from a replica T2K graphite target in the NA61/SHINE experiment at CERN [101].

3.2 INGRID

The Interactive Neutrino GRID (INGRID) detector is an on-axis iron scintillator detector designed to measure beam intensity and direction during run time. It consists of 14 identical modules (figure 3.7) arranged in vertical and horizontal stacks of seven modules forming a cross as shown in figure 3.6, allowing for beam direction information to be extracted from fits to the relative intensity of the beam in each module. Each module is constructed from alternating layers of scintillator bars and iron plates. An additional two modules have been used in different configurations over the years to provide measurements of the beam intensity off the vertical and horizontal axes.

INGRID gives the beam centre measurement a precision of 10 cm, or 0.4 mrad at the ND280 complex. This beam direction and intensity measurement is crucial for correctly modelling the off-axis beam at different angles. The total height and width of the detector modules spans 10m $\times$ 10m and the centre of the two overlapping modules is defined as 0° from the direction of the primary proton beamline. By fitting a Gaussian distribution to the interaction rates in each of the modules for both the vertical and horizontal stacks as shown in figure 3.8, the beam centre can be accurately determined, and used to produce proper inputs for the neutrino oscillation analysis.

3.2.1 Proton Module

In addition to the standard modules which make up the INGRID detector, there is a similar module where the iron plates are replaced with additional layers of hydrocarbon scintillator bars. This is to allow for measurements of on-axis rates of neutrino interactions which produce shorter tracks such as those from protons or pions, which can not be reconstructed in the other modules due to

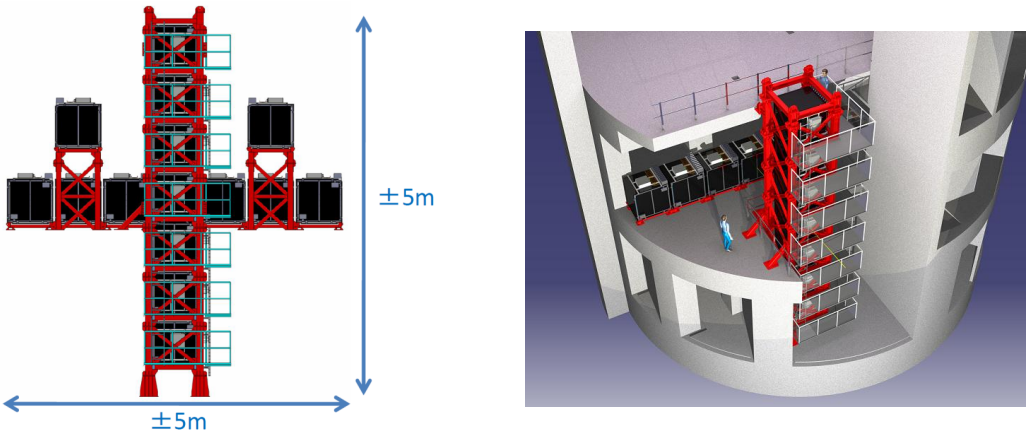

Figure 3.6: Left: INGRID module cross arrangement allows for precise measurement of the beam direction by measuring the number of interactions in each of its modules. The beam centre is through the two overlapping middle modules, one in the vertical stack of modules, and one in the horizontal modules. Right: The position of the INGRID modules in the ND280 complex. A hole in the floor plan of the higher levels of the ND280 pit is made to make way for the vertical stack of INGRID modules.

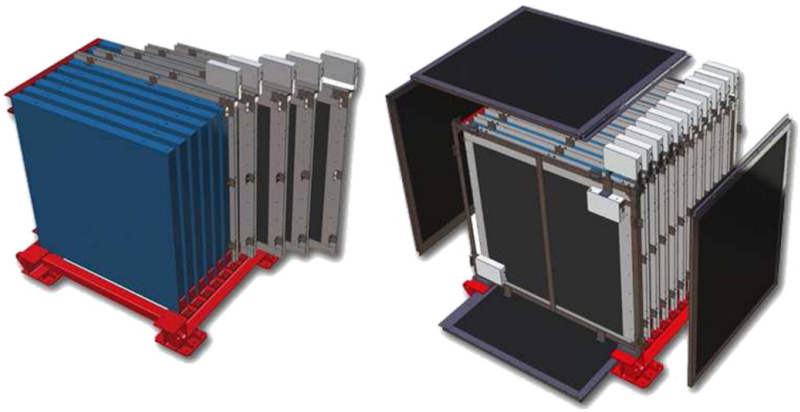

Figure 3.7: Exploded design of an INGRID module showing scintillator target (blue) and iron plates (grey) as well as black veto planes.

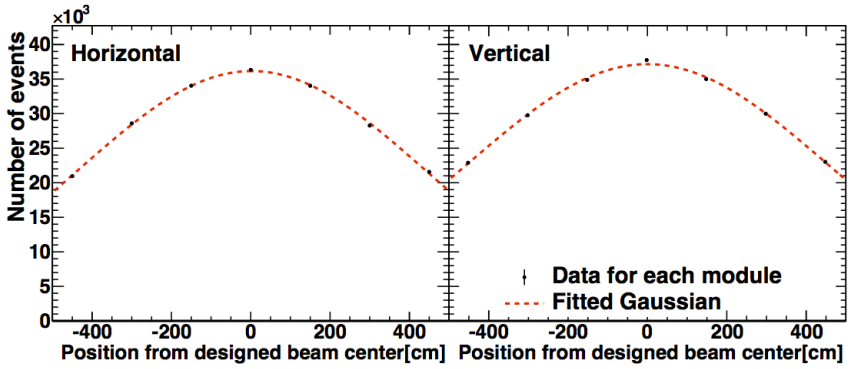

Figure 3.8: Neutrino intensity profile at INGRID in FHC beam mode with a nominal 250 kA horn current. Fits to the beam intensity in each of the modules allows for an accurate measurement of the beam direction.

the iron plates reducing the resolution of the detector. The proton module no longer is used as part of INGRID, but now functions as a subdetector of the WAGASCI/BabyMIND detector.

3.3 ND280

The ND280 detector is T2K's primary off-axis near detector which studies neutrino interactions at the same energies expected at SK. Following a similar detector design to those used in particle collider experiments, the ND280 is a magnetised tracking detector allowing for discrimination between particles of opposite charges and different energy loss ($\mathrm{d}E/\mathrm{d}x$) values for particle identification (PID) according to the characteristic Bethe-Bloch curves for each particle in a given medium. The charge discrimination at ND280 allows for a measurement of the wrong-sign contamination in the beam, and the PID allows for measurements of pion production rates as well as the electron neutrino contamination prior to oscillations. Constraining the wrong-sign component is important as it is a significant contamination in the antineutrino beam and SK cannot distinguish the sign of the outgoing charged lepton and therefore cannot distinguish between neutrino and antineutrino interactions. Thus this is an essential constraint on the measurement of CP violation. The ND280 and its subsystems are discussed in detail in chapter 4.

3.4 WAGASCI and BabyMIND

The WAterGrid-SCIintilator-Detector WAGASCI is an off-axis near detector in T2K's ND280 suite which sits at a 1.5° angle to the beam rather than the 2.5° angle of ND280 and SK, giving it access to slightly higher energy neutrinos than these other detectors. WAGASCI aims to help

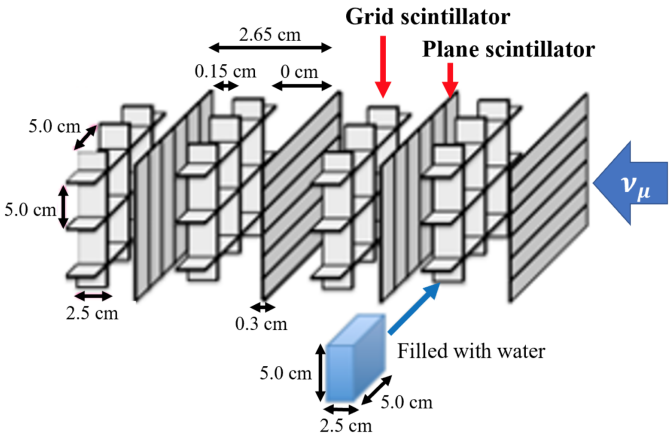

Figure 3.9: A 3D grid of scintillator strips make up the instrumentation of the WAGASCI modules, between which either water or scintillator fills the gaps, allowing for more of the detector mass to be made up of target material.

reduce the uncertainties in neutrino interaction cross-section differences on CH and water whilst providing a larger angular acceptance than that of ND280. Carbon and oxygen cross-section systematic uncertainties are non-cancellable between ND280 and SK so constraining this with better C and O interaction separation will benefit the oscillation analysis. A 3D plastic scintillator grid structure (figure 3.9) is filled with the desired target material, either plastic scintillator (CH) or water, maximising the ratio of target material to instrumentation (80%) whilst allowing for tracking and reconstruction even at high angles to the beam, something ND280's other detectors the PØD and tracker region struggle to do, as discussed in chapter 4

The central water and scintillator modules of WAGASCI are surrounded on two sides by wall muon range detectors (Wall-MRDs). The MRDs are magnetised steel scintillator sampling detectors which are placed at a small (20-65 cm) distance away from the target section to allow for time of flight measurements. The grid modules can be positioned either side of the repurposed INGRID proton module and NINJA experiment as shown in figure 3.10. The NINJA experiment is an emulsion-based experiment which is not part of WAGASCI or T2K.

3.4.1 BabyMIND

Downstream of the main WAGASCI detector is the Baby Magnetised Iron Neutrino Detector (BabyMIND), a muon spectrometer used to accurately measure the charge and momentum of muons exiting the main WAGASCI modules. BabyMIND is made of magnetised plates of iron separated by plastic scintillator, with the separation of the iron plates (shown in figure 3.11) optimised to measure the momentum of muons at the energies expected at 1.5° off axis. Early data results show a POT collection efficiency of 97% [103].

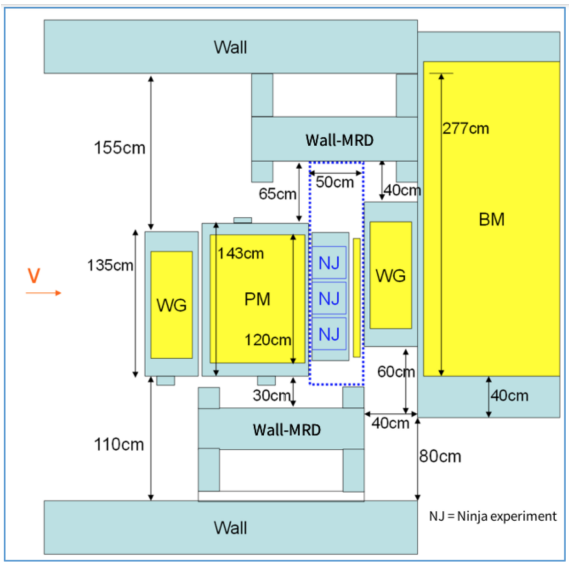

Figure 3.10: Layout of the WAGASCI (WG) and BabyMIND (BM) detectors during operation as well as the NINJA experiment (NJ) which is not a part of WAGASCI or T2K.

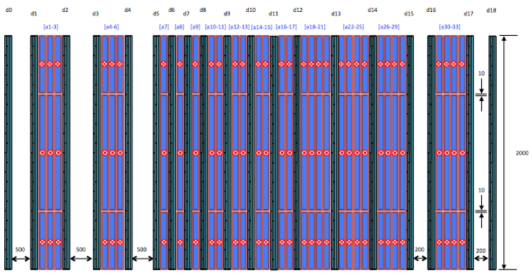

Figure 3.11: The arrangement of the BabyMIND iron plates (red and blue) and plastic scintillator planes (grey) are positioned to optimise muon momentum resolution down to 500 MeV/c

3.5 Super-Kamiokande

The Super-KamikokaNDE (SK) experiment is a 50-kton water-Cherenkov neutrino telescope in the Kamioka mountain region of Japan [104]. An experiment in and of its own right, SK has measured neutrino oscillation physics from atmospheric and solar neutrinos. It has also conducted searches for proton decay and is part of the SuperNova Early Warning System (SNEWS). Whilst SK has made these measurements as its own experiment, the detector serves as the far detector for T2K and is situated 295 km downstream of the T2K target and positioned at an off-axis angle of 2.5°. The detector consists of a stainless steel tank measuring 39.3 m diameter and 41.4 m tall filled with ultrapure water with 1000 m of rock overburden to block out the vast majority of cosmic ray muons. The tank is split into Inner Detector (ID) and Outer Detector (OD) volumes separated by an instrumented steel frame and the light-proof material Tyvek®, manufactured by Dupont, which provides an optical barrier between the two regions. The steel frame is instrumented with photomultiplier tubes (PMTs) produced by the Hamamatsu-Photonics company. The PMTs collect Cherenkov light emitted by charged particles moving above the speed at which light propagates through water. A cutaway diagram of the detector is shown in figure 3.12. The ID is lined with 11,129 inward-facing 50-cm diameter PMTs providing 40% coverage of the inner detector and serves as the target region of the detector for the oscillation analysis. The OD has 1,885 outward-facing 20-cm PMTs which are attached to 60-cm-square wavelength-shifting plates and acts as a veto for muons produced by interactions in the surrounding rock and any unattenuated cosmic ray muons which make it through the rock overburden. The wavelength-shifting plates allow the smaller PMTs and sparser coverage to still collect enough light to reject any background muons. A fiducial volume cut is used to reduce backgrounds from neutrinos which enter or leave the ID in such a way as it would prevent full reconstruction of the neutrino energy. This is applied as a distance cut from the wall of the ID. The PMTs are arranged into “super-modules” which contain six of the larger ID PMTs and two of the smaller outer PMTs. These are then stacked in towers which have readouts fed up to the top of the detector. Above the detector there is a low-energy linear accelerator (LINAC) for detector calibration, and electronics huts where the PMT signals are read out.

3.5.1 Photomultiplier tubes

A photomultiplier tube or PMT is a photosensor designed to collect photons over a large surface area and amplify the electronic signal produced. An incident photon on a photo-cathode under voltage releases an electron of equal energy to that of the photon minus the work function of the cathode. The cathode of one of the ID PMTs is visible in figure 3.13. The primary electron is

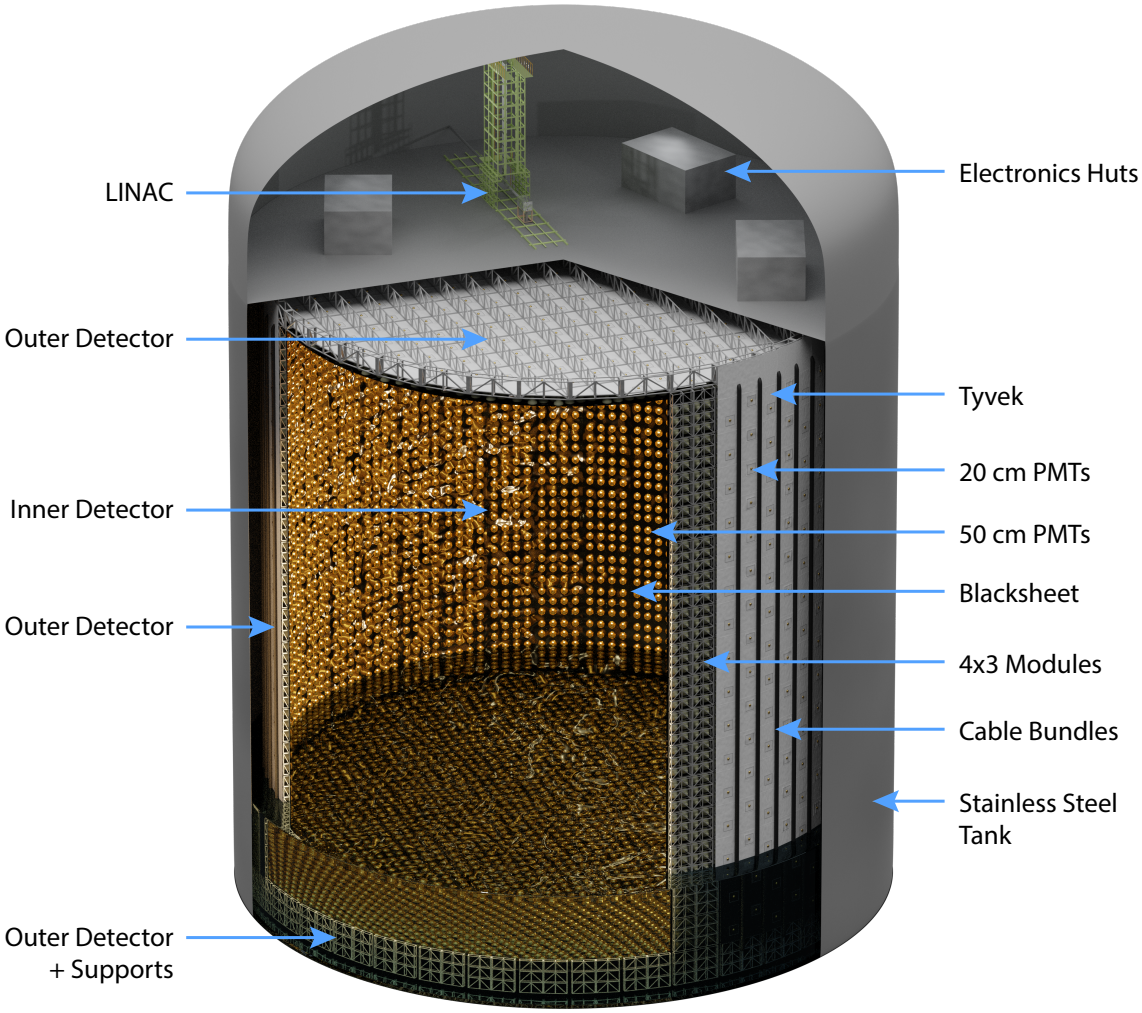

Figure 3.12: A cutaway view of the Super-K detector. The detector is split into two regions, the Inner Detector (ID) and the Outer Detector (OD). The ID contains the fiducial volume of the detector and the OD is used as an active veto on backgrounds from interactions in the rock. The PMT signals are read out by electronics in the dome above the detector. A steel support structure holds the PMTs and their readout cables and a lightproof building material, Tyvek®, separates the ID and OD forming a light proof barrier. The ID facing side is black to absorb any light leaving the ID, and the OD side is white to best reflect any light back in the OD. Water can be drained from the detector through the pipes at the base of the tank. Double sided white Tyvek®is also used to cover the walls of the steel tank to ensure as much of the light in the veto region is reflected and captured by the sparser OD PMTs. Additional semi-reflective Blacksheet surrounds the PMTs on the ID side of the structure to reflect light towards the PMTs and prevent it reaching the OD. During draining, the PMTs at the surface of the water can be accessed via a gondola from the top of the detector lowering workers down to a floating platform. *Rendering courtesy of Alex Goldsack* [105].

Figure 3.13: Left: The PMTs of the inner detector as photographed from the steel structure during 2018 open-tank cleaning. Right: Detail of the PMT inner photo-cathode.

focused by electrodes and accelerated onto a dynode which produces secondary electrons which are in turn accelerated producing a cascade of electrons. This allows a charge to flow in a circuit, which is the signal. The signal can then be digitised and collected with signals from other PMTs to reconstruct an event based on timing, photon energy and spacial position of hits.

3.5.2 Cherenkov Radiation

When a charged particle moves through a dielectric medium (such as water) at a greater velocity than the phase velocity at which light can propagate through that medium, it radiates energy in the form of photons in a similar manner to sonic booms produced by objects travelling faster than the phase velocity of sound. As the charged particle passes through the medium the charge polarises the dielectric and excites molecules into higher energy states. The molecules de-excite and the dielectric returns to its depolarised state releasing photons in a spherically symmetric wavefront. For a particle with a speed below the phase velocity of light, successive wavefronts will always be within each other; however, if the particle is able to move faster than the wavefront in some direction, successive wavefronts combine into a plane wave producing a cone of radiation at an angle to the direction of travel determined by the speed of the particle and the phase velocity of light in the dielectric.

3.5.3 Particle identification

Unlike ND280 in which track curvature in a magnetic field and track length are used to determine particle type and charge sign, SK does not have a strong magnetic field[1] or the ability to accurately

[1] Although SK does not have a field capable of noticeably impacting the track of any charged particles in the detector, a series of magnetic coils are used to ensure inhomogeneity from the Earth's magnetic field does not interfere with the calibration of the PMTs

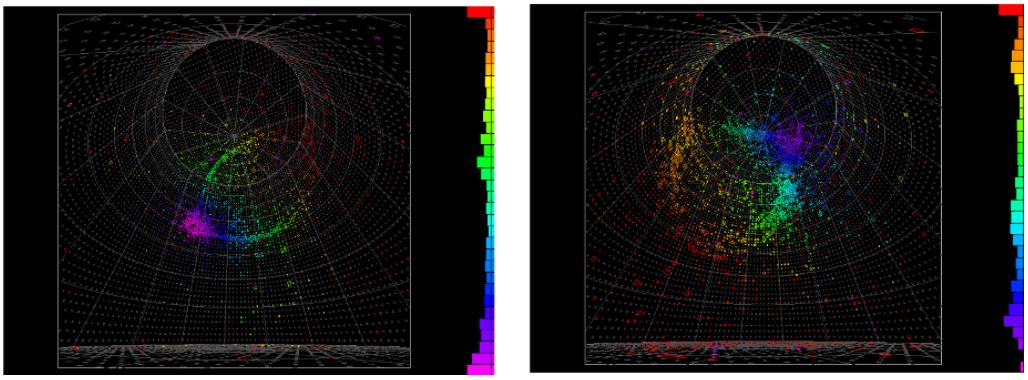

Figure 3.14: Muon (left) and electron (right) neutrino event candidates in Super-K during a T2K beam spill. The electron scatters as it travels producing a diffuse or fuzzy ring of PMT hits, whereas a muon is minimally ionising and will produce a sharp distinct ring.

measure curvature of the path travelled by particles in the water. SK can, however, distinguish between electron-like and muon-like particles from their signature Cherenkov light rings. By distinguishing between the shape of the ring pattern left on the PMTs by different particles as shown in figure 3.14, minimally ionising particles like muons which leave well defined rings can be separated from showering or scattering particles like photons and electrons which leave more diffuse or fuzzy rings.

3.5.4 Gadolinium doping and the SK 2018 open-tank cleaning

In an effort to search for supernova relic neutrinos in the 10-30 MeV range which is dominated by reactor and atmospheric backgrounds, SK utilises neutron capture to better distinguish between electron neutrino and electron antineutrinos. A coincidence of an initial or "prompt" positron from the antineutrino on proton interaction via inverse beta decay followed by the 2.2 MeV gamma emission from the capture of the neutron produced by the initial interaction on another proton (hydrogen in water) is used. By adding Gadolinium (Gd) to the water in SK, the thermal-neutron capture cross section can be increased significantly ($\sim 5\times$ the cross section on protons), and the higher energy 8 MeV gamma cascade from the decay of the excited Gd nuclei can be exploited for better detection.

In preparation for this, the detector was opened in 2018 and drained to allow for cleaning of the inner tank and steel structure and the application of sealant to the tank to prevent any possible leaking of the gadolinium sulphate into the surrounding environment. Draining or filling the 50-kton tank is a significant undertaking which is not done often, and so this also provided an opportunity to replace broken PMTs and do general maintenance of the typically submerged parts of the detector.

Maintenance in the tank relies on lowering the water level and using floating platforms for researchers working on the maintenance to stand on. In the ID there is a square platform that can be moved around the inner volume, as well as several inflatable boats to access the inner PMTs. Most of the work, however, takes place in the OD, and so the full circumference of the OD has connected floating platforms allowing workers to navigate the entire OD. The water-dropping schedule allowed for a three day cycle of cleaning until the water was dropped to access a new depth of the tank.

In November 2001, during the filling of the detector, an implosion of one of the SK PMTs caused a shock wave triggering a chain reaction of PMT implosions, destroying over 6000 of the submerged PMTs . The remaining PMTs were rearranged to cover the detector volume and covered with acrylic shields to reduce pressure on the glass of the PMTs. Subsequently, the missing PMTs were replaced in a later tank opening. In the aftermath of the disaster the tank was pressure-washed to remove the broken glass and debris from the detector, though this was not 100% effective, whilst the PMTs were rearranged. In the 2018 open-tank period, significant glass debris was still in many parts of the tank sitting on the steel structure holding the PMTs. In addition, significant diesel fume residue from the initial excavation and construction of the detector covered much of the surfaces within the detector. Whilst pressure washing had removed much of these from the outer tank in the initial opening in 2001, the internal steel structure was much more difficult to clean due to its geometry and so much of the cleaning of this structure was done in the 2018 open-tank period.

Despite the ultra-pure water used in SK, over the decades of operation the steel structure and outer tank plates have begun to rust in places. To extend the life of the detector and reduce the possibility of leaks due to rusting, efforts were made to remove as much of the rust in the tank as possible during this period. The three day cleaning cycle involved a water drop of roughly 1 m, allowing for access to a new half of a PMT super-module of each tower. Old Tyvek was first removed from the tank walls and steel structure to access the PMTs. Due to the need for time to properly apply the sealing compound to the joints of the outer plates, the first job was the cleaning and rust removal on the outer tank plates. Whilst much of this could be done with an electrolysis method which quickly removes rust from large areas with ease, this was not a good method to apply to the welds of the plates, as it can weaken the weld, and so this was done with small steel brushes and toothbrushes. Sealing compound could then be applied to these welds and joining plates which required time to set and dry. For the inner structure, first vacuuming was done to remove many of the remaining glass fragments. This was followed by wiping the structure with damp cloths, before waterproof industrial tape was repeatedly patted down on the surfaces

to remove the diesel residue and remaining glass particulates. Finally the electrolysis machines were used remove rust on the steel structure supporting the PMTs. New Tyvek was attached to the structure and sealed using cable ties, plastic clothes tags, and stainless steel staples. Care was taken to make sure that the overlapping sheets are folded to make a light-proof seal before being stapled, as well as to avoid piercing the waterproof coating of any wires from the PMTs when using the clothes tag gun. Outer PMTs could be replaced and wired whilst the structure was accessible. Once all of this was complete, the water level could be lowered again. Finally the outer wall Tyvek was replaced after the tank was fully emptied.

Chapter 4

ND280

The ND280 is the primary off-axis near detector of the T2K experiment [15]. Its purpose is to conduct detailed studies of neutrino interactions as well as characterise the unoscillated neutrino beam. As a magnetised tracking detector ND280 is able to extract more information about neutrino events than possible in INGRID or SK through more sophisticated particle identification and $\mathrm{d}E/\mathrm{d}x$ measurements. As well as making neutrino interaction cross-section measurements at the energies of the T2K beam, the data from the ND280 is used to constrain T2K's nominal neutrino interaction and beam flux models as part of the long-baseline neutrino oscillation measurements and so this detector will be the focus of this thesis.

The ND280 detector (figure 4.1) is comprised of a dedicated π^0 detector (PØD) and a tracker region which are held in a steel frame or "basket", and are enclosed by an electromagnetic calorimeter (ECal) which sits just inside the solenoid magnet. One module of the ECal is inside the basket as the most downstream subdetector of ND280. The magnet yoke is instrumented with a Side Muon Range Detector (SMRD) which is primarily used for cosmic ray muon rejection and detector calibration. Due to the expense of manufacturing a magnet, in particular the large iron flux return yoke, the ND280 magnet was refurbished from magnet of the UA1 experiment [4]. This constrained many of the design choices for the detector based on the physical dimensions of the magnet. The magnet is constructed in two halves to allow access to the inner subdetectors and the attached ECal modules are also split to facilitate this.

The tracker region is made up of two plastic scintillator Fine-Grained Detectors (FGDs) which are sandwiched between the three gaseous Time Projection Chambers (TPCs). The FGDs provide instrumented target material and the second (more downstream) detector, FGD2, has water layers interleaved between layers of plastic scintillator bars to allow for measurements of neutrino interactions on water which is the target material of the far detector.

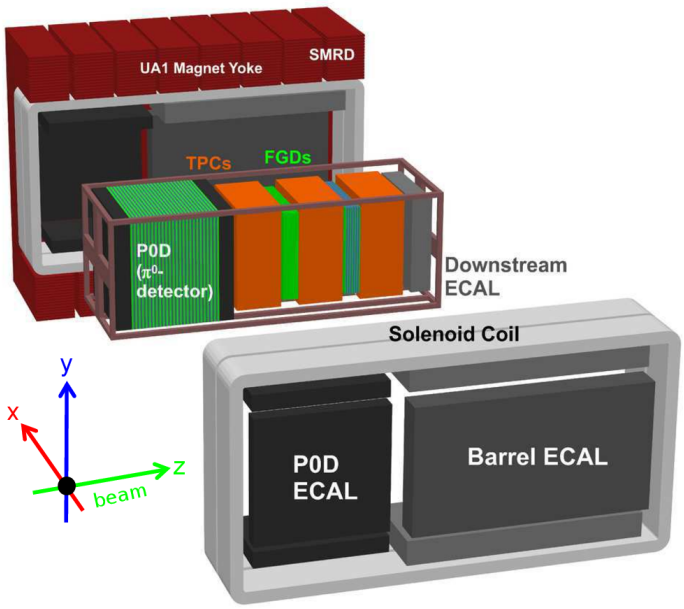

Figure 4.1: Exploded view of the ND280 showing its various subdetectors. The ND280 has three primary target subdetectors, the PØD and two FGDs, which provide both the target material mass (plastic scintillator, water, brass) and 3D tracking through scintillation light. Downstream of each of these subdetectors is a gaseous argon TPC, which provides detailed track curvature information for tracks which leave one of the target subdetectors and enters a TPC, which is typical of the muon tracks in CCQE interactions at ND280. The FGDs and the TPCs form the tracker region of the detector and are downstream of the PØD in the T2K beam. The analysis presented in chapters 6-8 uses primarily tracker information. Wrapped around each of these detectors, and downstream of the tracker region, are the thirteen modules of the ECal. The ECal is a lead-scintillator sandwich calorimeter designed to induce EM showering from photons and electrons, whilst being able to tag exiting muons as MIPs which leave only one track without showering. All of these subdetectors sit within the ND280 magnet which is comprised of an aluminium solenoid coil and an iron flux return yoke. The magnet provides a 0.2 T magnetic field for charge discrimination and dE/dx measurements in the TPCs. The flux return yoke is instrumented with the SMRD, a series of scintillator paddles which can tag muons passing through the magnet and is used for cosmic ray muon rejection [15].

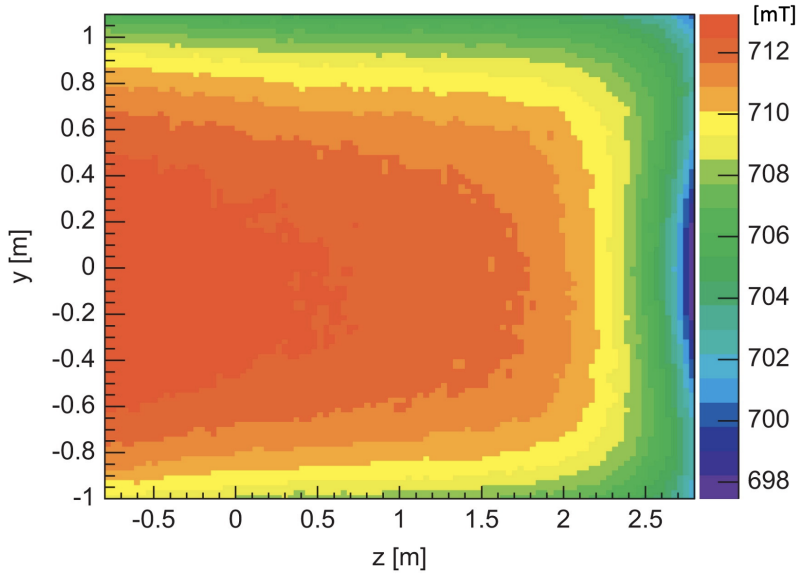

Figure 4.2: Measurements of the magnetic field within the central vertical slice of the tracker region of the ND280. The neutrino beam enters in from the negative z (left) and the DsECal is at z =2.8 m, on the right of the figure.

4.0.1 The ND280 magnet

All of the subdetectors of the ND280 detector are encased in the solenoid magnet. The magnet produces a uniform 0.2 T field across the tracking region of the detector, allowing for particle identification and momentum measurements from the radius of curvature of the particle tracks. The magnet is comprised of aluminium solenoid coils with a cross-section measuring 5.45-m square with a central 23-mm borehole for water cooling, enclosed within a 850-ton iron flux-return yoke. The magnet produces a dipole field within the basket region of the magnet and the flux external to the basket is directed back around the magnet to make the internal field as uniform as possible whilst magnetising as little of the region outside the ND280 as possible. The inner dimensions of the magnet measure 7.0 m $\times$ 3.5 m $\times$ 3.6 m and the outer dimensions measure 7.6 m $\times$ 5.6 m $\times$ 6.1 m. The magnet is made of two identical halves which are able to separate to allow access to the inner subdetectors. Each half consists of two coil loops and 8 “Cs” of the magnet yoke as can be seen in figure 4.1.

Field map

A computer-control-driven Hall probe system was used to make measurements of the field map in three dimensions within the basket region, which can be seen in in figure 4.2. These measurements were made with the field strength at a nominal value of 0.07 T and so must be scaled to the operational nominal value of 0.2 T taking into accounts hysteresis and saturation effects. An

analysis of this scaled map and fits to data were done to evaluate the uncertainty on this mapping using residuals between the data fit and the Hall probe measurements. The final uncertainty on the field value was 0.2 mT at the nominal value of 0.2 T. This allows for a 2% uncertainty on the momentum measurement of particles below 1 GeV/c.

Modelling inside the flux return yoke

Where the field cannot be directly measured such as within the iron of the flux return yoke, a field simulation is used. Prior to 2018 a simple uniform field with a total flux equal to that inside the detector was used to model the inside of the yoke, though a finite element analysis conducted in COMSOL Multiphysics [106] was available to more accurately simulate the field. This could impact the modelling of cosmic ray muons and sand muons which enter the detector from interactions outside ND280, as the finite element analysis model has far more varied field strengths, including regions near the internal corners of the yoke where strengths of over 1 T are predicted. The COMSOL simulation was validated in 2018, Comparisons of the COMSOL MC and the uniform-field MC on cosmic ray muon data showed that for most low-level observables in the ECal modules where the greatest differences were expected, the two MC predictions were closer to each other than to the data, and so the more sophisticated flux model has been used by default since 2018.

4.1 π^0 detector

The pi-zero detector (PØD) is a subdetector dedicated to the measurement of π^0 production in neutrino-nucleus interactions, in particular via the neutral current (NC) background for which there is no flavour-taggable charged lepton, which contributes a significant background to the ν_e appearance channel of the T2K oscillation analysis [107]. Whilst NC interactions have no charged lepton in their final state they may produce observable final-state particles if the momentum transfer to the target nuclei is sufficient to produce new particles coherently or resonantly, or induce deep inelastic scattering. The most common coherently- or resonantly-produced particles at T2K beam energies are pions. Neutral pions decay with a $(98.823 \pm 0.034)\%$ [108] branching ratio to two photons, and since photons can electromagnetically shower, they are indistinguishable from electrons in SK. Asymmetric decays of these neutral pions where one photon is not reconstructed can be mistakenly classified as a CCν_e-like event in SK. NC events may also offer a window into the total flux of the beam as the interaction of tau neutrinos is not restricted by the mass of the tau as in the CC case.

The PØD is a plastic scintillator tracking detector with water and brass layers in between the scintillator layers to measure interactions on water and induce electromagnetic (EM) conversion

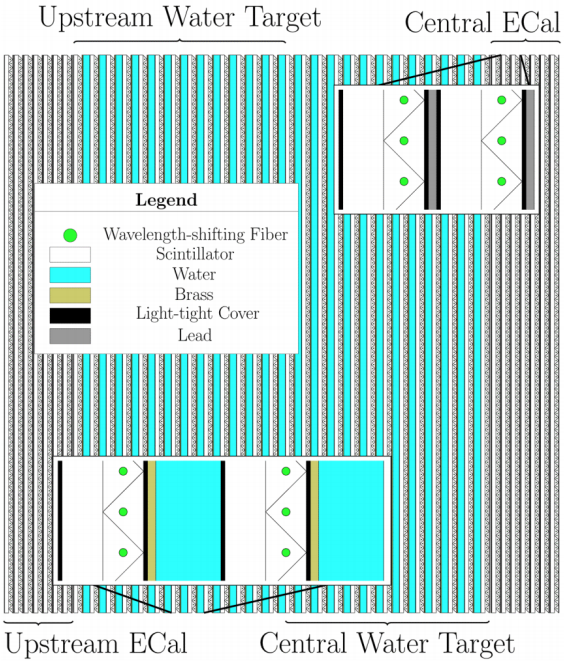

Figure 4.3: Schematic of the ND280 PØD or π^0 detector.

of the π^0 decay photons. Each of the scintillator bars in ND280 have a wavelength shifting (WLS) fibre running through the centre which feeds the light into the optical readouts at an optimal wavelength for detection efficiency. The plastic scintillator is made of polystyrene doped with 1% PPO and 0.03% POPOP, with a co-extruded thin layer of TiO_2 which reflects any light back into the bar to prevent loss and prevent light from bleeding between the bars. The PØD's scintillator bars are isosceles triangles in cross-section with a 33-mm base, 17-mm height with a 1.5-mm hole through the centre through which a wavelength shifting fibre is threaded. The scintillator bars have a 0.03-mm-TiO_2 coating. A discussion of the scintillator bar arrangement is given in section 4.2. Water layers provide target material to study neutrino interactions on oxygen, which is key to extrapolating model constraints to the far detector.

The PØD is enclosed on all sides by dedicated ECal modules. The upstream and central (central as it sits between the PØD and the tracker) ECals are similar in design to the PØD tracking area but with lead instead of brass and water, and cap each end of the PØD sitting inside the basket. The six other ECal modules are attached to the magnet similarly to the rest of the barrel ECal which will be discussed in section 4.4.

4.2 Fine-grained detectors

The two Fine-Grained Detectors [109], FGD1 and FGD2, are the primary targets for the flagship T2K analysis. Comprised of plastic scintillator bars arranged in planes with alternating orienta-

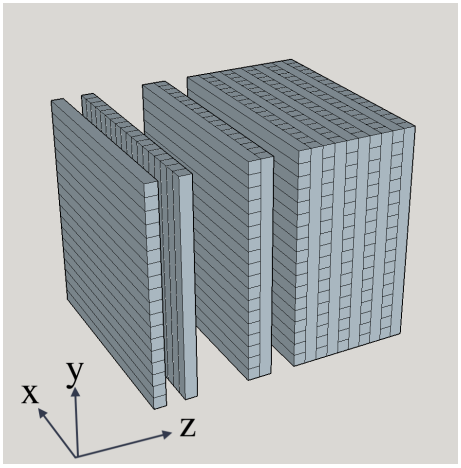

Figure 4.4: Partially exploded view of typical scintillator bar arrangement. Bars are aligned in the horizontal (X) or vertical (Y) directions, and pairs of XY alternated layers are stacked along the third axis (Z).

tions in a similar way to the PØD and WAGASCI, the FGDs provide active three-dimensional tracking within the target material of the detector. FGD2 contains water layers which can be filled or drained allowing for modelling of events on oxygen to be constrained at ND280.

Each FGD scintillator bar is made from extruded polystyrene in a square cross-section and measures 9.61 mm $\times$ 9.61 mm $\times$ 1864.3 mm with a WLS fiber, threaded through a 1.5-mm diameter hole along the bar length connected to a MPPC which collects the scintillation light and digitises the analogue light signal.

Typically the scintillator bars are arranged in layers such that the bars in successive layers are orthogonal to each other as shown in figure 4.4. For the FGD (as well as the PØD and DsECal) these layers are in planes of XY in detector coordinates (transverse to the beam direction) and stacked along the Z axis (parallel to the beam direction). This arrangement, coupled with the use of the TPCs, allows for excellent vertex and track reconstruction for events with forward going tracks. High angle tracks are more difficult to reconstruct because much of the length of a track may be restricted to the same layers, or even the same bars, and there is less chance of the particle entering a TPC. As such, for many ND280 analyses, including the fit to ND280 data for the oscillation analysis, the data samples are restricted to the forward going region where detector systematic effects are best understood.

4.2.1 Event reconstruction in the FGDs

In order to reconstruct events in the FGDs, hits are first digitised by fitting a polynomial to the waveform of the light pulse collected by each MPPC. The light pulse height is converted into

photo-electrons within some time period.

Hits with > 5 photo-electrons within 100 ns are then clustered using a radon transform. In a radon transform the Cartesian position of the scintillator bars in each (x,z) plane correspond to a line in polar coordinates (r,θ). Conversely, the hits which lie in a straight line in (x,y) will have lines in (r,θ) which intersect at a point. Points in (r,θ) coordinates with multiple hits can be used to identify straight line track segments. Hits which lie too far from this line are rejected to clean up the track.

The interaction vertex can be found by plotting the best fit line between points in the radon space which corresponds to the position of the vertex in Cartesian coordinates. Additional requirements such as proximity of hits between tracks are used to ensure hits truly are from the same track and tracks truly connect to a common vertex.

Tracks in each plane are then matched based on start and end positions in the Z plane as well as information about charge deposition and timing when multiple combinations are possible.

4.2.2 Iso-FGD particle identification

Whilst particle identification (PID) is preferably done with $\mathrm{d}E/\mathrm{d}x$ information from track curvature measurements in the TPCs, FGD-isolated tracks where charge deposited is entirely contained within the FGD, require different PID methods. Track length, shape, energy deposition, and timing can be used to distinguish between short tracks such as those made by protons, and longer tracks pions, electrons and Michel positrons from the decays of charged pions.

4.3 Time projection chambers

Three gaseous argon time-projection chambers or TPCs are used to provide accurate information about the direction and curvature of the particle tracks [110]. Applying a uniform magnetic field across the tracker allows particle identification and momentum measurements from radius of track curvature and energy loss characteristics for various particles.

The chamber is constructed as a double box, where the inner box forms a field cage, and the outer box is grounded. These are insulated by carbon dioxide as can be seen in figure 4.5. An electric field is generated by a high-voltage potential difference between a central cathode and the field cage. The chamber is filled with a mixture of argon, tetrafluoromethane and isobutane (Ar : CF_4 : iC4H10 (95 : 3 : 2)), or "TPC gas". A charged particle passing through the gas will ionise the TPC gas, and under the electric field the ions produced will drift from the cathode to the MicroMegas readout pads (discussed below) which collect the charge and timing information

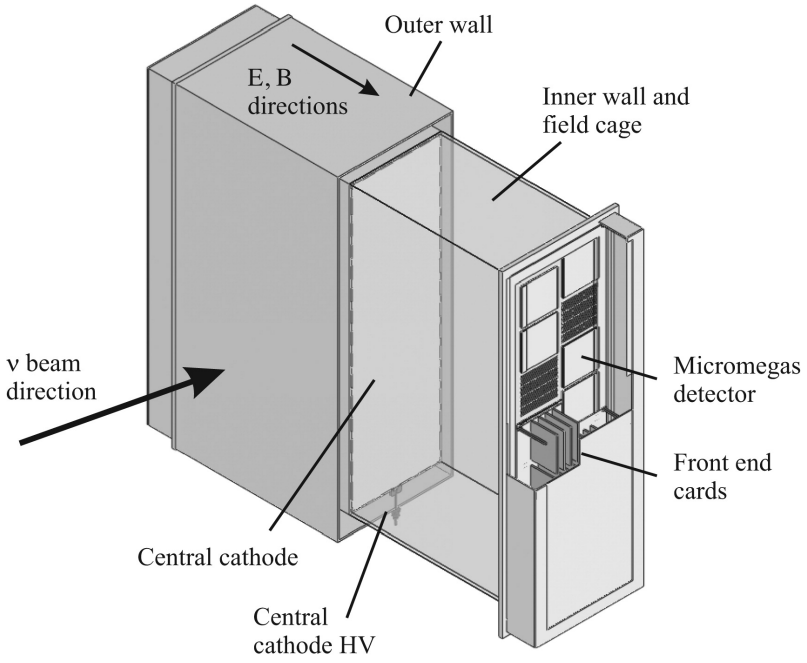

Figure 4.5: Diagram of the time-projection chamber design at ND280. A central cathode at high voltage generates an electric field under which ionised tracks of an argon gas mixture drift towards micromegas readout pads. High angular resolution can be achieved from the precise drift time difference measurements of each part of the track.

which can be used to reconstruct the 3D track direction and curvature to high precision from the time each part of the ionised track takes to reach the readout pad.

4.3.1 MicroMegas detector readouts

The charge collection and readout of the TPCs is done by Micro-Pattern Gas detectors or MicroMegas detectors which are finely segmented anode pads. The three TPCs have two readout planes each with 12 MicroMegas per module. Each of the 12 MicroMegas measures 342×359 mm^2 covering 9 m^2 of readouts for the 72 modules. The charge collected by the 1728 pad divisions is read out by printed circuit boards using time-binned switched capacitor arrays to determine the drift time for segments of the track which is used to reconstruct the position of the track transverse to the beam/detector axis.

4.3.2 Tracker reconstruction with TPC tracks

Whilst FGD-iso reconstruction is possible for tracks which do not exit the FGD into the TPC, the PID efficiency, charge discrimination and momentum measurement and measurement of the angle relative to the beam of tracks can be greatly improved if the FGD track is matched to a TPC track in the global reconstruction that collates the information from all of the subdetectors for a

given event time window. This is particularly important for the charged lepton exiting the event as the T2K beam and interaction models are constrained by fits to ND280 data binned in lepton kinematics p_μ and $\cos\theta_\mu$, as discussed in chapter 6.

4.4 Electromagnetic calorimeter

The tracker region of the ND280 detector is wrapped in a lead-scintillator sandwich calorimeter, or electromagnetic calorimeter (ECal), in order to distinguish between showering particles (electrons and photons) and minimally ionising particles or MIPs (e.g. muons and pions) [111]. The ECal is split into the PØD ECal sections that surround the PØD, the barrel ECal sections that surround the tracker region, and the downstream or DsECal[1]. The ECal modules sit just inside the magnet outside of the basket, with the exception of the DsECal which sits just downstream of TPC 3 within the basket. To allow for opening the detector, the top and bottom PØD and barrel ECals are made up of two modules each split down the centre so that the magnet and ECal can be separated from the inner subdetectors in the basket. As such there are 13 ECal modules in total: 6 PØD, 6 barrel, and the DsECal. The ECal modules are made of alternating layers of plastic scintillator bars for instrumentation and lead absorber to encourage EM showering of photons and electrons within the ECal. The ECal scintillator bars are wider than those of the FGD with a cross section measuring a 4.0 cm $\times$ 1.0 cm with a 0.25 mm layer of TiO_2. The central hole of the ECal scintillator bars for the WLS fibre is elliptical and measures roughly 2.0 mm $\times$ 3.0 mm with some variation due to the extrusion process. Where possible the WLS fibre is read out from both ends by MPPCs and is otherwise covered with an aluminium mirror coating to reflect light back towards the readout. The ECal is used to both help distinguish between particles with tracks leaving the tracker region and also to convert photons which could otherwise leave no signal in the tracker region. EM conversion is key to identifying NC backgrounds which may contribute to the electron like samples at SK such as NCπ^0 decays as well as background resonant processes to the CCQE-like samples.

4.4.1 Downstream ECal

The DsECal is made of 34 layers of 50 2.04-m-long scintillator bars and a 1.75-mm-thick layer of lead. This gives a total lead thickness of $10.6X_0$ radiation lengths. In the DsEcal the scintillator layers are arranged in the same orientation as in the PØD and FGDs as discussed in section 4.2. The majority of events in ND280 are forward going and the DsECal is used to aid particle ID and distinguish between muons and other tracks which exit the detector downstream of the final TPC.

[1] The PØD subdetector has two ECal components which are called the upstream and central ECals, though these are part of the PØD itself and generally not treated as part of the ECal in analyses.

4.4.2 PØD ECal

The PØD Ecal is split into six modules, two top modules, two bottom modules, and one module on each side allowing the top and bottom to split when the magnet is opened. Unlike the other ECal modules, the FGD and PØD, the PØD ECal modules only have bars arranged parallel to the beam axis. Since π^0 conversion takes place within the PØD itself, less sophisticated reconstruction is needed to capture the photons that do not fully convert in the active region of the PØD, or to distinguish between MIP-like and EM-like particles. Each module has six layers of scintillator bars and five 4-mm-thick layers of lead with for a total of thickness of 3.6 X_0.

4.4.3 Barrel ECal

The Barrel ECal, similarly to the PØD Ecal, is split into six modules: two top modules, two bottom modules, and one module on each side. The six Barrel ECal modules have 31 layers of scintillator bars, each with a 1.75-mm-thick layer of lead for a total of 9.7 X_0 with alternating layers arranged with bars parallel to the beam or transverse to it. Bars running parallel to the beam measure 3.84 m long; bars running transverse to the beam in the top and bottom modules measure 1.52 m long; and bars running transverse to the beam in the two side modules of the barrel ECal measure 2.36 m long. The scintillator bars in the barrel ECal have a mix of single- and double-ended MPPC readouts. The Barrel ECal wraps around the tracking region to aid in particle ID and reconstruction of high angle tracks, for which the FGD-Iso reconstruction is poorer due to the scintillator bar orientation.

4.5 Side muon range detector

The side muon range detector (SMRD) instruments the flux return yoke of the ND280 magnet [112]. The 440 mm wide paddles of scintillator which make up the SMRD sit in gaps in the yoke structure and can be used to discern between particles entering the detector from outside the main subdetector systems. The scintillator paddles have a curving groove on the surface in which the WLS fibre lies in a pattern which allows it to pass close to the whole area of one side of the paddle. The SMRD is used to tag cosmic ray muons which pass straight through the detector. Tracks including SMRD hits which are coincident in different sections of the SMRD are removed from the analyses. As ND280 sits on the surface, cosmic ray muons are a large background to analysis work; however, they do provide opportunities to calibrate detectors.

4.6 Electronics

For the detector signals to be processed into an analysable format they need to be read out and digitised. The MPPCs which collect the light and convert the light to electric charge produce a signal which is first processed by the front-end boards (FEBs) which digitise the analogue signal [113].

4.6.1 MPPCs

Each scintillator bar is read out from either one or both ends of the WLS fibre by a Multi-Pixel Photon Counter (MPPC) manufactured by Hamamatsu Photonics. MPPCs are small photosensors which make them easier to incorporate into a compact detector and they are unaffected by the strong magnetic field, unlike PMTs, which makes them a suitable choice for ND280. Each MPPC contains 667 independent pixels measuring 50 μm $\times$ 50 μm which act as Geiger micro-counters on an area measuring 1.3 mm $\times$ 1.3 mm. The photo-diode of the pixel has an applied voltage close to the breakdown voltage of the diode such that an incident photon produces a photoelectron and induces a Geiger avalanche as the potential difference across the diode is above the breakdown voltage, allowing current to flow. MPPC pixel gain is determined from the charge accumulated and is given by $Q_{\text{pixel}} = C_{\text{pixel}}\Delta V$ [15] where C_{pixel} is the pixel capacitance, ΔV is the overvoltage or difference between the breakdown voltage and the applied voltage. The typical MPPC operational voltage is 70 V with an overvoltage around $0.8 - 1.5$ V and a pixel capacitance of 90 fF. This gives a typical gain of the order of 10^6. Diode capacitance is temperature dependent and so the applied voltage must be calibrated in order to prevent pixels firing from electrical noise, without losing sensitivity to incoming photons from the scintillator bars. Whilst the avalanche technique means each pixel acts as a binary switch, the large number of pixels on each MPPC means the total readout acts as an analogue detector with the number of pixels which fire being proportional to the number of incident photons, though constrained by the total size of the chip.

4.6.2 Data acquisition system

This digitisation process also includes calibration steps to remove noise from the signal. For example, in the scintillator bar readouts a threshold or voltage gain must be set for the charge collected which will be accepted as a "hit". This both removes noise and compresses the data format from a series of complicated analogue light pulses to a series of hits the detector with their intensity and timing information.

The PØD, ECal and SMRD use the same Trip-T Front-end Boards or TFBs. Each of the TFBs controls the MPPC voltage gains to reduce temperature dependent noise. This calibration takes place on a weekly basis during beam-down time. The output of these boards is then collated by

Readout Merger Modules (RMMs) which build the events that are then stored in the MIDAS file format [114].

Collectively this is referred to as the Data AcQuisition system, or DAQ. The electronics are also used to coordinate a cosmic ray muon veto which uses coincidence signals from pairs of subdetector modules in the PØD, SMRD and ECal to reject events where a track enters and leaves the detector and does not come from in interaction within the detector. Data quality is actively monitored by DAQ shifters during operation.

4.7 Data

Before the data can be analysed it must be processed to package it into a form that can be related to the underlying parameters of a model. This begins with processing the most basic and low-level information produced by the detector such as charge deposition and timing before building higher-level reconstructed objects which represent the particles and interactions described by the model.

4.7.1 Event building

In contrast to collider experiment detectors such as ATLAS or CMS in which precision knowledge of the initial particle allows for precise reconstruction of the interaction vertex, neutrino interactions in ND280 occur throughout the detector in both instrumented and uninstrumented regions. As a result it is not always possible to fully reconstruct the primary neutrino interaction and the subsequent interactions within the detector may not be correlated to the neutrino interaction. To account for this, an ND280 event is defined by the time window in which a spill from the beam is expected. This excludes much of the background noise from processes such as cosmic rays (although there is also an active cosmic ray trigger veto), as well as capturing information about neutrino interactions where the initial neutrino interaction may not be reconstructed, such as in NC interactions. This can result both in multiple neutrino interactions or no interactions taking place in one event; however, for the analysis discussed in this thesis, a single primary interaction vertex is selected and is treated as the neutrino interaction vertex for the event. When building events in the data, inter-detector timing information, and a signal from the beam control indicating the timing of the beam spills, are used to package information from each of the subdetectors in which a signal is detected into an “event”. In a given uninterrupted period of data taking at ND280, or subrun, the information for each event in a given subrun is then stored in the file for that detector run period in the MIDAS file format [114].

4.8 Data acquisition

During running various checks on the quality of the data are made based on the operative status of the different subdetectors. A flag can be set for a given run that indicates which of the ND280 subdetectors and systems are running nominally, and which may need to be excluded from the reconstruction. Different analyses require different levels of data quality. The fit to the ND280 data in the oscillation analysis requires that each of the subdetectors used in the selections are running smoothly and the magnet is running at nominal field strength. Some analyses are less dependent on measurements of track curvature from the magnetic field and so data taken without the magnet, or in "magnet-open" configuration are still of use to these analyses. Much of T2K run 9 was conducted in magnet-open mode due to an electronics fault in the FGDs to allow for PØD π^0 production analyses to continue. Data quality is provisionally tagged by the DAQ shifter who details the subdetectors in operation when starting and stopping the run periods, and then calibration data prepared by the various detector experts and run coordinators are used to confirm the quality of the data.

4.9 Data distribution

To make the data collected available to T2K collaborators all over the world, copies of the raw data files are stored at multiple data centres on the CERN LHC computing grid, commonly referred to as "the grid". T2K data must be uploaded from the KEKCC to the closest grid storage element (SE) at KEK (KEKSE) before being transferred to the RAL and TRIUMF SEs in the UK and Canada. At each of these data centres the files are stored on combination hard drives and tape storage for long term backups. During beam running periods, scripts are used to upload files to the grid using supported grid tools, interfaced by an in-house wrapper called the T2K Data Manager (T2KDM). These scripts load the files onto the grid and perform various checks to ensure the data has been successfully loaded before logging the successfully-transferred files. The scripts check the existing files at KEKCC against the transfer logs and will periodically attempt to load any files not included in the logs. This means the scripts can continuously transfer the data collected at J-PARC to the grid, as well as being able to cope with short-term connection problems. Files which are still not successfully transferred after the automated checks are stored in a temporary storage area which is manually checked on a roughly daily basis. Files can then be manually uploaded to the grid, or if a file is corrupted, this can be logged and input into the data quality information.

4.9.1 Simulation and analysis

The distributed data can be processed through the ND280 software chain, discussed in Appendix A, and compared to simulated data produced using Monte-Carlo methods, which is referred to as the Monte-Carlo data, or more simply "the Monte Carlo" (MC). The MC is produced in multiple stages outlined in Appendix A which are based on neutrino interactions simulated in `NEUT` [115], an in-house MC generator developed and maintained by the SK and T2K collaborations, and propagated through a simulation of the detector response in `GEANT4` [116]. The MC must also be propagated through the ND280 software chain and undergo the same reconstruction process as the real ND280 data. Once reconstruction has been conducted, the `Highland2` package is used to develop selections and systematics for data samples. These selections can then be ported to the more lightweight `psyche` framework which produces the input files for the fit to the ND280 data in the oscillation analysis (Chapter 6).

4.10 Event classification in ND280

Once events have been reconstructed into tracks and vertices, particle identification (PID) methods are used to best estimate what type of particle left a track in a detector, and to determine the topology of the interaction. PID can involve simple cuts on physics observables or more complex likelihood functions with comparisons to known control samples of tracks.

Neutrino interactions in ND280 typically have a primary interaction vertex from which tracks associated with the final-state products of the interaction begin. Additional vertices or reconstructed objects in the detector may come from the reinteraction of products of the initial interaction. These additional vertices may be disconnected from the visible products of the initial interaction if one of the products is neutral or not reconstructed. The initial neutrino interaction will happen in the densest parts of the detector and many interactions from the beam take place in the magnet which cannot be reconstructed. Additionally interactions is within the ECals are difficult to reconstruct and so the ECals are typically only used for distinguishing between EM-showering and MIP like products of interactions in the other subdetectors. This leaves the PØD and the two FGDs as the primary-target subdetectors. This section will focus on interactions in the FGDs but the same description can, in general, be applied to interactions in the PØD. To ensure the interaction vertex truly is within the material of the FGDs, a fiducial volume (FV) is defined which excludes vertices constructed at the edges of the FGDs, where they may in fact be secondary interactions of products of interactions in other parts of the detector. These excluded backgrounds are called out of FV (OOFV) backgrounds.

CC interactions can be identified by a muon-like or electron-like track produced at the interaction vertex. These will typically carry most of the momentum of the neutrino, and so the best candidate for the lepton is usually the highest momentum track. Muons are minimally ionising particles (MIPs) and so will leave long tracks through the detector often exiting the tracker through the ECal without stopping. As most events in ND280 have very forward going final state leptons, this is often the DsECal. Electrons, on the other hand, are more ionising due to their lower mass and deposit more energy as they travel and will induce electromagnetic (EM) showering within the dense ECals. Photons also will induce EM showering, but can be distinguished from electrons by the lack of an associated charged track in the tracker region of the detector. The intrinsic ν_e content of the beam is small, and although there are dedicated analyses to study their interactions, these events are not included in the main T2K analysis. If a proton is ejected from the nucleus it may leave a short track exiting the primary vertex. Neutrons, however, do not leave tracks as they are neutral, but may undergo secondary interactions the products of which leave tracks not directly connected to the interaction vertex.

The signal topology of the T2K oscillation analysis is that of a CC interaction with no final-state pions or hadrons (other than the ejected nuclei), and is referred to as CC-zero-pions or CC0π. This is enriched in the signal interaction mode, CCQE, but will also include 2p2h interactions as well as SPP processes in which the pion is adsorbed during FSI within the nuclear remnant as an irreducible background. These interactions can be identified by a reconstructed topology of a vertex within the FV of the FGDs, with no disconnected tracks upstream of the vertex, a single lepton-like track and no hadron-like tracks. There may or may not be a short proton track associated with the vertex left by the ejected proton. If the lepton track is muon-like, this will typically exit the tracker region into the ECal and pass out of the detector without showering. An example of a ν_μ CC0π candidate data event can be seen in figure 4.6. If the track is electron-like, it will proceed through the tracker until converting by showering in the ECal, or stopping within the detector. Currently only events with forward going charged leptons are included as a cut on backwards going tracks which enter the TPC directly upstream of the target FGD is used to reduce misreconstructed OOFV backgrounds. In addition to the non-CCQE interactions with a true CC0π final-state topology, interactions with other true topologies may be misreconstructed where other final state products are not identified due to reconstruction inefficiencies.

As the dominant background interaction mode to the true CC0π topology is CC resonant pion production, a dedicated CC-one-pion or CC1π topology is defined as any CC interaction with a single charged pion in the final state with the opposite charge of the final state lepton. An example of a ν_μ CC1π candidate data event can be seen in figure 4.7. Pions, like muons, are MIPs, and so

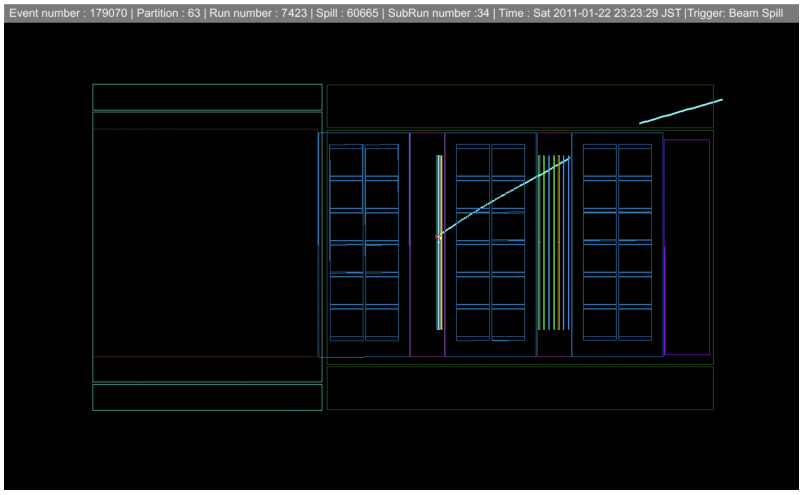

Figure 4.6: Event display showing a candidate for the CC0π topology. A long, muon-like track is produced at the primary vertex in FGD1, which passes through TPC2, FGD2, and into the Barrel ECal without showering. A very small track near the vertex indicates the ejection of a proton from the nucleus.

can be difficult to distinguish from muons in ND280, though pions will often undergo secondary interactions in the detector and may shower in the ECal at lower energies. Charged pions may decay within the detector to a muon and muon neutrino. This muon will be lower in momentum and quickly decay to a Michel electron. This muon track may not be observed in the detector due to its short lifetime. If the pion is produced at low energies, the pion itself may only be inferred from the Michel electron as neither pion nor muon live long enough to be seen in the detector. A cut on the time delay of the electron like track can also be used to confirm an electron-like track is from the decay of the pion due to the 2.2 μs lifetime of the intermediate muon.

High energy neutrinos which interact in an inelastic manner producing many hadrons or undergoing DIS will often leave many tracks in the detector leaving the interaction vertex in the FGD. A charged pion track of the same sign as the leading lepton is also a good indicator of DIS. High track multiplicity primary vertices and vertices where heavier mesons such as kaons or etas are present are classified as CC-Other. An example of a ν_μ CC-Other candidate data event with more than one final-state hadron can be seen in figure 4.8.

Neutral pions may be produced in interactions in a the same manner as charged pions, but must be inferred from their decay products. These are typically a pair of photons which may either convert in the ECals by undergoing EM showering, or by pair production in the tracker (as shown in figure 4.8). An event with two photon tagged showers in the ECal, or a photon tagged shower in the ECal and an electron-positron pair in the tracker, or two electron-positron pairs in the tracker may indicate the intermediate neutral pion. These interactions are also included in the

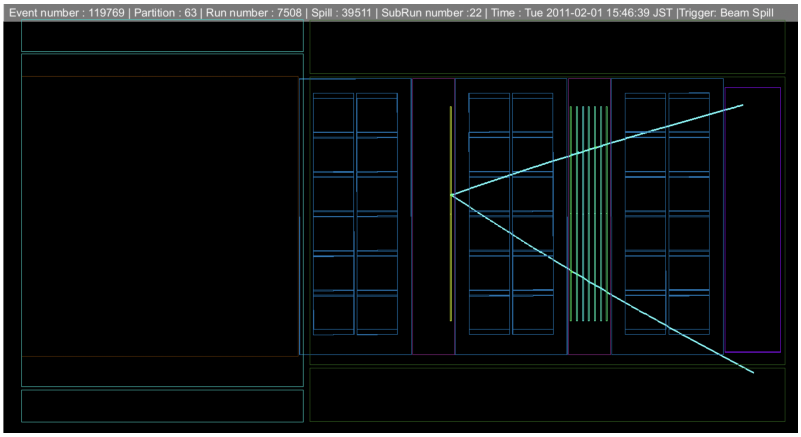

Figure 4.7: Event display showing a candidate for the CC1π topology. A long, muon-like track is produced at the primary vertex in FGD1, plus another long MIP-like track curving in the opposite direction indicates a final-state pion. Both the muon candidate and the pion candidate leave the detector through the ECal through the side of the detector.

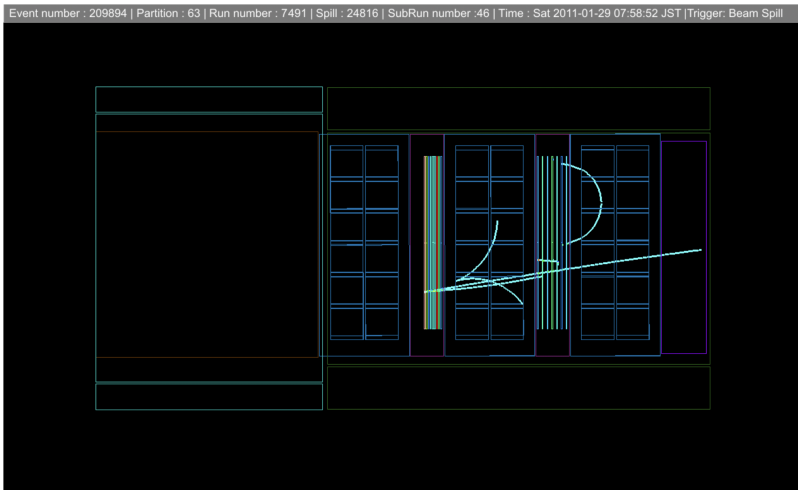

Figure 4.8: Event display showing a candidate for the CC0-Other topology. In addition to a long muon-like track starting in FGD1 which exits the detector through the ECal, an additional track at the primary vertex indicates a charged pion. Two oppositely curving and oppositely charged electron-like tracks can be seen in the TPC suggesting pair production from a photon. This photon may be from the decay of a hadron such as the decay of a π^0 in which the second photon is not reconstructed. Additional tracks in FGD2 may also be from the decay of a neutral particle.

CC-Other topology.

While T2K is most interested in the CC interactions, NC interactions contribute backgrounds to the analysis and must be studied. NC interactions in ND280 will not (directly) produce a charged lepton from the interaction vertex, but may through the decay of other particles produced through the momentum transfer of the Z^0 boson. NCE-like vertices may be identified by an ejected proton, or the secondary interactions of an ejected neutron. The former will appear as a short proton track in the FGD, but the latter may not be visible depending on the energy of the ejected neutron. It may also not be possible to accurately determine the position of the interaction vertex and this may contribute to out out of FV (OOFV) backgrounds.

Chapter 5

The T2K oscillation analysis

The T2K flagship analysis is the extraction of the accelerator-PMNS matrix terms, θ_{13}, θ_{23}, δ_{CP} and neutrino-mass splitting Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) from the differences between the unoscillated and oscillated energy spectra of the T2K neutrino beam. This is commonly referred to as "the oscillation analysis" or "the OA". The OA is a highly dimensional analysis which requires careful treatment of systematic uncertainties supplemented by many internal and external measurements of neutrino interactions.

The oscillation parameters are extracted from the T2K data from their impact on the predicted neutrino energy spectrum of the beam. By changing the values both of the oscillation parameters and the nuisance parameters (flux, cross section, detector response) in the MC the combination which best predicts the data can be found, and Bayesian credible intervals and contours or frequentist confidence intervals and contours for the oscillation parameters can be constructed after marginalising over the nuisance parameters. The effect of varying the oscillation parameters on the probability of muon (anti)neutrino survival for the T2K baseline and T2K energies (shown in figure 5.1) can be convolved with the T2K flux to predict the composition of the beam at SK which is shown in figure 5.2 for the CP-conserving case.

For the global best-fit value of θ_{23}, the deficit of muon neutrinos in the oscillated flux prediction is large, allowing for precise measurements of $\sin^2\theta_{23}$ and the magnitude of the mass splitting Δm^2. However, the sign of the mass splitting is obscured by the square of the sine term in which the mass splitting features in the oscillation probability equations (2.25 and 2.26) to first order. This sign is dependent on the order of the neutrino masses and so the mass splitting term is defined as Δm^2_{32} for the normal ordering ($\nu_1 < \nu_2 < \nu_3$) and as $-\Delta m^2_{31}$ or $|\Delta m^2_{31}|$ for the inverted ordering ($\nu_3 < \nu_1 < \nu_2$). The largest component of the neutrino beam in SK at these distances and energies is of tau flavour and so the electron neutrino component of the beam is still small when compared to the total flux of all neutrino types. This is due to the relatively small value of

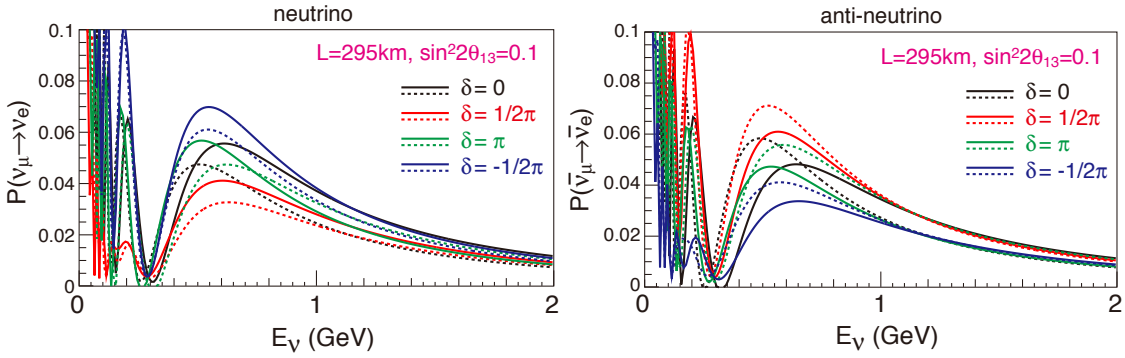

Figure 5.1: Left (right): Probability of a muon (anti)neutrino at a given energy E_ν, produced at the T2K target, oscillating to an electron (anti)neutrino after traversing the T2K baseline L of 295 km for different values of δ_{CP} and a fixed value of $\sin\theta_{13} = 0.1$. Normal ordering probabilities are shown in solid lines and inverted ordering probabilities are shown in dashed lines.

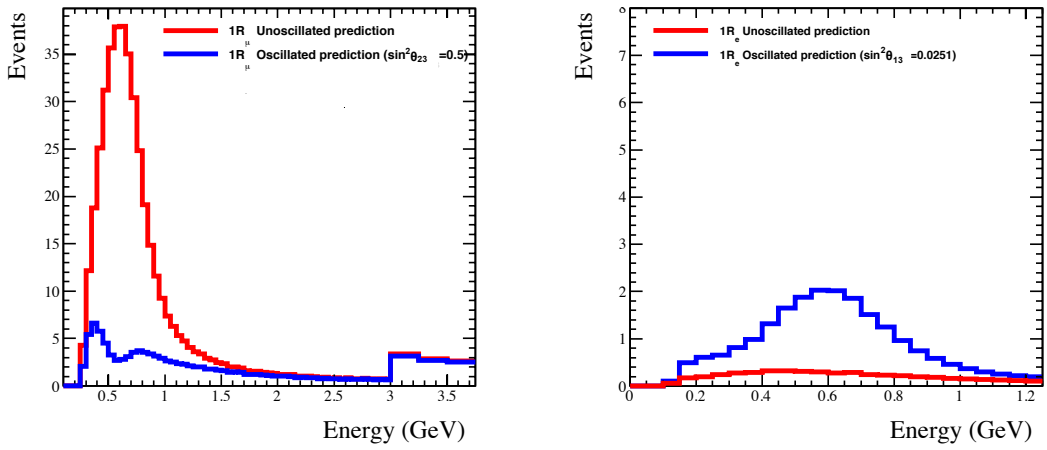

Figure 5.2: Comparison of the unoscillated (red) and oscillated (blue) neutrino spectra for the T2K neutrino beam at an off-axis angle of 2.5° assuming $\sin^2\theta_{13} = 0.0251$, $\sin^2\theta_{23} = 0.5$, $\Delta m^2_{32} = 2.4 \times 10^{-3}$ eV^2, $\delta_{CP} = 0$ and normally ordered neutrino masses for the muon (left) and electron (right) neutrino species.

$\sin^2\theta_{13}$ in the leading term of the $P(\nu_\mu \to \nu_e)$ expression in equation 2.26, which whilst making the electron neutrino appearance lower in statistics and so more difficult to measure, does expose the trailing terms with their dependence on the value of $\delta_{\rm CP}$ and the bare Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO). By determining the size of these second-order terms the sign of the splitting, and thus the order of the neutrino masses, can be determined. T2K has some sensitivity to these terms, but not enough to make a precise measurement of $\delta_{\rm CP}$ or exclude a mass ordering to the 5σ standard required for a discovery in particle physics. As well as the composition and flux of the beam, the rate at which the beam can be sampled must be known; i.e. the physics governing neutrino interactions in matter and the efficacy of the detectors to extract information and classify the events reliably must be understood. The expected observed rate $R_\mu(E_\nu, L)$ of ν_μ events at the far detector at distance L from the target for neutrinos of energy E_ν is given by

$$R_\mu(E_\nu, L) = \Phi_{\nu_\mu}(E_\nu) \times \sigma_{\nu_\mu A}(E_\nu) \times \varepsilon_{\rm det}(E_\nu) \times \underset{\nu_\mu \to \nu_\mu}{P}(E_\nu, L) \tag{5.1}$$

where $\Phi_{\nu_\mu}(E_\nu)$ is the beam flux, $\sigma_{\nu_\mu A}(E_\nu)$ is the ν_μ interaction cross section, $\varepsilon_{\rm det}(E_\nu)$ is the detector efficiency, and $\underset{\nu_\mu \to \nu_\mu}{P}(E_\nu, L)$ is the ν_μ survival probability. The largest source of systematic uncertainty on this measurement is the measurement of the interaction cross section.

Some neutrino experiments such as NOνA [77][117] use the same technology for near and far detectors to exploit the cancellation of the nuisance parameter uncertainties using the correlations between detector responses at near and far detectors. In contrast, T2K uses very different detector technology for its near detectors than used at SK. This means that, whilst more information about the neutrino interactions can be extracted at ND280 than SK, more care is needed to understand any correlations between modelling of interactions and event selections at the different detectors. The modelling differences are not simply due to the different detection methods (scintillation vs Cherenkov) but also due to the complicated correlations between the modelling of different nuclear targets like carbon and oxygen.

5.1 The T2K analysis streams

Two broad analysis streams are followed at T2K outlined in figure 5.3, one stream consisting of two analyses sequentially fitting to the near detector and then the far detector (5.3(a)), and the other simultaneously fitting to data from both near and far detectors(5.3(b)). Splitting the likelihood into two halves serves several purposes: firstly, the ability to understand the fit results and do extensive validation of the model and fitter at the near detector acts as a blinding of the far detector data; secondly, splitting the likelihood up reduces the dimensionality of the fit which can become computationally cumbersome, particularly in any analytical navigation of the parameter

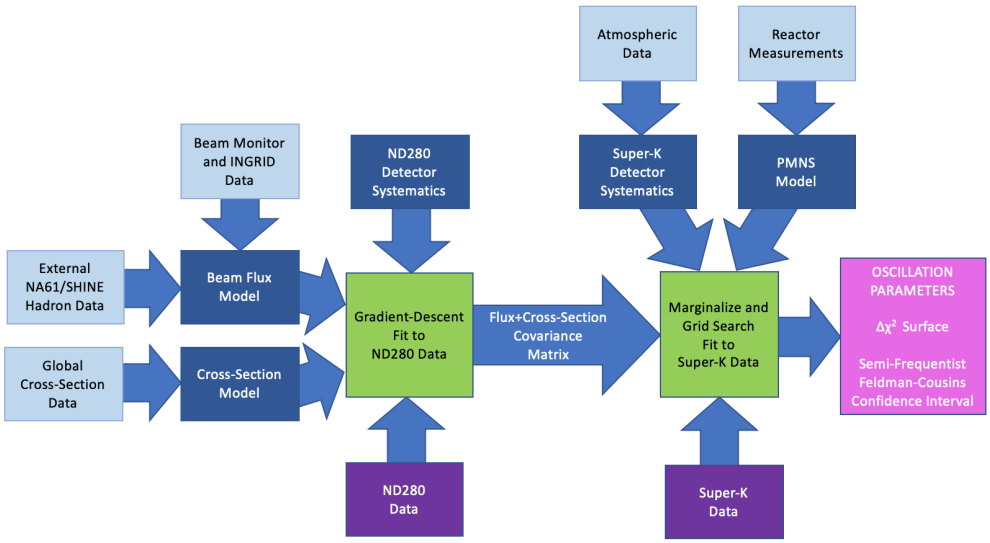

(a) Frequentist sequential analysis stream: BANFF and P-Theta

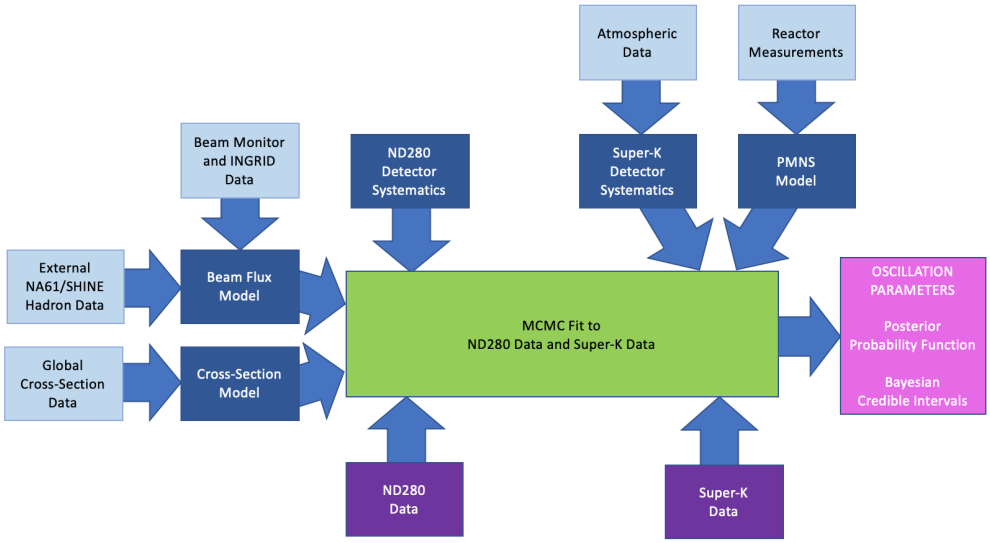

(b) Bayesian simultaneous analysis stream: MaCh3

Figure 5.3: Flow diagram outlining the two parallel analysis streams for the T2K oscillation analysis. Each analysis stream takes the same inputs at ND280 and SK (shown as Super-K in the figure) which come from analysis of both internal and external measurements. In addition to the sequential versus simultaneous differences, the fitting methods and interpretation of the results are different. In the sequential fit stream, the BANFF gradient descent fit produces a best-fit set of parameters and their covariance matrix, which are used as prior inputs to the P-Theta analysis which then marginalises over the flux, cross section and SK detector systematic uncertainties in order to make a grid search of the remaining oscillation parameters computationally feasible. This produces a $\Delta\chi^2$ surface which is interpreted in a frequentist manner. In the simultaneous fit stream, the MaCh3 MCMC fitter marginalises over all of the nuisance parameters from both detector likelihoods and samples the oscillation parameter space to build a posterior probability density function which is interpreted in a Bayesian manner.

space such as in a gradient descent fit; thirdly, whilst the simultaneous fit analyses the data of both detectors, the two different fits can be cross-validated at the near detector in order to ensure that they have consistent results from each fit prior to the extraction of the oscillation parameters.

In the first stream, the BANFF near detector fit minimises a χ^2 surface by following the gradient to find the combination of the parameters for which the MC best predicts the data. Once the minimum is found, the covariance matrix of the ND280 parameters is calculated from the inverse of the Hessian matrix as described in section 5.5. The ND280 detector parameters are then marginalised over and this matrix is then passed onto two far-detector fitter groups, P-Theta and VALOR. The results shown in this thesis are from the P-Theta fitting group. The far-detector fitters then marginalise over all of the parameters except for the oscillation parameters, before doing a grid search of the remaining one- or two-dimensional likelihood surface to find the best-fit value of the POI(s).

In the second stream, the MaCh3 fitter uses a MCMC method to semi-randomly step through the entire parameter space of both the ND280 and SK likelihoods according to a probability determined by the likelihood. This populates a posterior probability density function which can be marginalised over to extract the posterior probabilities of the oscillation parameters.

The differing fitter mechanics and statistical interpretations of these analyses are important to consider whenever making comparisons of the results. The necessity of a smooth likelihood to differentiate in the gradient descent method prevents the inclusion of any parameter which would shift the events into different bins the fitted observable space. Therefore, an alternative effective parameter is used that changes the likelihood in a smooth and continuous manner, which has advantages as well as shortcomings but may not be able to match the variations in the MCMC method which can fit these directly. The results of the sequential analyses are interpreted in a semi-frequentist manner, constructing confidence intervals. The results of the simultaneous fit are interpreted in a Bayesian manner constructing credible intervals.

5.1.1 Interpretation of Bayesian and Frequentist inferences

Statistics can be split into two schools of thought, Bayesian and frequentist. Each approach asks different questions and whilst they produce analogous metrics and quantities, an awareness of the difference in interpretation is necessary to understand the results.

Bayesian statistics is an attempt to quantify a degree of belief about a system. For example, if one assumes a die is fair, then one can assume that the probability of a die landing on any given side is equal, to some degree of certainty. In the case of the die, the prior assumption may be that the die is perfectly fair and has six sides and so a six will be rolled one sixth of the time.

When the die is thrown, the outcomes of each side may not be perfectly equal, but unless they significantly deviate from equal rates of occurrence, a Bayesian statistician may still believe that the observations are consistent with the prior assumption of fairness. The degree of certainty can be quantified by the prior uncertainty or prior probability distribution. Often this is a symmetric two-sided uncertainty in which the probability distribution is Gaussian, but this distribution can take other forms, and even when there is no prior knowledge of a quantity, a choice may need to be made of the shape of the prior distribution. An unknown quantity may have a "flat" prior probability of equal probability at all values, but if this lack of knowledge extends over multiple orders of magnitude, a flat prior distribution in the logarithm of the quantity may be more suitable. The prior probability is the degree of belief that the true value of the quantity, in the case of the die this may be the rate of rolling a 6, is within some range of values. After comparison to observation, a posterior distribution or uncertainty can be ascertained. This may now deviate from the assumption of fairness, and upon collection of a sufficiently large number of observations, the probability of the die being fair may be low enough to be discounted. In the frequentist approach, the die must be thrown a large number of times, and the probability of each number being cast is the fraction of times in which that given number is cast. A purely frequentist approach would be completely empirically driven, and the probability of an outcome is measured to be the rate at which that outcome occurs. It cannot be assumed that the die will roll a six, one sixth of the time; the die must be rolled to determine the probability of rolling a six. Often, a purely frequentist approach is not feasible for some experimental or computational reasons. In these cases a semi-frequentist approach can be used which attempts to incorporate some aspects of Bayesian statistics such as inclusion of prior knowledge. The BANFF analysis framework described in chapter 6 is a semi-frequentist framework which attempts to find the most likely set of parameters and their covariances which would describe the ND280 data, but begins with a prior covariance informed by previous internal and external measurements.

5.2 Choice of PMNS parameters in the SK Asimov fits

The Asimov fits at ND280 simply fit the prior model to itself, though a choice of the oscillation parameter values in the MC at SK for the P-Theta and full MaCh3 analyses. The Asimov fits to the far detector are prepared using MC generated or weighted to the the "Asimov A" PMNS values listed in table 5.1. Other Asimov data sets are used, but these are the values used in this work unless stated otherwise.

Parameter	Asimov A Value
Δm^2_{21}	$7.53 \times 10^{-5}\text{eV}^2$
$\Delta m^2_{32}/\lvert\Delta m^2_{31}\rvert$	$2.509 \times 10^{-3}\text{eV}^2$
$\sin^2\theta_{23}$	0.528
$\sin^2\theta_{12}$ ($\sin^2 2\theta_{12}$)	0.307 (0.851)
$\sin^2\theta_{13}$ ($\sin^2 2\theta_{13}$)	0.0218 (0.0853)
δ_{CP}	-1.601
Mass Ordering	Normal

Table 5.1: The selected values of the oscillation parameters used in the preparation of the primary Asimov MC simulated data set, or Asimov A. These parameters are taken from the best-fit 2014 result [118] of the fit to T2K runs 1-4, with $\sin^2\theta_{13}$ taken from the 2019 PDG value from reactor experiments [34].

5.3 Reactor constraint

Although when performing the fits, T2K places an independent constraint on the value of θ_{13} from its data alone, the 2019 Particle Data Group (PDG) value of θ_{13} from reactor experiments [34] such as Daya Bay and RENO is much more precise and therefore this θ_{13} result can be much stronger constraint in the fits. Exploiting the PMNS oscillation framework and this reactor constraint on θ_{13}, T2K can obtain a much stronger constraint on the value of δ_{CP} during the fit than would otherwise be possible. This constraint is applied as a Gaussian penalty on the χ^2 as a function of θ_{13} only, though there are ongoing efforts to include more sophisticated 2D likelihood contributions from experiments such as Daya Bay.

5.4 Likelihoods

The likelihood function, or the likelihood, of a modelling problem is an n-dimensional surface for n variable model parameters which describes the "goodness of fit" of a model's prediction to an observed sample of data at each possible combination of parameter values. This can be constructed by calculating the difference between the prediction and the data using a metric such as the χ^2 per degrees of freedom (DoF) of the fit to the data. Particle physics experiments are typically a type of counting experiment, and so the probability of successive events are discrete and independent random variables. For a simple counting experiment, the Poisson probability of counting k events for a predicted λ events is

$$p(k|\lambda) = \frac{\lambda^k \exp(-\lambda)}{k!} \tag{5.2}$$

where λ is the prediction according to the model and is dependent on some set of parameters θ. Most physics problems, however, cannot be explained by an observation described by a single number, and so this can be multiplied by the probabilities of a set of observations such as different bins of a measured quantity or observable. This may be continuous, such as an energy spectrum, or discrete, such as topological classification with some number of observable bins N. Extending the probability of observing k_i events in the ith bin for λ_i predicted events, we can find the likelihood $\mathscr{L}$ of observing all values of k_i across these N bins:

$$\mathscr{L} = \prod_{i=1}^{N} \frac{\lambda_i^{k_i} \exp(-\lambda_i)}{k_i!} \tag{5.3}$$

where λ_i all depend on some common set of parameters θ.

This likelihood expression is unnormalised but by taking the ratio to the most likely outcome of the nominal model, terms can be cancelled and the expression for the relative likelihood, or the likelihood ratio, is

$$L_{\text{ratio}} = \prod_{i=1}^{N} \frac{\lambda_i^{k_i} \exp(-\lambda_i) k_i!}{k_i^{k_i} \exp(-k_i) k_i!}. \tag{5.4}$$

Using the relationship between the Poisson likelihood and the simple χ^2

$$-2 \ln \mathscr{L} = \chi^2, \tag{5.5}$$

the difference between the χ^2 of a general set of parameter values to the χ^2 of the nominal set of parameter values, $\Delta\chi^2 = \chi^2 - \chi^2_{\text{nom}}$, can be expressed as

$$\Delta\chi^2 = 2 \sum_{i}^{N} \left(\lambda_i - k_i + k_i \log \frac{k_i}{\lambda_i} \right) \tag{5.6}$$

which is easy to calculate for a binned likelihood.

In some cases, where statistics are limited, there may be benefits to using an unbinned likelihood in which case a function is used to estimate the underlying probability distribution and normalisation of the data. Though a purely frequentist analysis would include no prior uncertainty constraint, prior constraints, $\pi(\theta)$, can be included as penalty terms in the expression of the likelihood ratio

$$L_{\text{ratio}} = \frac{\pi(\vec{\theta})}{\pi(\vec{\theta}_{nom})} \prod_{i=1}^{N} \frac{\lambda_i^{k_i} \exp(-\lambda_i) k_i!}{k_i^{k_i} \exp(-k_i) k_i!} \tag{5.7}$$

and if the priors are Gaussian, with Gaussian covariances, they take the form

$$\pi(\vec{\theta}) = (2\pi)^{-\frac{k}{2}} |V_\theta|^{-\frac{1}{2}} e^{-\frac{1}{2} \Delta\theta (V_\theta^{-1}) \Delta\theta^T} \tag{5.8}$$

where V_θ is the prior covariance matrix. The penalty terms therefore increase the chances that minimising $\Delta\chi^2$ results in a set of $\vec{\theta}$ that do not deviate more than a few standard deviations from their nominal or prior central values. After taking the logarithm a more complete expression for

the -2logL_{ratio} is given by summing over the covariances of parameters, p and q for the number of parameters N_{pars}

$$\Delta\chi^2 = 2\sum_{i}^{N}\lambda_i(\vec{\theta}) - k_i + k_i\ln[k_i/\lambda_i(\vec{\theta})] + \sum_{p}^{N_{\text{pars}}}\sum_{q}^{N_{\text{pars}}}\Delta\theta_p(V_\theta^{-1})_{p,q}\Delta\theta_q, \tag{5.9}$$

and if the parameter prior distributions are roughly Gaussian, then they can be approximated as quadratic functions around the minimum of the χ^2 surface, χ^2_{min}, where the ratio is redefined relative to the most likely set of parameters is then

$$\Delta\chi^2_{min} = \sum_{p}^{N_{\text{pars}}}\sum_{q}^{N_{\text{pars}}}\frac{\partial^2\Delta\chi^2(\vec{\theta}^{min})}{\partial\theta_p\partial\theta_q}(\theta_p - \theta_q^{min})(\theta_q - \theta_p^{min}) \tag{5.10}$$

where $\vec{\theta}$ are all of the parameters, and $\vec{\theta}^{min}$ are their values at the minimum of the of the χ^2 surface, χ^2_{min}. The "Hessian matrix",

$$\frac{\partial^2\Delta\chi^2(\vec{\theta}^{min})}{\partial\theta_i\partial\theta_j} \tag{5.11}$$

can then be inverted to find the covariance matrix at the best fit point.

5.4.1 Profiling and marginalisation of nuisance parameters

In statistical modelling, rarely are all of the parameters of the model which govern the prediction of some data equally interesting. Although some parameters in the modelling of an experiment may not be of physical significance, they must still be included to ensure appropriate freedom or uncertainty is included in the model.

These parameters are called nuisance parameters. For example, particle physics experiments aim to measure some physical quantity like the mass of a particle or the cross section of an interaction; however, there are many parameters describing the systematic uncertainties of the measurement, such as the momentum resolution of the detector, or the mass of the detector, which are of little interest. In other cases, physical parameters which data does not have the power to significantly constrain may have some correlations with the parameter of interest (POI), and so the extra freedom here must be taken into account.

Typically, frequentist analyses use profiling (defined below) to evaluate the minimum possible $\Delta\chi^2$ and so maximise the log$\mathscr{L}$ at some value of the POI, though doing so assumes that the likelihood is symmetric in shape around this best-fit point in all of the nuisance parameters, and may lose information about their true behaviour. The likelihood across the POI is therefore not being evaluated at the global best-fit point of the other parameters.

Profiling takes the values of the nuisance parameters which maximise the likelihood for some given value of the parameters of interest; whereas, marginalisation integrates the likelihood over the nuisance parameters to arrive at the value of the probability for each value of the parameter of

interest. The profiled likelihood $\mathscr{L}_p$, for a POI x and nuisance parameter $theta$ can be expressed as

$$\mathscr{L}_p(x) \equiv \mathscr{L}(x|\hat{\theta}_x) \tag{5.12}$$

where $\hat{\theta}_x$ is the value of the θ which maximises the $\mathscr{L}$ for the given value of x. The marginalised likelihood $\mathscr{L}_m$ can be expressed as an integral of the nuisance parameters

$$\mathscr{L}_m(x) = \int_\theta p(\theta|x)\mathscr{L}(x|\theta)d\theta \tag{5.13}$$

over the range of values they can take.

Bayesian analyses use marginalisation, integrating over the nuisance parameters and their likelihoods according to their priors and covariances. By sampling the parameter space according to the likelihood equation with all of the nuisance parameters free to vary, a posterior probability density function is constructed. When values of the POI are taken from this distribution, this is equivalent to sampling the posterior distribution of the POI with all of the nuisance parameters integrated over.

Finding the maximum likelihood point for every possible set of values for the parameters of interest can be incredibly difficult - especially in disfavoured regions where the excluded region is of interest, and so marginalisation is often preferable to profiling in high dimensions. Marginalisation also better handles non-Gaussian uncertainties which profiling cannot take into account. As such, often analyses will be called “semi-frequentist” as they use the Bayesian marginalisation method to deal with non-Gaussian nuisance parameters whilst being interpreted in a frequentist manner.

5.5 Fitting methods in the T2K analysis

T2K utilises different fitting methods which come with different strengths and weaknesses in terms of applicability and interpretation. A common fitting method is gradient descent, used in the BANFF fit to the ND280 data (see chapters 6-8), in which the likelihood surface is navigated by using its gradient until the maximum likelihood value, or minimum χ^2 value, is found. This method is easy to interpret in a frequentist manner with a best-fit set of parameter values, and the covariance matrix can be found using much of the same machinery. Where detail of a large range of values of a POI is needed, as in the p-theta fit to SK data (The results of which are presented in chapter 9), a grid search method may be preferred. Here the marginalised or profiled likelihood is sampled at a large range of values for the POI. This is more beneficial when calculating detailed shape of exclusion boundaries rather than assuming a Gaussian uncertainty. Another increasingly

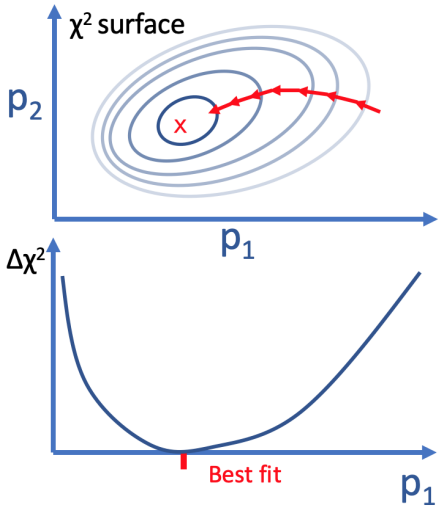

Figure 5.4: Simplified diagram demonstrating how the gradient descent algorithm incrementally navigates the likelihood surface in the direction of decreasing χ^2 values. Normalising the likelihood function by taking the subtracting the χ^2_{min} produces the $\Delta\chi^2$ which can be used to calculate confidence levels. Contours of equal $\Delta\chi^2$ are shown for the parameter space p_1, p_2, which can be profiled or marginalised to give a one-dimensional $\Delta\chi^2$ for each parameter. The best-fit point of the parameter is then the lowest value found, and the uncertainty can be evaluated through the Hessian calculation or by finding the parameter values with a $\Delta\chi^2$ corresponding to the 68% or 1σ interval.

popular method is the Markov Chain Monte Carlo method used in many Bayesian analyses, including the MaCh3 simultaneous fit to ND280 and SK data (Discussion of and comparisons of both the BANFF and p-theta fits to MaCh3 are shown in chapters 6-9), which samples the parameter space according to rules governed by the likelihood. Each of these is discussed in more detail in this section.

5.5.1 Gradient Descent

Gradient descent [119] uses the derivatives and coderivatives of the likelihood with respect to the model parameters to navigate the likelihood surface, moving incrementally to lower values of $\Delta\chi^2$ until a minimum value is found as shown in figure 5.4. Gradient descent requires a smooth and continuous surface to differentiate and so may run into problems with parameters which cause discrete jumps in the likelihood. Parameters which migrate events between observable bins in the fit may cause this to happen and so directly fitting these parameters may not be possible.

5.5.2 Grid Search

A grid search minimisation steps through the likelihood at regular intervals across the allowed range of the parameters and samples the likelihood. This is impractical in large likelihoods as the computational cost increases exponentially with dimensions, and unlike gradient descent in which each step can move in all dimensions, the grid search samples the same value in one parameter multiple times. This method may still be attractive where detailed knowledge of a small subset of the likelihood is wanted. All but the parameters of interest can be marginalised over leaving a manageable number of parameters to sample with a grid search method.

5.5.3 Markov Chain Monte Carlo

A numerical method of sampling the likelihood uses a Markov Chain Monte Carlo (MCMC) method which steps semi-randomly through the parameter space according to a probability dictated by the likelihood, as shown in figure 5.5. In the most simple MCMC method, parameters are initialised to some random set of parameter values, then random changes of the parameters within some predefined step size are used to select a new point in the parameter space. If the likelihood evaluated at the new set of parameters is more likely than that of the current set of parameters, the new parameter values are accepted and the process repeats. If the likelihood at the new parameter values is lower than the initial likelihood, the new values are rejected and are randomly selected again. As the new values are accepted, a chain is built of these values, and as the chain is allowed to run longer and longer it slowly covers the most probable values of the parameters and builds a posterior probability distribution. By sampling values for a parameter from the chain of accepted values, the posterior probability distribution for a parameter can be drawn, marginalising over all of the other parameters. This version of the MCMC method may find it difficult to populate high-dimensional likelihoods due to the high rate of rejection, and so for high-dimensional problems more sophisticated methods such as the Metropolis-Hastings (MH) algorithm are used [120]. In the MH-MCMC method, an acceptance ratio is defined such that even some fraction of the disfavoured parameter values are still accepted. This can prevent the fitter from sampling too small a region of parameter space or getting stuck in local minima. As the method can take some time to find the region of the likelihood which will be most sampled, initial values from the chain before the method can stabilise are rejected. This rejected part of the chain is called the “burn in”. In addition to this, whilst values in Markov Chains should only be dependent on the previous values, MCMC methods have correlations between successive steps beyond a single step, and so chains may be rejected in order to reduce the autocorrelations between the steps included in the posterior probability distribution.

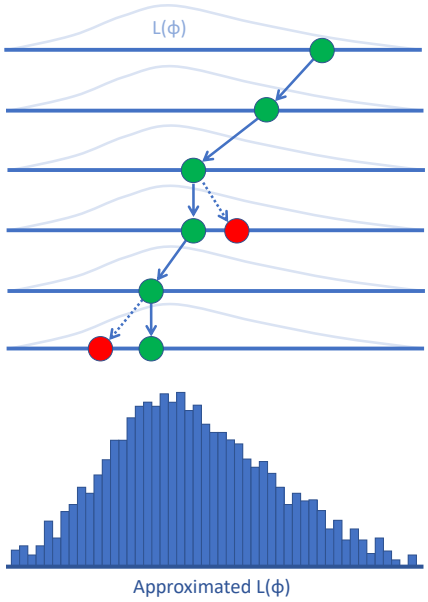

Figure 5.5: Simplified diagram demonstrating how the MCMC method steps through parameter space according to the likelihood and slowly builds up a posterior density distribution from the sampled values of the parameter ϕ whilst marginalising over the other parameters. Each of the green dots shows an accepted value of the parameter ϕ and each of the red dots shows a rejected value.

5.6 The T2K likelihood

As the number of model parameters increases and the shape of a likelihood becomes more complex, the chance of a gradient descent fitting method converging decreases, and the computational costs increase. It is therefore desirable to reduce the number of parameters in the fit. To avoid this problem with the high dimensionality of the total T2K likelihood, Bayes theorem can be used to split the likelihood into two parts: the part of the likelihood which is dependent on the oscillation parameters, and the part which is not. By first constraining the oscillation-independent part of the model, this can then be used as a new prior for the fits to the oscillation-dependent part of the likelihood. The total likelihood is given by

$$\mathscr{L}_{Tot}(\vec{b},\vec{x},\vec{o}) = \mathscr{L}_{Beam}(\vec{b}) * \mathscr{L}_{NA61}(\vec{b}) * \mathscr{L}_{Ext-xsec}(\vec{x}) * \mathscr{L}_{ND280}(\vec{b},\vec{x},\vec{d}_{ND280}) * \mathscr{L}_{SK}(\vec{b},\vec{x},\vec{d}_{SK},\vec{o}) \tag{5.14}$$

in which the beam parameters ($\vec{b}$), cross-section parameters ($\vec{x}$), detector parameters ($\vec{d}_{ND280}$) and ($\vec{d}_{SK}$), and the PMNS oscillation parameters ($\vec{o}$), describe the T2K data. The likelihood includes terms for the T2K beam monitor and INGRID data ($\mathscr{L}_{Beam}$), external data sets used to constrain both the beam ($\mathscr{L}_{NA61}$) and cross-section ($\mathscr{L}_{Ext-xsec}$) modelling, the ND280 data ($\mathscr{L}_{ND280}$), and the SK data ($\mathscr{L}_{SK}$). The beam and external measurements can be used as prior constraints to a near-detector only fit and so equation 5.14 then be expressed as

$$\mathscr{L}_{tot}(\vec{b},\vec{x},\vec{o}) = \mathscr{L}_{BANFF}(\vec{b},\vec{x},\vec{d}_{ND280}) * \mathscr{L}_{SK}(\vec{b},\vec{x},\vec{d}_{SK},\vec{o}) \tag{5.15}$$

where the oscillation parameters $\vec{o}$ have been isolated to the SK likelihood term. By exploiting this separation, the common flux and cross-section parameters can be constrained by the ND280 data with its larger data set prior to fitting to the SK data, in which flux and cross-section effects are convoluted with the oscillation parameters which have broad uncertainties. The ND280 detector parameters are effectively marginalised over in the BANFF ND280 fit and are not propagated to the far detector analysis, P-Theta. The MaCh3 fit runs parallel to the BANFF (and P-Theta) fits and is able to fit both detectors simultaneously due to the advantages of the MCMC method in high-dimensional likelihood spaces. A discussion of the production of the MC, and the reweighting procedure applied during the fitting procedure is given in appendix A.

5.7 Neutrino beam-flux model

The neutrino beam is simulated by first simulating interactions between protons and carbon nuclei in the graphite target using FLUKA 2011 [102][121] before the hadronic products of these interactions are propagated through the magnetic horns and decay volume with T2K's beam MC

generator JNUBEAM which uses GEANT3 [116] to simulate the geometry of the beam production hall and the semileptonic decays of the hadrons within it. The hadronic products from the proton interactions in the target are tuned to results from the NA61/SHINE experiment [101]. For the 2020 oscillation analysis, this tuning was updated from thin-target data [122] to data from a T2K-replica target run at NA61/SHINE [123][124]. The tuning is applied as a weight

$$w(p,\theta,z,i) = \frac{\mathrm{d}n^{\mathrm{NA61}}(p,\theta,z,i)}{\mathrm{d}n^{\mathrm{MC}}(p,\theta,z,i)}$$

to the hadron production rate in bins of exiting hadron momentum, p; angle of exit from the surface of the target, θ; position along the target, z; and particle type, i.

The tuned flux model is then fit to internal T2K data from beam monitor systems [125] and INGRID [126] and then a covariance matrix, in bins of neutrino flavour and energy, is produced based on throws of the beam simulation parameters. The nominal prediction of the T2K beam is shown in figure 5.6 which shows the impact of updating from the thin-target NA61/SHINE data to the replica-target NA61/SHINE data. The size of each component of the tuned flux relative to the untuned flux is given by table 5.2. The total contribution of each flavour to the unoscillated flux prediction is given in table 5.3.

The T2K-replica-target flux tuning reduces the total systematic uncertainty on the flux of the muon neutrinos in the peak of the FHC beam energy distribution from 8% to 5%, as shown in figure 5.7. The uncertainties in the flux covariance are then implemented as normalisations on the neutrino flux in bins of beam mode, neutrino sign and flavour, and neutrino energy E_ν.

Flux tune				
	ν_μ	$\bar{\nu}_\mu$	ν_e	$\bar{\nu}_e$
FHC	96.5%	87.6%	90.5%	77.8%
RHC	87.8%	96.2%	78.3%	91.1%

Table 5.2: Overall tune of each flavour component of the beam relative to pre-tuned total flux, after the hadronic parent production has been weighted to the NA61/SHINE T2K-replica-target data set [124].

Below are the bin boundaries from 0-30 GeV in E_ν for the flux normalisation parameters. The same binning in E_ν is used at ND280 and SK and for each beam mode with the following binning in GeV

- (FHC ν_μ) & (RHC $\bar{\nu}_\mu$): 0.0, 0.4, 0.5, 0.6, 0.7, 1.0, 1.5, 2.5, 3.5, 5.0, 7.0, 30.0
- (FHC $\bar{\nu}_\mu$) & (RHC ν_μ): 0.0, 0.7, 1.0, 1.5, 2.5, 30.0

Flux composition				
	ν_μ	$\bar{\nu}_\mu$	ν_e	$\bar{\nu}_e$
FHC	92.5%	6.2%	1.1%	0.2%
RHC	9.9%	88.9%	0.3%	0.9%

Table 5.3: Neutrino flavour composition of the beam for the FHC and RHC modes after the hadronic parent production has been weighted to the NA61/SHINE T2K-replica-target data set [124]. The RHC mode has a higher wrong-sign contamination than the FHC. The overall flux is decreased for all flavours, but the most significant change is in the wrong-sign components of the beam, particularly for the ν_e component.

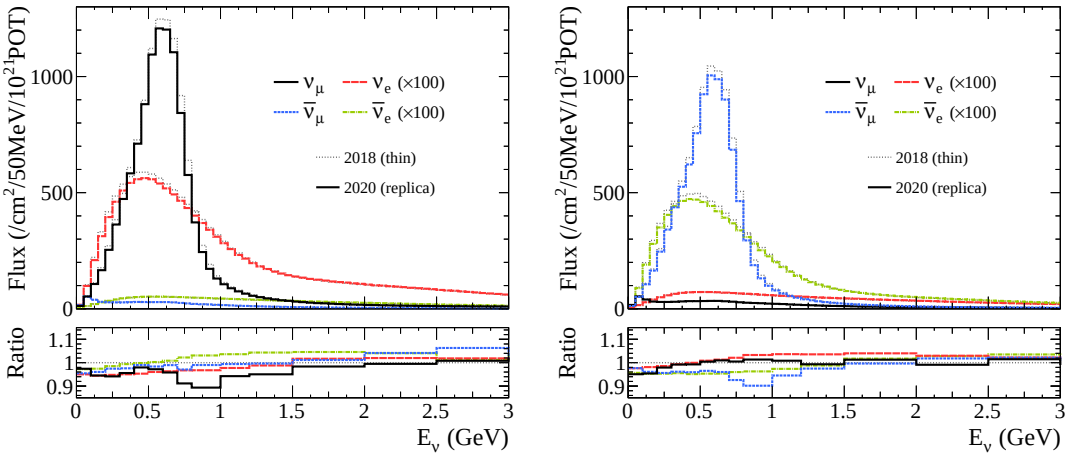

(a) FHC ν-mode beam

(b) RHC $\bar{\nu}$-mode beam

Figure 5.6: Predicted unoscillated SK flux composition for both ν-beam (FHC) and $\bar{\nu}$-beam (RHC) modes with the thin-target tune [122] used in the previous T2K analysis [20] in the grey lines, and the T2K-replica-target tune [124] in the solid lines. ν_e and $\bar{\nu}_e$ components are scaled up 100 times for shape comparison. The bottom panel of each subfigure shows the ratio of the replica-target tune to the thin-target tune. The most notable feature is a suppression of the ν_μ flux between 650 MeV and 1 GeV, visible in the replica-target tune, which may impact the measurement of the magnitude and shape of the $\nu_\mu(\bar{\nu}_\mu)$ disappearance.

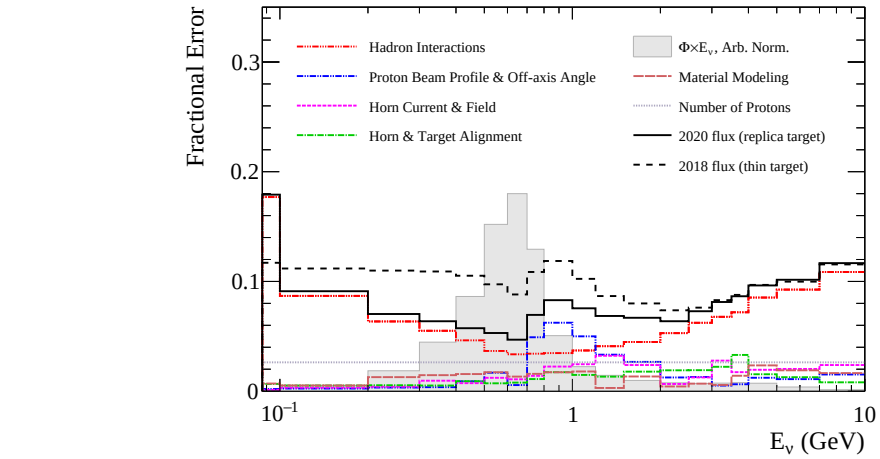

(a) FHC ν_μ flux

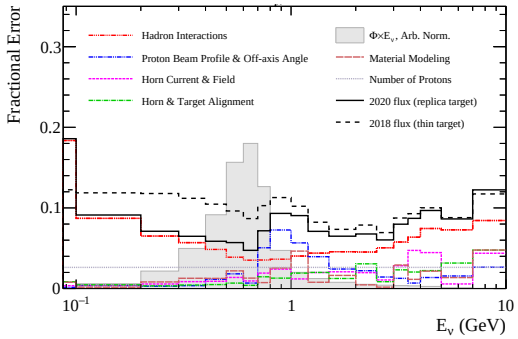

(b) RHC $\bar{\nu}_\mu$ flux

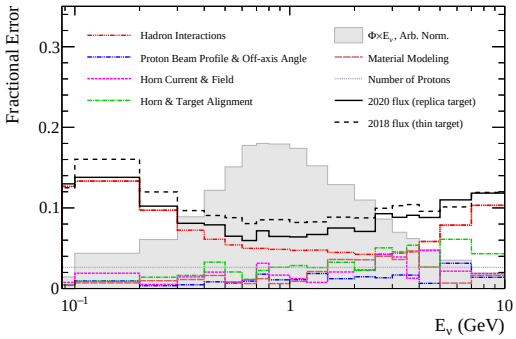

(c) RHC ν-background flux

Figure 5.7: Contributions to the uncertainty of the FHC ν_μ (a), and RHC $\bar{\nu}_\mu$ (b) and ν_μ (c) flux. The dashed line shows the total systematic uncertainty when using the thin-target tune [122] used in the previous T2K analysis [20], and the solid black line shows the total systematic uncertainty when using the T2K-replica-target tune [124]. The largest source of systematic uncertainty is from the hadronic interaction modelling, though there is a significant contribution in the right-sign (FHC ν_μ and RHC $\bar{\nu}_\mu$) flux to energy bins just above the peak of the flux spectrum (visible in the shaded grey histogram) from the impact of the off-axis angle of the detectors on the position of the peak.

- (FHC ν_e) & (RHC $\bar{\nu}_e$): 0.0, 0.5, 0.7, 0.8, 1.5, 2.5, 4.0, 30.0
- (FHC $\bar{\nu}_e$) & (RHC ν_e): 0.0, 2.5, 30.0

with eleven bins for the right-sign ν_μ-type component of each beam mode, five for the wrong-sign ν_μ-type component of each beam mode, seven for the right-sign ν_e-type component and two for the wrong-sign ν_e-type component.

5.8 Cross-section model

Neutrino-nuclear interaction cross-section modelling is the largest source of uncertainty in modern long-baseline neutrino-oscillation measurements in experiments like T2K or NOvA. The parameters allowed to vary in the fit are intended to give appropriate freedom to cover uncertainty associated with the various models. These parameters may be directly associated with the model parameterisation as implemented in the MC generator, or provide freedom to cover a known but unmodelled possible physical effect from alternative hypotheses. The relative contributions to the nominal T2K model prediction of the ND280 samples is discussed in section 6.2. A full list of the cross-section systematic parameters as implemented in the T2K oscillation analysis is listed in appendix C.

5.8.1 Charged-current quasi-elastic interactions

The T2K oscillation analysis relies on an assumption of the quasi-elasticity of interactions in order to reconstruct an unbiased estimator of the neutrino energy E_{QE}^{rec} to sample the neutrino flux. As such, CCQE is the target interaction mode, and the signal topology is defined as a single charged lepton with no mesons in the final state in order to enrich the CCQE component of the data collected. It is crucial to the analysis then, to accurately be able to predict the CCQE interaction rate in the T2K beam. The strength and final-state kinematic shape of the CCQE interaction is largely determined by vector and axial-vector nucleon form factors. These form factors are tuned to reproduce electron-nuclear scattering data [37].

The current T2K QE implementation uses the spectral function (SF) discussed in subsection 2.4.3 to determine nucleon kinematics, though previously the relativistic Fermi gas (RFG) discussed in subsection 2.4.2) was used in conjunction with a random phase approximation (RPA) to account for suppression of interactions with a low four-momentum transfer, Q^2, and forward going lepton kinematics due to nucleon correlations. This update of the T2K interaction model to incorporate the spectral-function model has allowed for a significant reduction in the systematic

uncertainty on the reconstructed neutrino energy. The equivalent elastic NC (NCE) interactions, are modelled using the same 1p1h formalism as the CCQE interaction.

A splined response to variations of the nucleon axial mass term in CCQE interactions, M_A^{QE}, is applied to each event, along with a normalisation in eight bins of Q^2 to give additional freedom to the SF which may not be covered by the dipole model.

The central value of $M_A^{QE} = 1.03$ is tuned to a global fit of external bubble chamber experiment data, with a strong constraint of ± 60 MeV applied. The normalisation parameters for the five lowest Q^2 bins ($Q^2 < 0.25$ GeV) have flat priors with no uncertainties. The normalisation parameters of the highest three Q^2 bins have prior constraints chosen based on differences between the shape of the dipole and z-expansion models [127]. These Q^2 normalisation parameters give freedom to the CCQE model which previously was covered by a series of RPA-related parameters to provide a low Q^2 suppression where the RFG model over-predicts the cross section.

Uncertainty in the nucleon removal energy and shell structure of the SF must also be accounted for. In the SF model, the parameter, ΔE_{rmv}, represents a translation of the SF along the E_{rmv} axis (see section 5.9), shifting the values of the removal energies for each shell. Notably, this parameter is not applied as a weight to the event as it does not change the cross-section value, but as a change to the shift in the momentum of the outgoing lepton, which has implications for the BANFF implementation. There are four correlated removal energy parameters, for each combination of $\nu/\bar{\nu}$ interacting on carbon or oxygen. ν_μ and ν_e are tied together on the assumption that lepton flavour has no effect on the nuclear structure. This is discussed in more detail in section 5.9 due to its unique complications for the BANFF fitter.

5.8.2 2p2h interactions

Multinucleon-quasielastic interactions are an irreducible background to the CCQE-enriched samples at T2K based solely on the final-state topology since the ejected nucleons may be difficult to reconstruct. The largest of these are for interactions with bound pairs of nuclei, or "two-proton two-hole" (2p2h) interactions. These interactions do have different final-state lepton kinematics to 1p1h CCQE interactions, and so E_{QE}^{rec} is a poorer estimate of the true neutrino energy, E_ν. T2K models these interactions using the "Valencia" model described by Nieves et al. [38].

The uncertainty parameterisation of the 2p2h model is split into three parts: shape, normalisation, and energy dependency. There are three normalisation parameters: one describing the normalisation of neutrino 2p2h interactions, one describing the normalisation of antineutrino interactions and one describing the normalisation of the ratio between the rate of events on carbon and the rate of events on oxygen. There are also two shape-controlling parameters, one for carbon

and one for oxygen, intended to cover differences between the Valencia and Martini [128] models, and their resulting final-state lepton kinematics. Then there are four energy-dependent 2p2h parameters with non-linear responses corresponding to low- and high-neutrino energies, E_ν, for neutrinos and antineutrinos. The ND280 data do not have the power to constrain these parameters and so they are not fit in the BANFF fit; however, they are allowed to freely vary in the MaCh3 fits to the ND280 data.

5.8.3 Single pion production

Single pion production (SPP) is the second most dominant interaction topology at T2K beam neutrino energies, and so modelling of single meson and gamma production is important for predicting the inclusive rate of neutrino interactions as well as constraining the dominant backgrounds in the measurement of CCQE-like events, since final-state interactions can cause the mesons to be absorbed by the nucleus.

SPP is the dominant interaction in the secondary signal topologies defined at the near and far detectors to control the dominant background to the the CCQE-like signal topology when the resultant pion is absorbed in final-state interactions or is not properly reconstructed. In SK this reconstruction is largely dependent on the final-state kinematics of the pion, whether it is above or below Cherenkov threshold, and whether it decays to either a Michel electron (for charged pions) or a pair of photons (for neutral pions).

For the T2K 2020 OA, the modelling of the resonant production of mesons is done by first modelling the excitation of the nucleon in the initial interaction and then modelling its subsequent decay into a meson or a photon, and a baryon. The excitation and decay of the nucleus in NEUT (see section A.2) are modelled using the Rein-Sehghal model [39] for events where the intermediate hadronic mass, W, is below 2 GeV, and for higher values of W they are modelled using the default model in the Pythia 5.72 event generator [129]. Kaon, eta and gamma production are modelled in the same framework.

Coherent pion production is modelled using the Berger-Seghal model [40] as well as the Rein diffractive model [41]. The coherent interaction does not proceed according to the impulse approximation, as the neutrino scatters off the nucleus as a whole, and final-state interactions are not modelled.

Whilst the T2K CCQE interaction model has been updated to use the SF nuclear model, the pion production interaction model is still based on the RFG nuclear model. For the hard scatter, two parameters give freedom to the form factors in the Rein-Seghal CC resonant pion production model: the resonant axial mass M_A^{RES} and the vector form-factor term C_5^A. Two splined response

parameters are used to control the non-resonant $1/2$-Isospin background component of single pion production. One of these is specifically for antineutrino events that produce low-momentum pions, although this parameter is not fit at the near detector due to lack of sensitivity. Two normalisation parameters control the rate of CC-Coherent pion production for carbon and oxygen.

5.8.4 Multi-pion production and DIS

Inelastic scattering and multiple-meson production, whilst subdominant at T2K's beam energies, become dominant in the high-energy tail of the flux. In the inelastic regime the interactions are dependent on the underlying quark structure and so the cross section is calculated using Parton Distribution Functions (PDFs) [42] and extended below the standard momentum-transfer regime of DIS to more typical T2K values of $Q^2 < 1.5$ GeV^2 using "Bodek-Yang modifications" [130][131]. DIS is restricted to hadronic system invariant masses, W, above 1.3 GeV and only multi-pion interactions are considered in order to avoid double counting with the SPP models. For values of W below 2 GeV an interpolation between the resonant production of $\Delta(1232)$ and the DIS model is used. For values of W above 2 GeV only DIS events are modelled, using the Lund string model implemented in PYTHIA [132].

Two parameters control the Bodek-Yang corrections as applied to either DIS or MultiPi events. These parameters are "on" in the prior model and can be varied between on and off continuously.

One parameter reweights between the NEUT nominal model and the model obtained if using the values used in the AGKY model for model parameters which are correlated between the two models [133] as implemented in GENIE [134]. Finally, two overall normalisation parameters are applied for neutrino and antineutrino interactions, respectively.

5.8.5 Final-state interactions

Pion final-state interaction parameters control the rates of pion interactions within the nucleus in the fit. Before exiting the nucleus, pions can undergo charge exchange, production and absorption processes, as well inducing further hadron production. These parameters have correlations from fits to T2K cross-section analysis data at ND280.

Other parameters

Two normalisation parameters are included to give freedom to the uncertainty on the Coulomb correction to the nuclear potential. Normalisation parameters varying the ratio of ν_e to ν_μ and varying the ratio of $\bar{\nu}_e$ to $\bar{\nu}_\mu$ in order to control the inherent wrong-flavour contamination of the beam. A series of NC parameters give a crude normalisation to the rate of NC interactions, in-

cluding specific parameters for coherent pion production and gamma production. Two separate normalisation parameters are used for other NC interactions at ND280 and SK, the former is not propagated directly and the latter is not fit but a constraint is placed on the SK parameter through imposed prior correlations. The ND280 samples have efficient NC background rejection from the requirement of a muon track and so the ND280 fit does not have a strong ability to constrain these parameters given the small rate of events; thus, the NC gamma production parameter is fixed in the BANFF fit but left free to vary in the MaCh3 fit at ND280. Finally a CC-Misc. parameter controls normalisations of other CC processes such as kaon and eta production.

5.9 Nucleon removal energy

In electron-nuclear scattering experiments, when accounting for the known kinematics for initial- and final-state products in QE 1p1h scattering, there is an energy deficit in the final-state products [50]. If we assume that these interactions happen on-shell, in most models this deficit is accounted for by the energy required to overcome the nuclear potential and by allowing the electron to interact on an unbound nucleon. This energy is referred to as the "binding", E_b, or "nucleon removal energy", E_{rmv}, and is understood to be the energy required to remove one nucleon from the nuclear ground state, leaving one hole. Note that for protons there is also an additional correction in order to take into account the Coulomb potential. Similarly, in neutrino 1p1h interactions these same nuclear effects come into play and are vital to understand in the reconstruction of the initial neutrino energy E_ν^{rec}. As such, a binding energy term E_b features in the calculation of E_ν^{rec},

$$E_\nu^{\text{rec}} = \frac{m_p^2 - (m_n - E_b)^2 - m_l^2 + 2(m_n - E_b)E_l}{2(m_n - E_b - E_l + p_l \cos\theta_l)}, \tag{5.16}$$

along with the proton, neutron and lepton masses, m_p, m_n and m_l, and so an incorrect value of E_b or an incorrect modelling of the lepton kinematics p_l, θ_l will lead to a bias in the extracted value of Δm^2_{32} in the oscillation analysis. The m_p and m_n terms are swapped for $\bar{\nu}$ events. In order to distinguish between the term used in the reconstruction of the neutrino energy and the modelled, true missing energy in the interaction, the term in equation 5.16 will be referred to as E_b and the energy required to remove the nucleon from its bound state in the modelling of the neutrino interaction will be referred to as E_{rmv}[1].

In previous T2K analyses, the uncertainty on E_{rmv} had a significant impact on the results, in part because of a less sophisticated nuclear model than is currently used and in part because of

[1] The terms E_b and E_{rmv} are sometimes used interchangeably in the literature, and so previous T2K publications may use E_b for this quantity as the RFG model has a single-value interaction energy which is analogous to the E_b term in the energy reconstruction equation, but this is not appropriate in the Spectral Function description which has a spread of values of E_{rmv} for a given nucleon momentum which are dependent on the underlying multipeak shell structure.

an inability to give appropriate freedom to this uncertainty in the fit. Since the value of E_{rmv} is understood not to primarily impact the cross section of an event, but rather the kinematics of the outgoing lepton in an event, an event reweighting to change the probability of an event occurring will not provide the freedom needed to vary the MC correctly. To properly assign uncertainty to this systematic effect in the fit, variations of this parameter must correspond to variations in the fitted observables used to calculate E_ν^{rec}. As uncertainties are applied to the MC and not the data, the value of E_b used in equation 5.16 is left fixed. Instead, variations are applied in a way which changes the charged lepton momentum p_l of the MC events based on the assessed impact of varying E_{rmv} in the MC generator, NEUT. This uncertainty treatment does not directly alter the interaction cross section, but does change the available final-state phase space. It is grouped with the other interaction parameters and referred to as part of the cross-section model. The different methods used to constrain the E_{rmv} uncertainty parameter in the T2K fitters are discussed below.

At the generator level

To reduce the uncertainty in E_{rmv} the impact that changing the relevant underlying model parameters in the MC generator has on the events must be assessed. As discussed in section 2.4, there are many different nuclear models and the binding energy term in the calculation of E_ν^{rec} is related to different underlying parameters. This affects the way in which we must implement the uncertainty on E_{rmv} must be implemented. Until the 2020 analysis, T2K used a RFG nuclear model in which the removal energy is constant for a given initial nucleon momentum, k, and is determined by a fixed "interaction energy" relating to the Fermi surface of the nucleus. In this model, the E_{rmv} parameter in the fit corresponds to this single fixed interaction energy value and can simply be varied around some central value defined in the generator as

$$E_{\text{rmv}} = \sqrt{p_F^2 + M_N^2} - M_N,$$

where p_F is the Fermi momentum, M_N is the on-shell nuclear mass. This model's simplicity, along with the large bias impact expected on the oscillation parameters, required that a conservative estimate of the uncertainty was used. The RFG model is unable to predict any of the nucleus' underlying shell structure that is expected to impact the removal energy for nucleons in different orbitals. In addition, the momentum variation approach had not been implemented and so significant bias was found when trying to implement this parameter. As a result of this, E_{rmv} was not included in the near detector fit, and an heuristic parameter was applied at the far detector constructed from unoscillated SK spectrum prediction studies with E_{rmv} set to its maximum and minimum values according to a large uncertainty.

For the 2020 oscillation analysis, the move to using the Benhar Spectral Function (SF) nuclear model for 1p1h events meant that a tighter prior uncertainty could be used based on electron scattering experiment data. The SF model incorporates the nuclear shell structure as well as allowing for random values of the removal energy in the event generation for a given nucleon momentum drawn from the SF. No longer a single value, the uncertainty implementation was updated to suit, and so a variation of a systematic uncertainty parameter, ΔE_{rmv}, in this analysis, corresponds to a translation of the whole SF distribution along the axis corresponding to the generated E_{rmv} in the interaction as indicated by the red arrow in figure 5.8(b). A comparison of the different possible combinations of E_{rmv} and initial nucleon momentum, k, is shown in figure 5.8 where the difference in available phase space for the RFG and SF models is very clear. The distinct peaks for different nuclear shells can be seen in the SF, showing up as dark bands of increased cross-section at particular values of E_{rmv} in figure 5.8(b).

Heuristic parameter in previous analyses

During the late stage of T2K's 2017 analysis, it was found that fitting a typical cross-section uncertainty comprised of response weights in bins of E_ν and Q^2 to fake data (MC) generated with a different true (RFG) removal energy led to a significant bias in the measurement of Δm^2_{32}. The impact of these biases on the lepton kinematics can be seen in figure 5.9. To avoid this bias, an ad-hoc parameter was included to ensure that this effect was covered in the total uncertainty on the extraction of Δm^2_{32} in the analyses between 2017 and 2020s. This was constructed by fitting the T2K model, without the removal energy parameter, to fake data at the near detector corresponding to the minimum and maximum values of the E_{rmv} parameter. For each of these fake-data fit results, and the BANFF Asimov fit result, the postfit covariance matrices were used to make predictions of the far detector E_ν spectrum. From these predictions, reaction and kinematic breakdowns at these three values of the removal energy parameter were used to create reweighting splines for each of the far detector bins. This still led to some bias, though it was reduced. A correction to the oscillation parameter $\Delta\chi^2$ surface was made in order to appropriately broaden, or smear, the uncertainty contours.

Event by event variations

After establishing how the ΔE_{rmv} parameter is applied at the generator level, a method to apply these changes during a fit to data for some arbitrary value of the parameter which does not require regeneration of the MC was developed. The fitted E_{rmv} parameter relates the true neutrino energy E_ν to the charged lepton kinematics p_l, θ_l, therefore MC events are generated and binned in three-

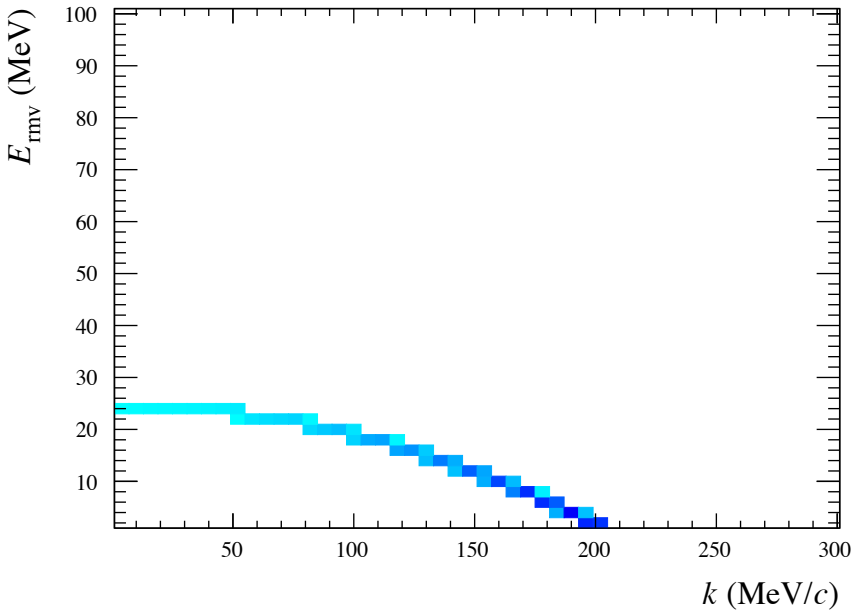

(a) Relativistic Fermi gas (RFG)

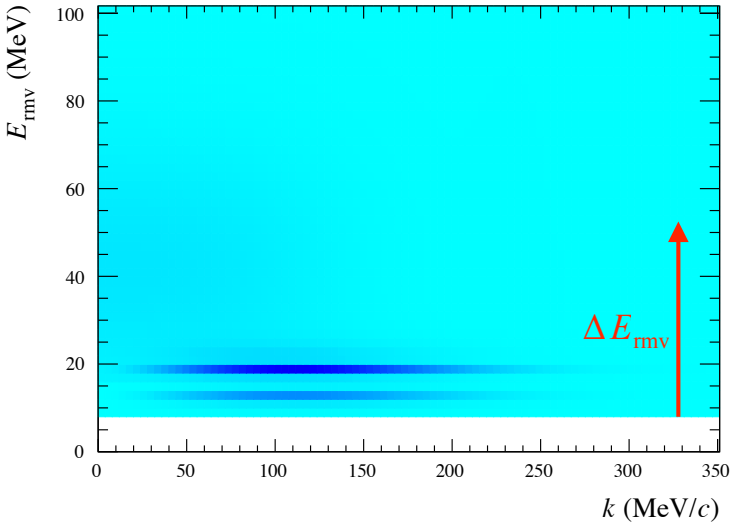

(b) Benhar spectral function (SF)

Figure 5.8: MC generated distributions of removal energy E_{rmv} and initial nucleon momentum k for the RFG (a) and Benhar SF (b) models where darker shades of blue represent values of higher probability. Whilst for RFG there is a single value of E_{rmv} for a given value of k, the spectral function takes a range of values with different probabilities with peaks resulting from the shell structure of the nucleus. The uncertainty on the removal energy in the SF, ΔE_{rmv}, is implemented as a shift of the entire spectrum along the E_{rmv} axis, as indicated by the red arrow.

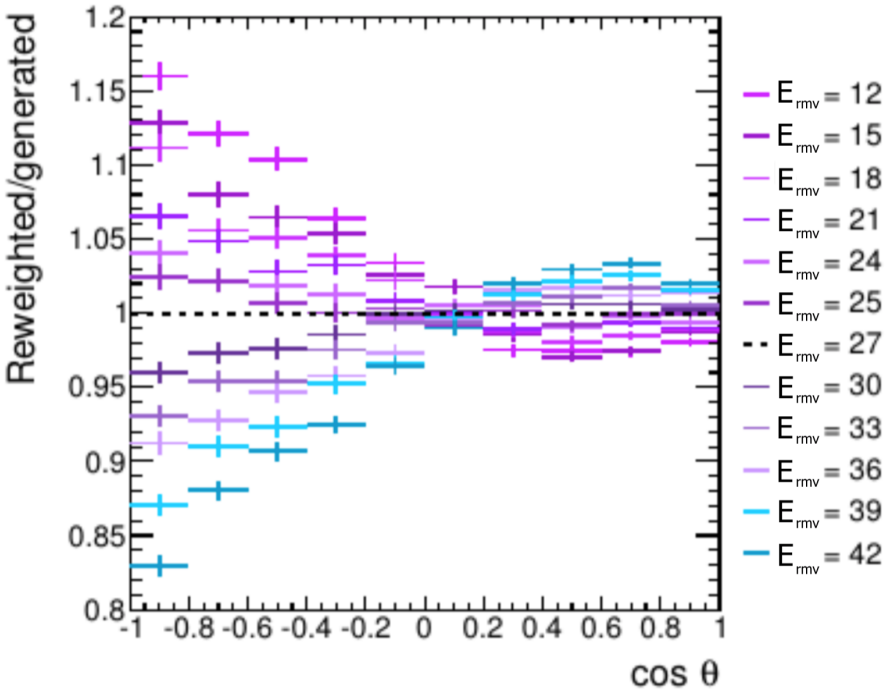

Figure 5.9: Ratios of the nominal MC reweighted to different values of E_{rmv} in the old reweighting scheme (in bins of E_ν and Q^2) to MC generated at that value of E_{rmv} for events binned in cosine of the lepton angle, $\cos\theta_l$. In the RFG, E_{rmv} is a single value (25 MeV for C and 27 MeV for O) and the parameter corresponds to this directly rather than the shift in the template as in the SF implementation. A lack of phase space coverage in the nominal MC prevents reweighting in bins of E_ν and Q^2 and prevents accurate reproduction the desired distribution in lepton kinematics.

dimensional histograms of E_ν, p_l and θ_l for a range of E_{rmv} values including the NEUT-generated value of $E_{\text{rmv}} = 0$. Histograms are generated for each ν flavour and target nuclei combination, and at integer multiples of the uncertainty of the parameter $\sigma_{E_{\text{rmv}}}$ within the physically allowed range. These histograms are then profiled into templates of average p_l for E_ν and θ_l as shown in figure 5.10. Taking the difference in average momentum values for each E_ν, θ_l bin between templates at one of the varied values and the generated value of $E_{\text{rmv}} = 0$, the momentum shift Δp_l was calculated, which was then applied to an event in order to change the expected momentum at the generated E_{rmv} value to the momentum corresponding to the desired E_{rmv} value. Using these values of E_{rmv} and Δp_l for knots, interpolation splines were constructed to continuously vary E_{rmv} and calculate an appropriate momentum variation for an event of a given E_ν and θ_l.

Whilst this variation parameter is not intended to change the interaction cross-section value, variations of E_{rmv} lead to changes in the likelihood of event generation in some regions of phase space, and so for some of the templates there may be very low numbers of, or even zero events. If an E_ν, θ_l bin in the nominal template is not populated, the spline is left null and not applied, as these templates are generated with large enough MC data sets to assume that there will not be selected events in this region of phase space in the production MC which is fit to the data in BANFF. If, however, the bin in the nominal template is populated and another template is not populated, one of two approaches is used based on population of this bin in the adjacent templates. If this knot is at the most extreme value of E_{rmv} used in the spline, or if none of the more extreme values away from the central value are filled, the nearest filled value or the nominal value is used. As such the variations cap at a momentum which has been calculated within templates with generated MC and the spline does not extrapolate beyond the properly generated MC values. If the knot is internal, such that there do exist values at more extreme values of E_{rmv}, the splines are constructed using a "dummy" knot to satisfy the need for equally spaced knots for computational efficiency, such that they have the same shape as a spline constructed without this knot. These interpolation splines then are used within the fit to calculate a momentum variation that is applied to the charged lepton's generated reconstructed momentum for a given combination of neutrino flavour, target nuclei, E_ν,p_l and ΔE_{rmv}.

Implementation in BANFF and avoiding a discontinuous likelihood

In the 2020 T2K ND280 fit, data is binned in the muon kinematic variables; thus, variations of the muon momentum result in events migrating between bins in the fit , leading to discrete changes to the likelihood surface when an event moves across a bin boundary. The resultant likelihood surface is therefore discontinuous and the gradient and covariance cannot be calculated because

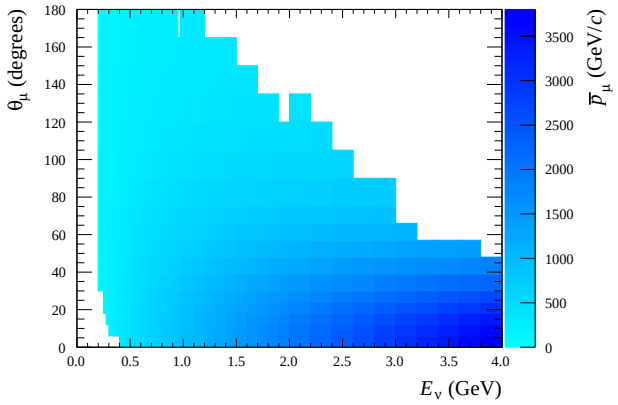

(a) $\Delta E^{\nu,\mathrm{C}}_{\mathrm{rmv}} = -10$ MeV

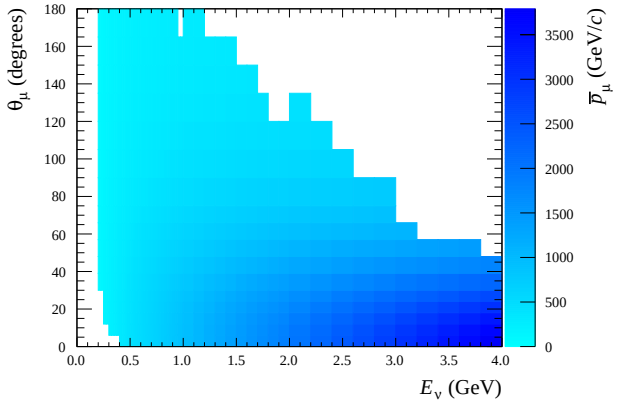

(b) $\Delta E^{\nu,\mathrm{C}}_{\mathrm{rmv}} = 0$ MeV

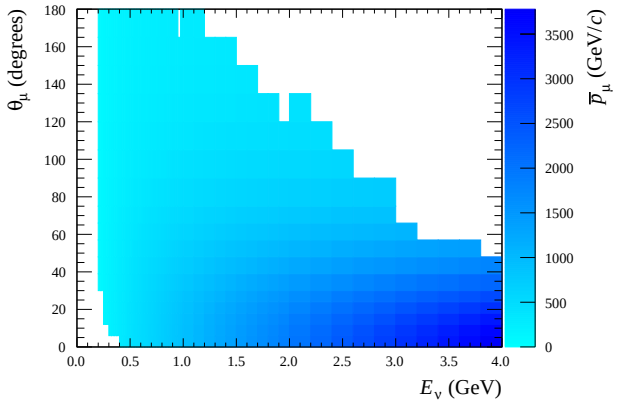

(c) $\Delta E^{\nu,\mathrm{C}}_{\mathrm{rmv}} = +10$ MeV

Figure 5.10: Average-momentum templates in E_ν, θ_μ for ν_μ CCQE interactions on carbon. Generated from profiling three-dimensional histograms in E_ν, θ_μ, p_μ, the average muon momentum represented by the z-axis is used to calculate a typical momentum shift for an event in a bin of E_ν, θ_μ by subtracting the value in the nominal template from the template corresponding to the desired $\Delta E_{\mathrm{rmv}^{\nu,\mathrm{C}}}$ value. Splines can be used to interpolate shifts, i.e. changes in p_μ, between the template values (knots) allowing the variation to be applied continuously during the shift. Care must be taken in bins in which there are too few events to pass a cut that is implemented in order to ensure sufficiently small relative uncertainty.

the derivatives cannot be found. Minuit, and by extension BANFF, are unable to fit this type of parameter. To fit these kinematically varying parameters in BANFF another approach is needed. Using the variations in ΔE_{rmv} to produce MC sample predictions with the ΔE_{rmv} parameter varied, events can be reweighted within a bin to reflect the expected number of events after a variation of the parameter has been applied and all of the events within the distribution have had their lepton momentum varied. These predictions can be used as templates (figure 5.11) for "binned splines", which then allow for continuous reweighting of the events. This procedure should reproduce shape changes of the sample exactly at the knot values used when compared to event-by-event momentum variations.

This approach, whilst protecting the fitter from pathological likelihood behaviour, does have some shortcomings. The intention of these momentum variations is to move events to regions of phase space which may be unpopulated in the generated MC; however, if the nominal MC in the bin that is used in the fit is completely empty, no reweight can be applied. This can also lead to a small change in the overall normalisation of the sample. Additionally, the templates should be split to capture as much of the original freedom of the parameter as possible, although introducing these breakdowns may also significantly affect available statistics. For example, for the E_{rmv} momentum variations, the event's true E_ν should be considered; however, since the fit binning uses lepton kinematics as a proxy for the neutrino energy, E_ν ranges within any given bin may be small which can lead to a dramatic drop-off in available statistics and too few bins with complete splines. Luckily, this very issue provides some constraint on the E_ν values within a range of θ_l or $\cos\theta_l$. Similarly, due to the low statistics of ν_e events at the near detector, many of the bins do not have enough statistics in the nominal E_{rmv} template to build splines for ν_e events. In T2K's fits, ν_e and ν_μ events are tied together to reduce the impact of this given the low statistics of the ν_e MC events in the ND280 samples. Thus, the four parameters included in each fitter are, $\Delta E^{\nu C}_{rmv}$, $\Delta E^{\bar{\nu} C}_{rmv}$, $\Delta E^{\nu O}_{rmv}$, and $\Delta E^{\bar{\nu} O}_{rmv}$, and their prior central values and uncertainties informed by electron-scattering data [48][50] are presented in table 5.4. Prior correlations based on theoretical considerations are given in table 5.5. Comparisons of the event-by-event variations and the binned reweighting splines used by MaCh3 and BANFF, respectively, and their compatibility are discussed in chapters 7 and 8.

5.10 Near detector likelihood

The near detector likelihood is calculated by comparing the ND280 data to the prediction of the combined flux, cross-section and ND280 detector models. The simulated MC events of neutrino interactions in the ND280 during neutrino beam operation are treated in the same manner as the

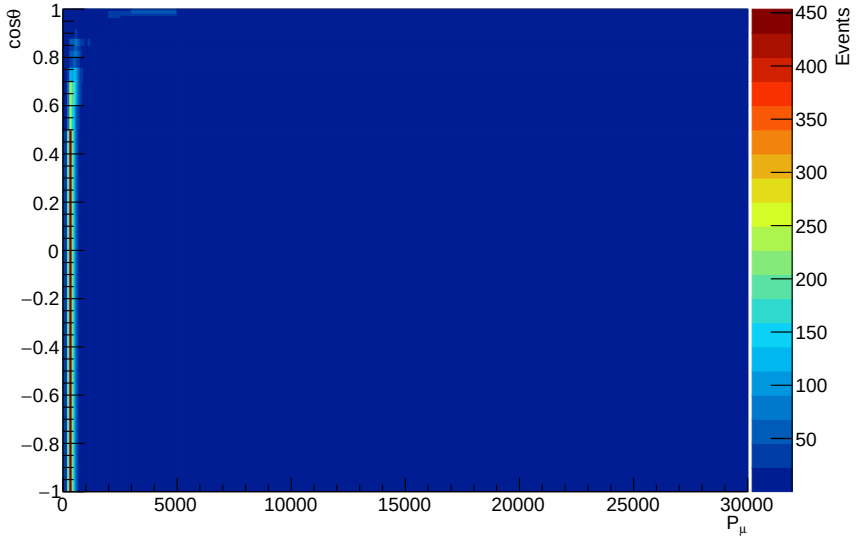

(a) Full template

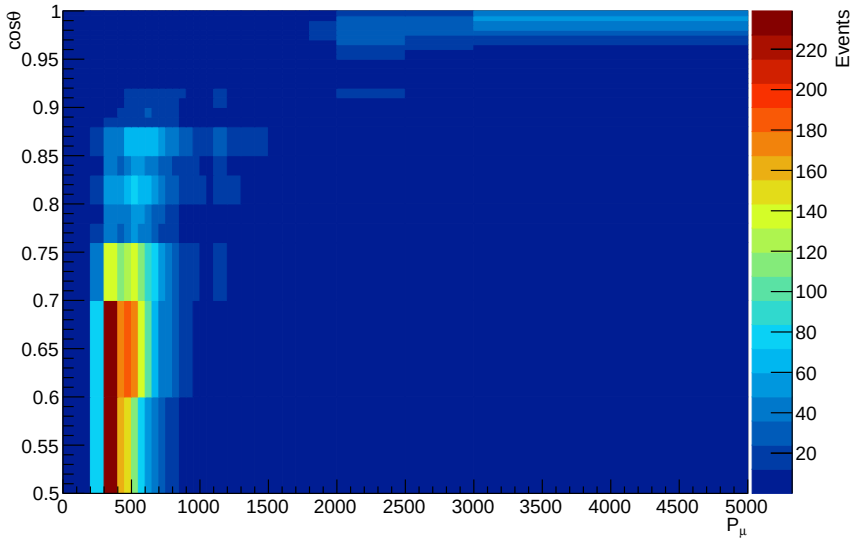

(b) Template in region of interest

Figure 5.11: Templates of event rates made from the sample distribution of specific subsets of CCQE events in the CC0π sample for a given value of E_{rmv}. Templates are split by neutrino flavour and nuclear target for each value of E_{rmv}. Weights are then constructed by taking the ratio of the event rate in each bin for each value of E_{rmv} to the event rate in that bin in the nominal ($E_{\text{rmv}} = 0$) template. An interpolation spline is then constructed from these weights for every bin in the sample.

Parameter	Central Value	Uncertainty
$\Delta E_{rmv}^{\nu C}$	+2 MeV	6 MeV
$\Delta E_{rmv}^{\bar{\nu} C}$	0 MeV	6 MeV
$\Delta E_{rmv}^{\nu O}$	+4 MeV	6 MeV
$\Delta E_{rmv}^{\bar{\nu} O}$	0 MeV	6 MeV

Table 5.4: Prior central values and uncertainties on the four CCQE removal energy parameters. A prior shift or weight is applied for MaCh3 and BANFF, respectively, based on a tuning to the electron scattering data [48][50].

Parameter	$E_{rmv}^{\nu C}$	$E_{rmv}^{\bar{\nu} C}$	$E_{rmv}^{\nu O}$	$E_{rmv}^{\bar{\nu} O}$
$E_{rmv}^{\bar{\nu} C}$	1	0.7	0.7777	0.6527
$E_{rmv}^{\nu C}$	0.7	1	0.6527	0.7777
$E_{rmv}^{\bar{\nu} O}$	0.7777	0.6527	1	0.7
$E_{rmv}^{\nu O}$	0.6527	0.7	0.7777	1

Table 5.5: Prefit correlations between the removal energy parameters in the BANFF and MaCh3 ND280 fits.

data when selecting events for different samples, but the predictions can be varied based on the "true" simulated properties of the events. Each parameter in the model can affect the likelihood of the combination of all of the parameter values through both the change in the sample prediction, and also through a penalty term applied as a prior uncertainty on the parameter.

5.10.1 ND280 detector systematic uncertainties

The ND280 detector systematic uncertainties are evaluated using the HighLAND 2 analysis package during the development of selections for both the ND280 fit and a dedicated program of cross-section measurements. The systematic uncertainty treatments are then ported to the PSYCHE package that BANFF uses to load events and to apply the relevant systematic weights. Currently there is no implementation to directly vary the detector systematic uncertainties in the fit, in part due to the non-Gaussian behaviour and non-multiplicative dependencies of the weights on the underlying parameters. In addition, detector systematic variations can affect the observables which are binned in the fit and can cause problematic behaviour in the likelihood. A covariance matrix is generated for variations on the number of events in the observable bins used in the fit based on variations of the underlying parameters in `psyche`. This matrix is then reduced in size from the

total number of bins used in the fit to a more computationally manageable number by merging adjacent bins that have an overall uncertainty within 7% of each other. This process is iterated until the matrix can be reduced no more. The uncertainty is then applied as a normalisation[2] on the events in each "detector bin" according to the correlation matrix.

5.11 Far detector likelihood

The far detector likelihood contains both the flux and cross-section models, but also the SK detector parameters and the modelling of the neutrino oscillations according to the PMNS description. The SK detector parameters are tuned using fits to atmospheric and cosmic ray muon data as well as dedicated calibration runs. A momentum scaling correction is applied which may shift events between bins in a similar way to the nuclear removal energy parameters. At SK, the momentum scaling and the nuclear removal energy parameters are implemented with a binned additive template spline method which differs from the binned multiplicative template splines applied at the near detector. This implementation avoids issues associated with much sparser regions of the samples in which the generated MC has no CCQE events, particularly of different neutrino flavours. The underlying nuclear removal energy parameter is the same, and so these parameters are fully correlated. The simultaneous fit uses the kinematic shifting implementation directly at both near and far detectors. Two large systematic uncertainties at SK are the pion secondary interactions, as well as the photonuclear effect in which photons can be absorbed by the nucleus which then can undergo subsequent decay. Both of these processes can affect the number of Cherenkov rings seen in the detector. The pion secondary interactions at SK are not correlated with those at ND280.

5.11.1 PMNS likelihood

The final and key part of the T2K model is the PMNS mixing model. T2K is not sensitive to the the two "solar" mixing parameters θ_{12} or Δm^2_{12} and so these are constrained using the 2019 global best-fit result reported by the particle data group (PDG) [34]. T2K does have sensitivity to the remaining four mixing parameters, θ_{13}, θ_{23}, δ_{CP} and the mass splitting Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO). As many of these mixing terms are enclosed in sine or cosine functions in the oscillation probability, the degrees of freedom given in the likelihood are in $\sin^2\theta_{13}$ and $\sin^2\theta_{23}$ rather than the bare mixing angle. The standard fits give freedom directly to δ_{CP} rather than $\sin\delta_{CP}$, though studies are also conducted with flat prior uncertainties in the Jarlskog invariant to check consistency with

[2] As these are binned in the observables of the fit, they are sometimes called "observable normalisation" paramters or "ObsNorm" parameters to distinguish them from the underlying detector parameters in `psyche`. In this thesis the "detector parameters" refers to these unless explicitly stated that it is referring to the `psyche` parameters.

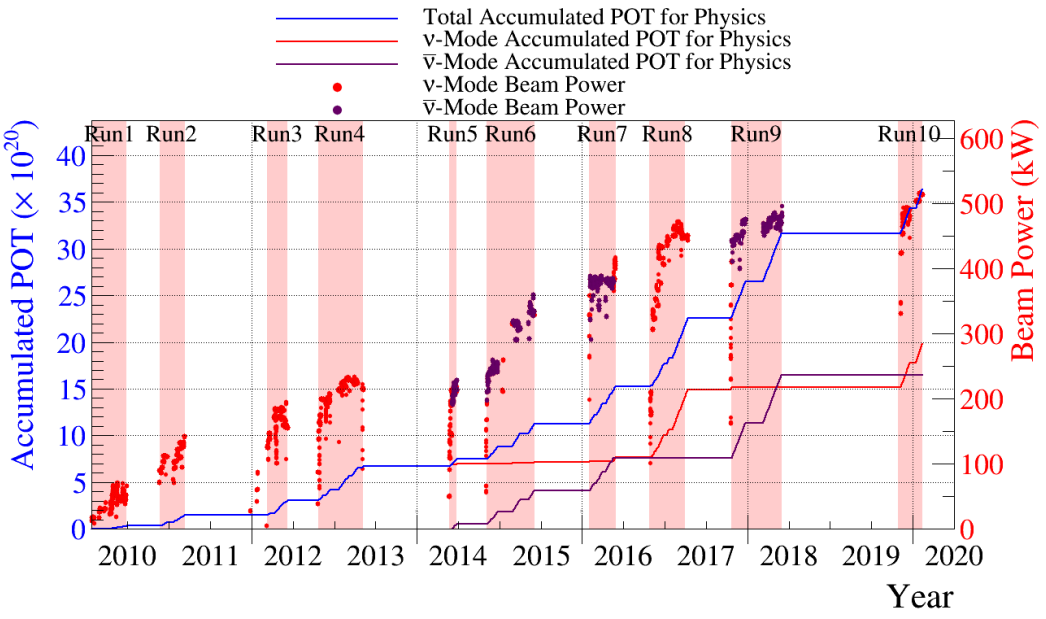

Figure 5.12: Accumulated protons on target (POT), overlaid with beam intensity, versus time for the T2K experiment. The ND280 analysis uses runs 2 to 9, and the INGRID and SK analyses use runs 1 to 10, with run-by-run POT listed in table 5.7.

less PMNS model-dependent parameterisations. An additional Gaussian constraint can be applied to the value of θ_{13} from reactor experiment measurements, as discussed in section 5.3, from the PDG [34].

5.12 Data collected

The data collected at ND280 and SK in terms of POT beam-spill equivalent used in the analysis is presented in table 5.7. Figure 5.12 shows the beam power over time, and total accumulated POT of the T2K experiment. The increase in beam power has allowed the accumulated POT to be more rapidly increased in recent years. The data collected at the near detector has doubled since the 2019 statistical update of the previous analysis [135][20], which used only runs 2-6, corresponding to an increase in beam exposure of 5.731×10^{20} POT (+99%) in FHC-beam mode and 4.476×10^{20} POT (+116%) in the RHC-beam mode for a total increase of 1.02×10^{21} POT (+106%).

This has allowed for the inclusion of new samples discussed below, particularly in RHC where the samples are now split by pion-multiplicity for both the $\bar{\nu}$ and intrinsic ν-background samples as in FHC. At the far detector, SK's beam exposure increased by 4.726×10^{20} POT (+33%) in FHC, and no more RHC data was taken, increasing the overall exposure by 15%. This resulted in 18 additional electron-like events and 75 more muon-like events included in the SK samples. The

	FHC			RHC	
	1Rμ	1Re	1Re1d.e.	1Rμ	1Re
events	318	137	14	94	16

Table 5.6: Total number of candidates collected over T2K runs 1-10 for each of the SK samples. The FHC samples have more collected data due to the for similar POT due to the higher rate of neutrino production possible in T2K's beam, and the higher interaction cross-section for neutrinos compared to antineutrinos.

total number of events in each sample over T2K runs 1-10 is presented in table 5.6.

5.13 Samples

The collected data is then categorised into samples using selections based on cuts to the data. A cut excludes data which does not meet the criteria of a selection. As the primary target interaction mode in the T2K analysis is CCQE, the cuts applied are designed to produce a sample of events with a high CCQE purity and selection efficiency. Purity is defined as the fraction of selected simulated events which are the true signal definition, and efficiency is the fraction of all true simulated signal events which are selected. The efficiency and purity of the ND280 selections used in the BANFF fit are given in table B.1 in appendix B. After cuts which are used to ensure particle identification and reconstruction purity and efficiency are high, the reconstructed topology of the event is used to select the desired interaction mode. A full list of the reconstructed topology cuts applied to the data at ND280 and SK are listed in tables 5.8 and 5.9, respectively. A comparison of the ND280 and SK reconstructed pion topology selections for the same desired true interaction mode is given in table 5.10. The CCQE-rich sample, CC0π, at the ND280 requires a muon track beginning in the fiducial volume of one of the FGDs with no reconstructed mesons in the final state. The CCQE-rich topologies at SK, 1Rμ and 1Re, are events with a single reconstructed Cherenkov ring within the beam spill time window, the sharpness of which is used to classify these as μ-like or e-like rings which separate the ν_μ disappearance channel sample and the ν_e appearance channel. As there are significant contributions of CC-Resonant pion production to the total cross section at T2K energies, and the reconstruction efficiency of pions at SK, as well as final-state absorption and production can impact whether an event is selected as 1Rμ or 1Re, a second sample, CC1π, at the near detector selects for events with a single reconstructed muon and a single reconstructed pion topology. At SK the 1Re1.d.e. sample selects events with a prompt electron-like ring, followed by a distinct electron ring after an initial delay which is used to tag the decay of low-momentum

Run number	Run start	Run end	Beam mode	POT ($\times 10^{19}$) ND280	POT ($\times 10^{19}$) SK
1	Jan. 2010	Jun. 2010	ν	—	3.26
2	Nov. 2010	Mar. 2011	ν	7.93	11.22
3	Mar. 2012	Jun. 2012	ν	15.81	15.99
4	Oct. 2012	May 2013	ν	34.26	35.97
5	May 2014	Jun. 2014	$\bar{\nu}$	4.35	5.12
			ν	—	2.44
6	Oct. 2014	Jun. 2015	$\bar{\nu}$	34.09	35.46
			ν	—	1.92
7	Feb. 2016	May 2016	$\bar{\nu}$	24.38	34.98
			ν	—	4.84
8	Oct. 2016	Apr. 2017	ν	57.31	71.69
9	Oct. 2017	May 2018	$\bar{\nu}$	20.54	87.88
			ν	—	2.04
10	Oct. 2019	Feb. 2020	ν	—	47.26
	Total		ν	115.31	196.63
	Total		$\bar{\nu}$	83.36	163.44
	Total		$\nu + \bar{\nu}$	198.67	360.07

Table 5.7: Collected protons-on-target (POT) included in the analysis of T2K data at ND280 and SK. The recorded POT at INGRID is similar to that of SK. The ND280 collected POT is reduced due to detector down-time or bad-data flags at ND280.

pions to Michel electrons via a short-lived muon. Finally, a third topology at the near detector, CC-Other, is used to select any other CC interaction with a reconstructed muon. This allows for the total CC interaction rate to constrain models for background processes of the signal topologies for the CC0π and CC1π samples at the near detector.

As the near detector is able to discriminate between positive and negative muon-like tracks, three sets of samples are used based on the charge of the muon and the neutrino beam mode, or horn current as outlined in table 5.8. The FHC beam is almost all (94%) ν and so there are the three topology-based samples outlined above for the selection of ν_μ events. The RHC beam has a higher wrong-sign contamination (10%), which along with the higher ν-cross section leads to a larger wrong-sign event rate, and so there are six samples, three for the events with μ^+-like tracks which are from $\bar{\nu}_\mu$ interactions, and three for the events with μ^--like tracks which are from the irreducible ν_μ background in the $\bar{\nu}$-mode beam.

SK is unable to discriminate between the charge of the final-state leptons from ν or $\bar{\nu}$ interactions, and so there are only two sets of samples, which are listed in table 5.9. This is why the background samples at ND280 play an essential role in estimating the ν contamination of the $\bar{\nu}$ samples and thus the estimation of any CP-violating effects in the data. Currently, there is no 1Re1.d.e. sample for the $\bar{\nu}$-mode beam.

Horn Curent	Target Neutrino	Sample	Nμ^-	Nμ^+	Nπ
FHC	ν_μ	CC0pi	$\geqslant 1$	0	0
		CC1pi	$\geqslant 1$	0	1
		CCOther	$\geqslant 1$	0	>1
RHC	$\bar{\nu}_\mu$	CC0pi	0	$\geqslant 1$	0
		CC1pi	0	$\geqslant 1$	1
		CCOther	0	$\geqslant 1$	>1
	ν_μ-bkg	CC0pi	$\geqslant 1$	0	0
		CC1pi	$\geqslant 1$	0	1
		CCOther	$\geqslant 1$	0	>1

Table 5.8: Selected cuts to the reconstructed event topology of the samples used in the BANFF and MaCh3 fits to the ND280 data. All of the samples currently used at the near detector are for ν_μ-like events. There are not currently any well-developed ν_e samples due to the low intrinsic ν_e content of the beam. The charge of the reconstructed muon track is used to determine the sign of the neutrino, which is particularly important to constrain the size of the intrinsic wrong-sign component of the RHC beam. Finally, the number of reconstructed pions or hadrons is used to enrich each sample in the desired interaction modes: CCQE, CC-Res, and Mult-π/DIS.

Horn Curent	Target Neutrino	Sample	μ-like rings	e-like rings	Decay Electrons
FHC	ν_μ	1Rμ	1	0	0
	ν_e	1Re	0	1	0
	ν_e	1Re	0	1	1
RHC	$\bar{\nu}_\mu + \nu_\mu$	1Rμ	1	0	0
	$\bar{\nu}_e + \nu_e$	1Re	0	1	0

Table 5.9: Selected cuts to the reconstructed event topology of the samples used in the P-Theta and MaCh3 fits to SK data. The charge of the final-state lepton cannot be determined and so the RHC samples have an irreducible ν background which cannot be separated out as in the ND280 samples. The single-ring samples are CCQE-enriched, and the one multi-ring sample has a large CC-Resonant component, which is the dominant background to the single ring samples.

ND280	SK	Target Interactions
CC0π	1Rμ or 1Re	CCQE
CC1π	1Re1.d.e	CC-Res
CC-Other	-	Multi-π, DIS

Table 5.10: Signal topologies of the selections at both ND280 and SK. The CC0π samples at ND280 provide the dominant constraint on the CCQE interaction model and on the single ring samples at SK used in the oscillation analysis. The CC1π selection at ND280 constrains the CC-Resonant pion production background of the CC0π and 1-Ring samples. An additional sample in the FHC at SK uses decay electrons from pions to constrain the total resonant pion cross section at SK where failure to reconstruct a pion, or pion absorption in the nucleus, may result in a single ring topology. The CC-Other selection at ND280 captures any remaining reconstructed CC interactions at ND280 and so can constrain backgrounds at higher than neutrino energies to the other two samples and constrain their signal topologies through correlations.

Chapter 6

The BANFF near detector fit

The BANFF (Beam And ND280 Flux extrapolation task Force) group is the T2K working group responsible for the near detector fit of the sequential analysis stream of the oscillation analysis. The group maintains and runs a fitter, known as the BANFF fitter, built around ROOT's Minuit2 package [119]. The MIGRAD gradient descent algorithm is used to maximise the high-dimensional ND280 likelihood to find a best-fit set of values for the T2K model parameters. Once the best-fit point in the parameter space is found, the Hessian matrix is calculated from each of the the second-order partial derivatives of the likelihood with respect to each of the parameter combinations. The Hessian matrix can then be inverted to determine the covariance matrix at the best-fit point and the appropriate postfit uncertainties for each of the parameters. As this is a maximum-likelihood estimation with a set of best-fit values for the model parameters, it is interpreted as a frequentist result, although prior constraints are still used. The BANFF fit constrains the parts of the T2K model which apply to ND280 and then the ND280-SK common parameters and their covariance matrix are propagated to the far detector fitter group, P-Theta, who extract the oscillation parameter constraints as discussed in chapter 5.

6.1 The BANFF test statistic

The ND280 likelihood is calculated from comparisons of the selected data to model predictions for each bin in the ND280 samples discussed in section 6.2. The relationship between the likelihood, $\mathscr{L}$, and the χ^2 of the data and the MC,

$$\chi^2 \equiv -2\ln\mathscr{L}, \tag{6.1}$$

is used to find the maximum likelihood by minimising the corresponding χ^2 surface. This log likelihood can be separated out as

$$\ln \mathscr{L}_{\text{ND280}} = \ln \mathscr{L}_{\text{stat}} + \ln \mathscr{L}_{\text{syst}} \tag{6.2}$$

where $\ln \mathscr{L}_{\text{stat}}$ and $\ln \mathscr{L}_{\text{syst}}$ are the statistical and systematic components of the likelihood, respectively. The systematic component can be further broken down into the different parts of the model

$$\ln \mathscr{L}_{\text{syst}} = \ln \mathscr{L}_{\text{flux}} + \ln \mathscr{L}_{\text{int}} + \ln \mathscr{L}_{\text{det}} \tag{6.3}$$

where $\ln \mathscr{L}_{\text{flux}}$, $\ln \mathscr{L}_{\text{int}}$ and $\ln \mathscr{L}_{\text{det}}$ are the beam-flux model, neutrino-nuclear interaction cross-section model, and detector model likelihood components, respectively. To reduce the computational cost and time of the fit, sections of the covariance matrix with highly-correlated parameters (flux, detector, FSI) can be decomposed to exploit these correlations.

In past T2K analyses, the $\mathscr{L}_{\text{stat}}$ term only took into account the data-MC difference as given by

$$-2\ln \mathscr{L}_{stat} = 2 \sum_{j}^{\text{samples}} \sum_{i}^{\text{bins}} \left[N_{\text{MC}} - N_{\text{Data}} + N_{\text{Data}} \ln \frac{N_{\text{Data}}}{N_{\text{MC}}} \right], \tag{6.4}$$

and any statistical variation was treated as part of this difference with an additional uncertainty term added to the diagonal of the detector covariance matrix to inflate the overall uncertainty on the number of events in each bin accordingly. For T2K's 2020 oscillation analysis (OA), the method outlined by Barlow and Beeston [136] was used, which separates out the MC statistics term $\mathscr{L}_{\text{MC stat}}$ from the standard Poisson likelihood $\mathscr{L}_{\text{Poisson}}$ for each bin such that

$$-2\ln \mathscr{L}_{stat} = -2\ln \mathscr{L}_{\text{Poisson}} - 2\ln \mathscr{L}_{\text{MC stat}} \tag{6.5}$$

and explicitly fits a scaling parameter β_i for the ith bin which accounts for statistical fluctuations in the generated MC on the true (infinite statistics) MC predicted by the model, such that $N_{\text{MC}}^{true} = \beta \times N_{\text{MC}}^{gen}$. This, however, introduces a very large number of parameters which would be unfeasible to fit in the same manner as the other parameters. Using an updated version of the Barlow-Beeston method proposed by Conway [137]. These parameters are allowed to behave in a Gaussian and independent manner such that when the χ^2 is minimised with respect to β, equation 6.6 can be solved for β analytically,

$$\beta_i^2 + (\mu\sigma_{\beta_i}^2 - 1)\beta_i - n\sigma_{\beta_i}^2 = 0 \tag{6.6}$$

where μ is the number of simulated events ($\mu = w_i N_{\text{MC}}^{gen}$) including all of the weights applied, w_i, to the events in that bin, and σ_{β_i} is the associated relative uncertainty on the number of events in the bin and can be expressed as a normalised sum of the square of the weights, $\sigma_{\beta_i} = \sqrt{w_i^2}/N_{\text{MC}}^{gen}$, where $\sigma_{\beta_i} = \sqrt{w_i^2}$ is the sum of the square of the weights in bin i.

By substituting in $N_{\mathrm{MC}}^{true} = \beta \times N_{\mathrm{MC}}^{gen}$ for N_{MC} in equation 6.4 we can arrive at a new expression for the likelihood

$$-2\ln\mathscr{L}_{stat} = 2\sum_{j}^{\text{samples}}\sum_{i}^{\text{bins}}\left[\left(N_{\mathrm{MC}} - N_{\mathrm{Data}} + N_{\mathrm{Data}}\ln\frac{N_{\mathrm{Data}}}{N_{\mathrm{MC}}}\right) + \frac{(\beta_i - 1)^2}{2\sigma_{\beta_i}^2}\right] \tag{6.7}$$

and so the MC statistical penalty in the likelihood can be applied quickly and at low computational cost since the only additional information that the fit needs to keep track of is the sum of the square of the weights.

6.2 ND280 Samples

The ND280 data in the BANFF fit is split into 18 samples with selections which best constrain the models that describe signal and background contributions to the ND280 data. Samples are split by beam mode, i.e. horn current direction (FHC or RHC), reconstructed pion multiplicity (0π, 1π, or Other), and the subdetector in which the neutrino interaction vertex is reconstructed (FGD1 or FGD2).

Samples are separated by reconstructed pion multiplicity to target specific interaction modes discussed in sections 2.3 and 5.8 as shown in figure 6.1: the CC-0π samples target CCQE-like interactions which are the primary signal events at SK; the CC-1π samples target CC-resonant pion production which is the second dominant interaction mode at peak energies of the T2K beam; and the CC-Other samples constrain multi-π and DIS interactions which dominate interactions from neutrinos in the high-energy tail of the beam flux. Each of these contributions not only has different interaction mode compositions but also contributes to events in different parts of the p_μ and E_ν spectra. The CC0π sample has the largest statistics and dominates the peak of the momentum spectrum, which largely corresponds to the peak of the E_ν spectrum from which the oscillation parameters are extracted at SK. Whilst largely made up of CCQE interactions, a significant background in the CC0π sample comes from CC-Resonant single-pion production (SPP). This may be due to a failure to reconstruct a pion, or from pion absorption within the nucleus resulting in a pionless final-state topology. A dedicated CC1π sample can be used to constrain this background in the CC0π sample and by extension the single-ring CCQE-rich samples at SK. The shape of this sample is generally broader in p_μ than the CC0π sample, with a larger relative contribution from higher p_μ events. The CC1π sample has significant contributions from the multi-π or DIS events which may similarly undergo FSI or mis-reconstruction as a single pion event which are backgrounds to the SPP processes. The final CC-Other reconstructed topology is intended to provide some constraint on higher pion multiplicities and DIS events which become dominant in the high E_ν tail of the T2K beam. There are some NC contributions to the CC1π and CC-Other samples;

however, there is not a large enough contribution to provide a significant constraint on the NC interaction models. A joint fit to all three topologies allows the constraints on these backgrounds to help to further constrain the CCQE interaction model, which is the signal definition for the SK single-ring samples.

Six FHC samples constrain ν interactions, and six RHC samples constrain $\bar{\nu}$ interactions, with an additional six RHC samples to constrain the significant wrong-sign ν contamination in the $\bar{\nu}$-beam. The nominal MC predictions and the data-MC ratio binned in p_μ for each of the FGD2 CC0π samples is shown in figure 6.2. Neutrino sign determination is generally not possible in SK, and whilst a prediction of the SK RHC sample could be made from a sign-agnostic treatment of the ν_μ and $\bar{\nu}_\mu$ events at ND280, this separation allows for a better constraint on the value of δ_{CP}. It is important to note that the reaction breakdown and E_ν (and so the lepton kinematics) differ between each of these samples, particularly between the two ν_μ samples due to the different kinematics of the mesonic parent of the neutrino.

For the 2020 analysis each of the samples are binned in muon momentum, p_μ, and cosine of the angle between the muon and the detector axis (parallel to the beam axis) as the muon leaves the interaction vertex, $\cos\theta_\mu$. The binning of each of these samples is adjusted to ensure that at least one data event and a minimum of 20 raw (unweighted) MC events are in each bin in the fit. This prevents the fit from attempting to evaluate the χ^2 for zero MC events which would not be defined, and reduces the impact of the MC statistical uncertainty on the fit.

The samples are also split by subdetector in which the neutrino interactions vertex lies. Events in FGD1 occur primarily on carbon and events in FGD2 occur on a mixture of carbon and oxygen. The differences between interactions on carbon and oxygen are constrained by using these FGD2 samples, which is crucial for the propagation to SK. Though the FGDs differ in construction because of the inclusion of water targets in FGD2, the similar design and detection principle leads to significant correlations between the detector systematic uncertainties relevant to each sample of events. The boundaries of each bin used in BANFF fit for the samples used in the 2020 analysis are listed below, in which the FGD1 and FGD2 samples have equivalent binnings.

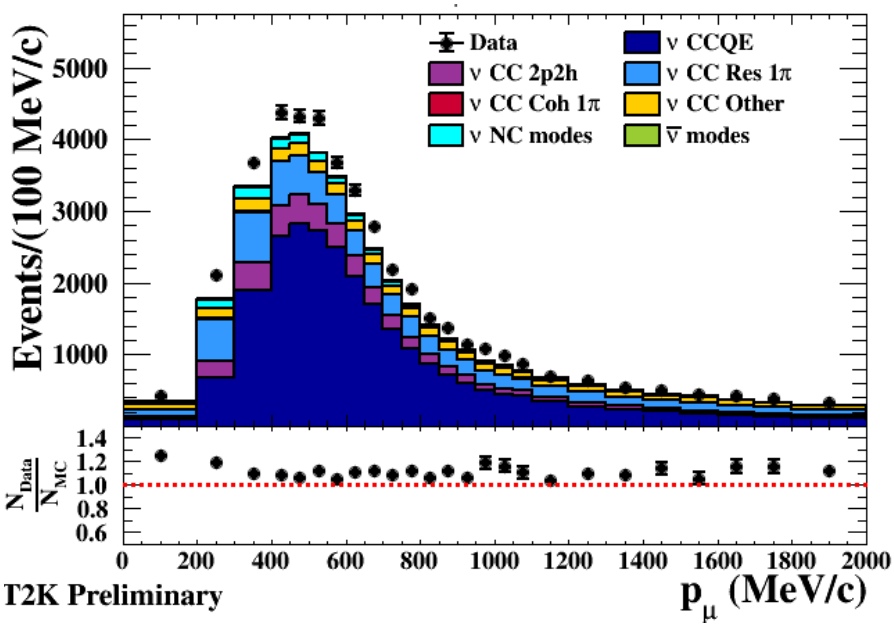

(a) CCQE-rich sample

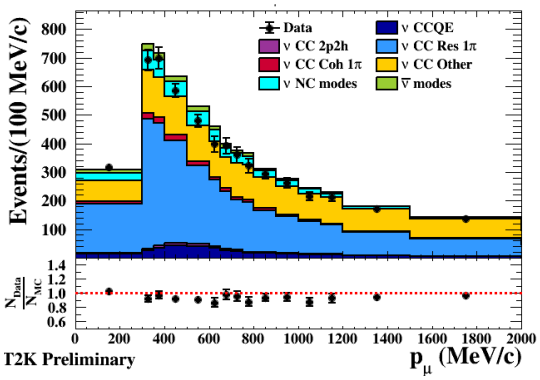

(b) CC-Res-rich sample

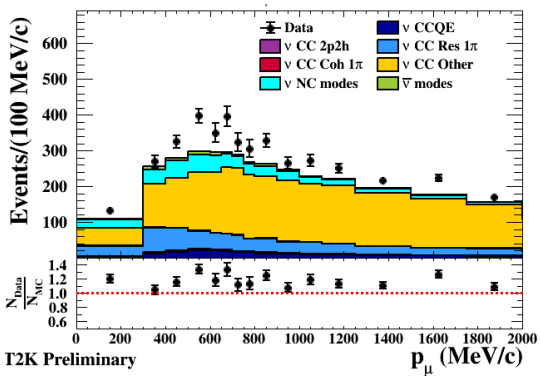

(c) Multi-π and DIS-rich sample

Figure 6.1: The stacked colour histograms in each of the upper panels of each subfigure show the nominal prefit reaction mode contributions to the MC prediction for the three FHC beam FGD1 samples ((a) CC0π, (b) CC1π, (c) CC-Other), as compared to data shown by the black points. The lower panel of each of the subfigures shows the ratio of the collected data at ND280 to the nominal MC prediction before the BANFF fit. The majority of the CC0π and CC1π sample MCs are primarily composed of their target interaction modes of CCQE and CC-Resonant pion production, with some background contributions due to mis-reconstruction or FSI which may change the final-state topology. The CC-Other sample acts as a catch-all side-band to constrain the backgrounds in the first two samples which become more significant at higher energies and lepton momenta.

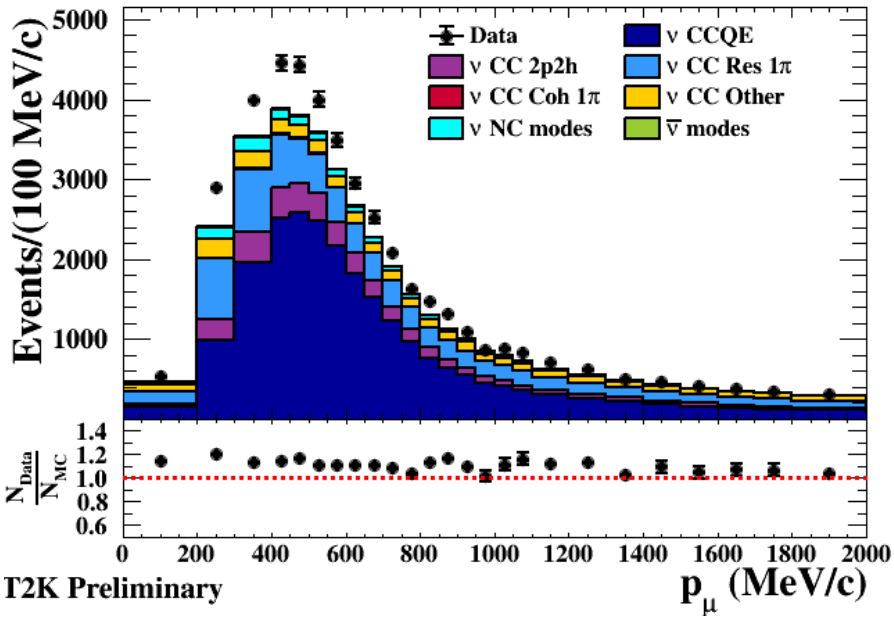

(a) ν_μ in ν-mode sample

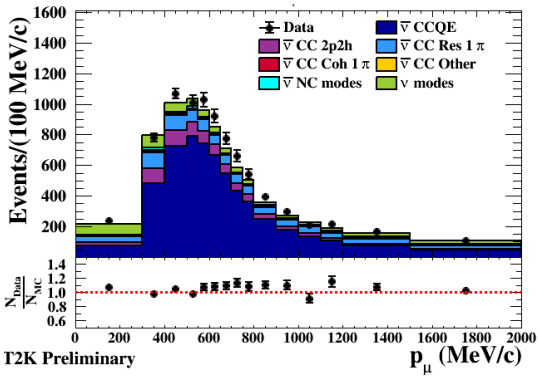

(b) $\bar{\nu}_\mu$ in $\bar{\nu}$-mode sample

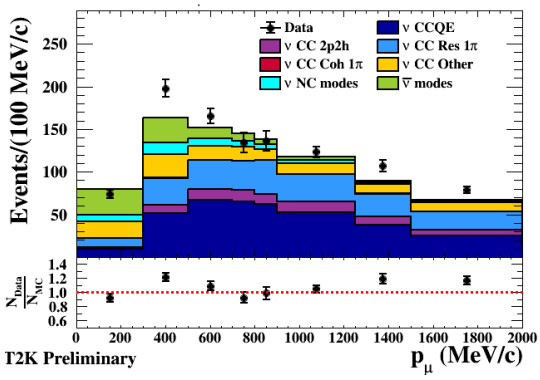

(c) ν_μ background in $\bar{\nu}$-mode sample

Figure 6.2: The stacked colour histograms in each of the upper panels of each subfigure show the nominal prefit reaction mode contributions to the MC prediction for the three FGD2 CC0π samples separated by neutrino sign and beam mode ((a) FHC ν_μ, (b) RHC $\bar{\nu}_\mu$, (c) RHC ν_μ background), as compared to data shown by the black points. The lower panel of each of the subfigures shows the ratio of the collected data at ND280 to the nominal MC prediction before the BANFF fit. Though the FHC beam mode is relatively right-sign pure (a), the RHC beam mode has a significant ν-in-$\bar{\nu}$-mode background for the collected statistics and so a third set of samples (c) is used to constrain the neutrino interactions in the antineutrino-mode beam. These background RHC-ν_μ events occupy a different part of the kinematic phase space to the FHC-ν_μ events, and the sample is much less-CCQE rich.

- **FHC ν_μ CC0π:**

  ***p* (MeV/c):** 0, 200, 300, 400, 450, 500, 550, 600, 650, 700, 750, 800, 850, 900, 950, 1000, 1050, 1100, 1200, 1300, 1400, 1500, 1600, 1700, 1800, 2000, 2500, 3000, 5000, 30000.

  cos θ: -1, 0.5, 0.6, 0.7, 0.76, 0.78, 0.8, 0.83, 0.85, 0.88, 0.89, 0.9, 0.91, 0.92, 0.925, 0.93, 0.935, 0.94, 0.945, 0.95, 0.955, 0.96, 0.965, 0.97, 0.975, 0.98, 0.985, 0.99, 0.995, 1.

- **FHC ν_μ CC1π:**

  ***p* (MeV/c):** 0, 300, 350, 400, 500, 600, 650, 700, 750, 800, 900, 1000, 1100, 1200, 1500, 2000, 3000, 5000, 30000.

  cos θ: -1, 0.6, 0.7, 0.8, 0.85, 0.88, 0.9, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99, 0.995, 1.

- **FHC ν_μ CCOther:**

  ***p* (MeV/c):** 0, 300, 400, 500, 600, 650, 700, 750, 800, 900, 1000, 1100, 1250, 1500, 1750, 2000, 3000, 5000, 30000.

  cos θ: -1, 0.6, 0.7, 0.76, 0.8, 0.85, 0.88, 0.89, 0.9, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99, 0.995, 1.

- **RHC $\bar{\nu}_\mu$ CC0π:**

  ***p* (MeV/c):** 0, 300, 400, 500, 550, 600, 650, 700, 750, 800, 900, 1000, 1100, 1200, 1500, 2000, 4000, 30000.

  cos θ: -1, 0.6, 0.7, 0.8, 0.85, 0.9, 0.92, 0.93, 0.94, 0.95, 0.96, 0.965, 0.97, 0.975, 0.98, 0.985, 0.99, 0.995, 1.

- **RHC $\bar{\nu}_\mu$ CC1π:**

  ***p* (MeV/c):** 0, 500, 700, 900, 1300, 2500, 30000.

  cos θ: -1, 0.7, 0.8, 0.9, 0.94, 0.96, 0.98, 0.99, 1.

- **RHC $\bar{\nu}_\mu$ CCOther: *p* (MeV/c):** 0, 600, 800, 1000, 1250, 1500, 2000, 4000, 30000.

  cos θ: -1, 0.7, 0.8, 0.85, 0.9, 0.93, 0.95, 0.97, 0.98, 0.99, 1.

- **RHC ν_μ CC0π:**

  ***p* (MeV/c):** 0, 300, 500, 700, 800, 900, 1250, 1500, 2000, 4000, 30000.

  cos θ: -1, 0.7, 0.8, 0.85, 0.88, 0.9, 0.92, 0.94, 0.96, 0.97, 0.98, 0.99, 1.

- **RHC ν_μ CC1π:**

  ***p* (MeV/c):** 0, 600, 800, 1500, 30000.

  cos θ: -1, 0.7, 0.8, 0.86, 0.9, 0.94, 0.96, 0.97, 0.98, 0.99, 1.

- **RHC ν_μ CCOther:**

  ***p* (MeV/c):** 0, 600, 1000, 1250, 2000, 4000, 30000.

  cos θ: -1, 0.7, 0.8, 0.86, 0.9, 0.93, 0.95, 0.97, 0.99, 1.

In the MCMC simultaneous fit to ND280 and SK data, MaCh3, different boundaries are used, with an irregular binning allowing for different p_μ binnings for different bins of $\cos\theta_\mu$, using the ROOT TH2Poly class [138]. This differs from the use of ROOT TH2Ds in BANFF which have the same p_μ binning for every bin in $\cos\theta_\mu$. This leads to some differences in sensitivity as the TH2D implementation may be much coarser in high statistics regions, due to the constraint on the number of events in other bins with low statistics in the sample which share the same p_μ bin boundaries. Conversely, sparser areas may have more bins than desirable, because they align in either p_μ or $\cos\theta_\mu$ with the peak of the distribution.

6.3 Parameters implemented in the BANFF fit

As discussed in chapter 5, the ND280 likelihood includes parameters which allow for variations of the event rates according to the uncertainties of the model. At ND280 these parameters fall into three categories, the flux, the cross-section and the detector parameters. The flux systematic parameters in the T2K likelihood are binned normalisations in neutrino energy applied to the events in bins of true neutrino energy and neutrino flavour. Half of these parameters apply to the near detector and half apply to the far detector with the corresponding bins between the near and far detector parameters being highly correlated. As a result, the 50 SK flux parameters, though not applied to any events at ND280, are still constrained through these correlations. Figure 6.3 shows the binning of the normalisation parameters for the FHC-beam ν_μ flux in E_ν. This figure also shows the prefit (red) and expected postfit (black) uncertainties on these parameters. The value of the flux normalisation parameters are ratios to the MC-generated value after tunes from the flux

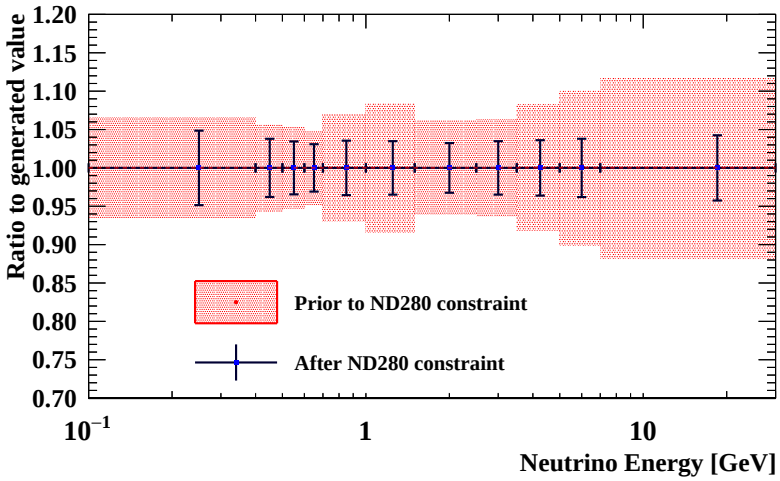

Figure 6.3: Prefit ND280 flux-parameter uncertainties in bins of E_ν and the expected constraint from the BANFF fit based on the collected POT. Parameter values are displayed relative to the generated value and so when fitting the MC to itself we expect the fit to return the input values. The reduction of the uncertainty is the expected sensitivity of the fit given the collected statistics.

group have been applied based on beam measurements and fits to external NA61/SHINE T2K replica-target data [123][124]. The flux binning is optimised to give greatest granularity in the region of the peak of the unoscillated spectrum, which is also where ν_μ disappearance is expected to be maximal.

Due to the very non-Gaussian behaviour of the detector parameters, directly fitting for these would be both computationally expensive and also risk convergence failures due to the assumption of quadratic behaviour of the likelihood around the minimum in the Hessian covariance calculation, so these are fit indirectly through an "observable normalisation" parameter for each of the detector bins in the fit. A total of 574 of these detector parameters are then included in the fit for the 2020 analysis.

The neutrino interaction cross-section (xsec) model parameters, more commonly referred to as the "cross-section parameters" often require more sophisticated treatment than the flux or detector systematics. Whilst some of the parameters are simple normalisations where a freedom is added to the fit to cover some uncertainty, such as the Q^2 CCQE which has previously been suppressed [135], others require more complicated response splines based on interpolation between externally-calculated weights in the T2K reweight packages detailed in section A.7. Whilst the majority of the xsec parameters are not decomposed during the fit, the final-state interaction (FSI) parameters are decomposed with the flux and detector systematics.

6.4 Event reweighting during the fit

As generating a complete set of MC events from scratch and propagating it through the entire ND280 analysis chain for every set of model parameter values explored by the fit would be impossible, the reweighting technique discussed in section A.7 is used to vary the MC prediction. First all of the "raw" MC is weighted to the collected POT equivalent as much more MC is generated than collected to reduce the MC statistical uncertainty and to cover any phase space in which data events are expected. Then the weight from the nominal flux tune to the NA61/Shine experiment data are applied [101]. For the first time in the 2020 analysis, the flux tune came from fits to the T2K-replica target data [124]. Then the cross section parameters are set to the prior-central values, which may differ from the generated or NEUT-nominal value, and a corresponding weight is applied. E.g. M_A^{QE} is weighted to bubble chamber experiment data [139] from ANL [140][141], BNL [142], BEBC [143], and FNAL [144]. Finally, the detector normalisation parameters, are set to the central value of the throws of the underlying `psyche` detector parameters used to generate the detector covariance, and a corresponding weight is applied. There is an assumption of multiplicative independence which holds close to the nominal prefit values of the parameters, though for this reason the detector covariance is generated with weights to the flux tune and cross-section prior values applied to the events. The flux and detector parameters are flat normalisations within some bin of E_ν or p_μ-$\cos\theta_\mu$, and so are some of the cross-section parameters; however, other cross-section parameters have more complicated non-linear responses and so the weights are precalculated at specific values and then a third-order polynomial spline is used to interpolate between the values. This allows the fitter to smoothly vary the weights as it explores parameter values between the precalculated values. The precalculated values are typically the integer standard deviations of the parameter between -3σ to 3σ, with special consideration for asymmetric parameters or parameters with physical bounds. This is interpolated using the ROOT TSpline3 class [138].

Chapter 7

Validation of the fitting framework

Before performing the BANFF and MaCh3 fits to the ND280 data, the fitting framework must be validated to ensure that the parameter responses behave in a well-defined manner and will return a known true set of parameters used to generate MC to an acceptable accuracy and precision. To do this, one can either fit the prior model to the MC produced by the prior model itself in an Asimov fit, or fit the prior model to alternative simulated data in a fake data study. This chapter will discuss the results of Asimov fits and fake data fits, and chapter 8 will discuss the results of fitting the MC to the data.

7.0.1 Event rate comparisons

The most basic validation of the fitter may be a simple calculation of the event rates for each bin in the fit to understand how close the nominally-weighted MC event rates are to the data. If a parameter over- or under-weights a particular region of the observable space, this may be apparent and can be used to diagnose initial problems with weight calculations or unphysical behaviour. In this work, event rates are calculated for the nominal model, and for specific variations of the individual parameters. As well as comparing to the data, these event rates can be compared across the different T2K fitting frameworks to ensure that the inputs to the likelihood are the same. A comparison of the BANFF and MaCh3 event rates for the ΔE_{rmv} parameters is given in appendix E.

7.1 Likelihood scans

To validate the parameter response, 1D-likelihood scans using the fitting machinery can show how a parameter is expected to behave during the fit. This is useful to highlight any issues in the implementation or behaviour which may be pathological to the fit. The values of all but the parameter of interest are fixed, and the $\Delta\chi^2$ or $-2\log\mathscr{L}$ value for variations of the parameter

are calculated. This shows how the event responses are being applied and what the expected sensitivity for a given parameter is. A flat response indicates that there is not enough sensitivity to constrain the parameter at the near detector, and so parameters with flat responses are not included in the BANFF fit, but may be allowed to freely vary at the near detector in the MaCh3 joint fit. A response which is ill-defined at some physical boundary may require special treatment to ensure smooth and continuous behaviour during the fit, such as mirroring of the spline. To mirror the spline additional knots are added above the boundary with values equal to those below the boundary in order to mirror the response at the boundary, allowing the calculation of the Hessian matrix to find appropriate values for the likelihood outside the bounds to use in the covariance calculation, even if these values are unphysical. Even if a parameter is well-behaved during a likelihood scan around the nominal parameter values, this may not highlight issues that arise when more than one parameter is allowed to vary, such as degenerate combinations of parameters leading to local minima that may not be observed in likelihood scans of the MC or the data. Variations of multiple parameters are studied in the Asimov fit and in the pull studies where they are they are moved in concert in a correlated manner. Whilst investigating the shape of the likelihood around a large ensemble of parameter combinations is not feasible, the fit failure rate as a function of parameter values in the pull study (section 7.4) and p-value (section 8.4) procedures can be used to identify problematic behaviour away from the nominal model.

Figure 7.1 shows how the flux normalisation parameters respond around their prior central values in a smooth and symmetric manner. This is expected as these parameters are normalisations with a Gaussian penalty around their prior value according to the uncertainty in the prefit flux covariance from fits to beam measurements, external NA61/SHINE data [124] and INGRID data.

More asymmetric behaviour can be seen in the interaction model shown in Figure 7.2. This is most evident in the FSI pion absorption parameter. Some asymmetry is to be expected and can be accommodated. It is important, however, that the LLH scan behaves quadratically around the minimum as the covariance calculation is reliant on the assumption that this is approximately true. Where a LLH scan is flatter or has multiple local minima, comparisons of the sample distribution predictions at these parameter values can be useful to assess the effect of this on the prediction, and to assess the impact of falling into a false local minimum on the analysis.

Figure 7.3 shows a comparison of the BANFF and MaCh3 LLH scans for the removal energy parameters. Using binned-template splines for the removal energy instead of a direct kinematic shift prevents discrete jumps which would result in a discontinuous likelihood, but the likelihood is otherwise compatible to within two units of $\Delta\chi^2$ or less. In the MaCh3 scans we can see these discrete jumps which correspond to event migrations between bins in the fit. Values of the scans

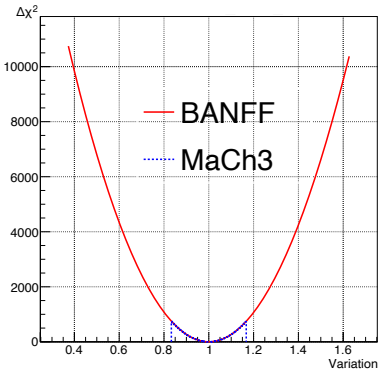

(a) ND280 FHC ν_μ 1-1.5 GeV

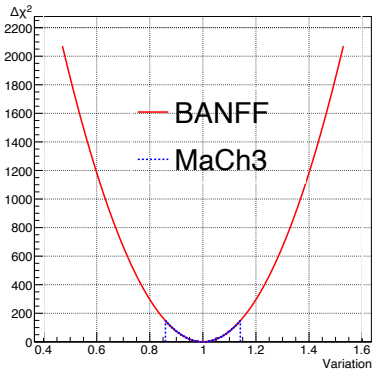

(b) ND280 FHC ν_e 0.5-0.7 GeV

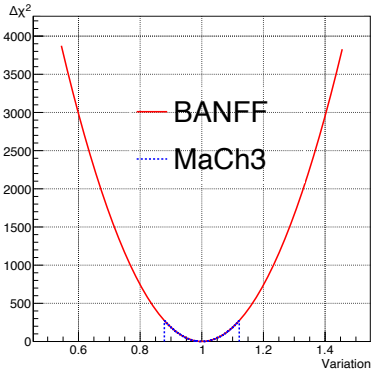

(c) ND280 RHC ν_e 2.5-30 GeV

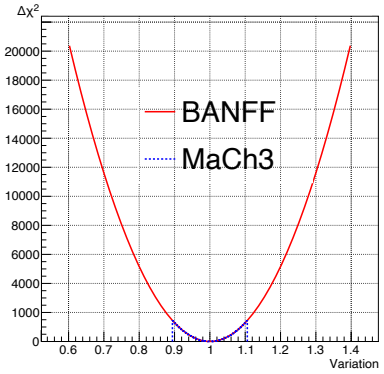

(d) SK RHC $\bar{\nu}_\mu$ 0.5-0.6 GeV

Figure 7.1: Log-likelihood scans (-2LLH$= \Delta\chi^2$) for selected flux parameters of the Asimov data set. The excellent agreement between the BANFF and MaCh3 fitters indicates consistent flux-parameter implementation between the two frameworks. The flux parameters are normalisations with a Gaussian penalty term and so return symmetric likelihood scans around their prior central value. A parameter value of 1 here indicates the NA61 replica flux tuned value of the flux for each E_ν bin. Variations are given relative to the NA61/tuned nominal value [124].

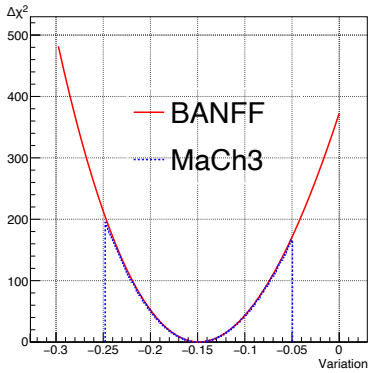

(a) M_A^{QE}

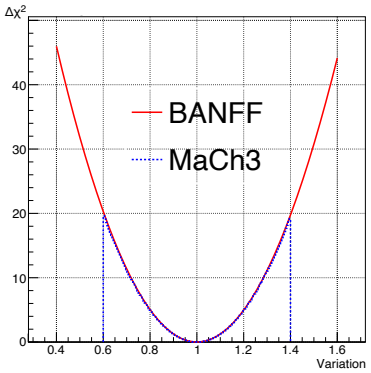

(b) 2p2h ^{12}C to ^{16}O normalisation

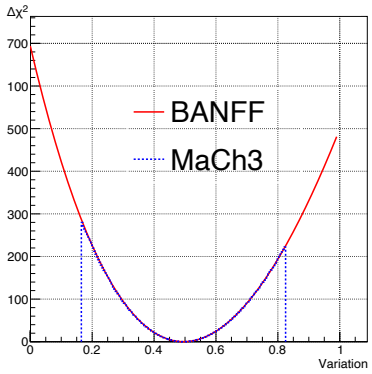

(c) $0.0 < Q^2 < 0.05$ normalisation

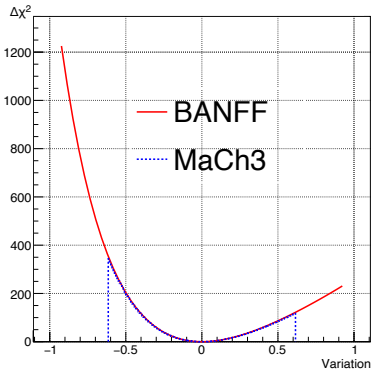

(d) FSI pion absorption

Figure 7.2: Log-likelihood scans (-2LLH$= \Delta\chi^2$) of selected interaction cross-section parameters for the Asimov data set for variations in units relative to the NEUT generated values. The offset of the minimum is due to the Asimov data being created at a prior central value other than the NEUT generated value.

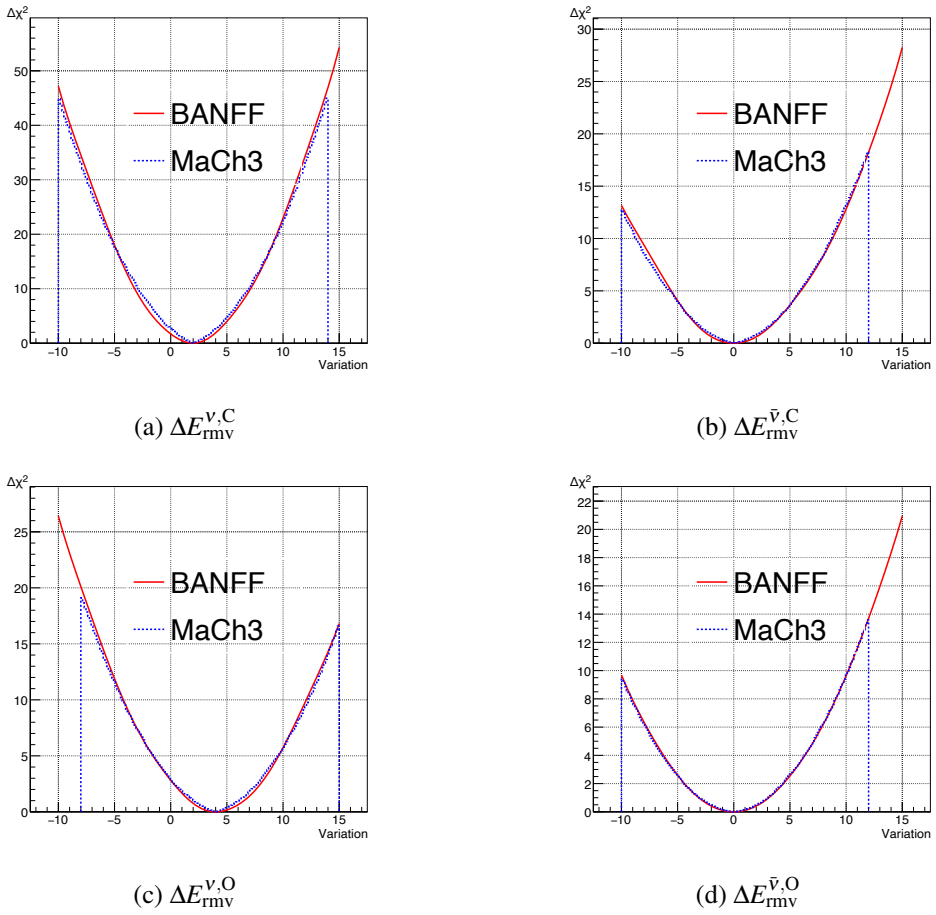

(a) $\Delta E_{\text{rmv}}^{\nu,\text{C}}$ (b) $\Delta E_{\text{rmv}}^{\bar{\nu},\text{C}}$

(c) $\Delta E_{\text{rmv}}^{\nu,\text{O}}$ (d) $\Delta E_{\text{rmv}}^{\bar{\nu},\text{O}}$

Figure 7.3: Log-likelihood scans ($-2\text{LLH}= \Delta\chi^2$) of the removal energy parameters calculated from the Asimov data set. Variations are given in units of MeV. For the Asimov data set there is good agreement between the kinematically-shifting parameter implemented in MaCh3, and the binned-spline implementation in BANFF. Some differences result from the discrete jumps produced by the momentum variations as events cross bin boundaries in MaCh3, as opposed to in BANFF where weights are interpolated from splines and so are smooth by construction.

are equivalent at the 5 MeV intervals for which the BANFF weight templates were produced but differ in between due to interpolation, binning and marginalisation effects. The compatibility of the fitters can also be seen in a direct comparison of the event rates used to calculate the likelihood. The BANFF and MaCh3 event rates for variations of the ΔE_{rmv} parameters can be found in appendix D.

7.2 Asimov fits

Asimov fits assess bias of the fitter around a known set of values of the parameters, as well as the sensitivity to each of the parameters included in the fit at those known values. The sensitivity of a

fit is the ability to exclude regions of parameter space given the amount of data collected. This will vary for different true values of the model parameters but an idea of the magnitude can be found by computing an "Asimov fit". At ND280, the Asimov data is defined as the nominal MC predicted by the prior model which is the same as the generated MC, but with all prior weights applied. By fitting the nominal MC to itself, the behaviour of the entire likelihood can be inspected before fitting to the data. This can be used to identify problems with convergence due to degenerate parameter responses producing multiple local minima. Pathological behaviour in the likelihood will cause the Asimov fit to fail to converge and will likely have the same effect on a fit to the data. This behaviour may only present itself due to the correlations between parameters which would not be indicated in likelihood scans. In the case of the Asimov fit, the fit should return all of the parameter values at their priors, and their postfit uncertainties should be smaller than their prefit uncertainties. The size of the reduction in these uncertainties is the sensitivity of the ND280 fit to this parameter for the statistical power of the sample, i.e. the POT collected.

Asimov fits are also conducted for the fits to SK data where it is important to assess the sensitivity of the fit for different true values of the PMNS parameters. The standard Asimov fit to the SK data uses the 2014 T2K fit result [118] for the values of those parameters to which T2K has some sensitivity, and the 2019 Particle Data Group global fit values for the other parameters [34]. Alternative Asimov data sets are also tested to check the T2K sensitivity at other values of physical interest such as CP-conserving and maximally CP-violating cases, or for comparisons with other experiments such as NOvA [77].

7.2.1 Results of the BANFF Asimov fit

Presented here are the results of the BANFF fit to the ND280 Asimov data. The parameters for the CC0π true topology interactions are shown in figure 7.4 where the expected sensitivity to the CCQE (1p1h) and 2p2h models can be seen. Each of the parameters returns its prior central value indicating the likelihood and parameter responses behave as expected around the prior values. New uncertainties are placed on the two 2p2h normalisation parameters and the five lower Q^2 normalisation parameters ($Q^2 < 0.25$ GeV^2). The uncertainties are reduced for the rest of the CC0π parameters which have prior constraints in the fit. The postfit uncertainty value should always be reduced relative to a prior constraint in the fit. An inflated uncertainty may be indicative of a poorly defined likelihood or fitter instability, as the prior constraint is added as a penalty to the χ^2 from the predicted sample comparison. The Asimov fit shows the expected sensitivity to the new ΔE_{rmv} parameterisation which behaves well during the fit, returning the prior values of each of the parameters, as expected. The carbon ΔE_{rmv} parameters are more strongly constrained than

their oxygen counterparts, and the ν parameters are more constrained than the $\bar{\nu}$ parameters, as expected. The $\Delta E_{\mathrm{rmv}}^{\nu,\mathrm{C}}$ parameter uncertainty is halved and the $\Delta E_{\mathrm{rmv}}^{\bar{\nu},\mathrm{C}}$ parameter uncertainty reduced by a third.

Figure 7.5 shows that the CC-resonant pion production parameters should be well constrained by the CC1π sample, and there is some sensitivity to the Isospin-1/2 non-resonant background parameter, although the second parameter which controls the rate of specifically $\bar{\nu}$ low p_π events was fixed for fitter stability reasons, and so there is no constraint.

Figure 7.6 shows a good expected sensitivity to the Bodek-Yang [130][131] corrections for both the high-multiplicity resonant pion production and DIS events, and some sensitivity to the AGKY-like model [133]. Though it is tightly constrained in the prefit, there is also a gain in sensitivity to the overall normalisation of ν Multi-π and DIS events. This gain is not as strong for the $\bar{\nu}$ parameter.

Finally, the sensitivity to pion FSI is shown in figure 7.7. The fit returns an uncertainty less than half of the prior constraint, with the hadron production parameter uncertainty reduced to 30% of its prior uncertainty. A strong constraint on the FSI should reduce the uncertainties on the prediction of the number of events which will be classified as single and multi-ring at SK, where effects such as the absorption of a pion from the hard scatter can change the reconstructed topology, moving events between samples.

7.3 Fits to fake data sets

In addition to fits to the Asimov MC, fits to MC generated with alternative models can be used to assess if the uncertainties are broad enough to capture any expected effects which may be missing from the T2K model. Often a more sophisticated model is available but a full systematic implementation may require significant work, and so if the difference between a simpler model and the more sophisticated model is smaller than the combined statistical and systematic uncertainty on the result there is no need to fully implement the more sophisticated model. As more data is collected and differences become more important, these studies are a good indicator of what freedoms T2K may need to include in the fits take advantage of the additional data. Fake data sets can also be used to explore the likelihood of a bias introduced by a single systematic parameter, such as the ΔE_{rmv} parameters discussed below. If either the true value of the parameter in the fake data is not converged on, or if it impacts the ability of the fitter to correctly evaluate the other parameters at their priors, this can indicate likely pitfalls in the data fit caused by degenerate behaviour or local minima which may not be evident at the prior central value. Whilst these studies are part of the near detector validation, they are also propagated to the far detector, where their

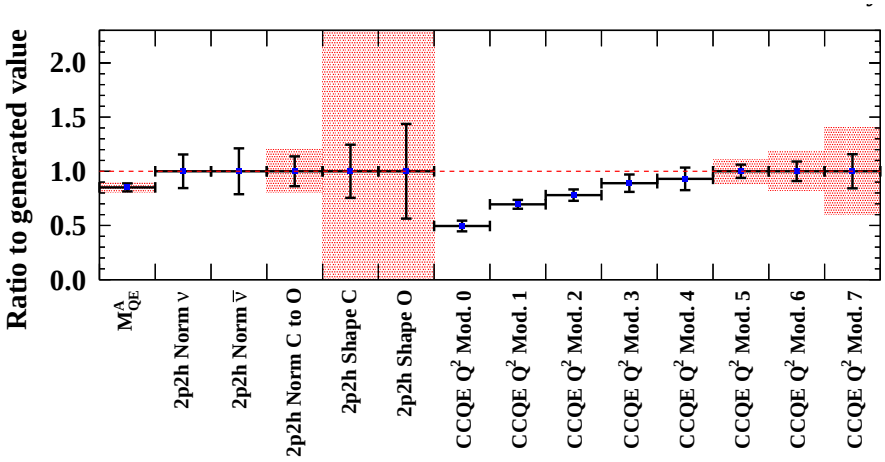

(a) CC0π cross-section parameters for both 1p1h and 2p2h models, excluding the ΔE_{rmv} parameters. M_A^{QE} is already tightly constrained by the bubble chamber data prior but is slightly further constrained by the fit. The 2p2- normalisation parameters are well constrained even without a prefit constraint. The fit shows a small sensitivity to constrain the 2p2h carbon-to-oxygen parameter but significant power to constrain the 2p2h-shape parameters which weight between the Valencia and Martini models as expected. The five low 1p1h Q^2 normalisation parameters are well constrained even without a prior constraint, and the high Q^2 normalisation uncertainties are approximately reduced by half.

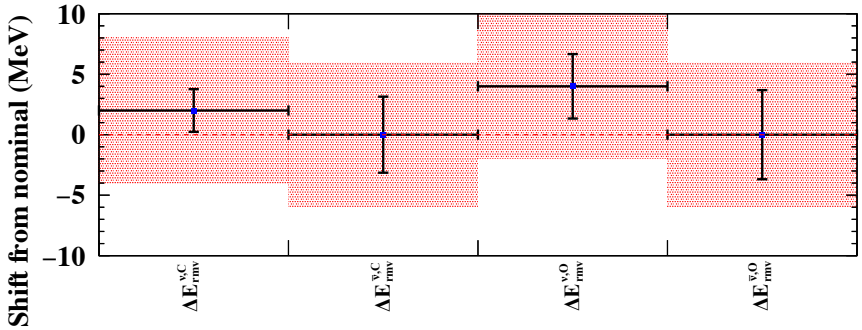

(b) ΔE_{rmv} parameters for the CCQE interaction. The Asimov fit suggests that the uncertainty on the neutrino oxygen nucleon removal energy uncertainty is expected to be reduced by half and the antineutrino oxygen removal energy uncertainty is expected to be reduced by a third. The Asimov fit returns the central value of the removal energy parameters suggesting the implementation is functioning as intended around the prefit model central values.

Figure 7.4: Prefit (red) and postfit (black cross) values and uncertainties for the true CC0π topology interaction models. Parameters with a flat or unbound prior value which do not have a Gaussian penalty term are presented without the red uncertainty band. The 2p2h normalisation and five lowest Q^2 normalisation parameters are unbound in the fit. Significant constraints are seen on the 2p2h shape parameters and the three high Q^2 normalisation parameters, as well as the four removal energy parameters, ΔE_{rmv}.

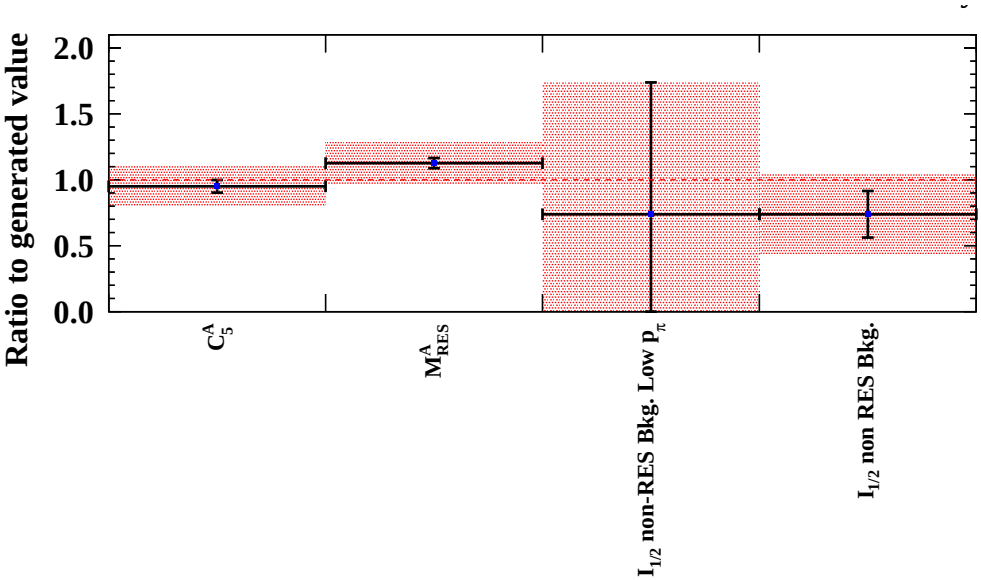

Figure 7.5: Prefit (red) and postfit (black cross) values and uncertainties for the true CC1π parameters. The resonant form factor parameters C_5^A and M_{RES}^A are well constrained indicating good sensitivity. Though the standard Isospin-1/2 non-resonant background parameter uncertainty is reduced by just under a half, the $\bar{\nu}$ low p_π parameter is fixed due to lack of sensitivity at the near detector.

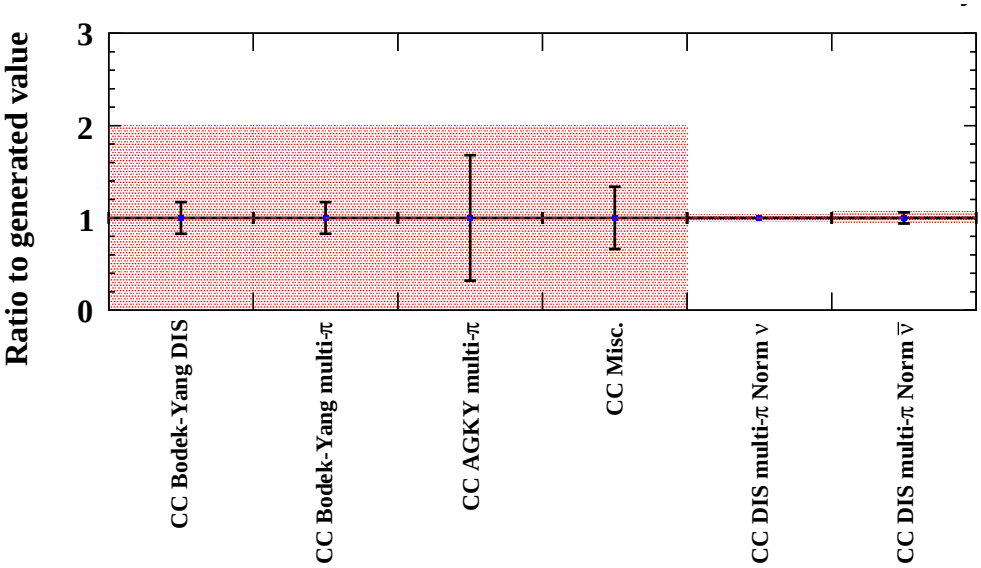

Figure 7.6: Prefit (red) and postfit (black cross) values and uncertainties for the true Multi-π and DIS parameters. The two Bodek-Yang correction dials are strongly constrained, and there is also sensitivty to the AGKY-like model. CC Misc., which controls a range of background processes such as eta decays, is constrained by more than half. Some sensitivity is gained to the overall normalisation parameters, though less for $\bar{\nu}$ events than ν events.

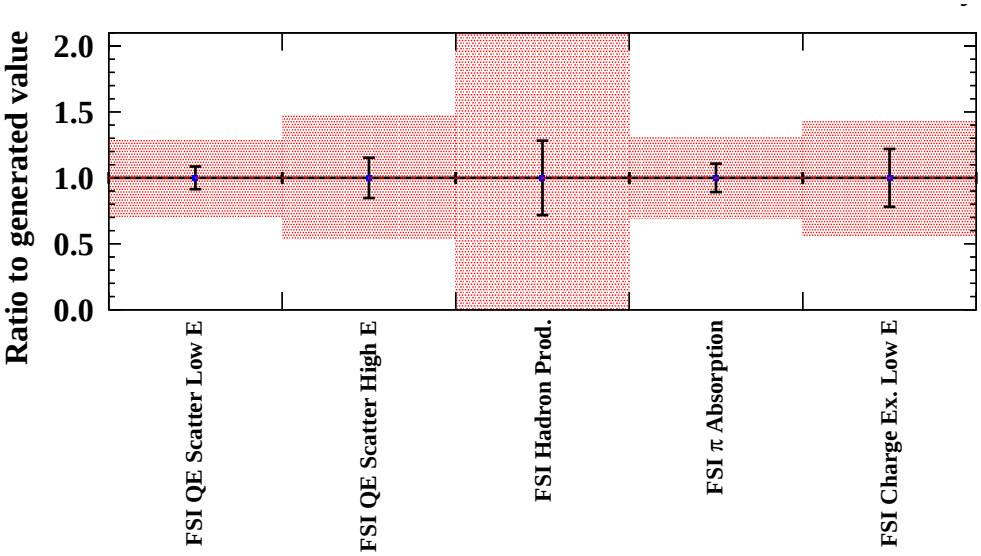

Figure 7.7: Prefit (red) and postfit (black cross) values and uncertainties for the pion FSI parameters. Each of the parameters are well constrained to within 50% of their prior uncertainty or less.

impact is used to provide additional uncertainty for these missing effects. The propagation of these studies, and the impact on the oscillation parameters is discussed in section 9.4

7.3.1 ΔE_{rmv} bias fake data set

A fake data study was conducted in which the true central values of the four ΔE_{rmv} parameters were all increased by 5 MeV above their prior central value to produce a fake data set. The ΔE_{rmv} parameters have their largest impact in the CC0π samples as these have the highest CCQE contribution, and so these are the samples for which the fake data will differ most to the standard ND280 Asimov or BANFF prefit model. Figure 7.8(a) shows the impact of this shift on the MC prediction compared to the prefit prediction. The higher typical E_{rmv} means more of the energy of the neutrino is used to liberate the nucleon and so the charged lepton has a lower momentum. This leads to a slight enhancement of events with $p_\mu < 500$ MeV and a suppression above this. The suppression above this momentum is more diffuse as there are generally fewer events in this region. The fitter is able to cover this 5 MeV bias well, as can be seen in the postfit ratio of the CC0π sample, and in the postfit parameter values shown in figure 7.9.

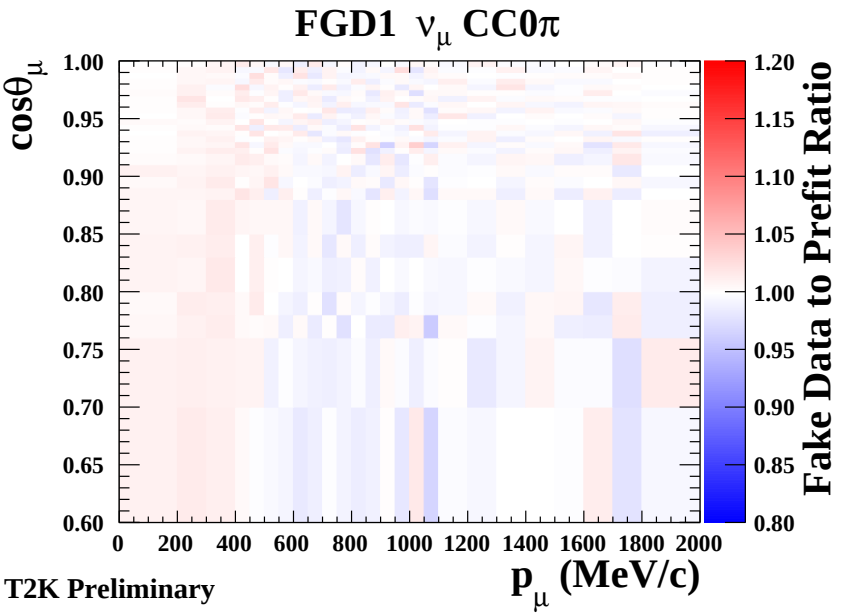

(a) Prefit to fake data ratio

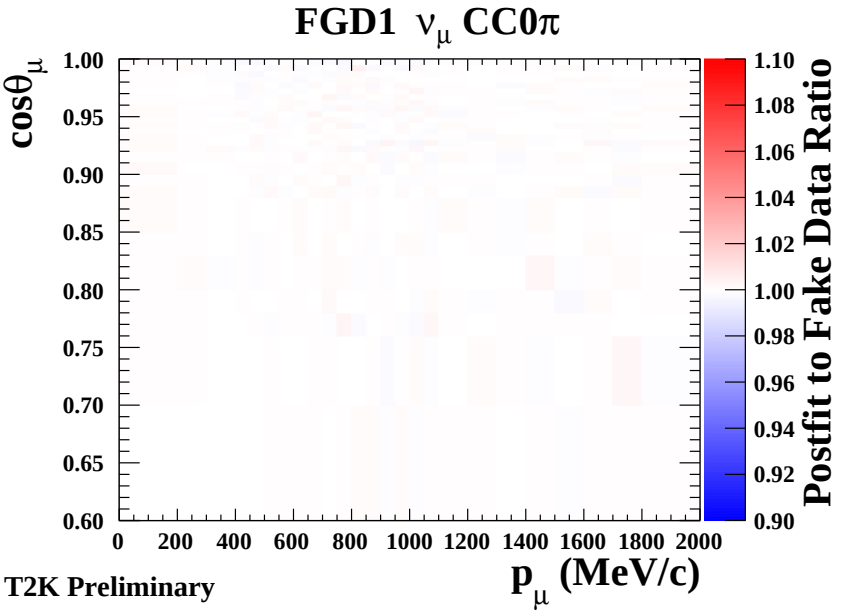

(b) Postfit to fake data ratio

Figure 7.8: Ratios of the prefit and postfit predictions to the +5 MeV-biased ΔE_{rmv} fake data set in p_μ and $\cos\theta_\mu$. The 5 MeV increase in ΔE_{rmv} enhances the low p_μ region slightly. In the higher p_μ and higher angle region, which typically has fewer events, the effect of event migration is more visible as discrete movement of events between adjacent bins where adjacent bins are enhanced and suppressed successively. The postfit ratio shows that the fitter is able to easily cover this +5 MeV bias evenly across the fitted observable space, as the ratio is reduced to subpercent differences, and appears flat across the distribution, indicating the parameterisation is able to respond to the shape changes produced by a true shift in the E_{rmv} spectrum.

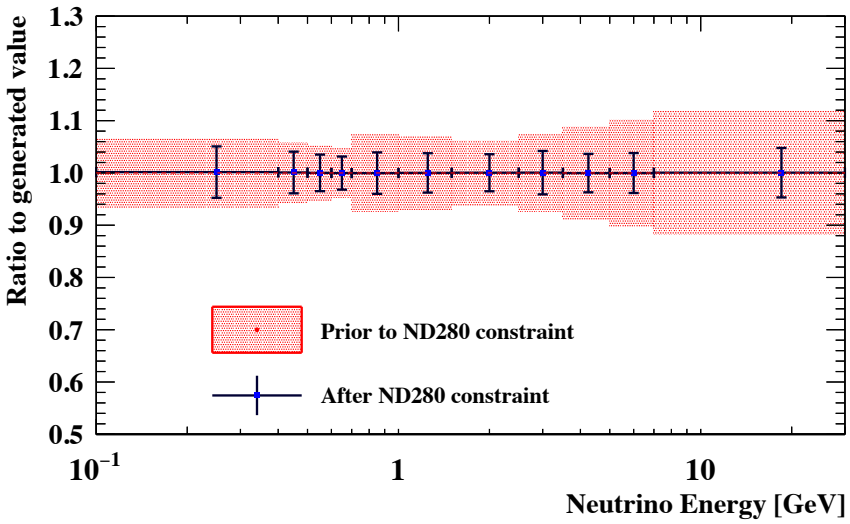

(a) SK FHC ν_μ flux

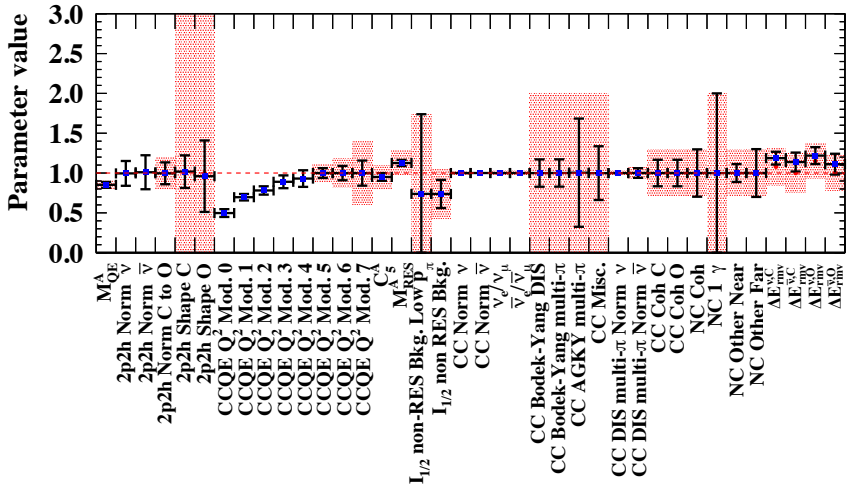

(b) Cross-section

Figure 7.9: Best-fit parameters from the BANFF fit to a fake data set where ΔE_{rmv} parameters were shifted upwards by 5 MeV. All other parameters are at their Asimov values. The SK flux parameters (a) return to their prior central values with a typical constraint when compared to the Asimov. The ΔE_{rmv} parameters, shown at the right hand side of figure (b), are given as ΔE_{rmv}/(25 MeV)+1 for ease of comparison. A slight deviation in the ΔE_{rmv} and 2p2h shape parameters is seen, but they are still very close to their prior central values when compared to their prefit and postfit uncertainties, all other cross-section parameters converge on their prior central value.

7.4 Pull studies and fit biases

Just as biases in the fitter around the nominal model can be assessed in an Asimov fit, a more generalised validation of fitter biases across the parameter space can be done by calculating the average "pull" of a parameter over the range of values expected according to the model. The pull of the parameter at a given true value is a measure of the accuracy of the postfit value of the fit given some statistical fluctuation of the events. A "throw" of the parameter values is a randomly generated set of values according to the probabilities dictated by the prior central values, uncertainties, and covariances. Then for each throw a "toy" data set is constructed from the prediction with a statistical fluctiation. For an ensemble of thrown true values of a parameter, and Poisson fluctuations of the MC statistics, the distribution of the pulls for that parameter should be Gaussianly distributed around zero. Assuming that the fitter is correctly assessing the postfit uncertainty in the fits to the toy data sets, the root-mean square width of the distribution should be that of the postfit uncertainty; i.e. the true value of the parameter should lie within 1σ of the postfit value that the fit converges on around 68% of the time. With this in mind, the pull of an individual throw of a parameter,p, is defined to be

$$\text{pull} = \frac{p_{\text{fit}} - p_{\text{true}}}{\sigma_{\text{fit}}} \tag{7.1}$$

where p_{true} and p_{fit} are the true thrown value and the postfit value of p, respectively, and the bias of the fit is

$$\text{bias} = \sum_{i}^{N} \text{pull}_i \tag{7.2}$$

where i is a single throw of all of the parameters according to the prior covariance matrix and N is the size of a sufficiently large ensemble of throws to thoroughly cover the parameter space expected by the prior model. Strictly speaking, equation 7.1 is the pull of an unconstrained parameter, and when a prior constraint is applied in the fit it is taken into account by including the prior uncertainty, using

$$\text{pull}_c = \frac{p_{\text{fit}} - p_{\text{prior}}}{\sqrt{\sigma_c^2 - \sigma_{\text{fit}}^2}} \tag{7.3}$$

where p_{prior} is the same as p_{true} since the thrown model is fit to itself, σ_c is the prior constraint and σ_{fit} is the postfit error on the parameter. However, this definition runs into problems in high-dimensional and highly-correlated parameter spaces as the denominator does not correctly assess the quadrature sum of the correlated parameter uncertainties. This distorts the width of many of these distributions, and so to have a consistent definition, the pulls are calculated with equation 7.1 rather than 7.3. This definition risks overestimating the biases in the parameters, but this may preferable when the biases are used to estimate the likely impact on the credible intervals of POIs in

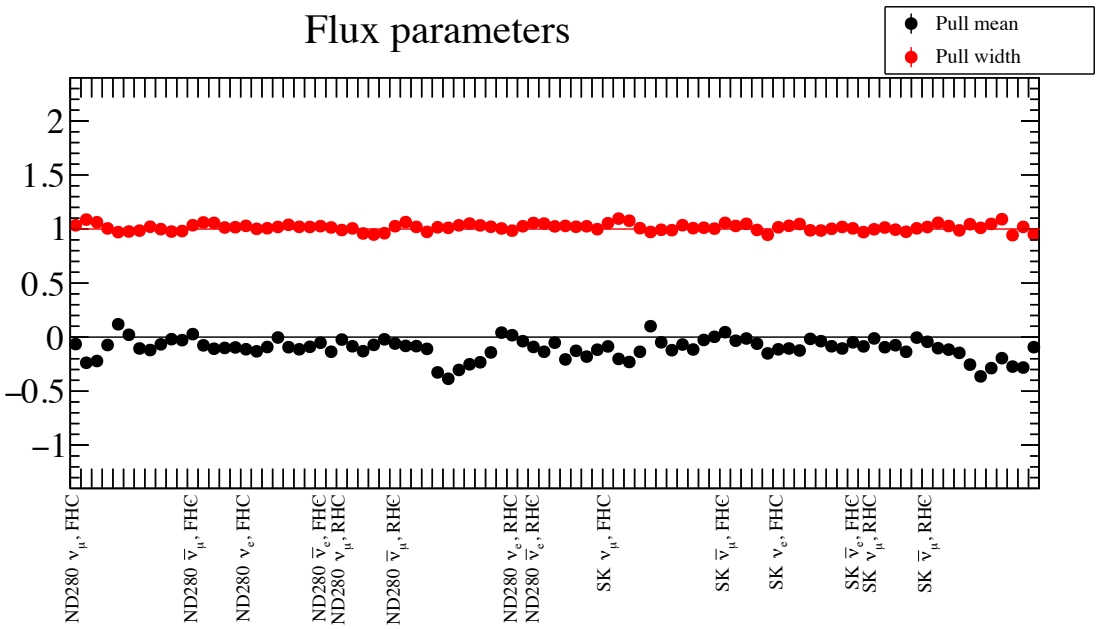

Figure 7.10: Flux parameter pull means and widths. The primary abscissa labels indicate the group of flux parameters for each detector, beam mode and neutrino flavour, each consecutive parameter, indicated by the secondary abscissa labels, are in corresponds to the flux normalisation parameters in increasing bins of E_ν for each group. Many of the flux parameters have an approximately 10% average bias towards an underestimate, with the exception the higher E_ν bins above the peak in the RHC $\bar{\nu}_\mu$ which have around a 20% bias.

the analysis, as a larger bias leads to a more inflated uncertainty and a more conservative statement about parameter space exclusion.

7.4.1 Results of the BANFF pull studies

For the BANFF pull studies, the flux, cross-section and effective bin-by-bin ND280 detector parameters are thrown in BANFF and their values stored along with the corresponding toy data set. Using the thrown values as the new prior central value of these parameters, the MC data set corresponding to the thrown parameter values is fit to the toy data, i.e. the statistically varied version of itself. This should return a value of the each parameter that matches that of the prior central value, within the postfit uncertainties returned. A bias of 50% in a SK-propagated parameter is considered significant enough to warrant further investigation; otherwise, the total bias impact study below in section 7.4.2 is done to assess any effect on the oscillation parameters.

The flux parameters in Figure 7.10 show little bias, except for the ND280 RHC $\bar{\nu}_\mu$ parameters. This was seen in the last iteration of the analysis, and the effect this has on the far detector analysis

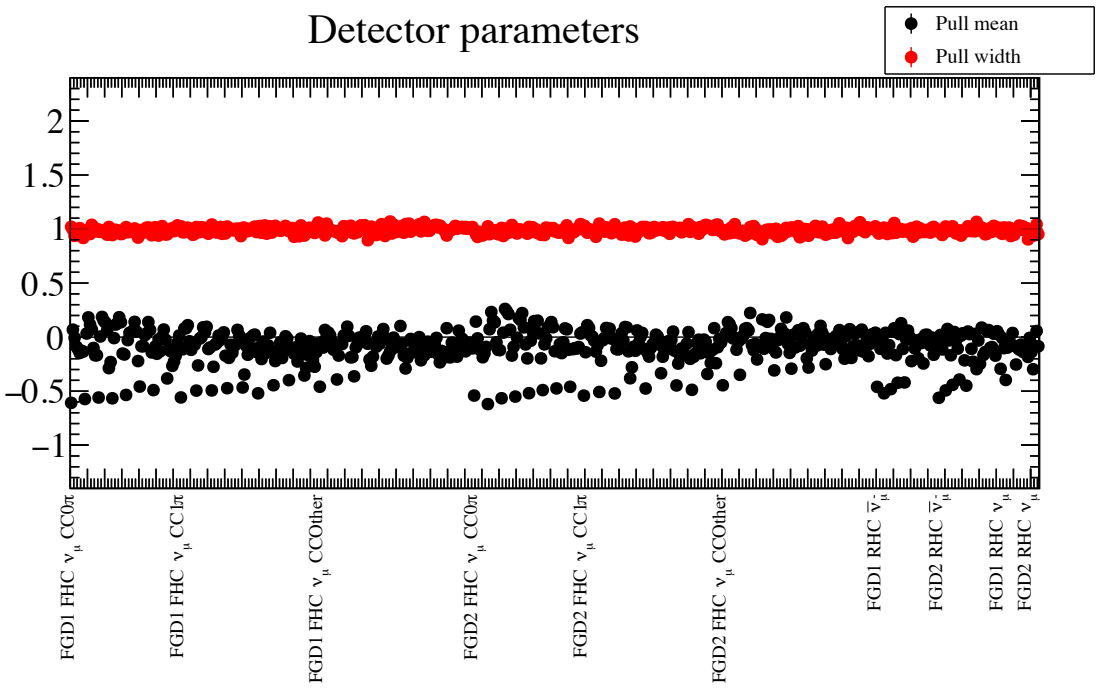

Figure 7.11: Binned detector normalisation parameter pull means and widths. The fitted detector parameters have a more significant and consistent bias than the flux or cross-section parameters in large part due to the strong correlations between bins. Groups of parameters can be seen with a clear trend of large (50%) biases which correspond to the most backwards-going angular bins in each slice of momentum. These bins have large systematic uncertainties due to poorer reconstruction in high-angle and backwards-going tracks and are highly correlated with each other.

has been tested and is discussed below. The widths of the pull distributions suggest that the uncertainties are well estimated by the fitter.

Similarly to the flux parameters, the ND280 detector parameters in Figure 7.11 show little bias in the majority of the detector bins. Some bins, however, show a larger bias, which seems to be correlated to the $p - \cos\theta$ structure of the parameters. Since the detector parameters are marginalised over and not propagated to the far detector fits, the effect of biases in a small number of these bins on the should not impact on the oscillation analysis.

In Figure 7.12 the cross-section parameters generally seem unbiased with the notable exception of the 2p2h Shape Carbon, Iso-$1/2$ non-resonant background, and the $\Delta E_{\mathrm{rmv}}^{\nu,\mathrm{C}}$ parameters. Pulls cannot be accurately calculated for parameters whose uncertainties allow them to reach a physical boundary for example, the 2p2h shape and Iso-$1/2$ background dial have physical boundaries at 0. The impact of the $\Delta E_{\mathrm{rmv}}^{\nu,\mathrm{C}}$ parameter biases on the far detector are likely to be covered by the fake data fit study, and the differences in the results of the BANFF and MaCh3 treatments for which

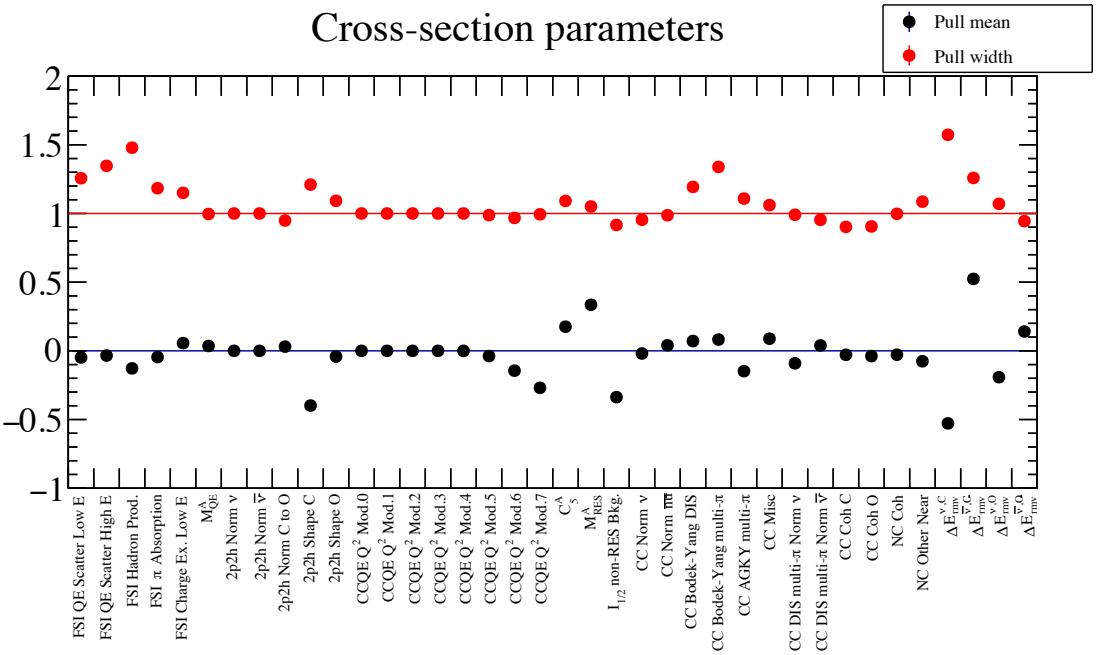

Figure 7.12: Cross-section parameter pull means and widths. A large bias can be seen in the 2p2h shape parameter which modifies the cross-section calculated using the Valencia model [38] to be more like the Martini model [128], that uses a mirrored spline at its boundaries and can prevent proper treatment of correlations in the fit when pushed to this extreme due to a zero second derivative in the likelihood. The high Q^2 and CC-Resonant, and ΔE_{rmv} parameters each have strong correlations with other parameters which can cause biases.

the impact on the oscillation parameter contours have been studied in detail.

7.4.2 Bias-corrected Asimov fit

To assess the impact of any bias at the near detector on the T2K results, a bias-corrected Asimov fit result is produced. The Asimov result is modified by the mean bias expected for each parameter multiplied by the Asimov postfit uncertainty for each parameter. Even though the corrected result is not the true value of the Asimov parameters, it is the best estimate for a bias at an unknown part of the parameter space. The effect of the bias corrections on the predicted event rates at the near detector is shown in table 7.1. The largest changes are in the CC1π samples with the RHC FGD2 ν_μ CC 1π dropping by 2.77% of the nominal prediction. The total change across all of the samples is 1.3%. The bias in the total event rate for the FHC FGD1 CC0π of 1.49% is larger than the total detector systematic uncertainty on the sample of 1.20%, though this is still smaller than the systematic uncertainty once the possible variations of the flux or cross-section parameters are included.

The FHC FGD1 CC0π and RHC ν_μ-background FGD2 CC1π samples have a bias larger than their nominal detector systematic uncertainties of 1.20% and 2.72% respectively, though these are within the total postfit systematic uncertainties with the uncertainties on the flux and cross-section systematics included. The total detector systematic uncertainty on the event rate for each sample is listed in table B.2 in appendix B.

To check that these biases do not significantly impact the results of the oscillation analysis, the bias-corrected Asimov postfit parameter values and covariance matrix were passed to the P-Theta fitting group for comparison to results from a nominal Asimov fit, shown in Figures 7.13 and 7.14. This demonstrates that the effect of the BANFF fitter bias on the T2K oscillation result is negligible, and as will be discussed in chapter 9, is smaller than differences seen between fitters due to implementation differences. Only a slight change is made in the best-fit value of δ_{CP}, though this region of the likelihood is known to be very flat, and the 68%, 90% and 99.7% confidence level contours are unchanged.

Beam	Topology	Target	Nominal	Bias Corrected	Change (%)
FHC ν_μ	0π	FGD1	27951.8	27534.05	1.49
		FGD2	27558.87	27258.62	1.09
	$1\pi^+$	FGD1	8358.62	8149.57	2.50
		FGD2	6723.79	6566.26	2.34
	Other	FGD1	7031.49	6912.58	1.69
		FGD2	6454.74	6345.15	1.70
RHC$\bar{\nu}_\mu$	0π	FGD1	7270.33	7297.51	0.37
		FGD2	7036.5	7091.23	0.78
	$1\pi^-$	FGD1	694.32	679.55	2.13
		FGD2	624.69	614.56	1.62
	Other	FGD1	1286.79	1269.30	1.36
		FGD2	1176.53	1159.12	1.47
RHC ν_μ	0π	FGD1	3035.59	2999.02	1.20
		FGD2	3012.4	2975.55	1.22
	$1\pi^+$	FGD1	1159.01	1130.24	2.48
		FGD2	930.64	904.85	2.77
	Other	FGD1	1073.13	1053.38	1.84
		FGD2	1000.39	980.62	1.98
Total			112379.63	110921.14	1.30

Table 7.1: The event rates for the nominal prediction compared to the bias-corrected prediction at ND280. The most significant changes are in the CC 1π samples which each change up to 2.77%. With the exception of the FHC FGD1 CC0π and RHC ν_μ-background FGD2 CC1π samples, which have a bias larger than their nominal detector systematic uncertainties of 1.20% and 2.72% respectively, all of the other sample event rates have bias smaller than their detector systematic uncertainties.

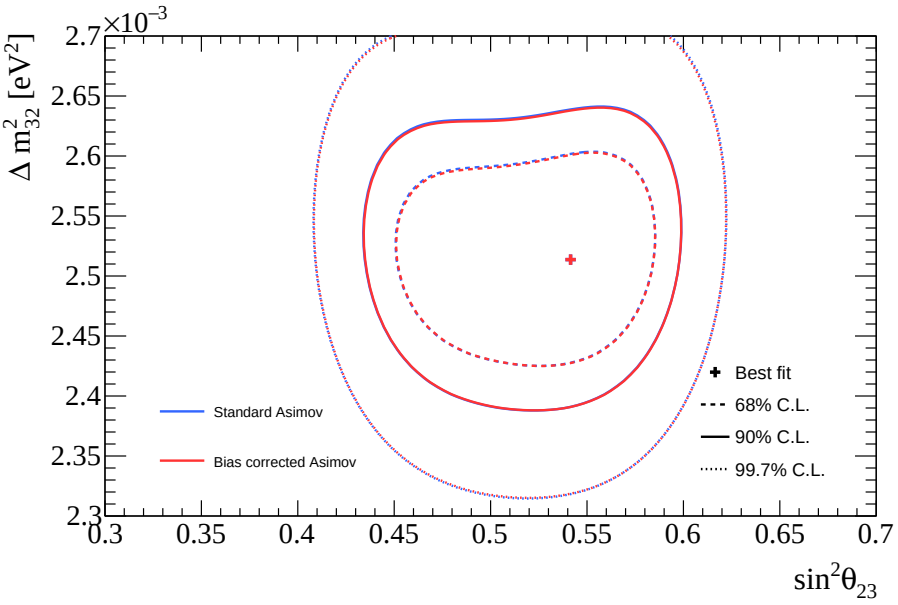

(a) Normal ordering

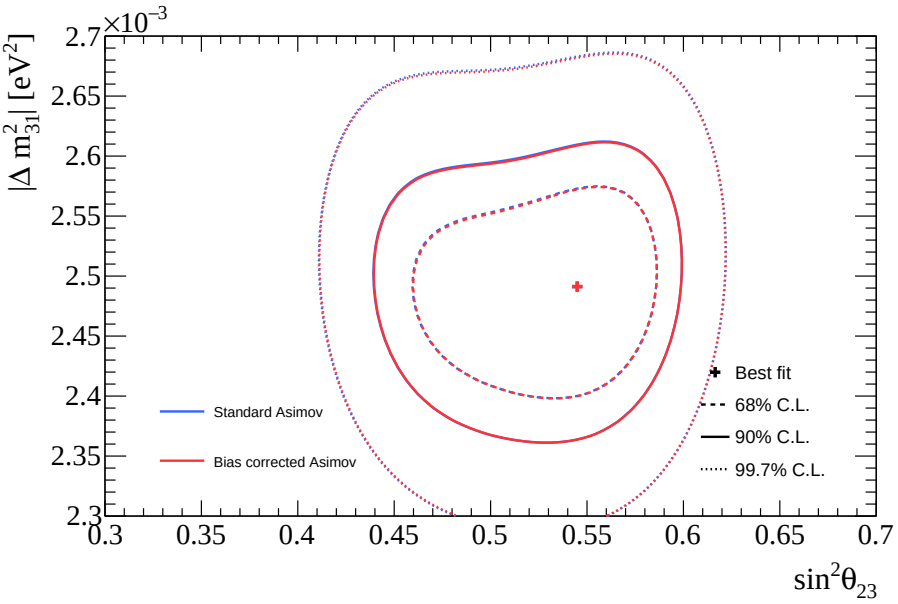

(b) Inverted ordering

Figure 7.13: Comparison of the bias-corrected Asimov fit results (red) to the nominal Asimov results (blue) for the disappearance channel parameters, Δm^2_{32} (Δm^2_{31}) and $\sin^2\theta_{23}$ for NO (a) and IO (b). Though there is a very slightly decreased Δm^2_{32} contour, this is far smaller than the impact from fake data studies and the additional uncertainty applied through smearing.

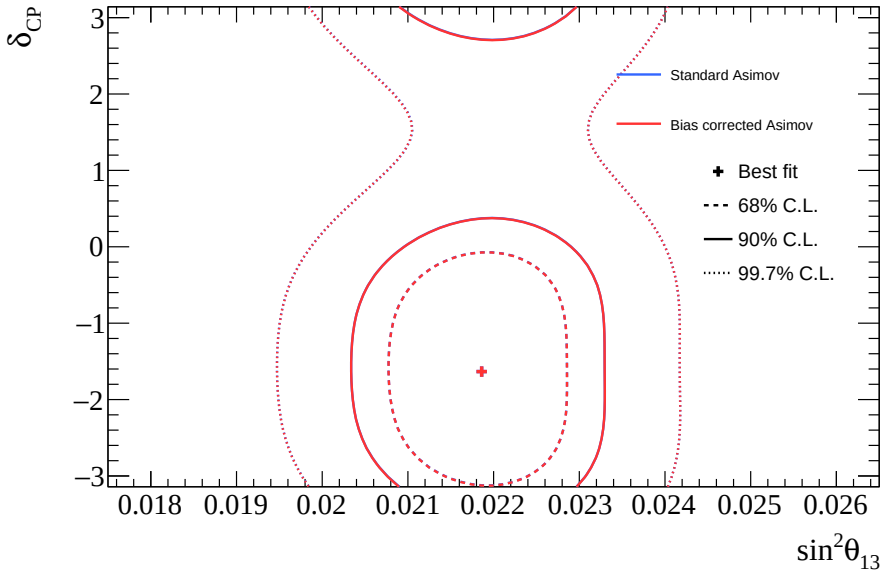

(a) Normal ordering

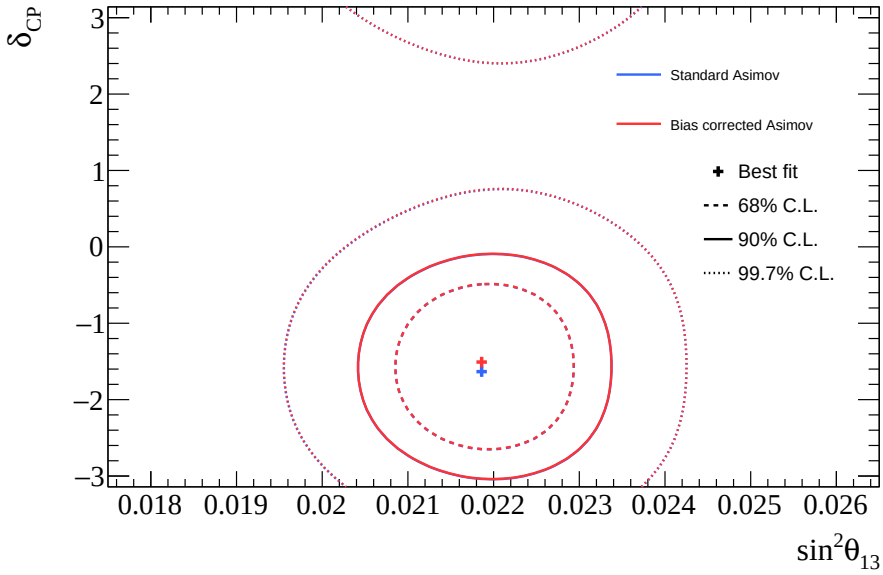

(b) Inverted ordering

Figure 7.14: Comparison of bias-corrected Asimov fit results (red) to the nominal Asimov results (blue) for the PMNS appearance channel parameters, δ_{CP} and $\sin^2\theta_{13}$ for NO (a) and IO (b). The largest impact of the ND280 biases on the best-fit point of any of the oscillation parameters is in inverted hierarchy δ_{CP}, where the likelihood is relatively flat. This impact, however, is still small, and the reported contours are essentially unchanged, when compared to the smearing applied to cover extra uncertainty from the fake data robustness studies.

Chapter 8

Results of the near detector fit

In this chapter the results of the BANFF near-detector fit to data, and data-model compatibility studies, are discussed. In the standard BANFF data fit, the full T2K model is used to fit the nominal MC to the data. To study the ability of the model to describe the region of parameter space in which the best-fit point lies, a p-value can be calculated from fits to toy data sets thrown from the model. A p-value is the probability of observing a result as extreme or more than that of the data according to the model. The toy data used for the p-value calculation can also indicate a bias in the failure rate of the fit in the region of the data which may also indicate that the model is a poor description of the data.

8.1 Fit to data

The BANFF postfit prediction shows improved agreement with the data when compared to the prefit prediction across all samples as expected. Total data event rates of each sample and the ratios of prefit and postfit MC predictions to data, shown in table 8.1, show a sub-5% disagreement between all of the postfit predictions and data, with all of the FGD1 CC0π samples within a subpercent difference, and the largest FGD2 CC0π sample difference being 3% between the postfit prediction and the data for the RHC ν_μ-background. The most significant change to the nominal MC is the increase in the CCQE (1p1h) cross section, which can be seen in figure 8.1, to correct the deficit in the prefit CC0π sample event rate prediction. This is also correlated with shape changes in both the flux and the momentum spectra, preferring low-E_ν and high-Q^2 CCQE events. The impact of the BANFF fit on the neutrino energy spectrum (convoluted with the cross section) can be inferred from the event rates in E_ν^{rec} shown in figure 8.2(b), and the momentum transfer distribution can be inferred from the event rates in Q^2_{rec}, shown in figure 8.3(b). The full set of prefit and postfit p_μ, $\cos\theta_\mu$ and p_μ-$\cos\theta_\mu$ distributions can be found in appendix E.

Beam	Topology	Target	Data	Prefit/Data	Postfit/Data
FHC	0π	FGD1	33443	0.91	1.00
		FGD2	33156	0.91	1.00
	$1\pi^+$	FGD1	7713	1.09	1.03
		FGD2	6281	1.09	1.03
	Other	FGD1	8026	0.88	0.99
		FGD2	7700	0.84	0.95
RHC	0π	FGD1	8388	0.97	1.00
		FGD2	8334	0.94	0.98
	$1\pi^-$	FGD1	698	1.00	0.98
		FGD2	650	0.96	0.98
	Other	FGD1	1472	0.88	1.00
		FGD2	1335	0.89	1.03
ν_μ in RHC	0π	FGD1	3594	0.89	1.00
		FGD2	3433	0.92	1.03
	$1\pi^+$	FGD1	1111	1.04	1.04
		FGD2	926	1.01	0.99
	Other	FGD1	1344	0.80	0.96
		FGD2	1245	0.81	0.96

Table 8.1: Event rates for each of the ND280 selections for data and the ratios of the prefit and postfit MC predictions to the data. The most significant changes are in the two ν_μ CC0π samples, which see an increase of around 10% each, and also across the CC-Other samples for each neutrino sign and beam mode, which have a 10-20% deficit in the prefit, which is reduced to less than 5%.

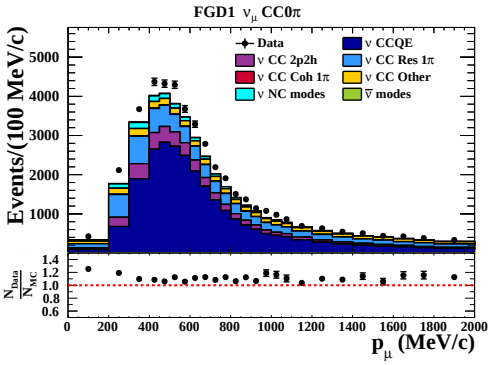

(a) Prefit

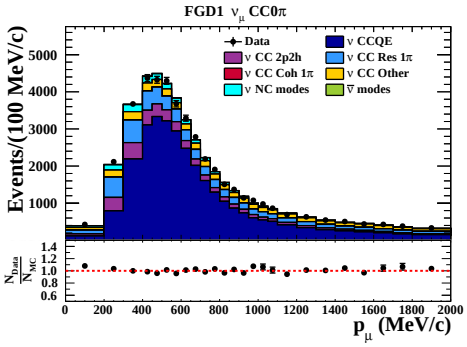

(b) Postfit

Figure 8.1: Comparison of the prefit and postfit event rate prediction in with data (black points) in reconstructed muon momentum, p_μ, broken down by true interaction channel, for the FHC FGD1 CC0π sample. The bottom panel shows the ratio of the data to the MC. The BANFF fit corrects for the 9% MC deficit in the event rate by increasing the CCQE cross-section at low E_ν and high Q^2.

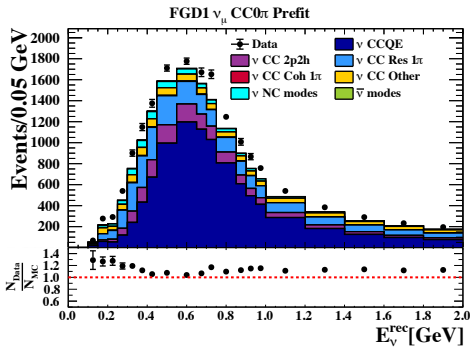

(a) Prefit

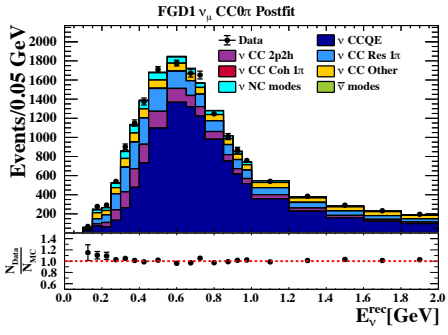

(b) Postfit

Figure 8.2: Comparison of the prefit and postfit event rate prediction with data (black points) in reconstructed neutrino energy, E_ν^{rec}, broken down by true interaction channel, for the FHC FGD1 CC0π sample. The bottom panel shows the ratio of the data to the MC. As with the p_μ event rates, the E_ν^{rec} MC prediction is enhanced across the spectrum.

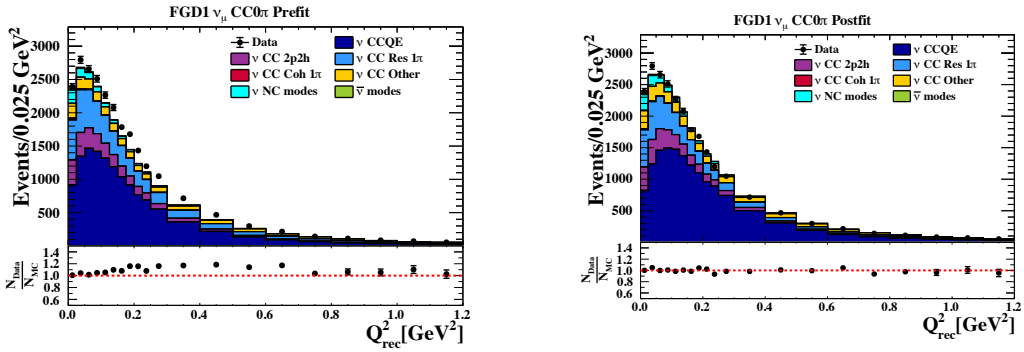

(a) Prefit

(b) Postfit

Figure 8.3: Comparison of the prefit and postfit event rate prediction with data (black points) in reconstructed four-momentum transfer, Q^2_{rec}, broken down by true interaction channel, for the FHC FGD1 CC0π sample. The bottom panel shows the ratio of the data to the MC. The region between 0.2 and 1.0 GeV^2 see the largest increase, which agrees with the postfit values of the CCQE Q^2 normalisation parameters. Below 0.1 GeV there is a slight suppression of the CCQE contribution, which also agrees with the postfit Q^2 normalisation parameter values.

8.1.1 Flux parameters

A change int the shape of the neutrino energy spectrum is indicated by the FHC ν_μ parameters, shown in figure 8.4. The normalisation parameter for the lowest E_ν bin (0-700 MeV) is increased by 11%, and the parameters gradually decreasing in value to 9% below the nominal for the highest E_ν bin (>2.5 GeV). This enhances low E_ν events whilst suppressing the highest E_ν events. This effect is not immediately obvious from the event rates shown in figure 8.2, in which the effect is also convoluted with an increase in the overall CCQE cross-section and a change in the shape of the underlying Q^2 distribution. The same effect is seen for the FHC ν_e flux parameters, as well as similar behaviour for the right-sign components of the RHC flux, though with a less gradual change between the low and high E_ν normalisations. The wrong sign components of each beam mode see an enhancement above the nominal prediction. Whilst only the ND280 parameters are fit in the BANFF fit, the SK flux parameters are also constrained by the BANFF fit due to their strong prefit correlations with the near detector flux parameters.

8.1.2 Detector Parameters

For the FHC samples, the general trend in the CC0π observable normalisation parameters (figure 8.5) is a decrease in the normalisation of the less forward, more high-angle parameters with a sharp increase towards the most forward-going bins of muon angle, with the exception of the

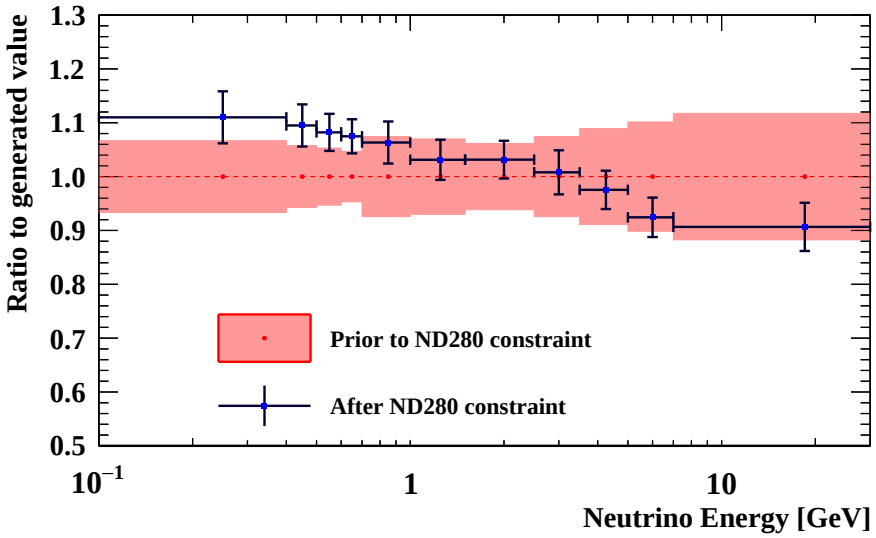

Figure 8.4: SK FHC ν_μ flux parameter values shown in bins of the neutrino energy, E_ν, they are applied to. E.g. The increase of 11% in the lowst bin corresponds to a weight of 1.11 applied to all FHC ν_μ events at SK. The prefit value shown (red) is post-NA61/SHINE tune. A shape-like effect is seen corresponding to a change in the shape of the distribution of events in E_ν is changed, favouring events with low neutrino energy over events with higher neutrino energies compared to the NA61/SHINE tuned NEUT-generated prediction. Whilst the lowest of these bins are all pulled around two standard deviations from the prior central value, the resultant penalty for this is not prohibitive due the strong correlations between adjacent flux parameters. This large number of degrees of freedom allows for significant changes within both the flux and the cross-section model through these correlations.

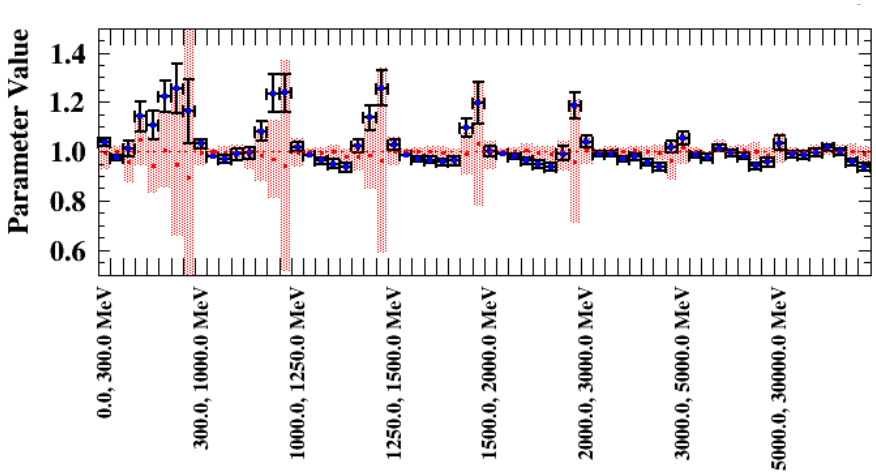

Figure 8.5: FGD1 FHC CC0π binned detector normalisation parameters. The detector parameters are applied as normalisations to the total number of events in each detector bin of the sample. Detector parameters are in order of lepton-angular bin from events with a backwards-going muon (left of each momentum bin) to events with a forward-going muon (right of each momentum bin) and then grouped in muon-momentum bins from lowest to highest as indicated by the axis labels. The prior central value and uncertainty, shown by a red dot and band, are calculated from event rates under variations of the underlying detector systematics. The postfit values and uncertainties, shown by a blue dot and black cross, are generally lower than their prefit values in the most backwards-going angular bins but higher in the most forward-going bins, with the exception of the lowest-momentum bin in which they are consistently higher, and the highest-momentum bin in which they are consistently lower.

lowest muon momentum range (0-300 MeV). The normalisations of the low-momentum bins are typically high compared to their prior central value whilst the highest-momentum bin (5-30 GeV) sees a general suppression across the angular range. The CC1π and CC-Other sample detector parameters have similar responses. These parameters give extra freedom to the model to cover the ND280 data, but are marginalised and so only affect the oscillation analysis through their correlations during the near detector fit.

8.1.3 Cross-section parameters

CCQE

Despite a strong constraint of ± 0.06 GeV placed on the prior value of $M_A^{QE} = 1.026$ GeV based upon fits to external bubble chamber data [139], the fit prefers values much closer to the generated value of 1.21 GeV, pulling the parameter significantly away from its prior central value, as can

be seen in the leftmost bin of figure 8.6(a). It is this change to the CCQE model which largely compensates for the under-prediction of events in the CC0π sample. The large penalty for this shift in M_A^{QE} is compensated for by the large number of degrees of freedom in the correlated flux and Q^2 normalisation parameters. There is a shape-like effect in bins of Q^2 with the fit preferring higher values of momentum transfer resulting in a low Q^2 suppression which was expected from previous cross-section measurements. This shape effect can be clearly seen in figure 8.6(b) in the true Q^2 normalisation parameters and also in figure 8.3 in the distribution of the reconstructed momentum transfer, Q^2_{rec}, which has a clear enhancement of high Q^2_{rec}, and a suppression of the lowest bins of Q^2_{rec} for CCQE events. These are both correlated with the shape effects seen in the FHC ν_μ flux parameters and this is the most clear interplay between the flux and cross-section parameters.

Each of the CCQE ΔE_{rmv} parameters are constrained but sit comfortably within their prior uncertainties as shown on the right-hand side of figure 8.6(a). Each of the carbon parameters are constrained to within ± 2 MeV, and the oxygen parameters propagated to SK are constrained to under ± 3 MeV, which is a significant reduction of the ± 6 MeV prior constraint. As these parameters affect the shape of the CCQE contribution to the event rate, they are weakly correlated with parameters controlling the strength of the dominant backgrounds to the CC0π topology, including the 2p2h normalisation parameters, and the CC-resonant pion production form factors. The decrease of the oxygen parameter values as compared to the prior value suggests a higher typical p_μ in the data than predicted by the unconstrained model. This can possibly change the shape of the single-ring samples at SK, through which there is sensitivity to the oscillation parameters. This will be propagated to the far detector as a higher typical momentum prediction, and so a higher typical calculated E_ν^{rec}. This means the position of the ν_μ disappearance signal in the data is shifted to lower true E_ν relative to the predicted peak of the spectrum, which can lead to a higher best-fit value of $\Delta m^2_{32}(\mathrm{NO})/\Delta m^2_{31}(\mathrm{IO})$.

2p2h

The rate of 2p2h events in neutrino interactions remains close to the nominal prediction, though with a slight enhancement; however, the antineutrino 2p2h cross section sees a suppression, as is evident from the decrease in the 2p2h Norm $\bar{\nu}$ parameter shown in figure 8.6(a). Whilst the carbon 2p2h cross section is increased, the oxygen 2p2h cross section remains at the prior value, though with a significant reduction in the uncertainty.

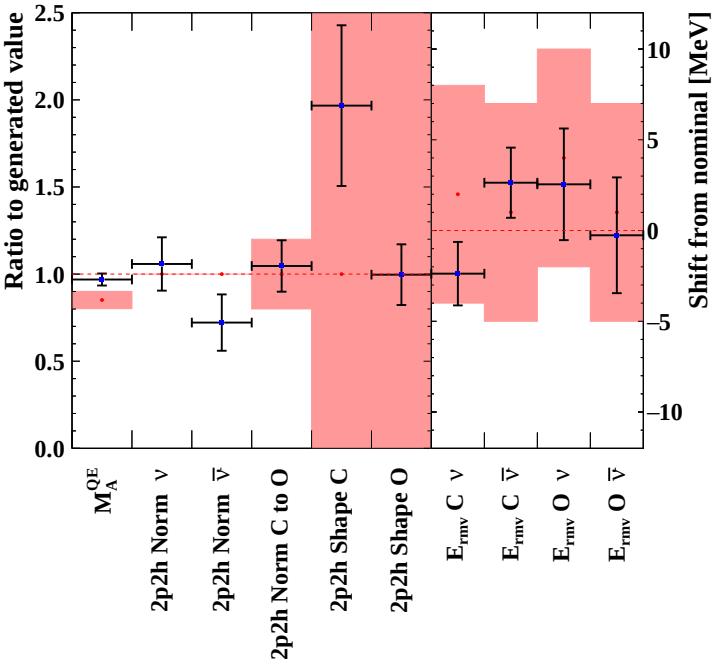

(a) CCQE and 2p2h cross-section and E_B parameters

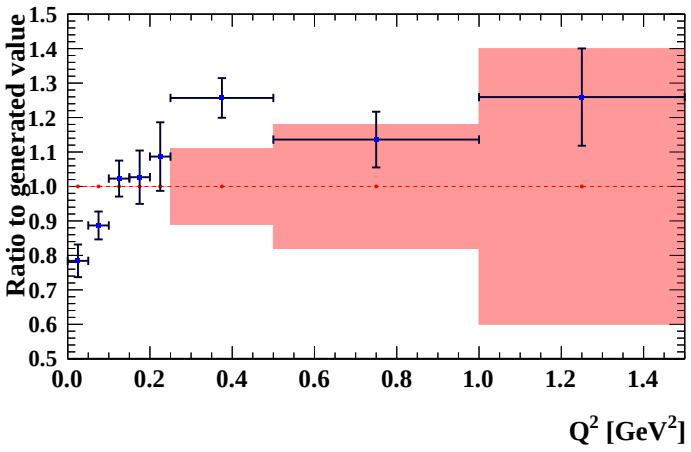

(b) CCQE Q^2 normalisation parameters

Figure 8.6: Normalised CC0π parameters controlling the CCQE and 2p2h models. The two 2p2h normalisation parameters are unbounded with flat prior uncertainties. M_A^{QE} is pulled significantly away from its central value of 1.026 GeV despite its strong constraint of $\pm$0.06 GeV towards the NEUT-generated value of 1.21 GeV. This is related to strong correlations with the eight Q^2 normalisation parameters and the flux parameters which both see shape-like effects, preferring low E_ν and high Q^2 CCQE events.

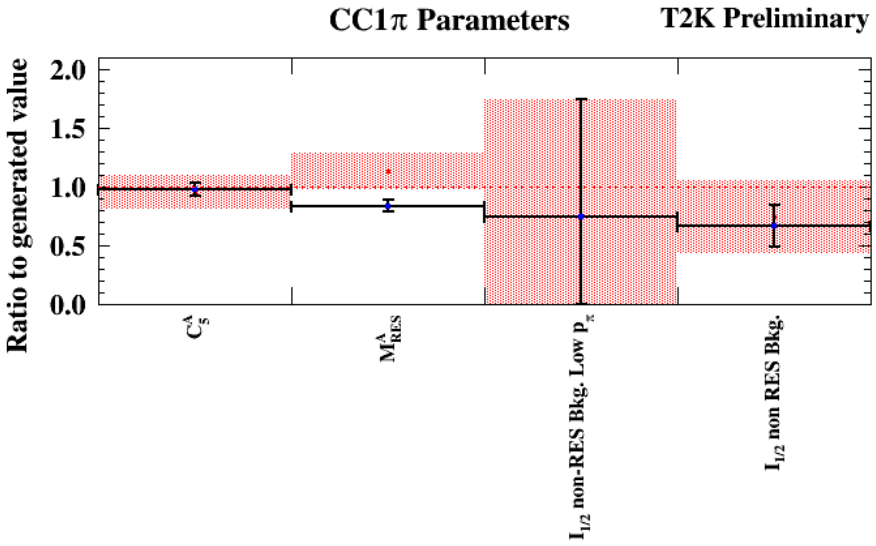

Figure 8.7: Normalised CC1π parameters controlling the CC-Resonant model and its backgrounds. The CC-Resonant form factor parameters, M_A^{RES} and C_5^A, and the standard non-resonant Iso-1/2 background parameters are constrained by the fit, though the specific Iso-1/2 background parameter for $\bar{\nu}$-low p_π events was fixed due to lack of sensitivity. C_5^A and the non-resonant background parameters are relatively unchanged from their priors, but the axial-nucleon mass, M_A^{RES}, is pulled significantly far from its prior value by approximately 2 σ, indicating significant mismodelling of the nuclear model in the resonant interaction.

CC-resonant pion production

In the CC-Resonant interaction model, the value of M_{Res}^A is pulled significantly away from its prior-central value despite a strong constraint (figure 8.7). This may be indicative of possible missing nuclear effects in the pion production model. The vector form factor, C_5^A, and the Iso-half background parameters are constrained though largely unchanged. Both prefit and postfit predictions for the rate of CC1π events (in which pions pass through the TPC) have poor agreement with data when binned in pion momentum. Although this has little impact on the lepton kinematics that T2K relies on most for its oscillation analysis, it is indicative of current failings of the pion production models.

CC-MultiPi and DIS

Figure 8.8 shows the postfit multi-π and DIS normalised parameter values and constraints. The Bodek-Yang DIS parameter is pulled upwards, just outside its prior uncertainty, whilst the Bodek-Yang Multi-π parameter is unchanged. The DIS-normalisation parameters are both pulled to the

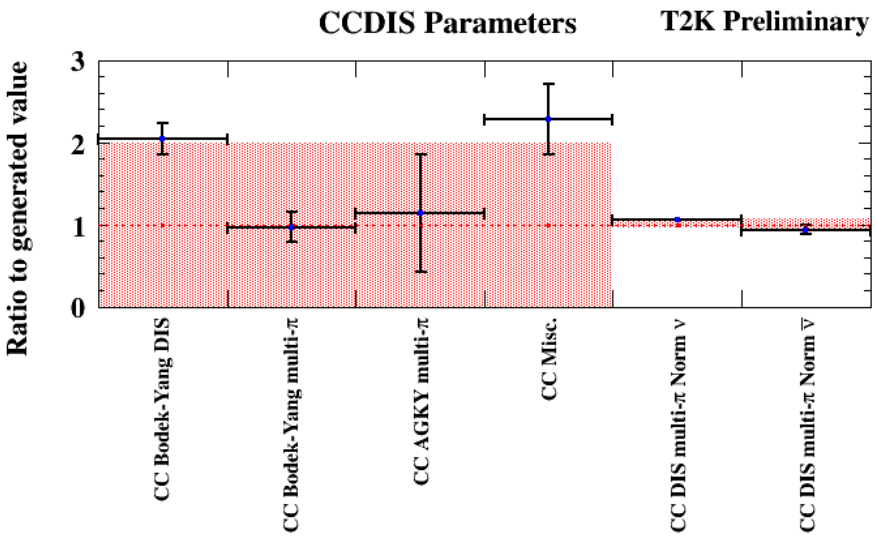

Figure 8.8: Normalised Multi-π and DIS cross-section parameters. The postfit Bodek-Yang parameters indicate a suppression (parameter at 1) is applied to the Multi-*pi* cross section ($W < 2$ GeV), but turned off (parameter at 2) for the DIS cross-section ($W > 2$ GeV). The CC Misc. parameter is increased with its central value lying outside its prior uncertainty suggesting an underprediction of the backround processes such as eta decays in the nominal model. The overall normalisations are pulled to the edge of their prior constraint favouring ν of $\bar{\nu}$ multi-π and DIS events.

edge of their uncertainty bands in opposite directions; the neutrino DIS normalisation sees an enhancement whilst the antineutrino DIS-Multiπ normalisation is suppressed and the parameter is pulled to the lower edge of its prior uncertainty. The CC-Miscellaneous parameter is pulled outside of its prior uncertainty with its value significantly increasing, suggesting larger contributions from background processes such as eta production and decays.

FSI

Whilst there is some movement in the FSI parameters (figure 8.9) and each of the FSI uncertainties is significantly reduced by the BANFF fit, each parameter value is consistent with its prior constraint. Quasielastic scattering is reduced in both the low and high E_ν bins, as is charge exchange in low E_ν events. Final-state hadron production and pion absorption parameters both see increases which can lead to greater background contributions to both the CC0π and CC1π samples from CCQE events which under go hadron production and resonant events which undergo absorption, respectively. A larger contamination of resonant events in the single-ring samples at SK could bias

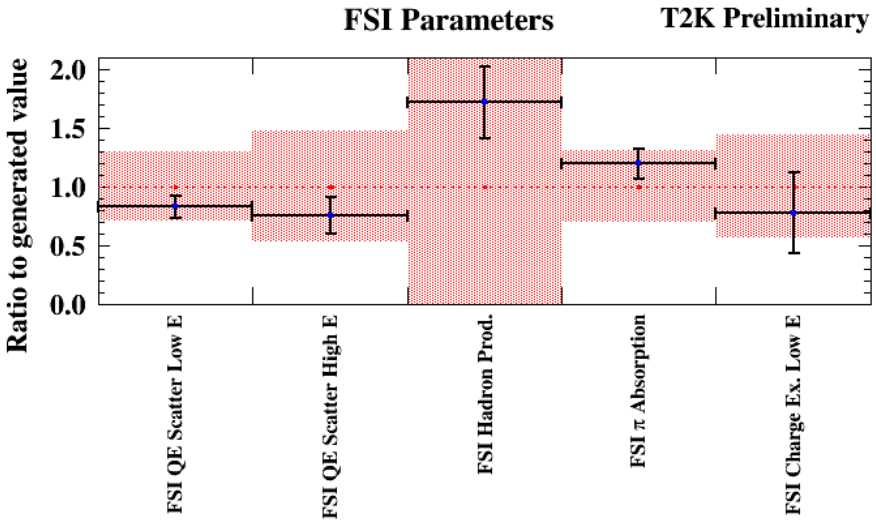

Figure 8.9: Normalised FSI parameters. Each of the postfit central values (blue points) lie comfortable within their prior constraints though differ somewhat from the prior central value (red dot). QE scattering is reduced across the neutrino energy spectrum, as is charge exchange. Hadron production and π-absorbtion are both increased, this may increase the number of 1p1h events with a CC1π topology, aslo increasing the number of resonant events in which the pion does not exit the nucleus, leading to a greater predicted CCQE background the CC1π sample as well as a greater predicted resonant background in the CC0π sample, which may bias the calculation of E_ν^{rec} at SK.

the calculation of $E_{\text{QE}}^{\text{rec}}$ and the constraint on Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO).

8.2 Correlations

The real power of the ND280 fit comes not simply from the evaluation of and reduction in the uncertainty of each of the parameters in the model, but more importantly through the evaluation of the correlations between parameters in different parts of the model in the absence of oscillations, in particular the anti-correlations introduced between the flux and cross-section parameters as shown in the difference between figures 8.10(a) and 8.10(b).

These large anti-correlations between the flux and cross-section parameters reduce the overall uncertainty on the event rate predictions at SK. A closer inspection of the cross-section correlations highlights the strong correlations amongst the Q^2 normalisation parameters and the very strong anti-correlations of the highest Q^2 parameters with M_A^{QE}. The strong correlations between the flux parameters, M_A^{QE}, and the Q^2 normalisation parameters can be seen in figure 8.11. The full postfit cross-section correlation matrix can be found in figure C.1 of appendix C. In the SPP

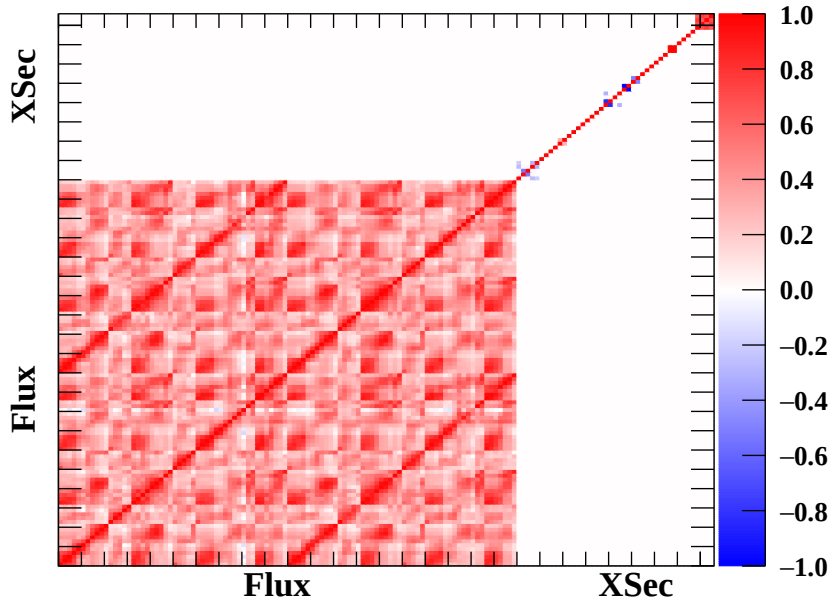

(a) Prefit flux and cross-section correlations

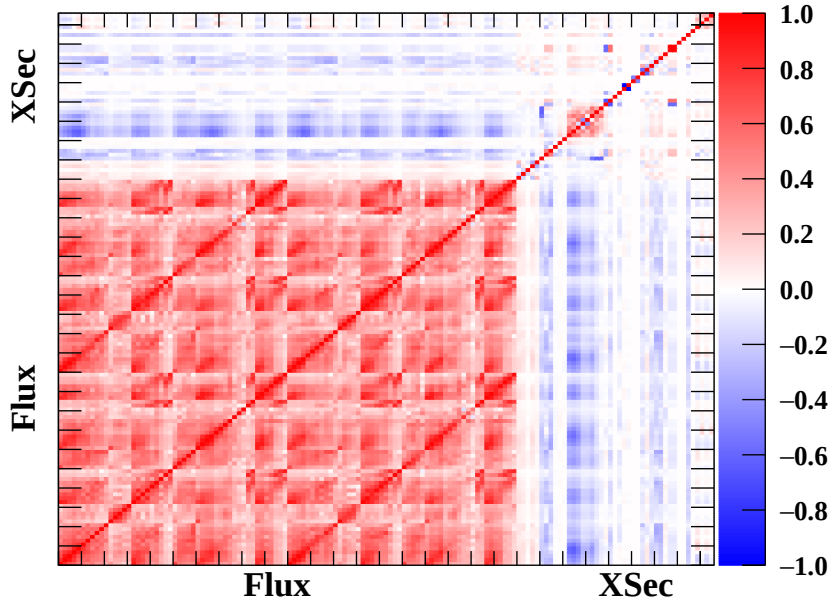

(b) Postfit flux and cross-section correlations

Figure 8.10: Correlations (red) and anticorrelations (blue) between the flux and cross-section parameters in the prefit model (a) and the BANFF postfit model (b). Strong anti-correlations are introduced by the BANFF ND280 fit between the flux and the cross-section parameters and is particular clear between the flux and the block of eight correlated Q^2 normalisation parameters, the detail of which is shown in figure 8.11.

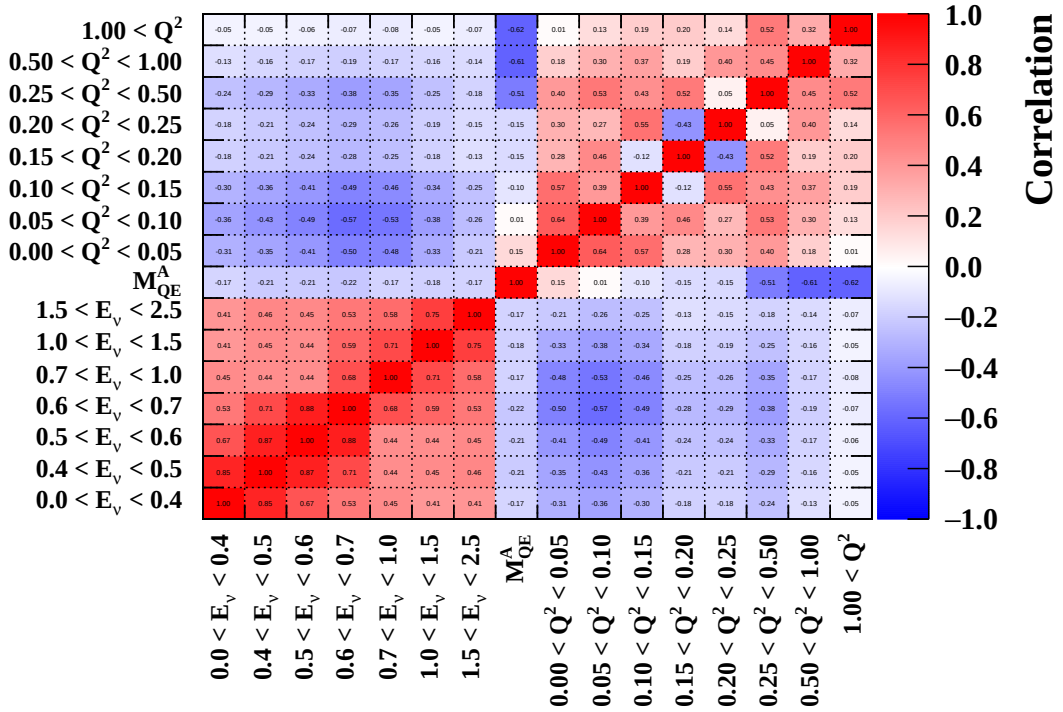

Figure 8.11: BANFF postfit correlations (red) and anticorrelations (blue) between SK ν_μ flux normalisation parameters, in units of GeV up to 1 GeV, M_A^{QE} and the Q^2 normalisation parameters in units of GeV^2 up to 1 GeV^2. These are the (anti)correlations which have the largest impact on the predicted CCQE event rate and the event rate of the single-ring SK samples. As such, these are the most significant (anti)correlations for the extraction of the oscillation parameters from the SK event rates, particularly due to their impact on constraining the shape of the interaction in Q^2 and E_ν.

model the nucleon form factors, M_A^{RES} and C_5^A, which are are strongly anti-correlated, see correlations introduced with the coherent pion production parameters. The four ΔE_{rmv} parameters see some weak correlations and anti-correlations introduced with the 2p2h normalisation, high Q^2 CCQE normalisation parameters and the SPP parameters. This is most likely due to their ability to respond to the effect of lepton kinematic-shape changes within the CCQE component of the samples, which the ΔE_{rmv} parameters produce.

8.3 Comparisons with the MCMC fitter

Comparisons of the BANFF fit results with the MaCh3 MCMC fitter at ND280 show consistent postfit flux and cross-section parameters as seen in figure 8.12. Some key differences arise due

to fitter and implementation differences. The most notable of these is the difference between the BANFF and MaCh3 in the change of the removal energy parameters from their prior values (figure 8.12(c)) which is in part due to the way the best fit value is extracted from the multi-modal structure of the MaCh3 E_B posterior probability density function, but also due to the mechanics of the implementation.

Examination of the posterior probability density distribution for the kinematically-shifting ΔE_{rmv} parameters implemented in MaCh3 shows clear binning effects with a multi-modal peak structure as shown in figure 8.13. The discrete likelihood changes cause the steps in the Markov Chain to bunch around specific values; without infinite statistics and/or an infinitely fine binning, the fit is unable to continuously sample the parameter values. This binning effect is dependent on the choice of bins, which can lead to quite varied posterior distributions. The reweighting parameter in BANFF, however, is designed to produce a continuous response by construction. Due to this difference in approaches, each fitter finds different best-fit values for the parameters. The BANFF and MaCh3 difference is also dependent on the way the postfit value and uncertainty are extracted from the posterior distribution. The MaCh3 values in figure 8.12(c) are taken from the highest posterior density, but the uncertainty comes from a Gaussian fit to the entire posterior distribution.

The impact of this on the oscillation parameter contours was studied and is discussed in chapter 9. Other differences arise between the fitters due to the binning differences between the fits. The BANFF fit uses a standard regular binning for its sample histograms, whilst the MaCh3 fit uses an irregular but rectangular binning with varying bin boundaries in p_μ for different bins in $\cos\theta_\mu$. The MaCh3 fitter can also produce results with the standard binning and the impact of these differences on the oscillation parameters is discussed in chapter 9.

8.4 Calculating a p-value

To evaluate the compatibility of the data with the model, the probability, according to the model, of making an observation as extreme or more when compared to the nominal prediction of the model is calculated. This is essentially the same as the widely-used standard hypothesis-rejection procedure but used in reverse. Instead of making the statement that data is inconsistent with some null hypothesis, and therefore the hypothesis must be rejected, this test is used to show whether the data is consistent with the model. The criterion set for compatibility is a p-value of $p > 0.05$.

By this criterion, if the model is correct, and if there were an infinite number of universes in which the true best-fit models describing these universes are distributed according to the prefit model, and if the fraction of universes in which the data is as different from the nominal prediction as the observed data is greater than 5%, then the result is accepted.

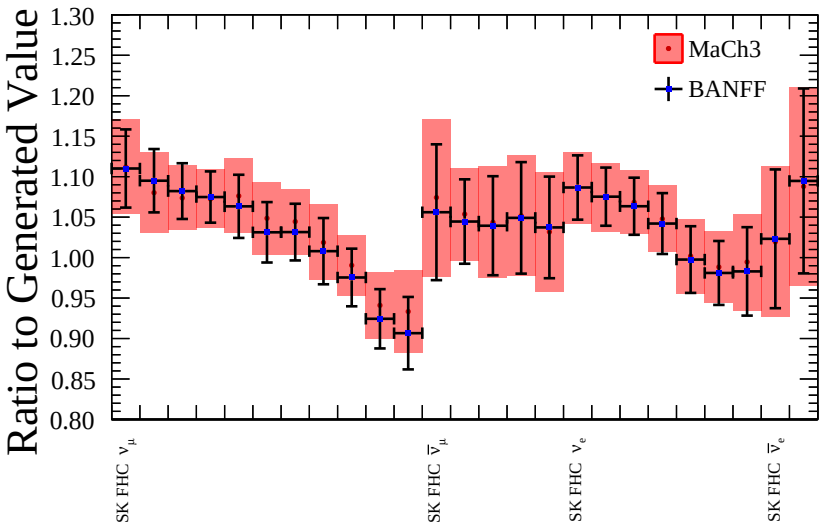

(a) SK FHC flux parameters

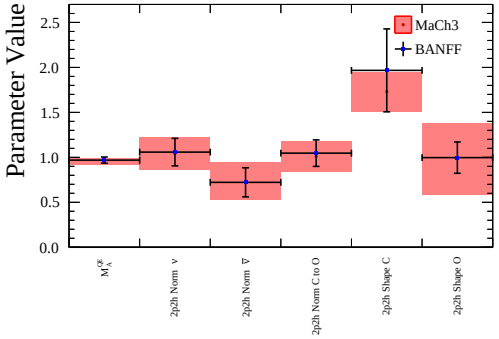

(b) The CC0π parameters

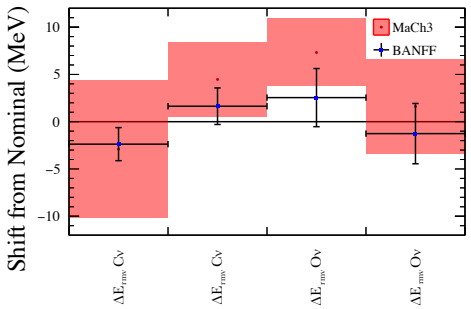

(c) The removal energy parameters

Figure 8.12: Comparison of the BANFF and MaCh3 fitter best-fit parameter values. MaCh3 best-fit values are taken from the highest value of the posterior probability density. Most parameters are consistent between fitters though the BANFF constraint on the ΔE_{rmv} parameters appears stronger due to the way the MaCh3 uncertainties are extracted from the multimodal posterior distribution. The BANFF and MaCh3 central value of ΔE_{rmv} differ by more than 1 σ from eachother. Differences in the BANFF and MaCh3 2p2h shape parameters resulting in the mirroring at the physical boundaries which can cause cusping in the likelihood which BANFF decorrelates parameters to correct for. The impact of these differences at SK is explored in section 9.3.2.

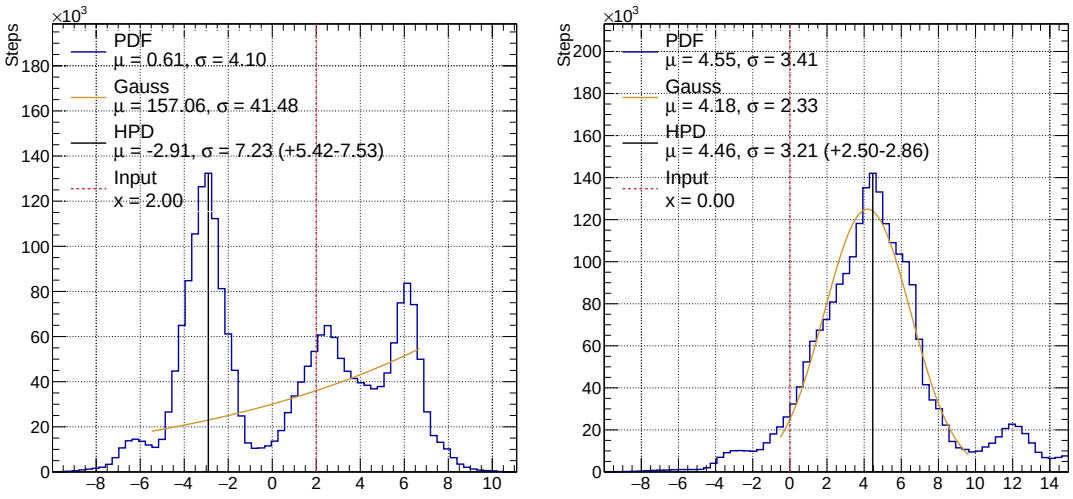

(a) $\Delta E_{\text{rmv}}^{\nu,\text{C}}$ (b) $\Delta E_{\text{rmv}}^{\bar{\nu},\text{C}}$

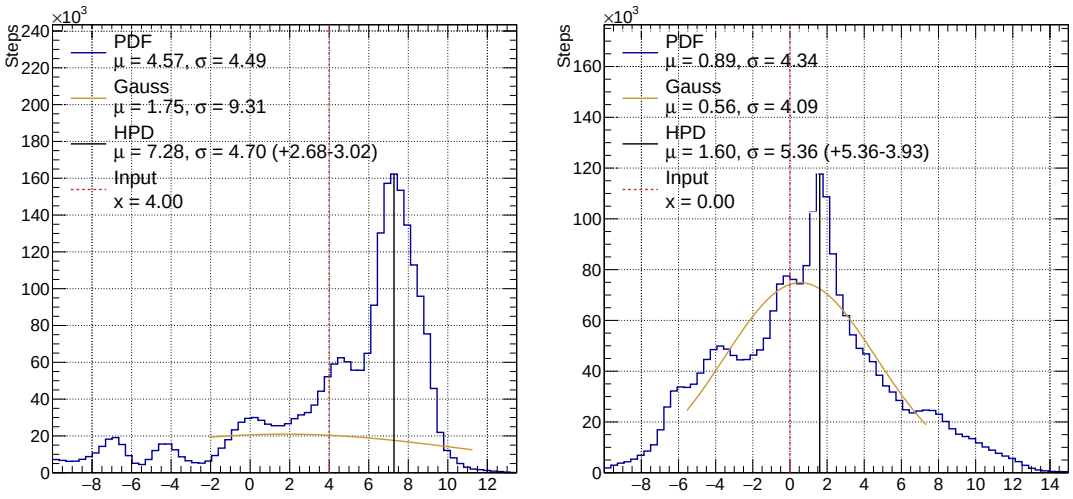

(c) $\Delta E_{\text{rmv}}^{\nu,\text{O}}$ (d) $\Delta E_{\text{rmv}}^{\bar{\nu},\text{O}}$

Figure 8.13: MaCh3 posterior distributions for the four ΔE_{rmv} parameters. A multimodal peak structure arises from marginalisation and binning effects as events crossing bin boundaries cause sudden changes to the likelihood. This results in a "bunching" together of steps in the Markov Chain. The BANFF fitter is shielded from this effect by the fact that the ΔE_{rmv} implementation is smooth by construction.

In order to calculate this, toy data sets are produced from throws of the parameter values from the prior covariance matrix. Fits of the nominal MC to these data sets are then used to rank them in terms of likelihood according to the model. The value of the χ^2_{min} found in the fit is used to rank the toys and a χ^2_{min} distribution can be found for the model. Comparisons of the χ^2_{min} found for the fit to the actual data to this distribution can be used to calculate the p-value by integrating the distribution of toys which lie at higher χ^2_{min} values than the data.

8.4.1 Toy throwing

To test the "alternative universes" described by the model, toy data sets are used. To ensure that these toys cover the parameter space predicted by the model, a large number of random toys is thrown according to the total prior covariance matrix. For this procedure, the true underlying detector parameters in `psyche` are used as opposed to the observable normalisation parameters which are fitted. From the throws of these detector parameters, throws of the flux and cross-section parameters are generated and the overall ensemble of toy data distributions to be fit is produced. For the calculation, 1387 toys were thrown; however, due to the demanding computational need involved in fitting even a single toy data set, limits on the computational resources used can impact the number of successful fits. No significant dependency or bias in the distribution of the failed toy fits when compared to the successful toy fits was found (figure 8.14(a)). This can be seen by the relatively flat distribution in the ratio of failed to successful fits in figure 8.14(b).

8.4.2 Prior p-value

For the 895 successful fits, a p-value of $p = 0.74$ was calculated demonstrating a good agreement between the model and the data. A comparison of the data χ^2_{min} and the distribution of χ^2_{min} values from the toy data fits is shown in figure 8.15. A prior p-value above 0.5 shows excellent agreement and the suggests that majority of likely combinations of parameter values, according to the prior model, are more unlikely than that which describes the data observed.

The p-value can be broken down into different contributions to the χ^2_{min} distribution. The samples are well described, with the statistical component evaluated at $p = 0.82$ ranging from 0.15 to 0.93 for the CC0π samples. The penalty contribution from the prior model means that the strong constraints give a poor cross-section-only p-value of $p = 0.01$ as fixing the flux and detector models removes most of the freedom in the model to cover the data. However, the flux-only p-values of $p = 0.74, 0.74, 0.31, 0.37$ for the FHC ν_μ, FHC $\bar{\nu}_\mu$, RHC ν_μ and RHC $\bar{\nu}_\mu$ ensure that the best-fit point is covered.

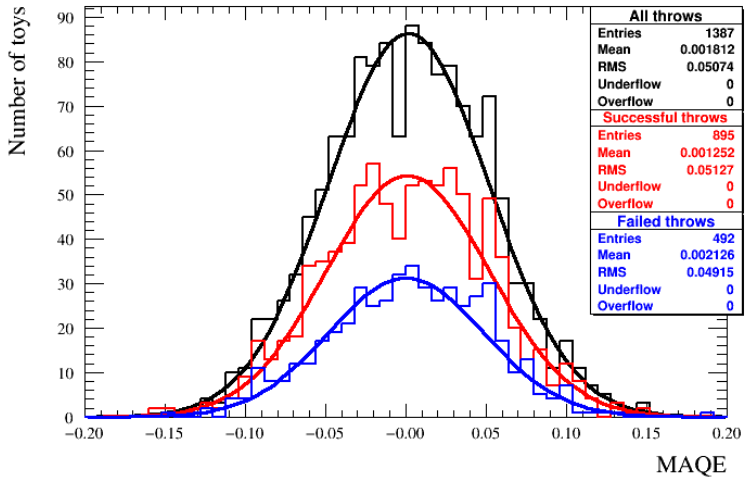

(a) Distribution of thrown values of M_A^{QE} used in the calculation of the p-value

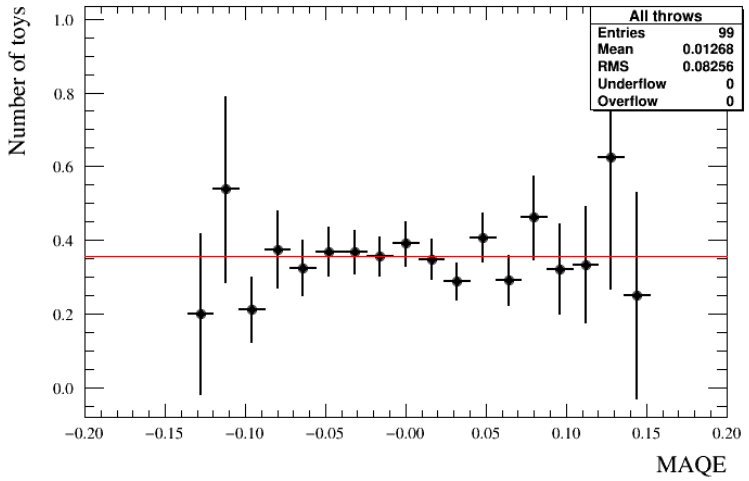

(b) Ratio of successfully fitted toys to failed fits across the range of thrown values of M_A^{QE}.

Figure 8.14: Toy throws and ratio of successful and failed fits to toys for values of M_A^{QE} in the calculation of the BANFF p-value. No bias is seen in any of the parameters.

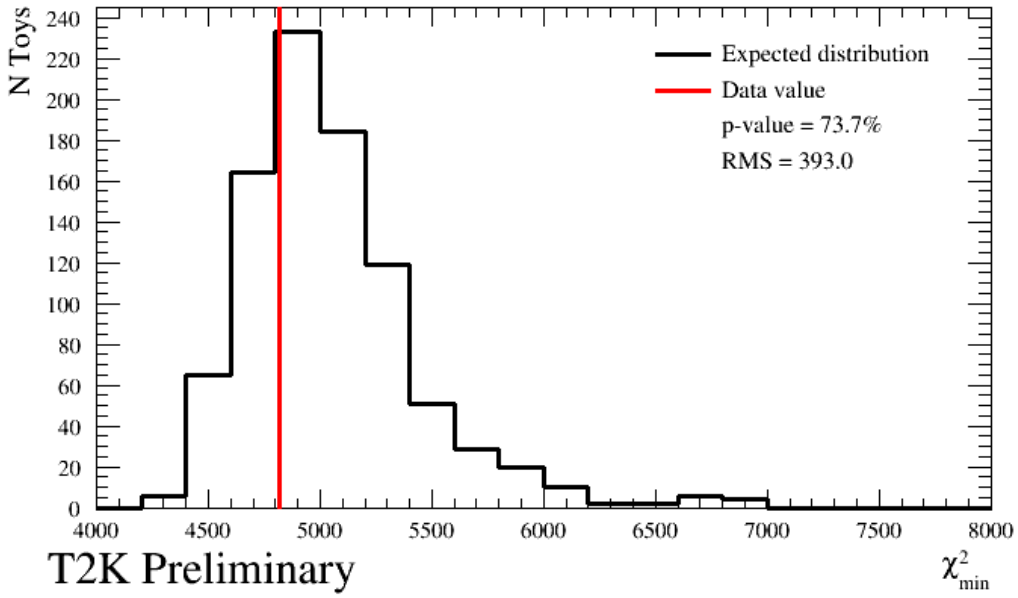

Figure 8.15: Distribution of χ^2_{min} values for fits to 895 toy data sets thrown from the T2K prefit model (black histogram). 73.4% of toy fits find a χ^2_{min} value greater than, and so less likely, than that found by the fit to the ND280 data (red line). This demonstrates that the prior model describes the data well.

8.4.3 Postfit p-value

It is also useful to test the compatibility of the post-BANFF model with our data. The test statistic used is still a fit of the nominal MC to each of the toys; however, the toys are thrown from the postfit covariance matrix. One would expect that if the best fit to the data lies at the central value of each of the parameters with the toys thrown Gaussianly around it, the p-value should approach 0.5 as there will be roughly an equal number of toys thrown with more-and less-likely χ^2_{min} values than the central value from which they are thrown. However, due to the correlations between each of the parameters, this may not be the case. A significant shortcoming of this study was the inability to throw the postfit detector parameters in the same way as for a the prior p-value, since there is no way to map the postfit observable normalisation parameter values back onto the underlying parameters in `psyche`. Due to this, the prefit values of the detector parameters were used, which due to the correlations between the different parts of the model, may have caused flux and cross-section parameters to be thrown in artificially disfavoured regions of phase space that would be avoided if the true postfit values of the detector parameters could be used. This may result in an overall inflated p-value, as many of the throws will be in disfavoured regions of phase space due to the combination of the prefit detector parameters and postfit flux and cross-section parameter correlations. This p-value is still indicative of any improvement of the different contributions to

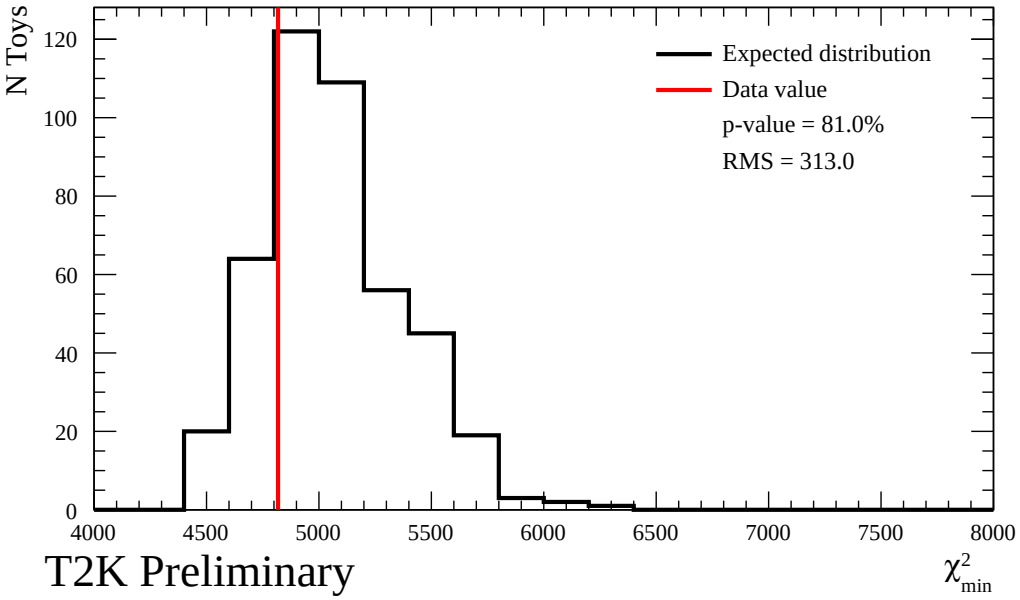

Figure 8.16: Distribution of χ^2_{min} values for fits to 895 toy data sets with the flux and cross-section parametrs thrown from the T2K postfit model (black histogram). The underlying detector parameters are thrown from the prefit model in Psyche and so this can cause some of the throws to populate regions of parameter space that are disfavoured by the narrower constraint on the flux and cross-section parameters. 81.0% of toy fits find a χ^2_{min} value greater, and so less likely, than that found by the fit to the ND280 data (red line). Notably here the RMS width of the χ^2 distribution is narrower, showing the effect of the ND280 constraint on the model.

the model.

The improved postfit p-value of $p = 0.81$ (figure 8.16) is mostly due to the improved cross-section only p-value of $p = 0.30$. Even though the cross-section parameters are more constrained in the postfit model than the prefit model, the shift in the central values of the flux (and detector) parameters allows the cross-section model to more thoroughly cover the parameter space which best describes the data.

Chapter 9

Results of the T2K 2020 oscillation analysis

This chapter will discuss the results of the T2K 2020 oscillation analysis (OA) as presented at Neutrino 2020 [16]. This analysis is still in the final stages of internal collaboration review before submission for publication. Results from both the sequential fit analysis in which the BANFF best-fit parameter values and their covariance are propagated to the P-Theta fit to the SK data, and the simultaneous fit to both using the MaCh3 MCMC fitter are presented. A comparison of these results is presented in subsection 9.3.1. The focus of this thesis is the results of the P-Theta framework for which the BANFF result is an input. The results of the P-Theta fit are interpreted in a Frequentist manner from $\Delta\chi^2$ surfaces and confidence levels for each of the parameters of interest; the results of the MaCh3 fit are interpreted in a Bayesian manner from posterior probability density functions and credible intervals for each of the parameters of interest. The BANFF near detector fit result changes the unoscillated prediction of the beam flux and the (anti)neutrino-oxygen interaction cross section and so changes the prediction of the event rates in each of the five SK samples. A comparison of the pre-BANFF and post-BANFF constrained event rates and uncertainties is discussed below. Once all of the nuisance parameters at SK have been marginalised, comparisons can be drawn between the collected data and the predictions of different PMNS parameters which can help in the interpretation of the results of the fit. The results of the P-Theta fit to the SK data, along with comparisons to the MaCh3 fit to both the ND280 and SK data, are presented in section 9.3. Finally, a procedure to include an additional uncertainty to cover unmodelled physical processes from the fake data studies discussed in section 7.3 is outlined, with a focus on a bias in the ΔE_{rmv} parameters, and its impact on the T2K results.

9.1 Impact of the ND280 fit on the SK sample predictions

In the T2K 2020 OA, events at SK are binned into five samples which separate events by beam mode and charged lepton flavour tag, as discussed in chapter 5.13. For each beam mode there is a sample of events with a single μ-like-Cherenkov ring, 1Rμ, and a sample of events with an e-like-Cherenkov ring, 1Re. In addition to these four samples, a fifth sample uses the decay of charged pions to muons to Michel electrons to attempt to tag charged pions produced in ν_e-like events. A secondary e-like ring occurring within a specific time window after a prompt e-like ring can be used as a tag to classify "one electron-like ring, plus one decay electron" (1Re1d.e.) events.

The impact of the ND280 constraint in terms of both the change of predicted event rate and reduction in the uncertainty on the event rate is presented table 9.1 and shown in figure 9.1. As a result of the under-prediction of the ND280 event rate by the nominal T2K model, the BANFF fit increases the size of M_A^{QE} and the cross section for low-E_ν high-Q^2 CCQE events, and so the overall event rate predictions of the four CCQE-enriched single-ring SK samples are increased. In contrast, the CC1π^+ enriched 1Re1d.e. sample event rate is decreased reflecting the significant decrease in the value of M_A^{RES} and an increase in final-state pion absorption.

The total event rate uncertainty on the muon (anti)neutrino disappearance is reduced from 11.1(11.3)% to 3.0(4.0)% by first fitting the T2K model to the ND280 data without yet fitting to the SK data as compared to the T2K nominal prediction. This is the largest reduction in the uncertainty in the SK event rates reflecting the larger CCQE-rich data set of the ND280 $\nu_\mu(\bar{\nu}_\mu)$ CC0π sample and reflects the efforts to better model the CCQE interaction and reduce the overall systematic uncertainty on the CCQE cross section.

The power of the ND280 fit in the OA is most evident in the reduction of the event rate uncertainty in the region of the 1Rμ deficit (subfigures 9.1(a) and 9.1(d)) as this is the primary part of the SK E_ν^{rec} spectrum from which the oscillation parameter values are extracted. Ensuring that these event-rate bins are well constrained is important for measuring the atmospheric oscillation parameters θ_{23} and Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO).

The increase in the prediction of the single e-like ring samples, particularly in the ν-mode (FHC) beam sample, has important ramifications for the extraction of δ_{CP}, the value of which is dependent on the differences between the ν-mode and $\bar{\nu}$-mode ν_e-like appearance channels. Due to the large wrong-sign contamination of the $\bar{\nu}$-beam, resolving a true asymmetry in the appearance of ν_e as compared to $\bar{\nu}_e$ neutrinos is somewhat obscured if both samples see an increase in events over a prediction. A significant unaccounted asymmetry in the flux or interaction model here would bias the T2K δ_{CP} result.

For the CCQE-like sample FHC(RHC) 1Re, the total event rate uncertainty is reduced from

13.8(12.7)% to 4.7(5.9)% and for the CC1π-like sample, FHC 1Re1d.e., the uncertainty is reduced from 18.7% to 14.3%. This is largely driven by the anticorrelations as indicated by the difference between the uncertainty that would be arrived at by the sum of the variances associated with the separate flux and interaction models, Flux^2 + Int.^{2}, and the variance calculated from the BANFF postfit covariance, (Flux $\bigotimes$ Int.)^2, as indicated in table 9.1.

9.1.1 ΔE_{rmv} parameters

The impact of the ND280 constraint of the two oxygen removal energy parameters ($\Delta E_{\text{rmv}}^{\nu,\text{O}}$ and $\Delta E_{\text{rmv}}^{\bar{\nu},\text{O}}$) on the SK prediction for the FHC 1Rμ sample is shown in figure 9.2. Since extensive work has been done to both reduce the intrinsic uncertainty in the model and improve the implementation of the systematic uncertainty at both ND280 and SK, the contribution to the uncertainty in the T2K nominal model (pre-BANFF) is small. As a result, the reduction in the uncertainty on the event rate due to the ND280 constraint on the ΔE_{rmv} parameters is not as dramatic as the overall reduction of the systematic uncertainty seen in figure 9.1.

The E_{rmv} systematic uncertainty still has a significant contribution to the event rate for the lower bins in E_ν^{rec}. Here, the ΔE_{rmv} systematic uncertainties can affect the shape of the spectrum, shifting the peak in E_ν^{rec} which has implications for the extraction of Δm^2_{32}(NO)/$|\Delta m^2_{31}|$)(IO). The size of the uncertainty in the region of the deficit, however, is very small and so the E_{rmv} systematic uncertainty should only impact the value of $\sin^2\theta_{23}$ through the degeneracy with the mass splitting term.

9.2 Comparison of the data with PMNS predictions

A comparison can be drawn between the data collected at SK and various predictions of the PMNS model after marginalising over all of the nuisance parameters. Table 9.2 indicates the samples which have the best (statistics-only) sensitivity to δ_{CP}, which have a broader spread in the event-rate predictions. The FHC 1Rμ sample has a total spread of 1.48 events but the statistical uncertainty on the number of collected data events is $\pm$17.8 events indicating the muon disappearance channel has little sensitivity to δ_{CP}. This spread is even smaller than that in the RHC 1Rμ sample which is 0.74 events. The FHC 1Re has the largest sensitivity as expected, as well a better separation of the CP-violating predictions (-π/2 and π/2) by 29.66 events for a statistical uncertainty on the data of $\pm$9.7. The RHC 1Re sample has much lower sensitivity with a total spread of 4.19 events due to its lower event rate for which the larger relative statistical uncertainty on the data of $\pm$4.0 events is roughly equivalent to the spread of the predictions. There is a large excess of events in the FHC 1Re1d.e. sample which should have some sensitivity to δ_{CP}. This excess pushes

Sample		Uncertainty source (%)			Flux⊗Int. (%)	Total (%)
		Flux	Int.	SK det.		
1Rμ	ν	2.9 (5.0)	3.1 (11.7)	2.1 (2.7)	2.2 (12.7)	3.0 (13.0)
	$\bar{\nu}$	2.8 (4.7)	3.0 (10.8)	1.9 (2.3)	3.4 (11.8)	4.0 (12.0)
1Re	ν	2.8 (4.8)	3.2 (12.6)	3.1 (3.2)	3.6 (13.5)	4.7 (13.8)
	$\bar{\nu}$	2.9 (4.7)	3.1 (11.1)	3.9 (4.2)	4.3 (12.1)	5.9 (12.7)
1Re1d.e.	ν	2.8 (4.9)	4.2 (12.1)	13.4 (13.4)	5.0 (13.1)	14.3 (18.7)

Table 9.1:

Uncertainty on the number of events in each of the five SK samples after (before) the BANFF constraint, broken down by source. "SK det." includes uncertainties from the SK detector, secondary particle interactions, and photo-nuclear effects. "Flux⊗Int." denotes the combined effect from the ND constrained flux and interaction parameters, and the unconstrained interaction parameters. The flux uncertainties see $\sim 40\%$ reduction on their own, whilst the ND280-constrained cross section (XSec) parameters see $\sim 70\%$ reduction. Once ND280 postfit correlations are taken into account, the total uncertainty is in fact smaller than the simple quadrature sum of these due to strong anticorrelations. There is a small contribution from the non-ND280-constrained parameters such as the 2p2h energy dependence dials or the non-resonant Iso-1/2 background normalisation. A large component of the final SK event-rate uncertainty comes from the SK detector-systematic uncertainties, proton- and pion-secondary interactions, and systematic uncertainties on the photonuclear effect. The thorough treatment of the CCQE model systematics and large CC0π sample data set at ND280 allows the near detector constraint to reduce the total systematic for each of the CCQE-rich single-ring samples by more than half. The largest reduction is in the FHC 1Rμ sample which drives the sensitivity to the atmospheric parameters $\sin^2\theta_{23}$ and Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO). The single e-like ring samples, from which the value of δ_{CP} is extracted, have larger uncertainties due to the lack of strong direct constraint on the intrinsic $\nu_e(\bar{\nu}_e)$ component of the beam at the ND280 which does not include dedicated ν_e samples. The least constrained sample is the CC1π-like FHC 1Re1d.e. sample as the SPP and Multi-π/DIS processes are less constrained at the near detector. The systematic error on this sample also has a much larger contribution from the pion SI.

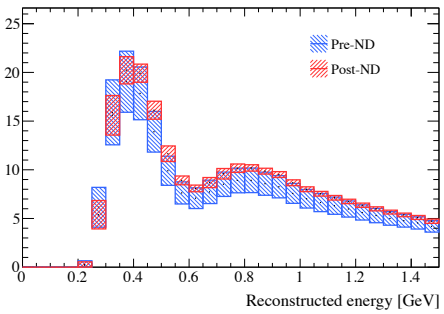

(a) FHC 1Rμ

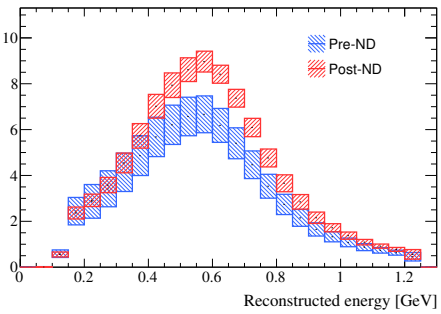

(b) FHC 1Re

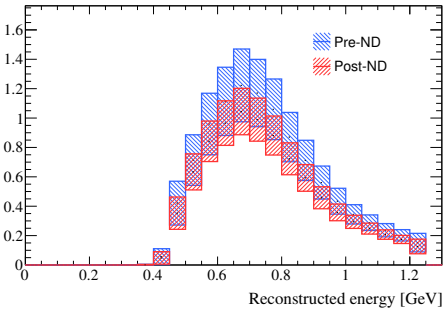

(c) FHC 1Re 1d.e.

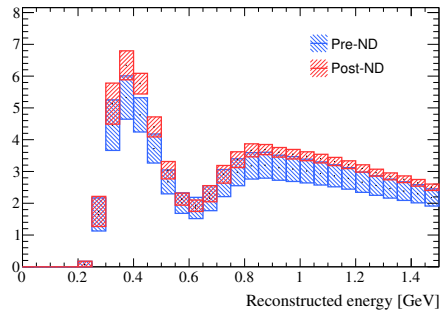

(d) RHC 1Rμ

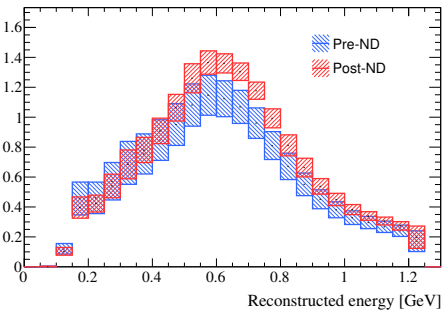

(e) RHC 1Re

Figure 9.1: Comparison of the total the T2K nominal model (pre-ND) and the BANFF postfit (post-ND) constraints (including variations of the SK detector parameters) on the predicted event rates of the five SK samples used in the T2K OA in FHC (ν_μ) beam mode and RHC ($\bar{\nu_\mu}$) beam mode. This effect is from the near detector fit only. The post-ND280 constrained event rates have been changed to match the new parameter values and the uncertainty on the event rates for each bin has been reduced as a result of the reduced parameter value uncertainties and the correlations introduced between the parameters at the near detector. A significant portion of the remaining uncertainty comes from the overall SK detector systematic uncertainty. The PMNS oscillation parameters are at their Asimov A values.

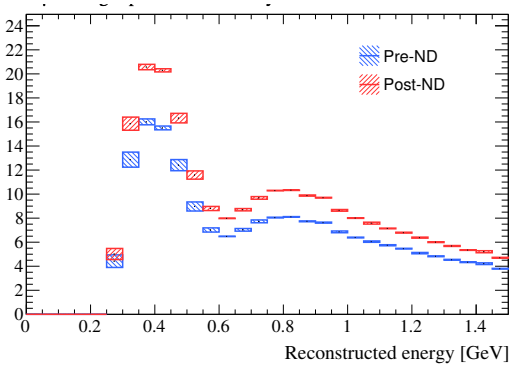

Figure 9.2: SK FHC 1Rμ event rate uncertainty from variations of the ΔE_{rmv} parameters only, with all other systematics fixed for the T2K nominal model (pre-ND) and the BANFF postfit (post-ND) constraints and PMNS parameters at their Asimov values. The ΔE_{rmv} variation contribution to the total systematic uncertainty is largest for the lower E_ν^{rec} bins below the dip, though it is still very small in the region of the dip from which the value of the PMNS parameters is extracted.

the T2K best-fit outside the PMNS-like predictions, given that the largest predicted event rate is for $\delta_{CP} = -\pi/2$ which still under predicts the data by more than the statistical uncertainty on the data collected. This excess has decreased since the previous analysis [20], thereby decreasing the sensitivity of T2K to δ_{CP}, bringing it more in line with the expected sensitivity.

The combined FHC 1Re+1Re1d.e. event rate and the RHC 1Re event rate are sensitive to $\sin\delta_{CP}$, the sign of the mass splitting term and the neutrino mass ordering, and the octant of θ_{23}, but the E_ν^{rec} distribution of the combined event rates of FHC 1Re and 1Re1d.e. also has some sensitivity to $\cos\delta_{CP}$. Figure 9.3 shows that the FHC data shows better consistency with negative values of δ_{CP} but the RHC sample is consistent with any value of δ_{CP}. This reflects the lower statistical power of the RHC sample which largely is due to the overall lower $\bar{\nu}$ nuclear interaction cross sections when compared to ν, as well as less separation of the δ_{CP} predictions due to the large contamination of ν_e events. A more detailed picture of the PMNS predictions can be seen in figure 9.4 which shows the predicted number of e-like ring candidates for both neutrino and antineutrino beam modes. Even though the predictions are highly degenerate for the two mass orderings, T2K's chosen energy and baseline provide good separation of maximally CP-violating δ_{CP} values ($\pm\pi/2$) due to their co-alignment between the mass orderings. T2K's mass ordering resolution, however, is poor and the predictions for each mass ordering largely overlap.

Similar conclusions can be drawn when looking at predictions for the combined FHC and RHC ν_e-like candidates against the combined FHC and RHC ν_μ-like candidates with a reconstructed energy below 1200 MeV in figure 9.5. The data favours $\delta_{CP} = -\pi/2$ and normal ordering, though for any non-maximally CP-violating values of δ_{CP}, the predictions of the different mass orderings

are highly degenerate. In fact, any combination other than $\delta_{CP} = -\pi/2$ and NO or $\delta_{CP} = +\pi/2$ and IO are highly degenerate with the other mass ordering, and between minimally and maximally violating values of δ_{CP}. There is also a small trend in the mass-ordering resolution for true values of δ_{CP} and the maximally positive CP-violating case, $\delta_{CP} = +\pi/2$, has weaker resolution of the mass ordering ellipses than if $\delta_{CP} = -\pi/2$. These combined factors give T2K higher sensitivity in the region in which T2K's data lies. The arrangement of these predictions are largely dependent on the L/E_ν characteristics of the experiment. In these predictions the data are most consistent with values of $\sin^2\theta_{23}$ between 0.5 and 0.55, and least consistent with very large or small values of $\sin^2\theta_{23}$ which would both correspond to less $\nu_\mu(\bar{\nu})$ disappearance.

In each of these figures, the T2K data is slightly outside the PMNS predictions which, if significant, would imply non-unitarity or non-standard mixing, though the T2K result is still consistent with a unitary PMNS parameterisation within the systematic and statistical uncertainties. This is largely driven by an excess of 1Re1d.e. events when compared to the PMNS predictions, which has in fact reduced upon the collection of more data since the previous T2K result [20] so is likely from statistical fluctuation.

Sample		True δ_{CP} (rad.)				Data
		$-\pi/2$	0	$\pi/2$	π	
1Rμ	FHC	346.61	345.90	346.57	347.38	318
	RHC	135.80	135.45	135.81	136.19	137
1Re	FHC	96.55	81.59	66.89	81.85	94
	RHC	16.56	18.81	20.75	18.49	16
1Re1d.e.	FHC	9.30	8.10	6.59	7.79	14

Table 9.2: Event rate prediction using oscillation parameters and systematic parameters at best-fit while varying δ_{CP}. Most of the sensitivity to δ_{CP} comes from the FHC 1Re sample, whereas sensitivity from the RHC 1Re sample is relatively weak due to the smaller number of events and larger statistical uncertainty.

9.3 T2K results

Whilst T2K can make independent measurements of the oscillation parameters and has the power to make independent constraints on both θ_{13} and δ_{CP}, the constraint on θ_{13} is about an order of magnitude weaker than that made by reactor experiments. The T2K constraint is consistent with the reactor experiments to within the 68% credible interval as indicated by the MaCh3 results

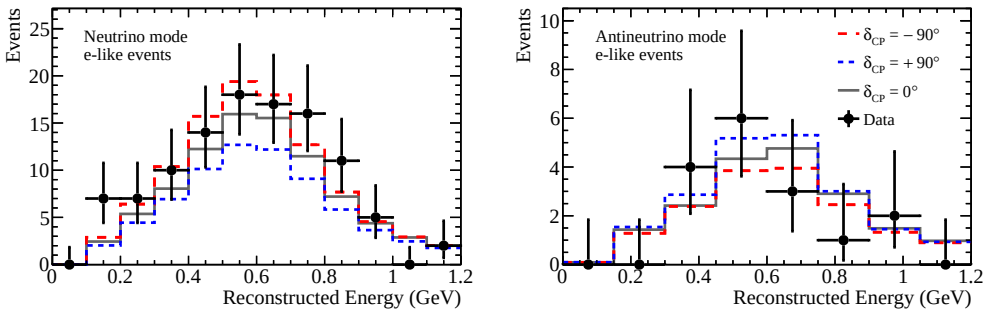

Figure 9.3: Predictions of the number of e-like candidates for ν-mode and $\bar{\nu}$-mode beam samples for different values of δ_{CP} with all other parameters fixed. Normal mass ordering is assumed, and $\sin^2\theta_{23}$ and Δm^2_{32} are at the best fit (SK) values. $\sin^2\theta_{13}$, $\sin^2\theta_{12}$ and Δm^2_{12} are at the PDG2019 values [34]. The flux and cross-section parameters are at the post-BANFF best fit values.

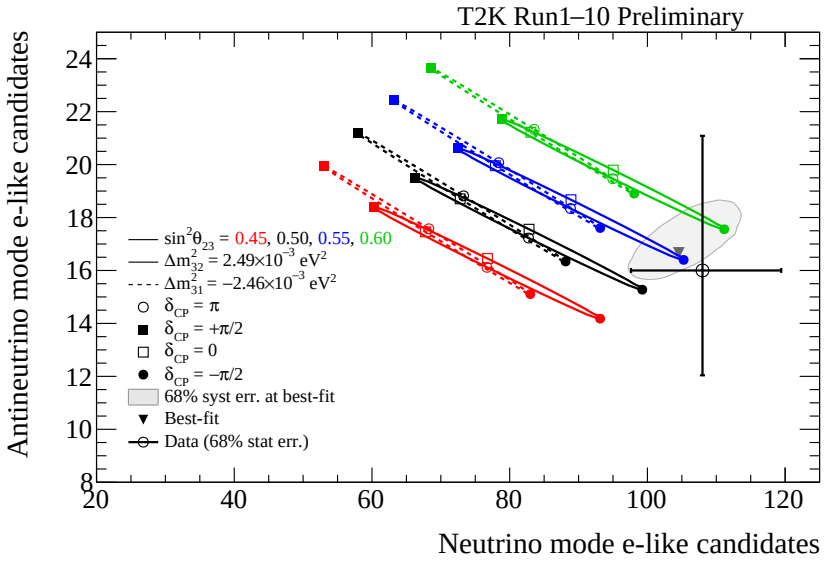

Figure 9.4: Biprobability plot showing a comparison of the T2K data collected for the number of $\bar{\nu}$-mode versus ν-mode-beam e-like ring candidates with the predictions given by different values of the PMNS oscillation parameters. T2K's leading resolution of δ_{CP} comes from the arrangement of these ellipses for the chosen value of L/E_ν. At T2K's L/E_ν, the maximally violating values of δ_{CP} are aligned in each hierarchy, allowing for δ_{CP}-resolution without the need for strong hierarchy resolution in the maximally CP-violating cases. An excess of ν_e candidates shows clear preference for negative values of δ_{CP} as well as some preference for the normal mass ordering. The best-fit point given uses the BANFF best-fit point for the flux and cross-section parameters and the p-theta best fit-point for the SK detector systematics and PMNS parameters. A 68% contour is drawn from predictions taken from throws of the systematic uncertainties.

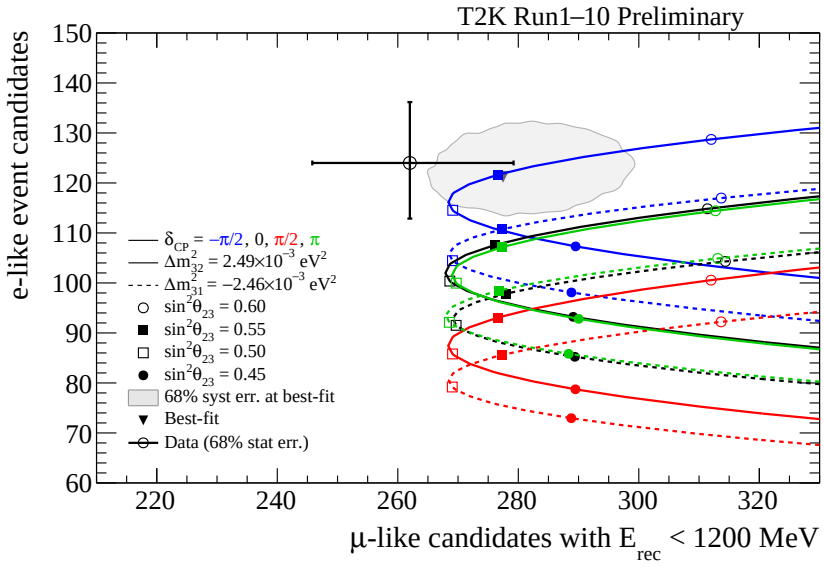

Figure 9.5: Biprobability plot showing a comparison of the T2K data collected for the number of e-like ring candidates versus μ-like ring candidates for muon-neutrino energies below 1200 MeV, with the predictions given by different values of the PMNS oscillation parameters. The data show preference for $\delta_{CP} = -\pi/2$ and a normal mass ordering with $\Delta m^2_{32} = 2.49 \times 10^{-3}\text{eV}^2$. The interdependence of these preferences can be seen by the overlapping ellipses for other possible values of δ_{CP}, and even the degeneracies between the inverted mass ordering for the same value of δ_{CP}, and the CP-conserving case for normal ordering. The best-fit point uses the BANFF best-fit point for the flux and cross-section parameters and the p-theta best-fit point for the SK detector systematics and PMNS parameters. A 68% contour is drawn from predictions taken from throws of the systematic uncertainties.

Parameter	Best fit			
Data	T2K Only		T2K + RC	
Mass Ordering	Normal	Inverted	Normal	Inverted
$\sin^2(2\theta_{13})$	0.109	0.120	0.0855	0.0860
$\sin^2\theta_{13}$	28.0×10^{-3}	31.0×10^{-3}	21.9×10^{-3}	22.0×10^{-3}
δ_{CP} (rads)	-2.22	-1.29	-1.97	-1.44
$\Delta m^2_{32}(\mathrm{NO})/\lvert\Delta m^2_{31}\rvert(\mathrm{IO})$ (eV^2)	2.495×10^{-3}	2.463×10^{-3}	2.494×10^{-3}	2.463×10^{-3}
$\sin^2\theta_{23}$	0.467	0.466	0.561	0.563
χ^2	597.72	598.56	598.05	600.49

Table 9.3: T2K's best-fit values of the oscillation parameters for each mass ordering, both with and without a reactor constraint (RC). The χ^2 of the best-fit point shows a slight preference for normal ordering of the neutrino masses both with and without the application of the reactor constraint. Values of δ_{CP} are consistently near to $-\pi/2$ (~ -1.57) but are not quite maximally CP-violating.

shown in figure 9.6. By using the global reactor constraint (RC) from the Particle Data Group (PDG) [34] as a prior uncertainty for the T2K OA, correlations between θ_{13} and δ_{CP} can be exploited to give a tighter constraint on δ_{CP}. Table 9.3 lists the best-fit values from the frequentist sequential (BANFF+P-Theta) fit to near and far detector data for the T2K-only and T2K+RC fits. The impact of this constraint on the 2D confidence levels is shown for the P-Theta fit result in figure 9.7. The direct impact of this on δ_{CP} when marginalising over all other parameters is shown in figure 9.8. The global (both mass orderings) $\Delta\chi^2$ values show a slight preference for the normal ordering with no regions of δ_{CP} within the 1σ confidence level for the inverted mass ordering when the reactor constraint is used.

This constraint also has an impact on the atmospheric parameters through the eight-fold degeneracy of the oscillation parameters. The impact of the reactor constraint on the mass squared splitting is largely the exclusion of more of the $|\Delta m^2_{31}|$ values and inverted ordering parameter space at the 2σ level, as shown in figure 9.9. Otherwise, the best-fit value in both orderings is unchanged by the reactor constraint. In contrast, the exclusion of $\sin^2\theta_{23}$ values is asymmetric, and removes the lower degenerate minima, moving the best fit in both orderings to the what was the higher and less favoured local minimum in the T2K only results as shown in figure 9.10. The best-fit point of $\sin^2\theta_{23}$ in the inverted mass ordering is not as disfavoured as the best-fit point of $|\Delta m^2 31|$ due to the marginalisation over the asymmetric structure of the 2-dimensional $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ and $\sin^2\theta_{23}$ confidence levels shown in figure 9.11.

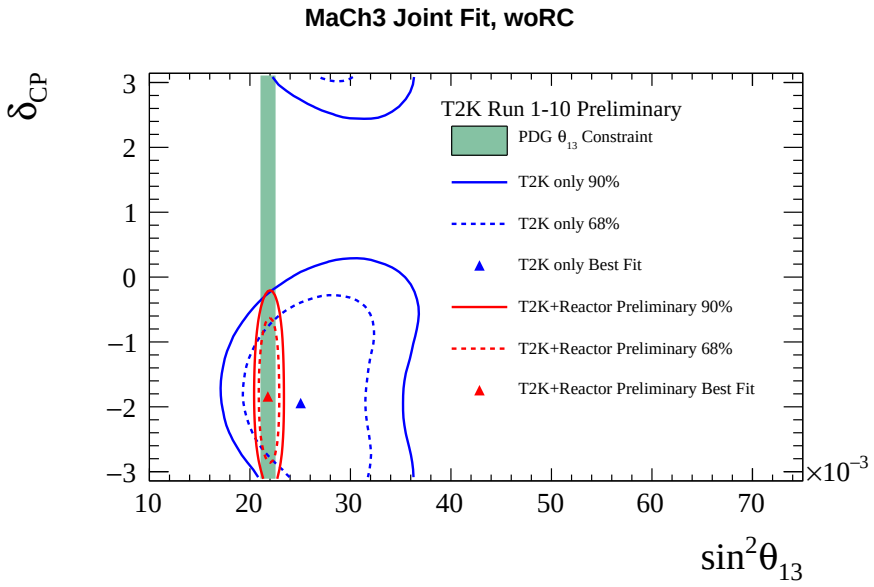

Figure 9.6: Two-dimensional 68% (dashed) and 90% (solid) credible-interval contours for δ_{CP} versus $\sin\theta_{13}$ constructed using chains from the joint ND280 and SK MCMC fit, MaCh3, marginalised over both neutrino mass orderings, shown both with (red) and without (blue) a constraint on θ_{13} from reactor experiments which is indicated by the green band.

9.3.1 Comparison of sequential and joint fit results for δ_{CP}

Despite the different the fitting methods outlined in chapter 5, and the different interpretations of the confidence and credible levels in $\Delta\chi^2$ surfaces and posterior density functions respectively, analogous statements can be made about favoured and excluded values of a parameter in each of the T2K analyses. This can be seen when examining favoured and disfavoured values of δ_{CP} in each framework. CP-conserving values can be excluded to a 90% confidence level in the frequentist analysis (P-Theta) as shown in figure 9.12, and lie outside the 90% credible intervals in the Bayesian analysis (MaCh3) as shown in figure 9.13. A large range of values around maximal positive CP-violation, $\delta_{\mathrm{CP}} = +\pi/2$, are excluded at over 3σ. The extent of this disfavouring can be seen in the MaCh3 posterior probability distribution in which large regions of the inverted mass ordering, and some regions of the normal mass ordering around $\delta_{\mathrm{CP}} = +\pi/2$, are unpopulated by the steps in the Markov Chain despite a chain length of 12 million steps. This is not to say that these values are impossible, but that they are so disfavoured that the chain does not sample the region. The preference for normal mass ordering also leads to the exclusion of CP-conserving values in the inverted mass ordering to 3σ; whereas, in the normal mass ordering $\delta_{\mathrm{CP}} = 0$ cannot be excluded at 3σ and $\delta_{\mathrm{CP}} = \pm\pi$ cannot be excluded at 2σ. The values of sequential fit δ_{CP} Feldman Cousins confidence level intervals are listed in table 9.4.

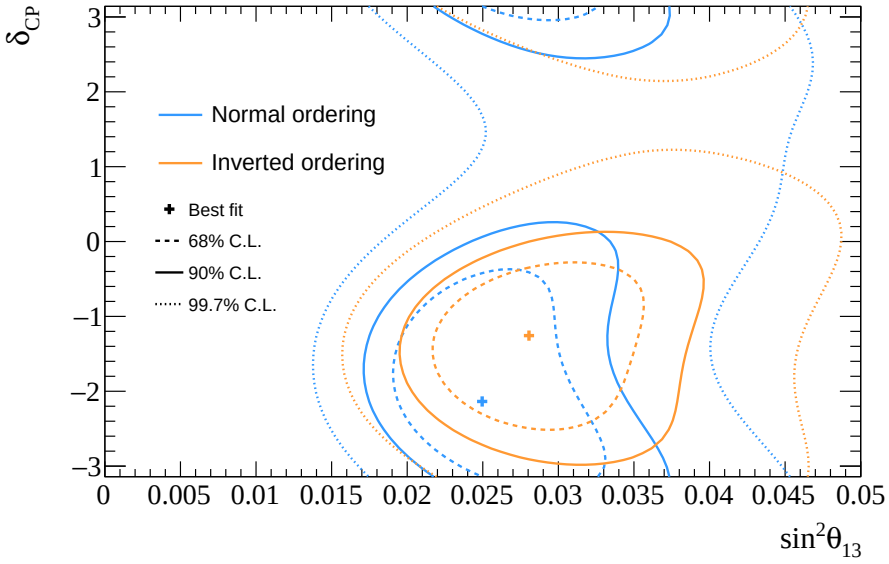

(a) T2K only

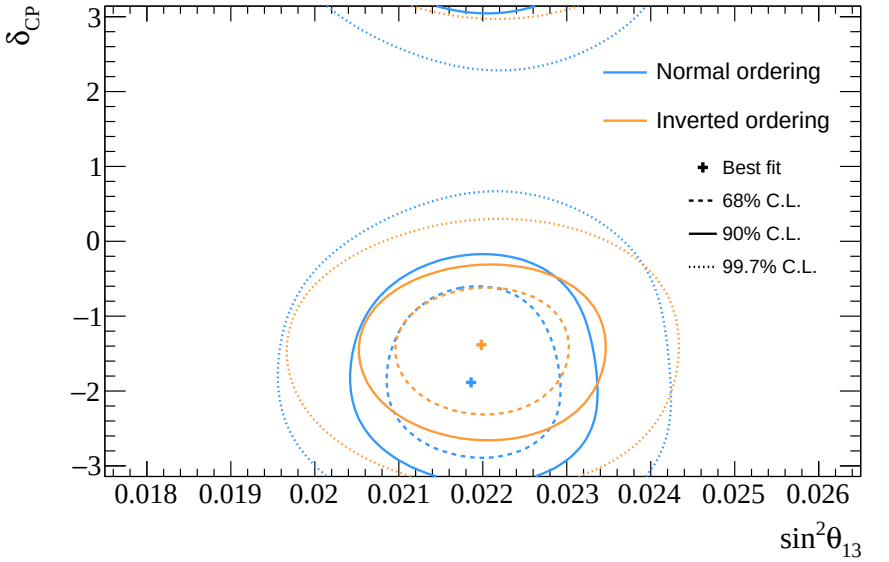

(b) T2K + R.C.

Figure 9.7: Two-dimensional 68% (dashed), 90% (solid) and 99.7% (dotted) confidence-level contours for δ_{CP} versus $\sin\theta_{13}$, computed using the best-fit point in each mass ordering and not the global best fit point, both without (9.7(a)) and with (9.7(b)) a constraint (R.C.) on θ_{13} from reactor experiments.

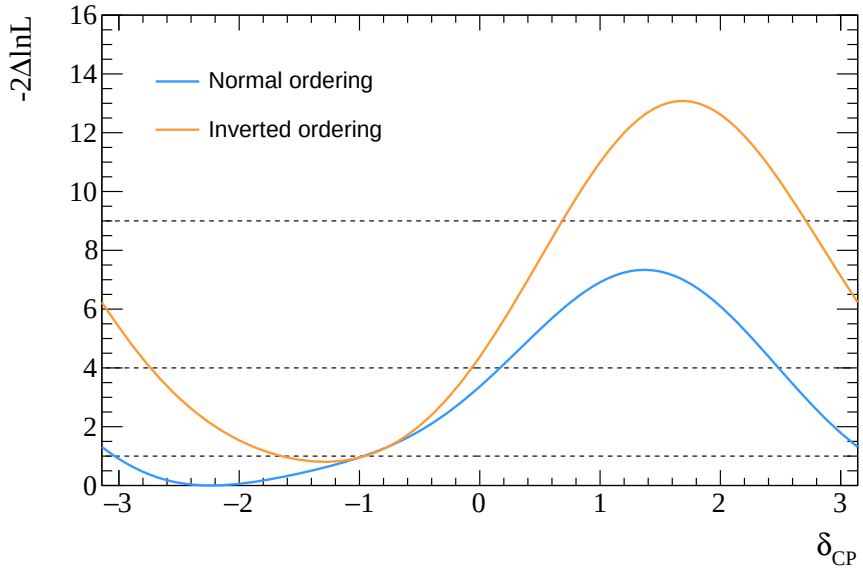

(a) T2K only

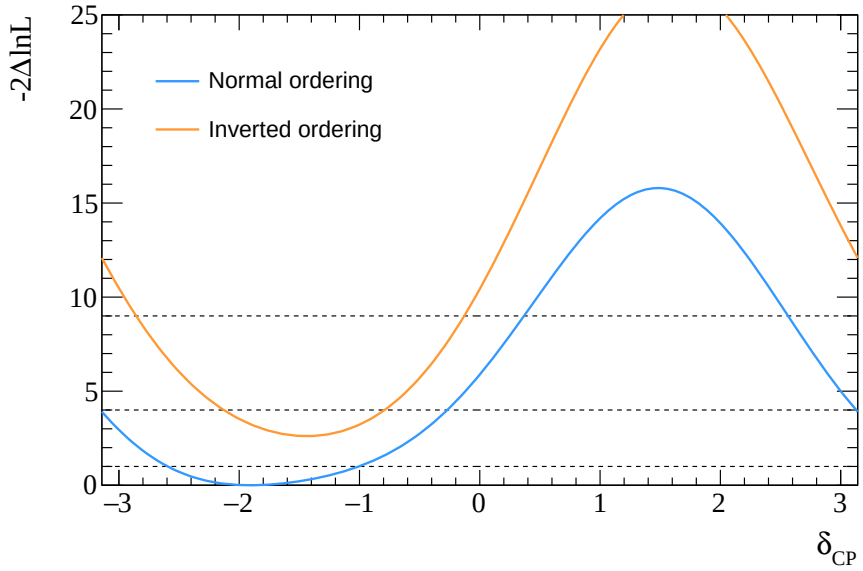

(b) T2K + R.C.

Figure 9.8: One-dimensional $\Delta\chi^2$ ($-2\Delta\ln L$) surface for δ_{CP} marginalised over all other parameters without (9.8(a)) and with (9.8(b)) a reactor constraint (R.C.) applied. The horizontal dashed lines correspond to the 1σ, 2σ, and 3σ confidence levels with higher of values of $\Delta\chi^2$ more disfavoured. A large region around $\pi/2$ has been excluded to over 3σ with the R.C. in both mass orderings. This region is not excluded in the normal ordering without the R.C.. The inverted ordering is more disfavoured in when the reactor constraint is applied for all values of δ_{CP}.

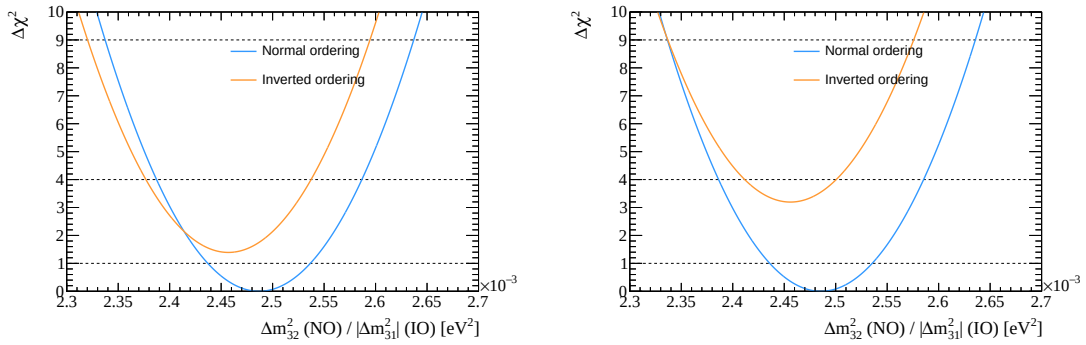

(a) T2K Only (b) T2K+RC

Figure 9.9: Data fit $\Delta\chi^2$ surface for the mass splitting term $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ for the T2K only and reactor-constrained T2K results. As well as a slight shift in the best-fit value of the $\Delta m^2_{32}(\mathrm{NO})$ mass splitting term, values of $|\Delta m^2_{31}|(\mathrm{IO})$ become disfavoured with the application of the reactor constraint on $\sin^2\theta_{23}$.

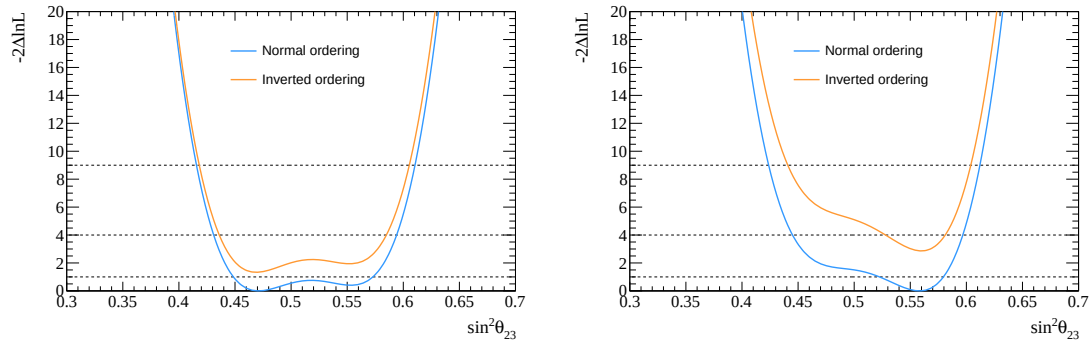

(a) T2K Only (b) T2K+RC

Figure 9.10: Data fit $\Delta\chi^2$ surface for the value of $\sin^2\theta_{23}$ for the T2K only and reactor-constrained T2K results. Applying the reactor constraint moves the best-fit value in both mass orderings from the lower to the higher minimum.

Confidence level	Interval (NO)	Interval (IO)
1σ	[-2.66,-0.97]	
90%	[-3.00,-0.49]	[-1.79,-1.09]
2σ	$[-\pi,-0.26]\cup[3.11,\pi]$	[2.20,-0.75]
3σ	$[-\pi,0.32]\cup[2.63,\pi]$	[-2.82,-0.14]

Table 9.4: Feldman-Cousins confidence intervals for δ_{CP} from the BANFF and P-Theta sequential fit framework for the T2K+RC fit. The inverted mass ordering is excluded at 1σ entirely. CP-conserving values (0 and π) are excluded at 90% in both orderings, and 37% of values around $+\pi/2$ are excluded at over 3σ.

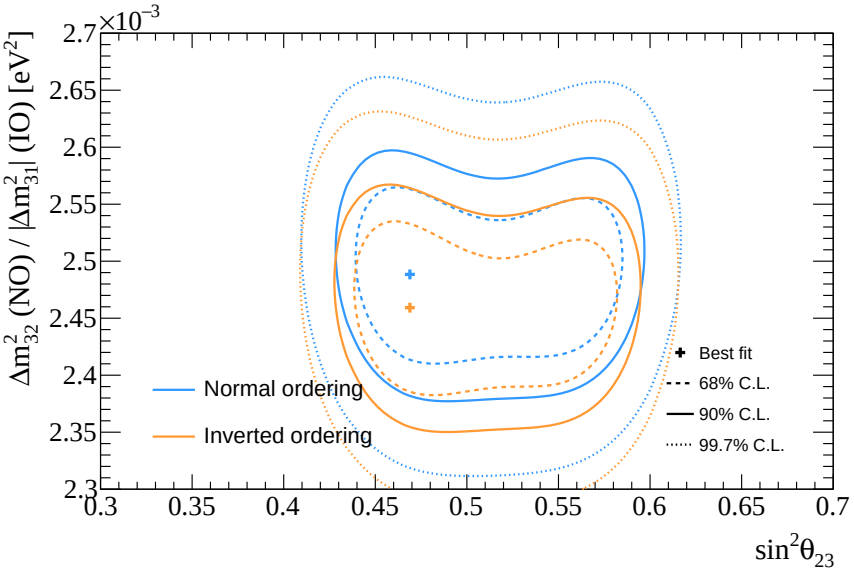

(a) T2K Only

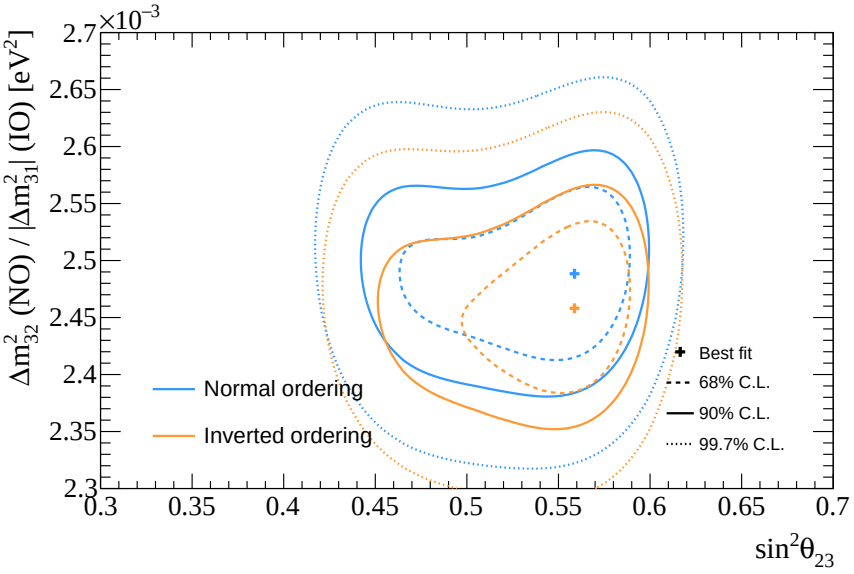

(b) T2K+RC

Figure 9.11: Contours of the 2-dimensional confidence level intervals for $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ and $\sin^2\theta_{23}$ for both the T2K-only result and the reactor-constrained T2K result. The $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ best-fit value is relatively unchanged by the application of the RC, but within the two-lobed contour structure the value of $\sin^2\theta_{23}$ has flipped to the higher local minima as the lower values are slightly more penalised, particularly in the inverted mass ordering.

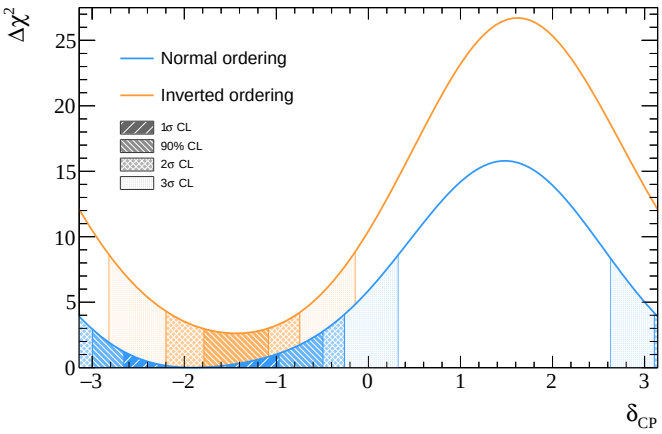

Figure 9.12: One-dimensional δ_{CP} Feldman-Cousins corrected $\Delta\chi^2$ ($-2\log\mathscr{L}$) surface and confidence intervals for the sequential fit (BANFF + P-Theta). Both hierarchies are shown, and all other parameters are marginalised over. CP-conservation can not be fully excluded at 2σ as $\pm\pi$ is still just within the confidence level in the normal ordering, though these values are both excluded at the 90% confidence level; however, a large range of values around $\delta_{CP} = +\pi/2$ can be excluded to 3σ in both orderings.

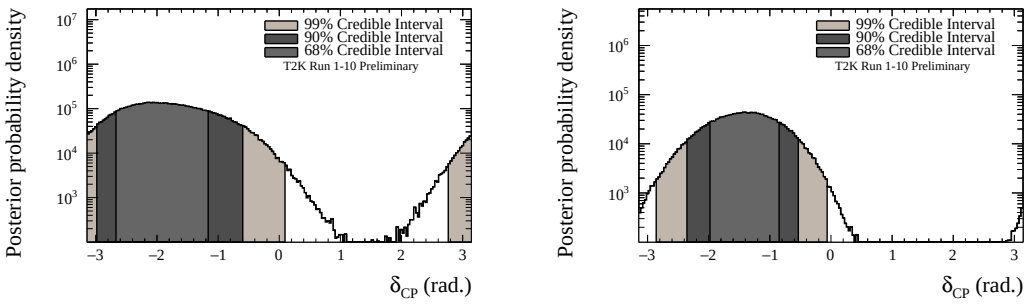

(a) Normal Ordering (b) Inverted Ordering

Figure 9.13: One-dimensional δ_{CP} posterior density for each mass ordering. Density is represented on a logarithmic scale. Whilst the region around $\delta_{CP} = +\pi/2$ is sparsely populated in both orderings, this is due to computational costs. The probability is still non-zero, though highly disfavoured and so very few steps in the Markov Chain are in this region as the Markov Chain is optimised to sample the most favoured values more frequently so the credible intervals can be determined.

9.3.2 Comparison of BANFF and MaCh3 ND280 fit constraints on the atmospheric parameters in the p-theta framework

To understand the magnitude of the differences between the BANFF fit and the ND280-Only MaCh3 fit result, the posterior probability distribution function obtained in a MaCh3 MCMC fit to the ND280 data for each parameter can be used as an input to the P-Theta framework. Figure 9.14 shows a comparison of the MaCh3 posterior probability distribution for a MCMC fit using standard MaCh3 binning (irregular), the posterior probability distribution for n MCMC fit using the BANFF binning (regular), and throws taken from the BANFF postfit covariance. Here, "regular" and "irregular" refer to the use of same (regular) or different (irregular) p_μ binnings for different $\cos\theta_\mu$ bins in a sample. These are implemented using the standard ROOT 2D histogram class (TH2D), and a ROOT histogram class which allows for the use of arbitrary 2D bin shapes or polygons (TH2Poly). The MaCh3 posteriors shown in figure 9.14 can be used as irregular prior uncertainties on the oxygen ΔE_{rmv} parameters in P-Theta which are marginalised over before a grid search of the oscillation parameters.

The impact of these differences on the oscillation parameter extraction can be seen in figure 9.15. Whilst differences between the fit results arise from differences in the fitting methods, gradient descent v.s. MCMC, or from differences in parameter implementation, such as with E_{rmv}, the differences between the P-Theta result using the MaCh3 regular and irregular binnings are larger than those between BANFF and MaCh3 when using the same regular binning. In fact, the MaCh3 regular binning result is generally more different to the BANFF result than the MaCh3 result when using the same binning as BANFF. This may be due to larger binning effects and marginalisation issues in MaCh3 when the bin boundaries are more aligned. The difference between the BANFF and MaCh3 irregular TH2Poly binning results are generally as close or closer than the unsmeared and smeared T2K+RC result in $\Delta m^2_{32}(\text{NO})/|\Delta m^2_{31}|(\text{IO})$ discussed in section 9.4.2 and so these differences are covered by the total systematic uncertainty quoted in the reported results.

9.4 Impact of fake data fits on the oscillation parameters

As discussed for the BANFF ND280 fit in section 7.3, fake data studies (FDS) are used to test the robustness of the analysis and assign any extra uncertainty is believed to be needed to cover unmodified effects or biases. This is propagated to the oscillation analysis by preparing data in a similar fashion to the ND280 fake data, and propagating the BANFF best-fit parameters and covariance as would be done in the standard fit. This procedure is outlined here for the ΔE_{rmv}

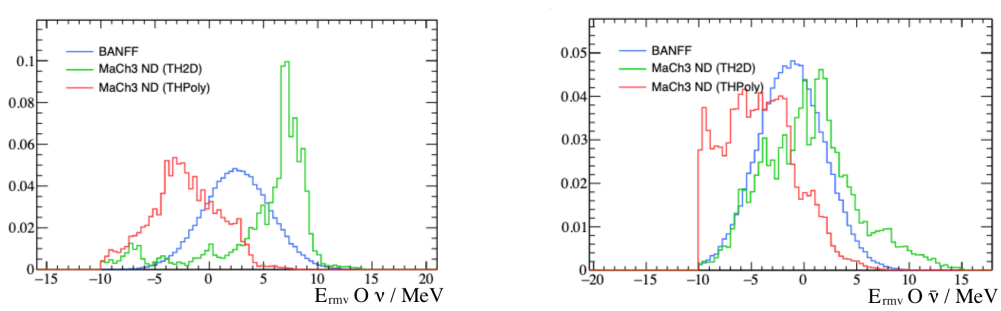

(a) ν on oxygen (b) $\bar{\nu}$ on oxygen

Figure 9.14: BANFF and MaCh3 ND280 fit posterior distributions, treated as prior distributions at SK, of the two oxygen nucleon removal energy parameters. The MaCh3 distribution is taken directly from the values sampled by the Markov chain, and the BANFF distribution is produced by Gaussianly throwing values from the postfit covariance. In addition to the differences between implementation of the parameter as a momentum varier or a binned reweight, the binning choice used by each near detector fitter has an impact on the ND280 postfit parameter values. The ΔE_{rmv} parameters in MaCh3 are very sensitive to this as they suffer from significant binning and marginalisation effects as events migrate across bin boundaries, which is evident from the multipeak structure in the posterior distribution. The difference between the use of a regular rectangular binning (TH2D) and the use of an irregular rectangular binning (THPoly) for the momentum varying ΔE_{rmv} in MaCh3 has a more significant difference than between the BANFF and MaCh3 implementation for regular binning (TH2D). A sharp cut off at -10 MeV can be seen where a boundary is enforced to prevent the SF peak structures from entering the unbound region with negative removal energy.

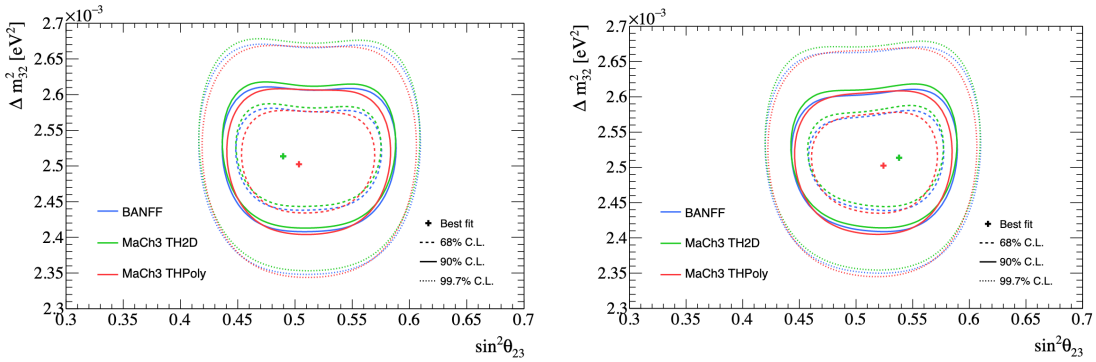

(a) T2K-Only N.O.

(b) T2K+RC N.O.

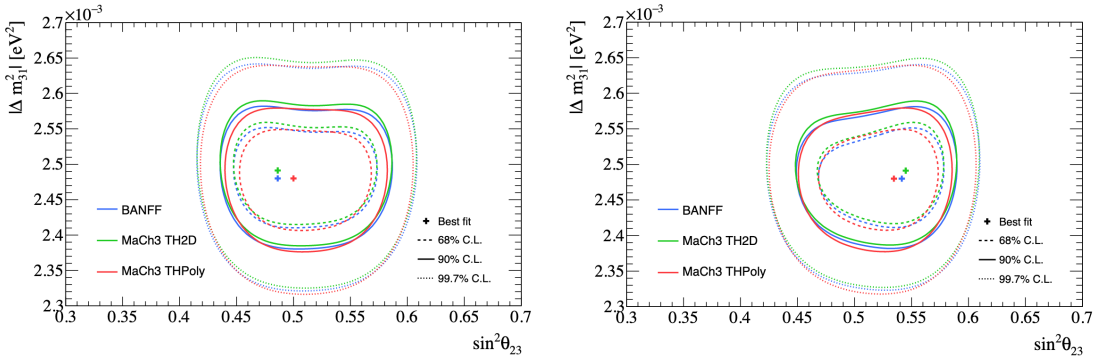

(c) T2K-Only I.O.

(d) T2K+RC I.O.

Figure 9.15: Comparison of the $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ and $\sin^2\theta_{23}$ contours in the frequentist P-Theta framework using the BANFF and MaCh3 ND280 fit results as inputs. Similarly to the ND280 fit posteriors for the removal energy parameters, the choice of binning for the MaCh3 fit has a greater impact on the spread of the contours and position of the best-fit point than the difference between the BANFF and MaCh3 ND280 constraints for the same regular TH2D binning. In the NO plots the BANFF and MaCh3 TH2D (regular) binning best fit results overlap.

biased fake data set. As fake data produced with this adjusted model was already fit to at ND280, and the postfit value of the parameter returned as expected, the ND280 fit should not introduce a large bias for this fake data set. As the nominal BANFF Asimov prediction of the SK spectra shown in figure 9.16 are still compatible with the 5 MeV biased ΔE_{rmv} fake data set within the uncertainty of the BANFF fit to the fake data.

9.4.1 Bias on the oscillation parameters in ΔE_{rmv} biased fake data fits

The result of fits from the P-Theta group to SK fake data prepared in the same way as the near detector +5 MeV biased ΔE_{rmv}, with the PMNS parameters at their Asimov A values assuming NO, are presented in table 9.5. The priors for the P-Theta fit to the SK fake data is taken from the postfit covariance of the BANFF fit to the ND280 5 MeV biased ΔE_{rmv} fake data set. The impact of this bias on the two-dimensional CL contours of the atmospheric and accelerator PMNS parameters are shown in figure 9.17 and one-dimensional log-likelihood-ratios of each of the parameters are shown in figures 9.19.

		θ_{23}	Δm^2_{32}	δ_{CP}
Asimov A	Middle of the 1σ interval	0.522	2.514e-03	-1.59
	1σ interval size	0.0514	5.733e-05	1.16
	1σ interval ratio to Asimov	1	1	1.01
	Change in the 1σ interval ratio to syst interval	0.999 %	-0.104 %	4.24 %
	Bias computed with 1σ middle and interval (lower octant θ_{23})	0.017 %	4.84 %	1.3 %
	Bias relative to σ_{Syst} computed with 1σ middle and interval	0.0419 %	13.4 %	5.2 %
	Middle of the 2σ interval	0.517	2.515e-03	-1.59
	2σ interval size	0.0791	1.156e-04	1.88
	2σ interval ratio to Asimov	1	0.998	1
	Change in the 2σ interval ratio to syst interval	0.65 %	-0.646 %	1.72 %
	Bias computed with 2σ middle and interval (lower octant θ_{23})	0.151 %	2.5 %	0.301 %
	Bias relative to σ_{Syst} computed with 2σ middle and interval	0.449 %	6.81 %	1.09 %

Table 9.5: Oscillation parameters extracted for a +5 MeV biased ΔE_{rmv} fake data set.

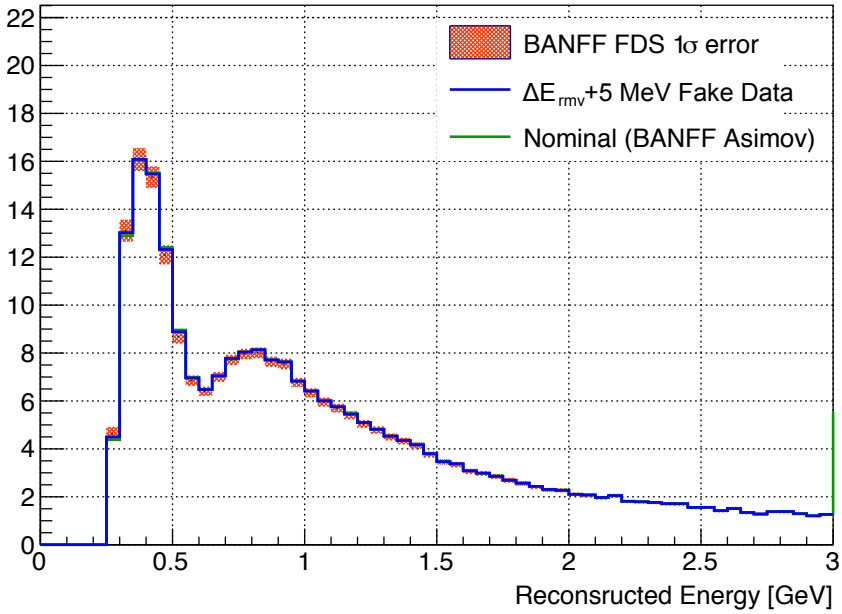

(a) ν_μ sample

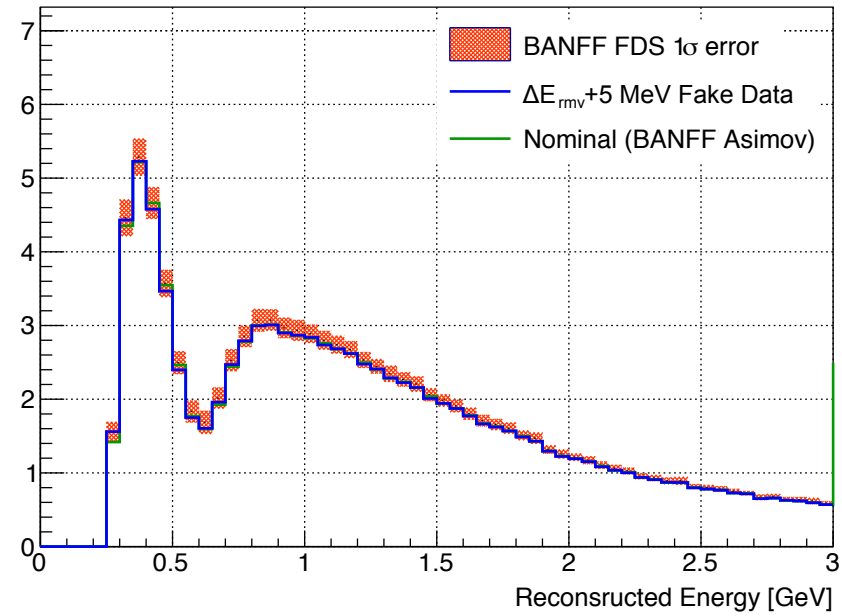

(b) $\bar{\nu}_\mu$ sample

Figure 9.16: ΔE_{rmv} bias fake data fit predictions of the SK event rates (continued on next page)

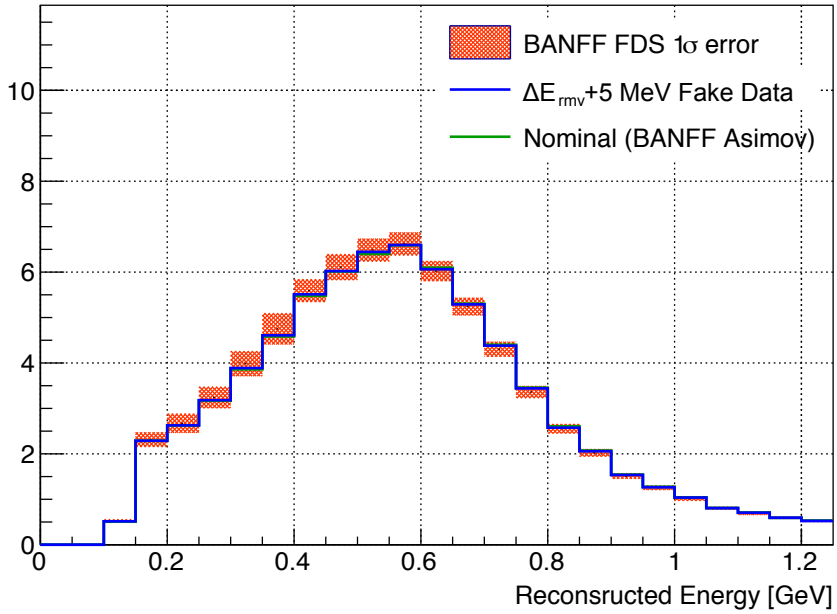

(c) ν_e sample

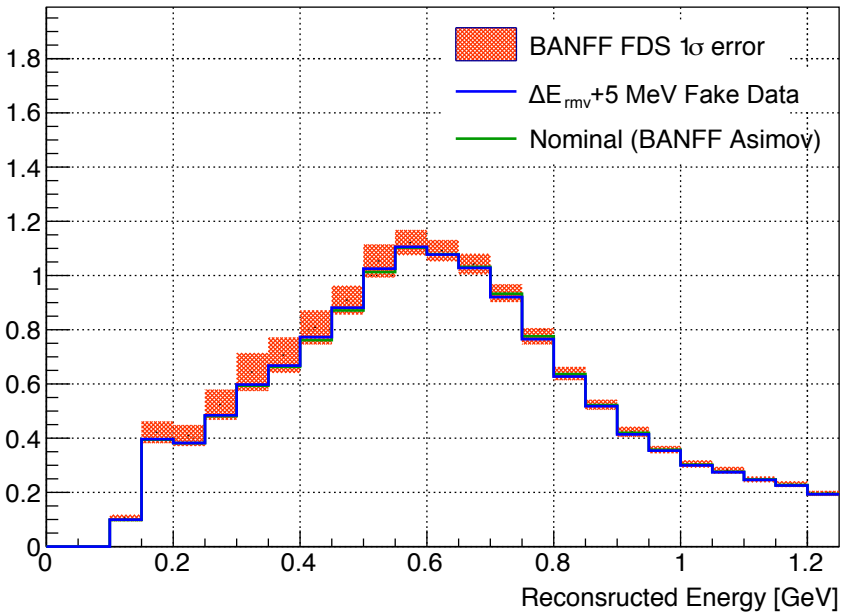

(d) $\bar{\nu}_e$ sample

Figure 9.16: ΔE_{rmv} bias fake data fit predictions of the SK event rates (continued on next page)

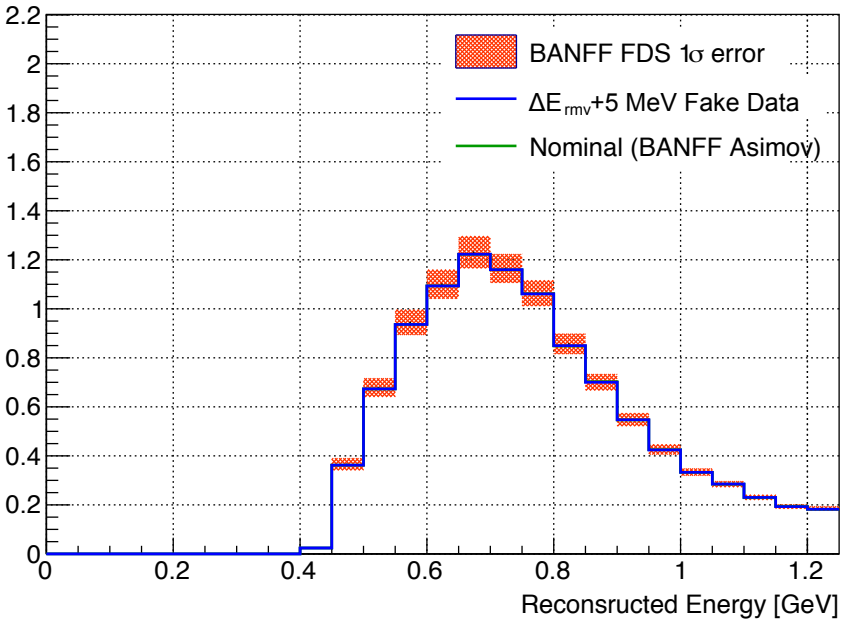

(e) ν_e 1 d.e. sample

Figure 9.16: Comparison between the nominal SK Asimov A event rate (green), the SK 5 MeV biased ΔE_{rmv} fake data event rate (blue), and the prediction and uncertainties from the BANFF FDS fit to the +5 MeV biased ΔE_{rmv} fake data (red band) made from throws to the BANFF FDS postfit covariance. Both the nominal T2K model and the fake data event rates lie within the BANFF postfit uncertainties, though the BANFF slightly overpredicts the rate of high E_ν^{rec} $\bar{\nu}_\mu$ and low E_ν^{rec} $\bar{\nu}_e$ events. The consistency of the true fake data event rate with the BANFF fake data fit prediction indicates that there is little bias from the ND280 fit when a bias of this size is introduced in the values of the ΔE_{rmv} parameters.

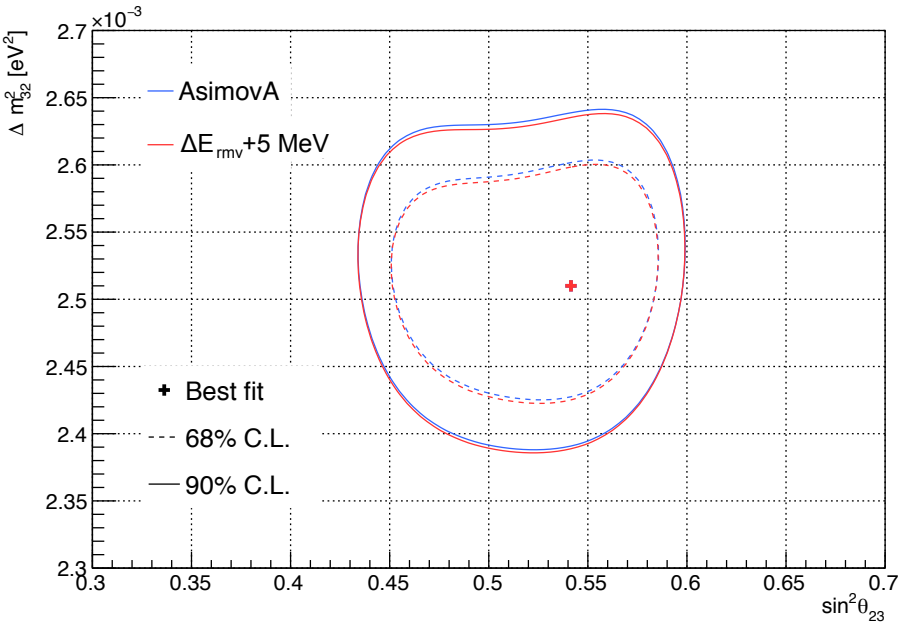

(a) $\sin^2\theta_{23}$–Δm^2_{23}

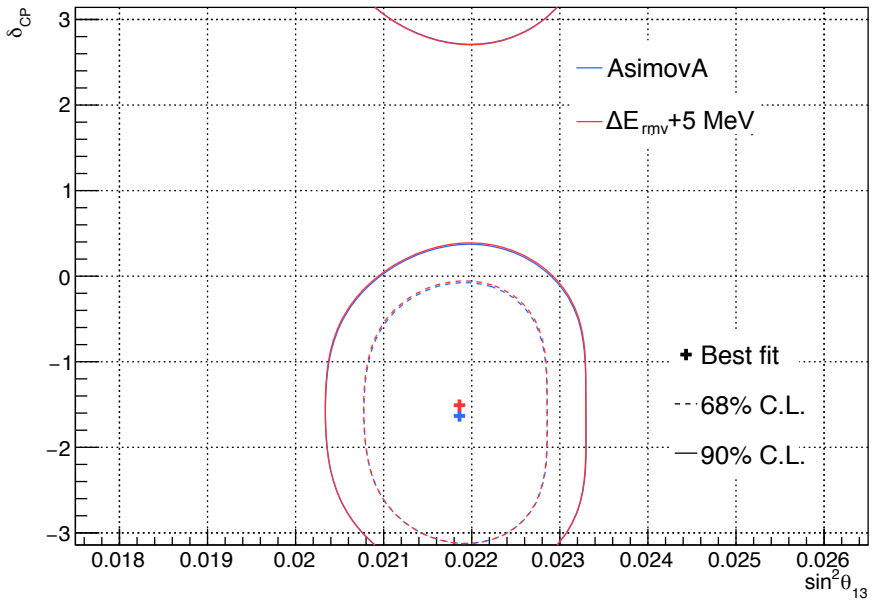

(b) $\sin^2\theta_{13}$–δ_{CP}

Figure 9.17: 2-D 68% and 90% contours for P-Theta fits including the reactor constraint on $\sin^2\theta_{13}$ and assuming normal ordering, for the +5 MeV biased ΔE_{rmv} SK fake data.

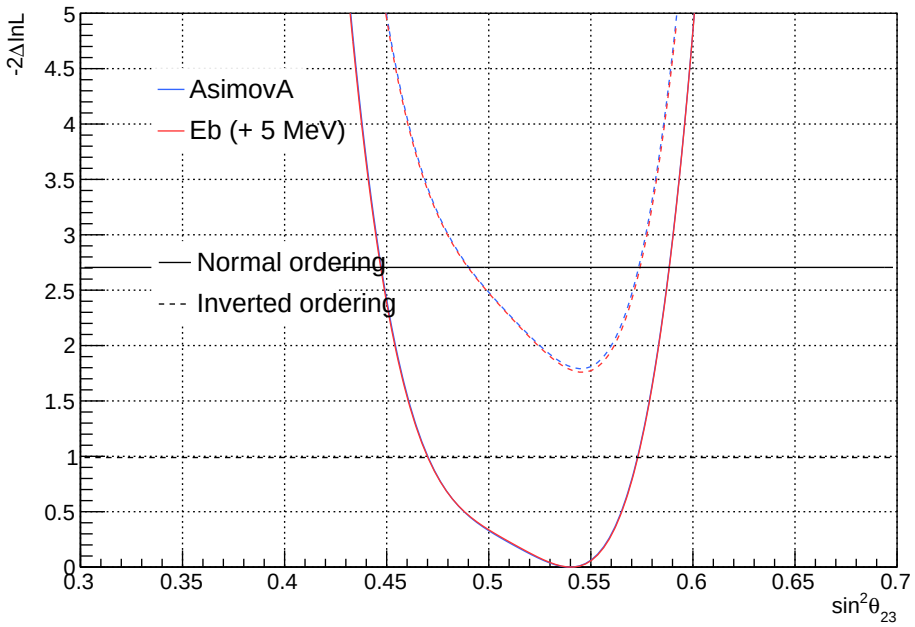

(a) $\sin^2\theta_{23}$ with reactor constraint

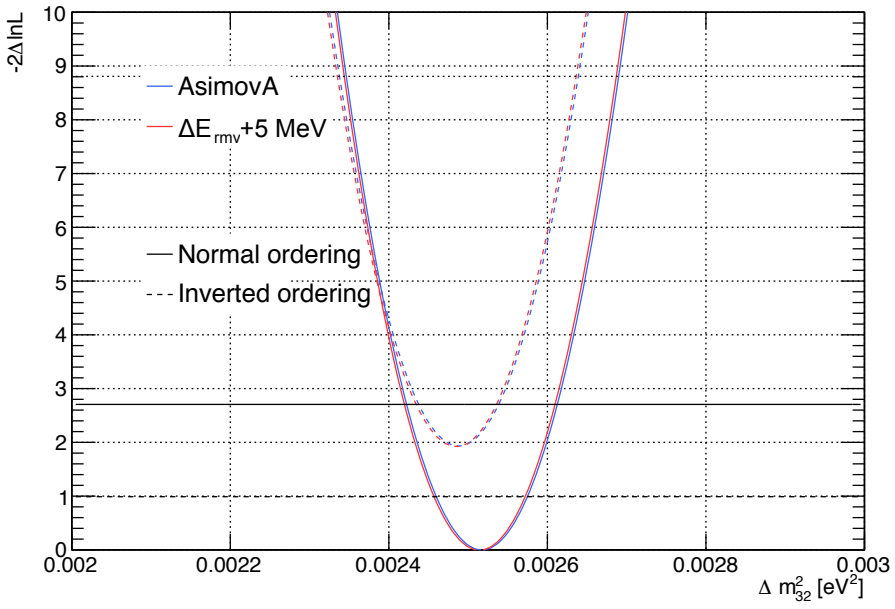

(b) Δm^2_{23} with reactor constraint

Figure 9.18: 1-D likelihood surfaces for all parameters, for P-Theta fits including the reactor constraint on $\sin^2\theta_{13}$ and assuming normal ordering, for the +5 MeV biased ΔE_{rmv} SK fake data. Continued on next page.

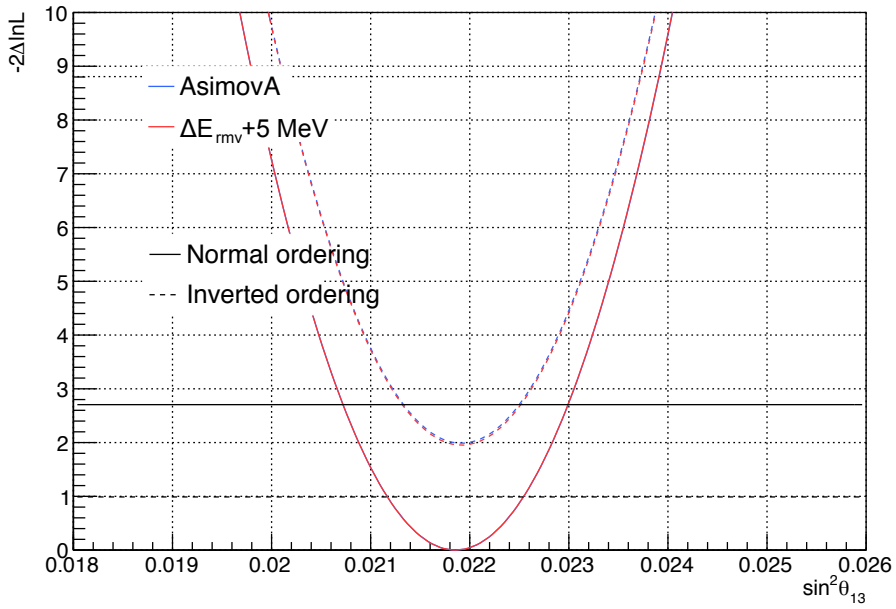

(a) $\sin^2\theta_{13}$ with reactor constraint

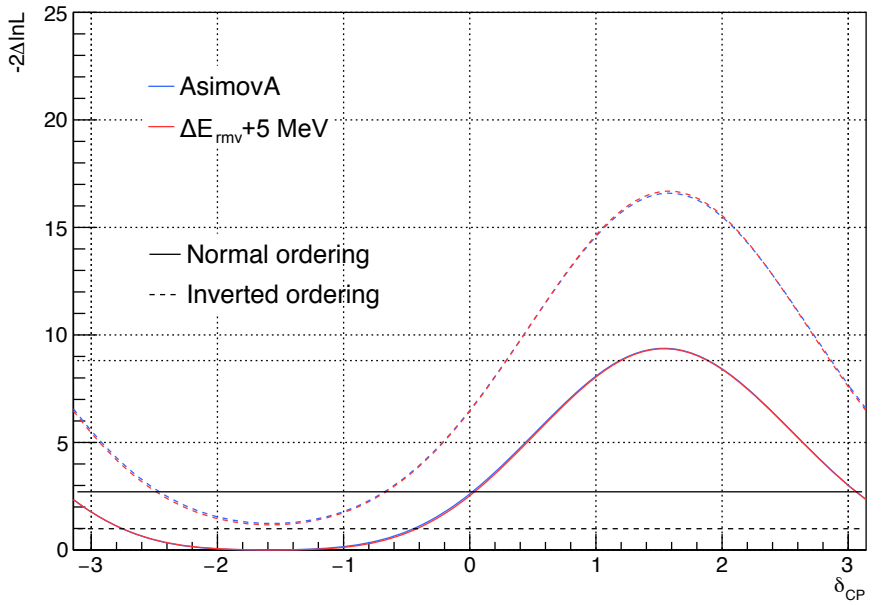

(b) δ_{CP} with reactor constraint

Figure 9.19: 1-D likelihood surfaces for all parameters, for P-Theta fits including the reactor constraint on $\sin^2\theta_{13}$ and assuming normal hierarchy, for the +5 MeV biased ΔE_{rmv} SK fake data.

9.4.2 Fake data smearing on Δm^2_{32}

In previous T2K analyses, the smearing applied to $\Delta m^2_{32}(\text{NO})/|\Delta m^2_{31}|(\text{IO})$ was large due to the large bias in the reconstructed neutrino energy from the poor removal energy uncertainty implementation and the RFG model's inability to accurately produce the effects of the nuclear shell structure. The total smearing applied to the $\Delta m^2_{32}\chi^2$ surface has been reduced by a factor of 2.8, largely due to the better removal energy uncertainty treatment which is no longer the dominant systematic uncertainty in the cross-section model. A comparison of the 2018 and the 2020 Δm^2_{32}fake-data smearing procedures as applied to the 2020 fit result is shown in figures 9.20 and 9.21. The 2020 BANFF fit biases calculated in the pull studies that are described in subsection 7.4, as applied to the Asimov fit result in figure 7.13 can clearly be seen to have a smaller effect than that of the applied smearing in the 2020 result, suggesting that the BANFF fit biases are covered by the additional uncertainty budget.

Six fake data sets were found to contribute a large enough bias to the value of $\Delta m^2_{32}$to be taken into account for the smearing. These were an enhancement of the non-QE contributions to the CC0π sample, an ND280 data-driven modification to the SK true pion momentum spectrum, a reweight to the prediction of the MINERvA parameterisation of SPP [145], the use of the z-expansion form factors for CCQE events [127], a reweight of 2p2h events from the Nieves model to the Martini [128] model, and the introduction of a bias in the value of ΔE_{rmv}. The size of the biases from each of these FDSs is listed in table 9.6. The largest of these effects is the enhancement of the non-QE contribution to to the CC0π sample, which is just under three times as strong as the bias from the value of ΔE_{rmv}. The data-driven pion momentum modification is also a larger bias than that of the +5 MeV shift in the ΔE_{rmv} parameters. These biases are then added in quadrature, and applied directly to the likelihood from the data fit as a Gaussian smear in the Δm^2-$\sin^2\theta_{23}$ contours.

9.4.3 Applying fake data uncertainty to δ_{CP}

Unlike $\Delta m^2_{32}(\text{NO})/|\Delta m^2_{32}|(\text{IO})$, the uncertainty on δ_{CP} is non-Gaussian, and asymmetric, and so a Gaussian smearing procedure is inappropriate. Instead, an additional term can be added to the $\Delta\chi^2$ surface of the oscillation parameters when a bias is observed. The impact on the δ_{CP} contours is assessed using

$$\Delta\chi^2_{\text{diff}}(\delta_{CP}) = \Delta\chi^2_{\text{nom}}(\delta_{CP}) - \Delta\chi^2_{\text{FD}}(\delta_{CP}) \tag{9.1}$$

to calculate the difference, $\Delta\chi^2_{\text{diff}}(\delta_{CP})$, between the nominal AsimovA fit result, $\Delta\chi^2_{\text{nom}}(\delta_{CP})$, and the Fake Data result, $\Delta\chi^2_{\text{FD}}(\delta_{CP})$, of the $\Delta\chi^2$ surfaces for δ_{CP}. The shift in the confidence intervals

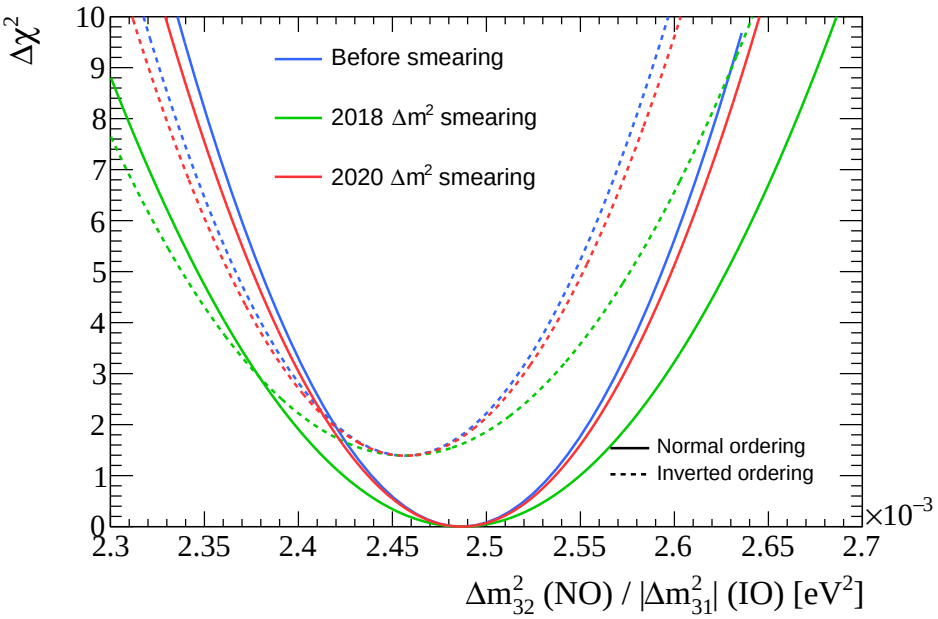

(a) T2K Only result

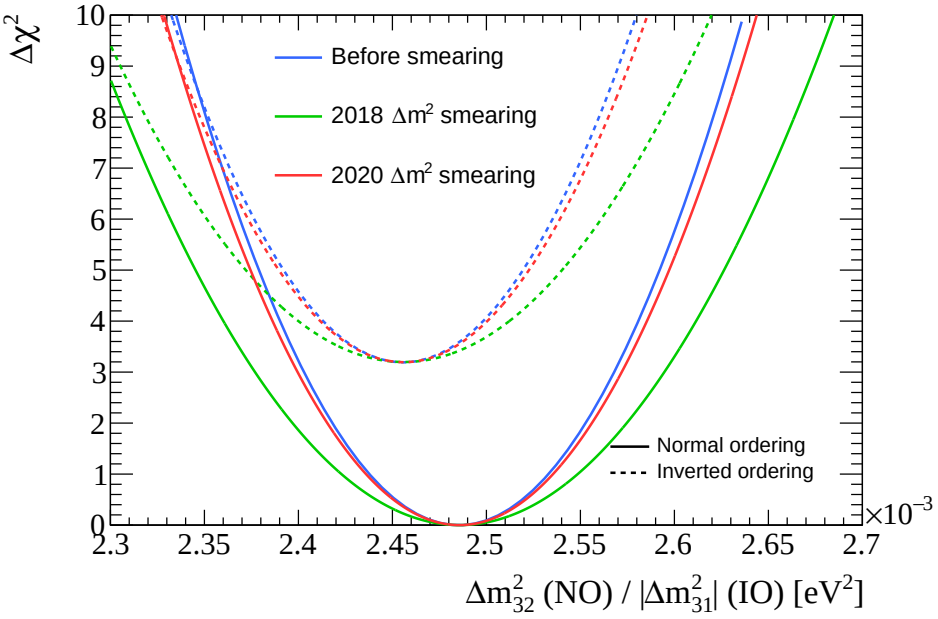

(b) T2K+RC result

Figure 9.20: Comparison of the 2018 and 2020 fake-data smearing applied to $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ for T2K only and reactor constrained (RC) results. A large reduction in the smearing in $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ can be seen due to the much smaller uncertainty budget required to cover biases in E_ν^{rec} due to the ΔE_{rmv} uncertainty implementation.

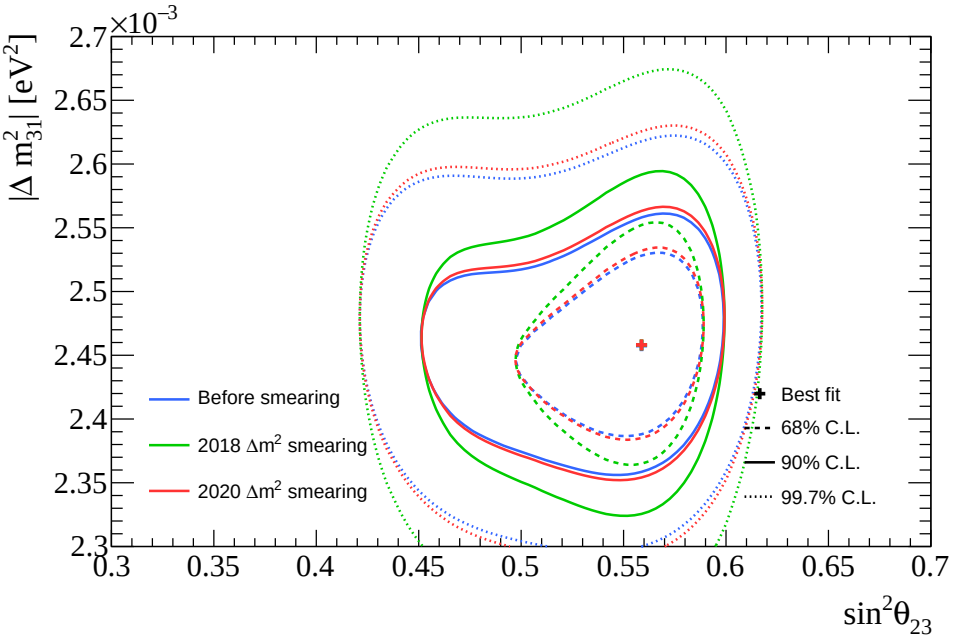

(a) Normal Ordering

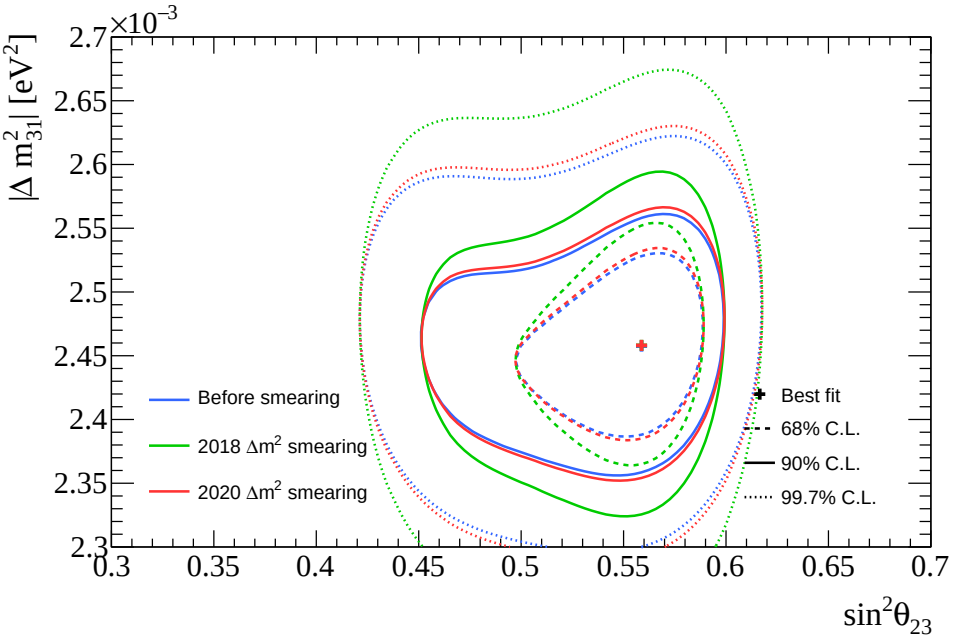

(b) Inverted Ordering

Figure 9.21: Comparison of the 2018 and 2020 fake-data smearing applied to the $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$ and $\sin^2\theta_{23}$ contours for the Normal and Inverted neutrino mass orderings (T2K+RC). The smearing as applied to $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$, whilst not directly applied to $\sin^2\theta_{23}$, can impact the $\sin^2\theta_{23}$ result through impacts to the marginalisation over $\Delta m^2_{32}(\mathrm{NO})/|\Delta m^2_{31}|(\mathrm{IO})$.

Fake data name	Bias
Non-QE CC0pi bias	32.662%
Pion data-driven	17.947%
Minerva	6.826%
Z-expansion	5.667%
Alternative Martini	7.312%
ΔE_{rmv} +5 MeV	13.388%

Table 9.6: The T2K fake data sets with significant biases on Δm^2_{32} used to calculate the smearing for Δm^2_{32} values. The biases are reported with more digits than their precision. The total size of the ΔE_{rmv}+5 MeV bias is just over a third of the largest bias which comes from an enhancement of the non-QE component of the CC0π topology.

can then be determined by applying this result to the data fit

$$\Delta\chi^2_{\text{shift}}(\delta_{CP}) = \Delta\chi^2_{\text{data}}(\delta_{CP}) + \Delta\chi^2_{\text{diff}}(\delta_{CP}) \tag{9.2}$$

where $\Delta\chi^2_{\text{data}}(\delta_{CP})$ is the result of the fit to data, and $\Delta\chi^2_{\text{shift}}(\delta_{CP})$ is the modified data $\Delta\chi^2$ from which the shifted contours can be calculated.

For the ΔE_{rmv} parameters, the difference between the nominal AsimovA fit $\Delta\chi^2$ and ΔE_{rmv} biased fake data set fit $\Delta\chi^2$ are shown in figure 9.22(a). The smeared data fit $\Delta\chi^2$, with shifted confidence levels, is shown in figure 9.22(b). The total change in the $\Delta\chi^2$ in δ_{CP} is small and the shift on the boundaries of the confidence levels is not large enough to change whether $\delta_{CP} = 0$ is or is not excluded to 2σ. Larger effects were seen in the other fake data studies such as the non-QE CC0π bias or the data-driven pion momentum modification as presented in table 9.6, which do impact this confidence level and so T2K does not report CP-conservation exclusion to 2σ. T2K does, however, report CP-conservation exclusion at the 90% confidence level.

9.5 Contribution of the removal energy parameters to the oscillation parameter uncertainty

A direct comparison of the 2018 and 2020 removal energy (or binding energy) uncertainty contributions to the oscillation parameters is non-trivial.

The nature of the heuristic parameter applied in previous analyses is best understood through the bias studies and not as a freely variable nuisance parameter with a well motivated physical response in the fit. The heuristic parameter had large impact on the event rate, which in turn caused

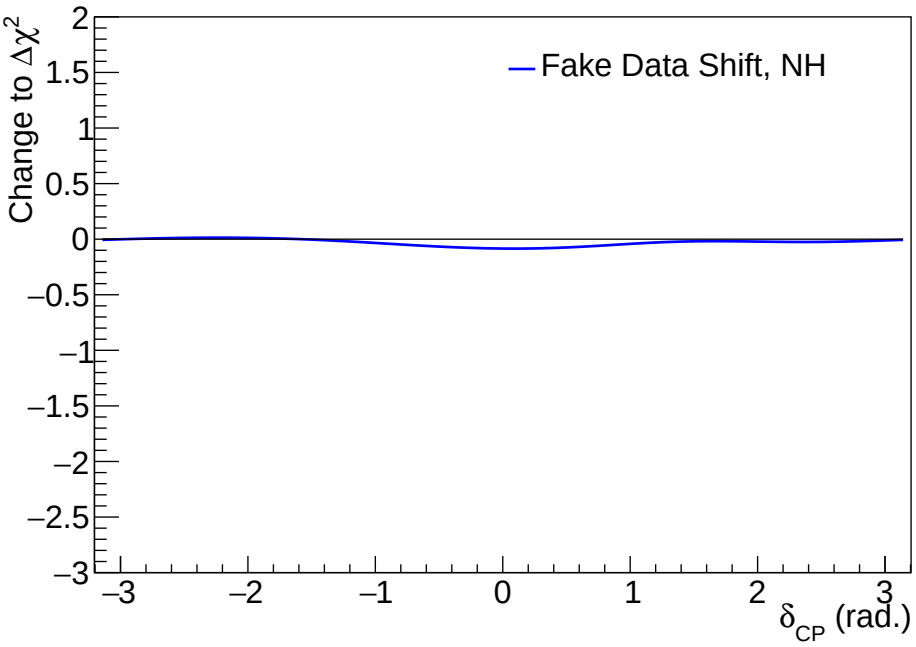

(a) 5 MeV $\Delta\chi^2_{\text{diff}}$

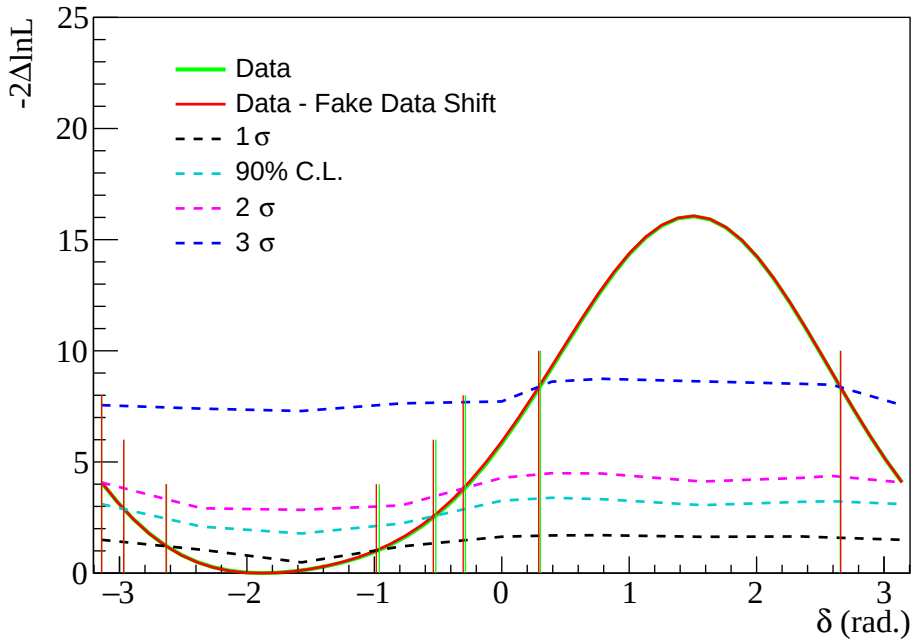

(b) 5 MeV $\Delta\chi^2_{\text{shift}}$ and CLs

Figure 9.22: (a) Difference between the one-dimensional δ_{CP} $\Delta\chi^2$ surface for the fit to the Asimov A fit and the fit to the +5 MeV biased ΔE_{rmv} data set. (b) The one-dimensional δ_{CP} $\Delta\chi^2$ surface and Feldman-Cousins confidence levels (CL) for the data before (green) and after (red) the additional uncertainty from the fake data fit is included. The vertical lines indicate the values of the CL-interval boundaries before (green) and after (red) the fake data study procedure. The impact of a bias in ΔE_{rmv} is small on the extraction of δ_{CP}, and a small bias of 5 MeV does not impact the exclusion of CP-conservation to 2σ on its own, though other T2K fake data studies do.

a bias in δ_{CP} to be seen. With the new uncertainty parameterisation this was reduced to an impact of less that 0.1 unit of $\Delta\chi^2$ near $\delta_{CP} = 0$, and less than this for values of δ_{CP} more than one radian away from 0. The impact of the heuristic parameter implementation on the Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) was large, and if applied to the 2020 result, the fake data smearing to account for the E_{rmv} uncertainty would account for $\sim 35\%$ of the total reported uncertainty. This is not entirely driven by the heuristic parameter but it can be assumed to be the largest contributing factor. This can be seen by the drastic reduction in the size of the smearing applied in the 2020 result in section 9.4.

Due to the computationally intensive calculation of the uncertainty contours using the Feldman-Cousins method, parameter-by-parameter contributions to the systematic budget are not calculated by T2K. However, an attempt to quantify the contribution was made through the calculation of the covariance matrix through the inversion of the Hessian matrix at the best-fit point found by a Minuit gradient descent fit to the SK Asimov B data set. This is not directly comparable to the the confidence levels calculated through the Feldman-Cousins method, and does not properly indicate the contributions due to correlations and anticorrelations with all of the other nuisance parameters. This was calculated not for the data fit but for fits to the Asimov B data set, in which the values of the oscillation parameters are less subject to boundary effects and are more symmetric, making the Hessian calculation more valid and so should only be used to estimate the relative size of the contribution, acknowledging that these assumptions may not hold near the best-fit to the data.

The procedure is as follows: The total uncertainty on the oscillation parameter is calculated using the Hessian procedure as in BANFF. The 2p2h energy dependence parameters are fixed in the fit due to their weak constraint and their nominal value is close to a physical boundary which can interfere with the covariance calculation. The postfit covariance is inverted, and the 2p2h energy dependence parameters are removed, before the matrix is reinverted to give a reference uncertainty for the oscillation parameter, σ_{tot}.

Beginning with the largest contribution to the uncertainty of the POI, the covariance is inverted, the rows corresponding to the systematic parameter are removed from the matrix and it is reinverted to produce σ_{fixed}. An estimate of the uncertainty for the systematic parameter is then calculated according to

$$\sigma^2_{\text{syst}} = \sigma_{\text{tot}} - \sigma_{\text{fixed}} \tag{9.3}$$

This process is iterated (setting $\sigma_{\text{tot}} = \sigma_{\text{fixed}}$ from the previous step) over for each contribution to the total uncertainty.

The estimated uncertainty corresponding to the leading ten contributions to the systematic uncertainty on Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) are presented in table 9.7. Using this method, an uncertainty contribution of $\sim 6\%$ of the variance on Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) can be estimated. In comparison to

Estimated uncertainty contributions to Δm^2_{32}			
Nuisance parameter	σ_{syst} [eV$^2 10^{-6}$]	$\sigma_{\text{syst}}/\sigma_{\text{tot}}$	$\sigma^2_{\text{syst}}/\sigma^2_{\text{tot}}$
$\sin^2\theta_{23}$	9.31312	0.20287	0.0411565
SK ν DIS NC	8.88506	0.193546	0.0374601
SK $\bar{\nu}$ NC	8.87629	0.193355	0.0373862
FHC flux ν_μ 0-0.4 GeV	4.58207	0.0998127	0.00996257
$\Delta E^{\text{O}\nu}_{\text{rmv}}$	4.56029	0.0993382	0.00986808
FHC Flux ν_μ 0.4-0.5 GeV	4.53739	0.0988394	0.00976922
RHC flux $\bar{\nu}_\mu$ 0-0.4 GeV	3.98853	0.0868834	0.00754872
RHC flux $\bar{\nu}_\mu$ 0.4-0.5 GeV	3.84865	0.0838363	0.00702853
FHC flux ν_e 0-0.5 GeV	3.78642	0.0824808	0.00680308
RHC flux $\bar{\nu}_e$ 0-0.5 GeV	3.72443	0.0811304	0.00658214
Sum of 10 dominant systs.	$\sim$18	$\sim$0.41	$\sim$0.17
Total	45.9067	-	1

Table 9.7: Estimated contributions to the systematic uncertainty on Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) in the fit to the Asimov B data set as calculated using the covariance subtraction strategy. The 10 largest contributing nuisance parameters are shown. The largest uncertainty comes from the uncertainty on $\sin^2\theta_{23}$, with which the mass squared splitting is highly correlated. The next largest contributions come from the SK only NC systematics. The neutrino removal energy parameter, $\Delta E^{\text{O}\nu}_{\text{rmv}}$, is only the fifth largest contribution to the systematic uncertainty on the mass squared splitting and is roughly 10% of the uncertainty, or more correctly 1% of the variance.

the $\sim$ 6% variance contribution of a 2018-like smearing factor this is a significant improvement. Note, T2K is still very much a statistics dominated experiment, particularly in the case of δ_{CP} and $\sin^2\theta_{23}$. The total uncertainty on Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO) has the largest systematic contribution of all the oscillation parameters.

Chapter 10

Beyond the 2020 oscillation analysis

The T2K 2020 oscillation analysis was a significant improvement on previous T2K analyses, both in terms of data samples used and models implemented, but there are known limitations and whilst currently these limitations do not seem to be a dominant source of uncertainty in the measurement of the oscillation parameters, they are important for modelling neutrino interactions and will become more significant in the future as experiments are able to collect more data and make more precise measurements. This chapter discusses ongoing and future work at T2K which aim to better constrain the oscillation parameters and prepare for the next generation of experiments.

The limitations of the previous analysis, and the improvements to the analysis made for this year (2021) are presented. This includes new data sample definitions and updates to the cross-section model which develop the SF nuclear model to a more complete shell-model description. Preliminary work on including pion kinematic information in the ND280 fit, which is intended to be included in the next analysis (2022) is discussed. An overview of upgrades to the ND280 that will improve sensitivity to physics and have the ability to constrain more sophisticated models is given. New data samples at SK will be included in the T2K analysis which take advantage of the Gadolinium doping discussed in section 3.5. An outline of improvements to the BANFF fitting framework is presented, which address current and foreseen complications to the analysis. Finally, joint fits with the Super-Kamiokande collaboration [104], and the NOvA collaboration [117], which aim to exploit complementary baselines and sensitivities in order to resolve the degeneracies between the PMNS parameters, and which provide better constraints than possible with any current single-experimental analysis, are discussed.

10.1 Limitations in the 2020 analysis

Although data-model compatibility studies and the ND280-fitter bias-impact studies show that the current T2K model well describes the data when binned in lepton kinematics, and that biases in the ND280 fit have minimal impact on the oscillation analysis, there are parts of the interaction model which are currently unsatisfactory for future planned samples and analyses. For example, neither the prefit nor postfit model accurately describes the reconstructed pion momentum peak for the CC1π sample as can be seen in figure 10.1. It is important to be able to properly predict this as in the near future (2022) separate samples for the 1Rμ plus one-decay-electron (1Rμ1d.e.) and two-decay-electron (1Rμ2d.e.) topologies will be included in the analysis. A major source of systematic uncertainty on the event rate for these samples is the pion momentum threshold at which pions are visible in Cherenkov detectors, as well as the rate of decay to Michel electrons which is highly correlated with the pion kinematics. The large difference between the 2020 BANFF prefit and postfit values of the M_A^{RES} term in the resonant nuclear form factor suggest that the model does not capture some nuclear effects.

As additional samples are added to the near detector fit, the current approach to the detector systematic variations risks convergence failure due to an increasingly high dimensional likelihood. The ability to fit the most significant sources of uncertainty directly in the fitter will allow the remaining binned detector normalisation parameters to be more strongly correlated, thus the detector-bin sizes can be expanded, reducing the number of parameters. Currently, the pion SI parameters are the dominant contributions to the total detector systematic uncertainty.

10.2 The 2021 analysis

The 2020 OA focused on a move away from an RFG nuclear model to a SF model with a better treatment of the removal energy systematic uncertainty; whereas, the 2021 analysis aims to include constraints from new observable information whilst adding additional uncertainties to make the SF model more shell model-like. The addition of new samples which go beyond the leptonic and mesonic topologies of the interactions will probe the ability of the ND280 data to constrain these new freedoms.

10.2.1 New selections at ND280

Two new event selections were validated since the 2020 analysis at the near detector and are combined to give the 2021 ND280 data samples. These new samples are currently only included for the ν-beam mode.

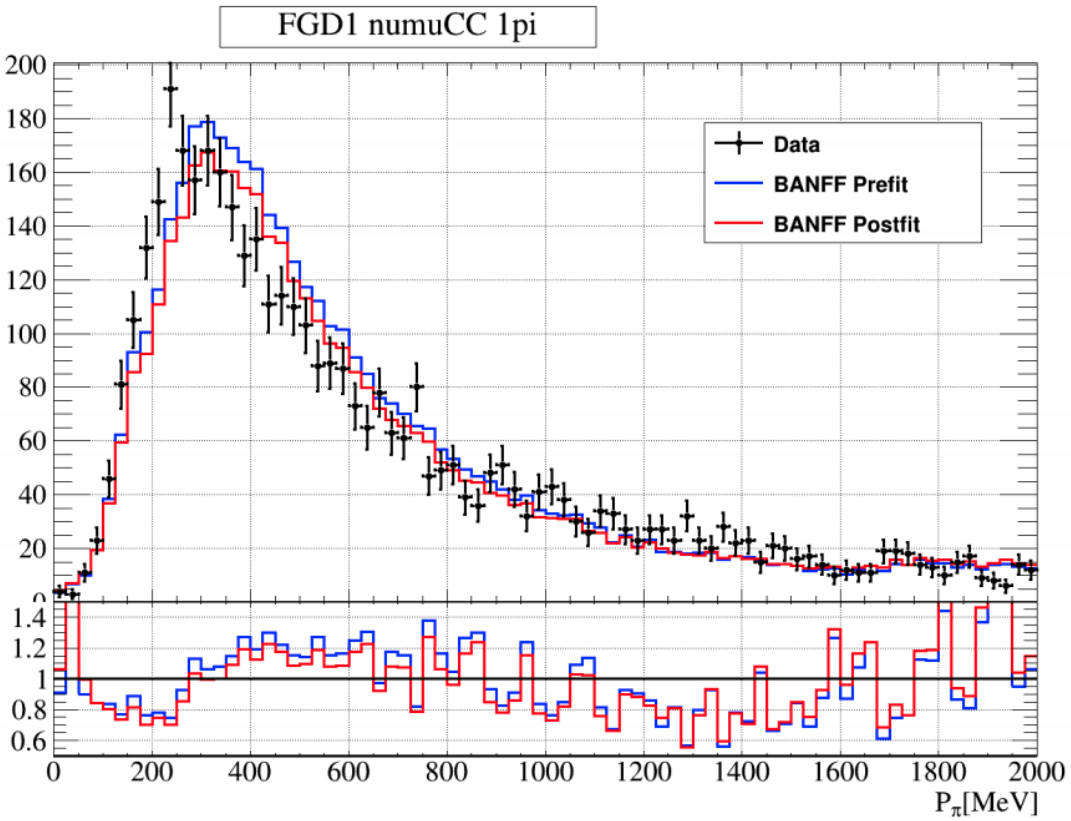

Figure 10.1: The top panel shows the reconstructed pion momentum, p_π, distribution for CC1π events in which the reconstructed pion track has a TPC component, for the prefit and postfit predictions and data. The bottom panel shows the ratio of each prediction to the data, showing where the current model consistently over or underpredicts the data. Neither the BANFF 2020 prefit nor postfit model accurately predicts the shape of the momentum distribution. The position of the peak is overestimated by approximately 100 MeV and the rate of events in the tail of the distribution is underestimated by around 20% between 1 to 1.5 GeV.

Proton Selections

Separation of events in the CC0π sample by the presence of a reconstructed protons from the initial interaction vertex gives sensitivity to final-state kinematics of the ejected nucleon. The presence or absence of a reconstructed proton is correlated with both the final-state lepton and initial-state nucleon kinematics; therefore, better sensitivity to the level of 2p2h contributions to the CCQE-like samples can be achieved. The NOvA experiment near detector tune accounts for a low hadronic energy E_{had} excess in the data compared to their prior model by enhancing the multinucleon contribution to their CC selection to 50% [117] at low E_{had}. This is significantly larger than the 2p2h contribution to the BANFF postfit reaction breakdown for the ND280 CC0π sample which is less than 20% at low E_ν^{rec}. Though the NOvA near detector tune is not a fit in the same sense as the BANFF, and relies more on systematic uncertainty cancellation, this is a key difference of the cross-section modelling between the two experiments that must be understood before a joint fit can be attempted.

Photon selections

Many different processes can produce photons in interactions at ND280, particularly the decays of mesonic products and of excited nucleon states. Common interaction modes at ND280 with final-state photons in order of event rate are

$$\pi^0 \to \gamma + \gamma$$

$$\eta \to \pi^0 + X \to \gamma + \gamma + X$$

$$\eta \to \gamma + \gamma$$

$$\Lambda \to \pi^0 + X \to \gamma + \gamma + X$$

$$K \to \pi^0 + X \to \gamma + \gamma + X$$

where X can be a variety of non-photonic final state particles. The photon topology of the event has not previously been used to separate events into different samples. In the 2020 analysis, the CC-Other sample captures a broad range of different events and so further distinction between these events can help improve models which are important for predicting the higher-energy interactions in the T2K beam. In addition to this, by first separating events by the presence or absence of a photon, the purity of the other samples can also be improved. In the 2020 analysis selections, events with a reconstructed muon and a photon in the final state, but no pion or meson, was not separated from other CC0π events, though this event is unlikely to be CCQE which is the target interaction mode of the CC0π sample. By adding this extra selection criterion, and removing

Analysis	Sample	Nμ	Nπ	Np	Nγ
2020	CC0π	$\geqslant 1$	0	-	-
	CC1π	$\geqslant 1$	1	-	-
	CCOther	$\geqslant 1$	$\geqslant 2$	-	-
2021	CC0π0p	$\geqslant 1$	0	0	0
	CC0πNp	$\geqslant 1$	0	$\geqslant 1$	0
	CC1π	$\geqslant 1$	1	-	0
	CCγ	$\geqslant 1$	-	-	$\geqslant 1$
	CC-Other	$\geqslant 1$	$\geqslant 2$	-	0

Table 10.1: Reconstructed topology definitions for the different samples included in the 2020 and 2021 ND280 fits in ν-beam mode. Each sample requires events to have at least one μ^- track from which the neutrino flavour and interaction vertex can be inferred.

events with a reconstructed photon from the CC0π sample, the QE purity of the sample can be increased. For the FGD1 (FGD2) CC0π sample there is a cost of -1.14% (-0.53%) in efficiency, but with a +5.04% (+4.62%) gain in purity. Similarly, for the CC1π samples there is a cost of -1.94% (-0.97%) in efficiency, but with a +7.75% (+7.25%) gain in purity.

The implemented ND280 samples

In the 2021 analysis, the proton and photon selections are combined by first applying a cut on the number of photons, then on the number of final-state pions, and then on the number of final-state protons. This turns the three FHC samples of the 2020 analysis into five samples as detailed in table 10.1. The existing CC1π and CC-Other samples are redefined as photonless equivalents.

10.2.2 New samples at SK

As with the 2020 1Re1d.e. sample in which a second, delayed e-like ring is used to infer the decay of a pion in ν_e-like events, there is a new sample in which either one or two decay e-like rings are tagged after an initial μ-like ring. These "multi-ring" samples target pion production. With the larger data set of the ν_μ disappearance channel than the ν_e appearance channel, there should be sufficient statistical power to separate this into two samples, 1Rμ1d.e. and 1Rμ2d.e.; however, this would require a better estimation of the final-state pion kinematics, which are a dominant source of the uncertainty on the reconstruction of a pion, and so these are left combined in the 2021 analysis. A better constraint on the pion kinematics from the near detector fit would help in the separation of the multi-ring events into two samples.

10.2.3 Updated cross-section model

New interaction parameters were included in the BANFF fit based on theoretical considerations in order to ensure that adequate freedom is included in the model to describe the new samples listed in this chapter, as well as to prepare the T2K interaction model to cover interactions at other experiments for joint fits.

In the CCQE parameterisation, the five lower Q^2 normalisation parameters that were designed to capture nuclear effects (except for the removal energy uncertainty) have been replaced with parameters that alter the shape of the SF and account for Pauli blocking effects, and with an optical potential correction to the SF. The shape of the SF is altered to reflect the change in the position and size of the peaks corresponding to the shells as predicted by relativistic mean field (RMF) models. An overall normalisation is applied to the region of the SF that corresponds to short-range correlations.

RMF calculations predict a strong linear dependence of the removal energy on the three-momentum transfer, q_3. A correction based on fits to electron scattering is applied to the 2020 ΔE_{rmv} systematic uncertainty. This is the first step to building a T2K model which goes beyond the impulse approximation in which the factorisation of the hard scatter from the SF ensures no dependence of the nuclear model on the incoming or outgoing lepton kinematics.

In the 2p2h model, the two 2p2h shape parameters included in the 2020 analysis, which vary the prediction from the Valencia to the Martini model, are split by the possible nucleon pairings. Neutrinos can interact either with a mixed pair *pn*, or a matched pair *nn* (or *pp* for $\bar{\nu}$). One of the new parameters included in the fit varies the ratio of these, which changes the shape of the interaction in the ejected nucleon kinematics. This ratio is highly model-dependent and there are significant differences between the predictions of the *pn* component in the Valencia model [38] and the SUSAv2 model [146]. The new $0p$ and $1p$ CCQE-like samples should have the power to constrain this new parameter and the now separated shape parameters as they have different peak structures in q_0 and q_3.

In the resonant pion production model, template normalisations in E_ν, Q^2, q_3 and W have been implemented for the variation of the resonant pion production removal energy Res-E_{rmv}. A parameter has been included to adjust the decay of the Δ resonance in the Rein-Seghal model which has a dependence on the true pion momentum and so should impact the postfit prediction for the reconstructed pion momentum. A simple normalisation parameter controls the ratio of $\pi^0/\pi^\pm$ production in CC events, which should be constrained by the new CCγ sample at ND280 allowing for a reduction in the uncertainty on the rate of backgrounds that result from asymmetric π^0 decays at SK.

Finally, in the FSI model, the high-energy charge-exchange parameter which was previously fixed due to lack of constraining power is now free to vary to ensure proper freedom for the proton-tagged split CCQE-like sample. The current FSI parameters only apply to events for which the pre-FSI and post-FSI have different kinematics in order to reduce memory usage and computation time. A new parameter which controls the amount of "FSI on" events which contribute to each sample is implemented to cover the uncertainty associated with this.

10.2.4 Removal energy q_3 dependency implementation

In fits to electron scattering data [48], there was an offset in the value of ΔE_{rmv} extracted from events with different values of the three-momentum transfer q_3. This is consistent with external measurements and predictions of relativistic mean-field models as discussed by González-Jiménez *et al.* [147]. T2K fits to the electron scattering data find that the dependency is approximately linear and so the current implementation can be modified using

$$\Delta E_{\text{rmv}} = \delta + \alpha(mq_3 + c) \tag{10.1}$$

where δ replaces the ΔE_{rmv} shifting parameter included in the 2020 fit, m and c are constants evaluated in these fits and α is a parameter that allows the strength of the correction to be varied. The momentum shift can then be calculating using this new value of $\Delta E_{\text{rmv}}(\delta, \alpha, q_3)$. This is the first inclusion of effects which go beyond the impulse approximation in the T2K interaction model.

Incorporating this correction into the kinematically-shifting removal energy parameter is trivial; however, the binned spline implementation in BANFF is somewhat dependent on the assumption of the universality of the ΔE_{rmv} parameter across all events. Two options were considered: one in which the weights are stored in a two-dimensional spline, and a second in which the relatively gradual change in q_3 both within a fit bin and between adjacent bins allows for the assumption that all events which would move in or out of a bin have a roughly similar value of ΔE_{rmv}. Due to the computational expense of two-dimensional interpolation, the latter was implemented.

10.3 Pion kinematics sensitivity studies

As the long-term plan for the SK multi-ring sample is to split the 1Rμ1d.e. and 1Rμ2d.e. samples, modelling of final-state pion kinematics is needed to accurately predict the number of events selected in each sample. A comparison of the prefit and postfit pion kinematics to data for the ND280 reconstructed momentum of pions in the CC1π sample in figure 10.1, shows that the 2020 model is unable to predict the pion momentum distribution peak position based on a constraint of the sample binned only in lepton kinematics. This is due to both the lack of direct constraint on

the pion kinematics, and also the lack of correlations between the lepton and pion kinematics in the T2K 2020 model.

There are three types of events selected with a CC1π (or CC1π0γ in 2021) topology at ND280: events with a reconstructed pion track with a component that passes through one of the TPCs (CC1π-TPC), events with a pion track that begins in one of the FGDs but does not enter a TPC (CC1π-FGD), and events in which the pion is inferred from its decay to a muon and subsequently to a Michel electron (CC1π-ME). Each of these event topologies broadly corresponds to different pion kinematic ranges with decreasing momentum. The most energetic pions are more likely to be forward going and leave the FGD in which they are produced and pass through a TPC. The lowest energy pions decay quickly via a muon to a Michel electron after a short time interval.

A study was undertaken to see whether any extra constraint on the SPP models could be gained from the inclusion of this pion tag information, and to explore whether additional freedoms would be needed for any future fits which would bin events in pion kinematics rather than purely lepton kinematics. The CC1π sample was split by tag into three subsamples. This was done in a single histogram method to avoid the need to fully prepare new samples in `psyche` by redefining the second observable as

$$\cos\theta_\mu \rightarrow \cos\theta_\mu + 2 \times \mathrm{Tag}_\pi$$

where Tag_π is an index indicating the pion tag, (TPC $= 1$, FGD $= 2$, M.E. $= 3$).

Initial comparisons for the 2020 model indicated that the pion tag did not provide any extra direct constraint on the existing SPP models, as shown in figure 10.2, but did show improved constraints on background processes including coherent pion production and FSI, shown in figure 10.3. This is in line with the relative contributions of the interaction processes shown in the lepton kinematics in figure 10.4, where the background composition of each sample differs more than the shape of the resonant pion production when binned in p_μ. A comparison applying a single selection (e.g. TPC vs not-TPC tagged pions) indicated that the largest contribution to this constraint came from the separation of the events with TPC tagged pions from events with M.E. tagged pions.

An auxiliary study is to be conducted to probe the sensitivity of the 2021 model to additional freedoms intended to allow for a better prediction of the pion kinematics. The Rein-Sehgal delta decay systematic uncertainty should weight events in bins of pion momentum, thus the extra sensitivity of the split pion tag samples to constrain this parameter would result in a better postfit prediction of the reconstructed pion momentum distribution. If this is confirmed by this study then there would be clear a benefit to binning the CC1π0γ sample in pion kinematics for future fits.

This study will be used to inform future decisions about samples to be included at the near detector and the observables in which to bin them. Whilst a full detector systematic uncertainty evaluation on the reconstructed pion kinematics was not yet prepared, the pion identification systematic uncertainty is already used and so splitting the sample into three based upon the detector tag (TPC, FGD, M.E.) could be done with only a regeneration of the detector covariance matrix after separating out the samples. An update to the fit and detector binning was done to ensure that each bin contained at least one data event and twenty MC events; and then, beginning at the new fit binning, iteratively reduce the detector binning until the minimum number of bins for which no adjacent bins have a correlation within 7% of each other.

10.4 ND280 Upgrade

In 2022 T2K will undergo a long shutdown, during which the beam power will be increased and the ND280 detector will undergo an upgrade [148] in which new subdetectors will be added in preparation for T2K-II and their eventual use in the Hyper-K experiment [43]. The PØD will be removed to make room for the Super Fine Grained Detector (Super-FGD), and two horizontally oriented "high angle" time projection chambers (HA-TPCs) above and below the Super-FGD as shown in figure 10.5, and a time of flight detector (TOF) will enclose these new detectors to provide better track direction determination. The purpose of these new detectors is to provide better angular resolution and tracking, allowing ND280 analyses to have the same solid-angle coverage as Super-K. Unlike the current FGDs in which scintillator bars allow for readouts in alternating planes, the SFDG is constructed from 1 cm $\times$ 1 cm $\times$ 1 cm scintillator cubes through which three WLS-fibres are threaded building layers in three dimensions allowing for readout in three planes as shown in figure 10.6. This improves the spacial and angular resolution of the detector. The Super-FGD is oriented flat, in contrast to the existing FGDs in which bars are vertical and transverse to the beam, with a HA-TPC above and below the Super-FGD. This, along with the existing TPCs means that a much greater solid angle around the target mass benefits from the high angular resolution of the TPCs. In addition to the upgrade of the ND280, the J-PARC main ring beam power is expected to be increased from its current maximum power of 515 kW to 1.3 MW, resulting in pile-up or larger numbers of neutrino interactions in each event window. The TOF detector system uses six panels of scintillator bars which enclose the Super-FGD and HA-TPCs entirely, precise enough timing to be able to infer track direction of charged particles exiting the Super-FGD.

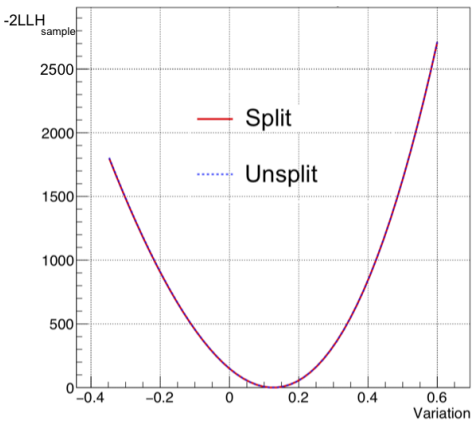

(a) M_A^{RES}

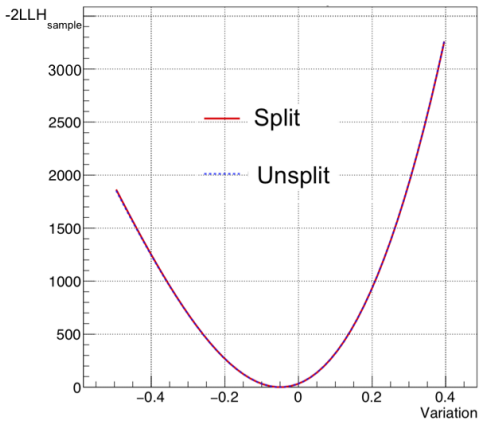

(b) C_5^A

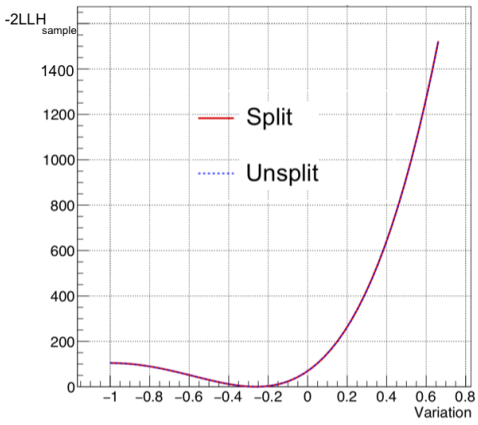

(c) $\mathrm{I}_{\frac{1}{2}}$-Background

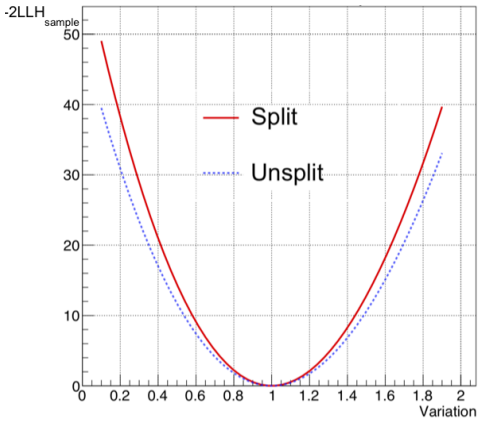

(d) CC-Coh. C

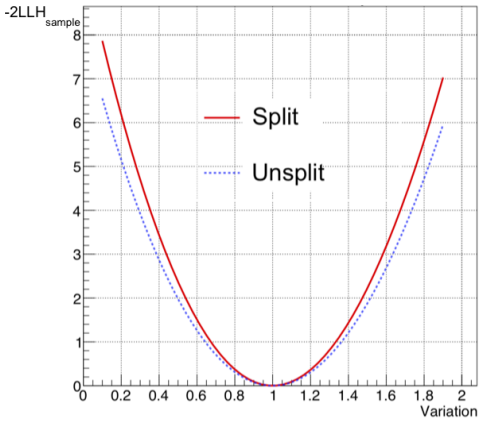

(e) CC-Coh. O

Figure 10.2: Comparison of LLH-scan sample contributions for the CC-resonant and CC-coherent pion production cross-section parameters for the 2020 ND280 fit with the standard CC1π sample (unsplit) and CC1π sample split by pion tag (split). The resonant parameters do not gain any extra constraint from using the split samples, though the coherent background, which contributes mainly to the TPC-π, does gain an extra constraint from this separation.

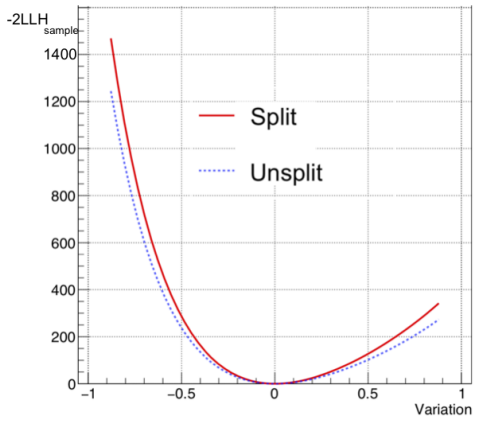

(a) Low QE

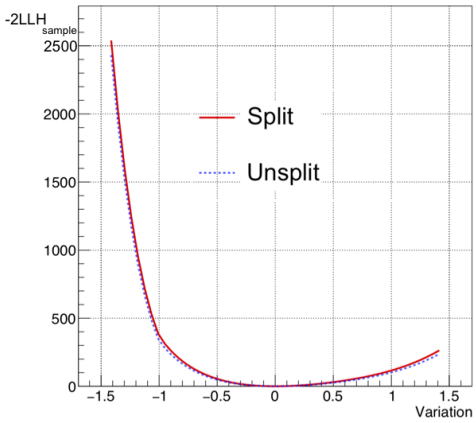

(b) High QE

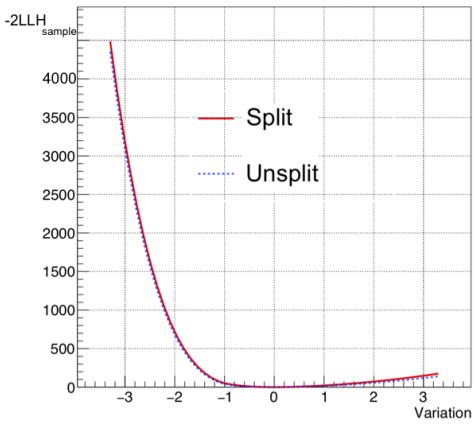

(c) Hadron production

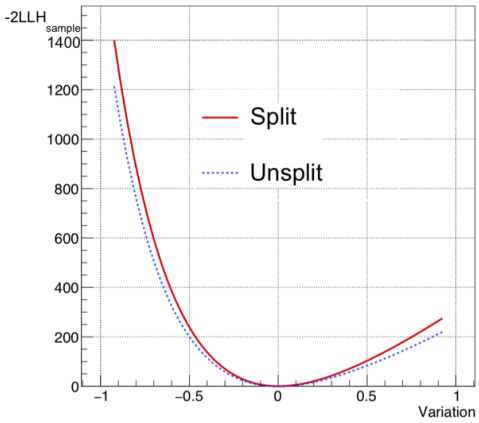

(d) π absorption

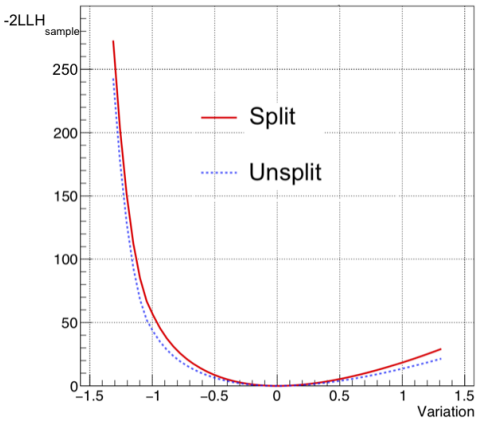

(e) Charge exchange

Figure 10.3: Comparison of LLH-scan sample contributions for the FSI parameters for the 2020 ND280 fit with the standard CC1π sample (unsplit) and CC1π sample split by pion tag (split). The FSI parameters, which can impact the final-state topology of pion events, do gain extra constraint. In particular, the low-QE parameter, and the π-absorption and charge-exchange parameters which can change the final state topology of CCQE and SPP events, gain some constraint.

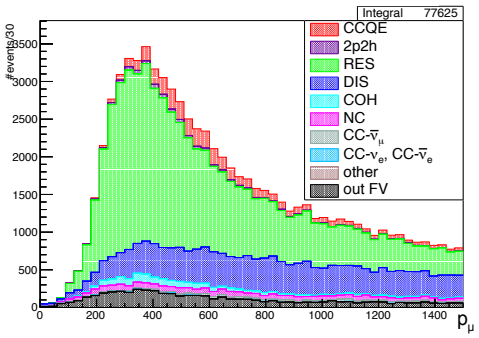

(a) Total CC1π

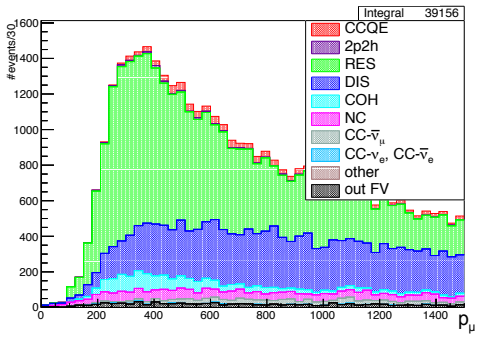

(b) TPC-π

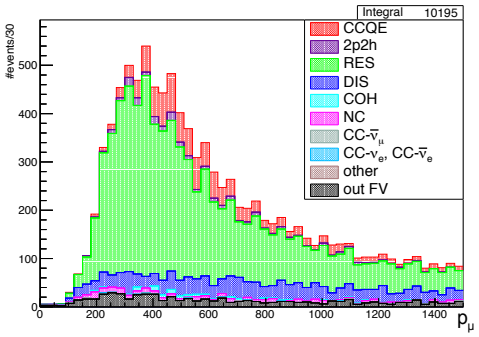

(c) FGD-π

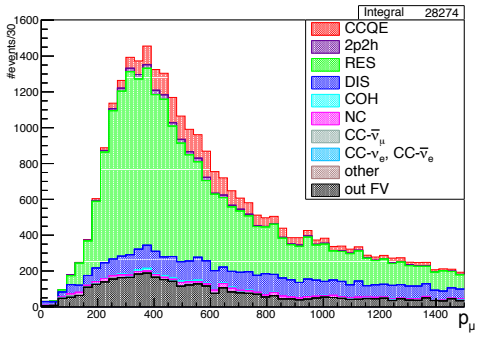

(d) ME-π

Figure 10.4: True MC interaction modes (coloured histograms) of the CC1π sample when divided by pion tag, binned in the true lepton momentum, p_μ. Whilst the shape of the resonant pion production contribution to each sample is roughly the same when binned in lepton kinematics, the contribution of different background processes varies. Most of the coherent and NC backgrounds are from the TPC-tagged pions whilst the FGD-and ME-tagged pions have a larger CCQE contamination. A significant background to the ME-tagged pion selection comes from out-of-fiducial-volume events. Higher pion multiplicity and DIS events contribute significantly to the TPC-tagged pion selection, in part due to unreconstructed pions and mesons.

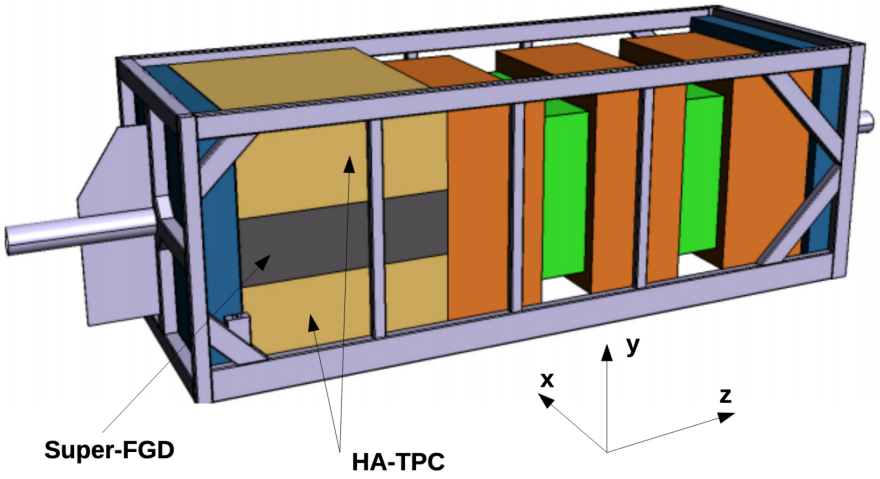

Figure 10.5: CAD model of the ND280 upgrade detectors inside the basket of the ND280 [148] (the TOF detectors are not shown). The Super-FGD and two HA-TPCs will replace the PØD in the region upstream of the tracker. The orientation and arrangement of these new subdetectors allows for better reconstruction of high angle tracks transverse to the beam which currently have lower reconstruction efficiencies due to the orientation of the FGD-scintillator bars and the position of the tracker TPCs.

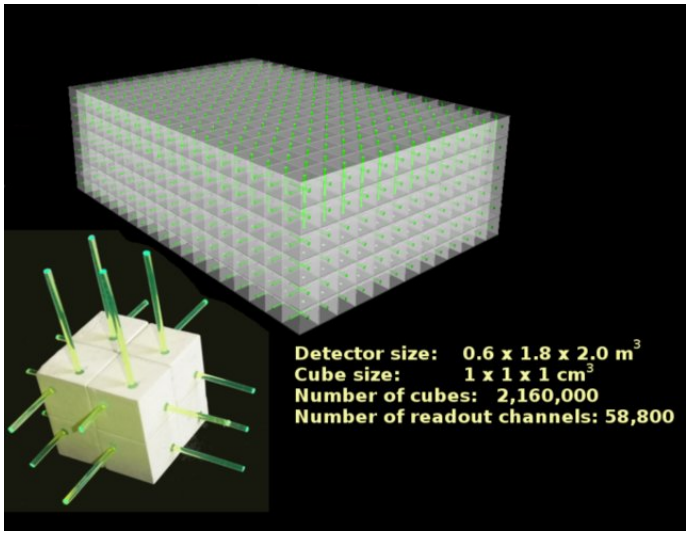

Figure 10.6: Schematic of the Super-FGD scintillator cube and WLS-fibres (centre) and a photograph of an individual scintillator cube showing the three readout fibres passing through it [148]. This arrangement gives the same angular resolution for tracks moving transverse to the beam as along it and so will allow better measurement of high-angle and backwards going events.

10.4.1 Physics potential of the Super-FGD

The higher resolution of the Super-FGD allows for better reconstruction of short tracks from particles such as ejected nucleons and low-momentum pions. Access to the kinematics of these particles can provide a probe into the nuclear physics and non-QE models which are becoming more important with the addition of new samples at Super-K which are not targeting CCQE events, and as T2K works on joint fits with other experiments with different interaction energy regimes.

Different observables such as transverse momentum, or the angle between final-state particles, may provide a better constraint on parts of the interaction model. For example, the sum of the transverse momentum of a lepton and ejected photon constrains the momentum which has been transferred to the nuclear remnant. From this a better estimation of the missing energy can be made and any bias in E_ν^{rec} can be reduced. Similarly, from the angle between ejected protons and pions in resonant production, information about the intermediate hadronic state can be inferred. The Super-FGD will have a better reconstruction transverse to the beam than the current FGDs, and so will have the precision necessary to make these measurements for a much wider range of solid angles than would be possible with the tracker FGDs and TPCs. The higher resolution of the Super-FGD will also allow for neutron tagging in $\bar{\nu}$-interactions, which will improve constraints on the 2p2h contributions to the RHC beam.

Neutrino interactions which produce high angle (and backwards) charged lepton tracks have a higher typical Q^2 than the events with forward going leptons that ND280 can best reconstruct. As seen in chapter 8, the three high Q^2 normalisation parameters of the CCQE model have strong anti-correlations with both the neutrino energy, E_ν, and the axial form factor mass M_A^{QE}. This is another aspect of the T2K model in which the Super-FGD will have clear benefits when it comes to interaction model constraints.

10.5 SK Gadolinium samples

The SK experiment now uses gadolinium-doped water as its target material [149] allowing for more efficient neutron-capture tagging. In $\bar{\nu}$ interactions in which a neutron is ejected from the nucleus, the neutron may be captured by another nucleus exciting it into a higher energy state as shown in figure 10.7. This excited nucleus, can then undergo decays corresponding to the discrete energy gaps in the nuclear shell structure. Gadolinium is efficient at neutron capture and the characteristics of its decay gammas is well understood. This addition was made for the SK collaboration's own analyses searching for super-nova relic neutrinos; however, it can be of benefit to the T2K analysis for wrong-sign background discrimination in the RHC as neutrons are ejected

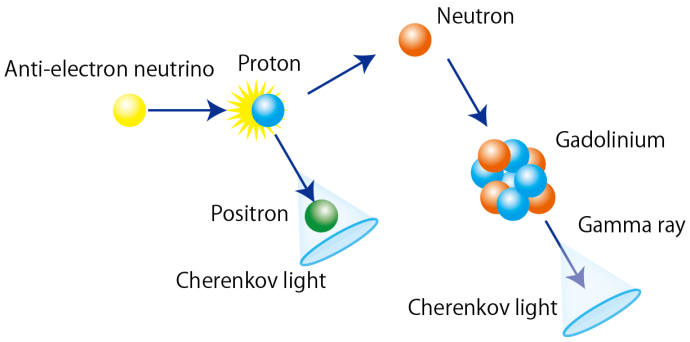

Figure 10.7: The capture on a gadolinium nucleus of an ejected neutron from antineutrino interactions with protons in hydrogen or oxygen nuclei produces a metastable isotope which decays via gamma emission. The characteristics of this gamma emission allow for a tagging of otherwise unobservable neutrons [150].

from oxygen nuclei in antineutrino interactions but not neutrino interactions.

10.6 BANFF fitting framework updates

As more sophisticated model and sample treatments are required to take advantage of new detector capabilities and the corresponding larger data sets, the fitting framework must be updated to handle this. Alternative fitting methods to the gradient descent algorithm are being explored, such as a simple implementation of the artificial bee colony algorithm [151], and detailed comparisons of different methods on the full ND280 likelihood are planned for the near future. Directly fitting the largest contributions to the detector systematic uncertainty may reduce some of the biases currently seen in the analysis. A study into the direct fitting of the pion SI systematic uncertainties is underway which should reduce the total uncertainty that must be covered by the detector covariance and the observable normalisation parameters, and in turn allows for more of the bins to be merged resulting in a reduced number of parameters. Another improvement to the detector parameter treatment is to use a principal component analysis to remove the smallest-contributing detector parameters to the variations described by the detector covariance. Both of these changes should allow for better treatment of the increasingly large set of detector parameters as more samples are included in the fit. Otherwise, the addition of new samples, potentially with three-dimensional fits, would entail on the order of hundreds or even thousands of extra parameters as the number of fit bins are increased.

10.7 Joint fits with other experiments

Due to the different neutrino energies and baselines of different neutrino oscillation experiments, each have different sensitivities to different oscillation parameters. Joint fits intend to exploit the strengths of the different experiments in complementary ways. Two joint T2K analyses are currently underway with the NOvA and Super-K collaborations. Whilst many theorists make global fits to neutrino data by combining published likelihoods in only the oscillation parameters after all nuisance systematics have been marginalised, that procedure may overlook important correlations between the different models implemented by different experiments. Identifying these correlations and fitting directly to the full likelihoods of multiple experiments may break any tension between the results and resolve the degeneracies in the PMNS parameterisation in complementary ways for each experiment.

10.7.1 T2K-Super-K

Given that the SK experiment uses the same detector as T2K as well as using the NEUT MC generator in many of its analyses, joint fits to T2K beam data and SK atmospheric and solar neutrino data benefit from shared detector systematics and models. Unlike T2K which has a narrow-band beam with a peak energy of 600 MeV, SK is exposed to atmospheric neutrinos with energies that range from the few MeV scale to the TeV scale. As such, SK sees many more non-QE-like events which undergo multi-pion production and DIS than T2K. These SK events may complement the multi-ring samples which are being added to the T2K analysis. Atmospheric neutrinos that pass through the earth are subject to significant matter effects that can enhance the oscillation probability dependence on the neutrino mass ordering. The degeneracy of the mass ordering and the value of $\sin^2\theta_{23}$ limits SK's ability to resolve the sign of Δm^2_{32}(NO)/$|\Delta m^2_{31}|$(IO). T2K has limited sensitivity to the mass ordering as it is designed to have minimal matter effects, probing only the first oscillation maximum and at lower neutrino energies; however, T2K does have the current world-leading measurement of $\sin^2\theta_{23}$, and so a joint fit may lift this degeneracy and allow for a better sensitivity to the determination of the mass ordering.

10.7.2 T2K-NOvA

The results of the 2020 T2K analysis detailed in this thesis seem to be in slight tension with the results presented by the NOvA collaboration at Neutrino 2020 [117]. Whilst there remains a small region of the δ_{CP}-$\sin^2\theta_{23}$ phase space still compatible at 1σ, the experiments clearly favour different regions of phase space, each excluding the other experiment's global best fit to 90% confidence as shown in figure 10.8. It must be stressed, however, that the "tension" between the

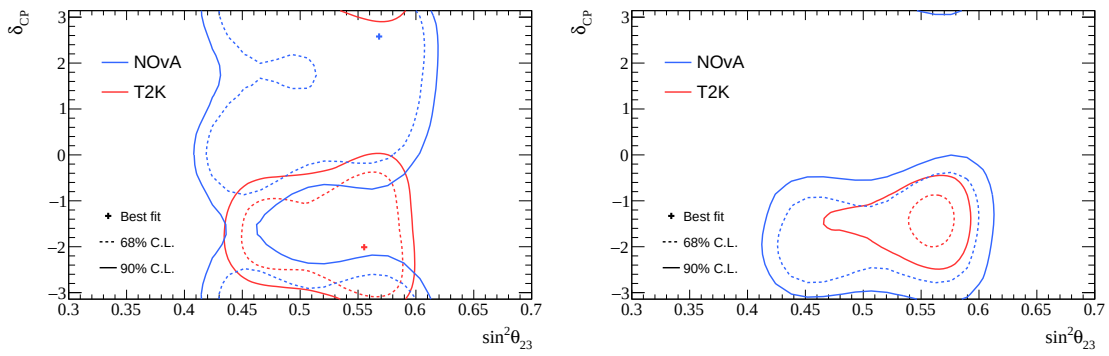

(a) Normal ordering

(b) Inverted ordering

Figure 10.8: Feldman-Cousins confidence level contours for δ_{CP} and $\sin^2\theta_{23}$ for both the T2K and NOvA [117] preliminary results presented at Neutrino 2020. Global $\Delta\chi^2$ values over both orderings are used for both experiments. Though small regions of phases space are still consistent at the 1σ C.L. in both hierarchies, an apparent tension exists between the best fit points and regions excluded above 90% C.L. by each experiment in the normal ordering.

results is not large and may resolve on the collection of more data as both experiments are still statistically limited. These fits do not take into account the full likelihoods of each experiment and cannot be appropriately assigned correlations between nuisance parameters. As a result, important effects may have been marginalised out which could impact the results. To resolve tension between the two experiments, a joint fit to the two data sets with the likelihood of each experiment is being prepared in which the two collaborations will attempt to assess these correlations and include them.

10.7.3 Role of the ND280 in joint fits

Many of the considerations in the near detector fit have been specified to benefit the T2K analysis; however, its cross-section measurements, as well as some of its fake data studies, are intended to probe effects seen in other experiments. As ND280 sees a much higher flux of neutrinos than SK, it still may see an appreciable number of interactions which are impacted by effects which, whilst not important for T2K, are significant sources of systematic uncertainty for NOvA or SK. Understanding these effects, or the ability of NEUT to predict these effects, may improve understanding of how the different models implemented at the different experiments are correlated. The difference in the BANFF fit and the NOvA near-detector tune predictions for the rate of CCQE (1p1h) and multinucleon (2p2h) events may contribute to differences in the results. An under-prediction of the CCQE rate at T2K could lead to an under-prediction of the number of neutrinos which oscillate from ν_μ, as well as to an under-prediction of the rate of ν_e QE-like interactions, resulting

in an enhancement of the ostensibly observed CP-violation.

Here is where studies such as the pion-production sensitivity will become important for joint analyses with experiments of typically higher neutrino energy. CCQE and multinucleon processes are subdominant compared to resonant pion production in the NOvA CC selection, and it is a significant contibution to the ND280 CC0π selection. Whilst there are good neutrino energy estimators for the QE-like range at T2K, E_{rec}, and the DIS/Multi-π-like range at NOvA, E_{had}, the estimation of neutrino energy for resonant interactions which bridge the gap between these neutrino energy domains is poorer. Better modelling of the SPP processes may therefore address important systematic uncertainties for which the correlations between the experiment likelihoods need to be understood in order to bridge the gap between T2K's reconstructed energy and NOvA's hadronic energy estimators of the true E_ν.

Chapter 11

Conclusions

The field of neutrino physics has proved to be a fertile ground for BSM discovery with existing theoretical avenues still to explore. Future experiments will likely yield results that have profound implications for our understanding of particle physics and cosmology, but as a ten-fold, or even a hundred-fold more data is collected with higher precision instruments, the systematic uncertainties on the modelling of neutrino-nuclear interactions will be of increasing importance when refining measurements of the oscillation parameters. The sign of Δm^2_{32}, the octant of θ_{23}, and the size of the CP-violating term, δ_{CP}, all require a detailed understanding of differences in the cross section for different neutrino flavours, energies and sign, as well as for different nuclear targets including new detector materials such as liquid argon which is becoming commonplace. Efforts to improve neutrino interaction models and their systematic uncertainty treatment have been fruitful in the modelling of final-state lepton kinematics, as in the case of the nucleon removal energy demonstrated in this thesis; however, clear evidence of mismodelling of other interaction observables, such as the kinematics of resonantly-produced pions, indicates that there is still room for improvement. These are aspects of the model which, whilst not currently of critical importance for T2K, are significant for future experiments as well as for assessing correlations for existing experiments at higher neutrino energies where these processes are dominant to the QE processes that are the primary focus for T2K. For the 2020 oscillation analysis of the T2K data, the removal energy systematic uncertainty is no longer the dominant source of uncertainty in the interaction cross-section model. The improved treatment of the nucleon removal energy, along with other improvements to the cross-section uncertainty treatment, has allowed for a 2.8 times reduction of the smearing applied to the contours in the extraction of the value of Δm^2_{32} in a way which allowed the associated parameter in the MC generator to be varied appropriately during fits at both ND280 and SK. The binned spline parameter developed in this work allowed for a smooth and continuous variation of the parameter at ND280, preventing pathological likelihood variations as a result of

finite MC statistics and binning. Comparison with the directly-fit kinematic variation parameters in the MCMC fitter shows that, whilst the gradient descent implementation used in the BANFF fit avoids other non-physical effects of binning and marginalisation, the results between the two fitter methods are consistent to within 1σ for both of the carbon parameters and the $\bar{\nu}$-on-oxygen parameter, and only differ by less than 1.5σ for ν-on-oxygen. The impact of these differences on the oscillation parameter contours was studied in depth and found to be smaller than both the effect of the binning choices in the MCMC method and the effects seen in robustness studies in which alternative models are fit. As such, the additional uncertainty added to the oscillation parameters through smearing should cover the possible differences in the confidence intervals or credible intervals obtained by using the ND280 best-fit values or posterior probability density distributions obtained through the different methods. It is estimated that the contribution of the removal energy uncertainty to the uncertainty on Δm^2_{32} was reduced from a $\sim 35\%$ to $\sim 10\%$, or more accurately the contribution to the variance was reduced from $\sim 6\%$ to $\sim 1\%$. This small difference is due mainly to the smaller uncertainty afforded to the nucleon removal energy by updating the nuclear model from the RFG to the Benhar SF which includes more physical information about the nuclear shell structure. The impact of this change can also clearly be seen in the improved p-value calculated for the 2020 T2K model at the near detector of 73.7% when compared to the previous T2K model with a p-value of 50%.

The near detector constraint on the flux and neutrino interaction models is a vital part of the oscillation parameter extraction and therefore it is carefully tested. Through thorough testing of fitter stability and bias it was concluded that the impact of fitter machinery bias was small at both the near and far detector when fitting to the region of parameter space dictated by the T2K prior model. Robustness studies indicate that alternative models which are not currently fit at T2K have small impacts on the oscillation parameters, and so the T2K results are adjusted according to these studies and report conservative confidence levels and credible intervals where appropriate.

Whilst T2K has expended significant person hours in the thorough investigation of its systematic uncertainties, and this will continue as considerations of future experiments and joint fits are catered to by modelling groups, the inclusion of new constraints on the data through more sophisticated selection development may help to guide model-building which will be of use in the future. An initial investigation of the ND280 sensitivity to pion kinematics and its possible impact at the far detector has shown that in the 2020 model, resonant pion production was not directly constrained by the addition of pion kinematic information to the fit. However, new uncertainties in the 2021 model on the impact on the final-state pion kinematics includes freedom via which the observable pion kinematics may better constrain the resonant production model.

Appendix A

T2K analysis software

To make these comparisons the data must be simulated in a form in which it can be treated in exactly same way as the data, but with the full "true" information about the simulated data available. These simulations are computed using "Monte Carlo" methods and the simulated data are often referred to as "the Monte Carlo simulated data" or more simply, "the Monte Carlo" (the MC), as detailed in section A.1. In the high-statistics limit MC-predicted distributions should approach the "true" model prediction for the data. Monte Carlo methods are useful for predicting complex but random processes, either truly random processes such as those in quantum mechanics, or chaotically random processes such as those from statistical mechanics.

For example, if nuclear beta decay were a neutrinoless two-body problem, the precise momentum of the β particle would be calculable and could be confirmed by measurement exactly within some measurement resolution. However, in real beta decay the momentum is split amongst three particles allowing energy to be given to an undetectable neutrino; hence, the momentum of the β particle can take a range of possible values according to some probability distribution. Whilst this distribution may be calculable, it is not possible to calculate with certainty the momentum of any individual beta particle. The calculable distribution can, however, be used to calculate the probability of an individual emitted beta particle to have some specific momentum, and a process known as "accept-reject" is used to generate an ensemble of MC events which will have the predicted distribution of momenta. This process is discussed in more detail in section A.1. Even in the first case where the final-state kinematics of the beta particle are completely calculable, a full modelling of detector apparatus and the physical extent of the radioactive source will introduce fluctuations to the measurement and smear the resulting distribution.

This appendix outlines the process of generating MC, and the ND280 software chain which is used to compare MC to data for analysis.

A.1 Monte Carlo

In many fields there is a need to simulate random or semi-random processes and here Monte Carlo[1] (MC) methods can be used to do this. Named for the Monte Carlo Casino in Monaco, MC methods aim to approach problems which involve a large degree of chance or combinatorics by randomly sampling possible solutions to a set of criteria. Typically these rely on some form of "accept-reject" method to simulate a discrete set of data according to some probability which is calculated from a model, using randomly selected values of event parameters, or "throws", which obey a set of rules. For example, in a simple neutrino-interaction model, the probability of interaction σ may be solely dependent on some parameter x. The accept-reject method draws two random numbers, one uniformly in the range of x, the parameter of interest, and one, y, in the range $[0,1]$, which will determine whether the event is accepted or rejected based on whether it is below or above $\sigma(x)$ which acts as the probability function. Therefore the criteria for a randomly generated event to be accepted is

$$P_{\text{accept}}(x_i, y_i) = \begin{cases} 1, & y_i <= \sigma(x_i) \\ 0, & y_i > \sigma(x_i) \end{cases} \tag{A.1}$$

which upon a large ensemble of throws will result in a distribution of events in x which, when normalised, approaches the true probability density function of x. The process of simulating data in this way is referred to as "generation", and a software package which is used to produce MC events in this way is a MC "generator".

Although this is a simple system, this method can be extended to any number of dimensions for which the probability of a discrete event occurring or not occurring can be calculated. In the T2K CCQE interaction generation, a Spectral Function which describes the event cross section at a given removal E_{rmv} (a function of the neutrino energy, final-state lepton energy, and bound nucleon energy) and the initial nucleon momentum, k, is used to draw a viable initial-state nucleon and is dependent on the underlying shell structure of the nucleus.

The law of conditional probability allows different parts of the model to be separated out and modelled separately. For instance, it is not necessary to simulate every stage of the T2K experiment from the proton bunches hitting the target and interacting, to the interaction of oscillated neutrinos at SK, in one long simulation.

The expected neutrino flux at each detector can be simulated from the simulation of the

[1] Often the term "the Monte Carlo" or more simply "the MC" is used as a shorthand for the simulated data produced using MC methods or MC generators. Though MC may refer to the method or the simulated data, generally when used as a noun it is referring to the data, and the methods will be explicitly referred to as "Monte Carlo methods" such as in a "Markov Chain Monte Carlo" (MCMC) methods.

hadronic production in the graphite target, and from the selection and decay of these hadrons to neutrinos. This simulated neutrino flux can then be used as one of the probability functions that acts as a criterion in the accept-reject process before throwing any of the other parameters. Accept-reject can then be used in stages, saving the accepted values of the parameters from one part of the simulation and throwing the remaining values instead, as below. The final probability of an event with some set of variables occurring is then given by

$$P(\text{event}) = P(E_\nu|\text{flux}) \times P(Q^2) \times P(\text{mode}) \times P(E_{\text{rmv}}, k|\text{SF}) \times P(\text{final state}|E_\nu, Q^2, E_{\text{rmv}}, k) \quad \text{(A.2)}$$

where: $P(E_\nu|\text{flux})$ is the probability of a neutrino having a given specific energy, E_ν, given the simulated flux spectrum; $P(Q^2)$ is the probability of specific value of 4-momentum transfer; $P(\text{mode})$ is the ratio of the specific interaction mode cross section to the total inclusive cross section defined in the generator; $P(E_{\text{rmv}}, k|\text{SF})$ is the probability of a specific combination of removal energy and nucleon momentum given the spectral function; and $P(\text{final state}|E_\nu, Q^2, E_{\text{rmv}}, k)$ is the probability of a specific set of final state kinematics given the other selected variables of the event, taking into account constraints on the parameters to be thrown. This allows the selection of specific event types of interest by setting some condition manually, such as the interaction mode, or a given neutrino energy. In these cases care must be taken to make sure the generator does not become stuck in an accept-reject loop. For example, if only CCQE events are generated at a specific value of E_{rmv}, but neutrino energies are allowed to be drawn from the regions of flux spectrum at energies which are lower than the sum of $E_{rmv} + m_l$, where m_l is the mass of the final state lepton, this will cause the generator will continually try to throw values of the other parameters until it can find an accepted event though this is not possible. This can be prevented with a loop counter and limit, but the limit should be sufficiently large as to not prevent modelling of events that, whilst unlikely, are still possible.

A.2 NEUT

Initially developed for the KamiokaNDE experiment, the NEUT MC event generator has been used by several neutrino experiments such as Super-K, K2K, SciBooNE, and T2K [115]. Currently maintained by Super-K and T2K collaborators, NEUT is the "in-house" generator in which T2K's Neutrino Interaction Working Group (NIWG) implements models. Written primarily in FORTRAN, NEUT is a library of generator programs which includes a range of interaction models for different neutrino energy regimes, target materials and interaction modes. NEUT underpins all of the T2K MC simulated data and is used to prepare many of the uncertainty treatments which are applied during fits to the data.

The T2K neutrino interaction model, as implemented in NEUT, is outlined in full by Hayato and Pickering in [115]. The model is factorised into stages and, with the exception of coherent interactions, treats the target nucleon and nuclear remnant separately. This gives some degree of independence between the nuclear model and the initial chosen kinematics of the event. These stages can broadly be separated into the following steps:

Interaction channel selection

After drawing a neutrino flavour l and energy E_ν from the beam flux spectra, and a target nucleus T from the detector geometry, an interaction channel, i, is assigned at random with a probability

$$P = \sigma_T^i(E_{\nu_l})/\sigma_T^{tot}(E_{\nu_l}) \tag{A.3}$$

equal to the ratio of the channel's cross section for the given neutrino flavour, nuclear target and neutrino energy, $\sigma_T^i(E_{\nu_l})$, to the total inclusive cross section, $\sigma_T^{tot}(E_{\nu_l})$, of that target. The total cross section of the nucleus is calculated as the sum of the cross sections for interactions on the constituent protons and neutrons

$$\sigma_T^i(E_{\nu_l}) = Z\sigma_p^i(E_{\nu_l}) + (A-Z)\sigma_n^i(E_{\nu_l}) \tag{A.4}$$

where A and Z are the nucleon and proton numbers of the target and σ_p^i and σ_n^i are the interaction cross sections on protons and neutrons, respectively.

Primary neutrino interaction

Once the neutrino, target and interaction mode have all been selected, the initial interaction of the neutrino, or the "hard scatter", can be simulated. Primary interactions are modelled using the impulse approximation, in which the neutrino interacts with an individual nucleon (or a bound pair of nucleons) rather than the nucleus as a whole; the nucleon and nuclear remnant are treated separately ignoring statistical correlations and final-state interactions [152]. A bound nucleus is chosen with initial-state and final-state kinematics drawn from the model if energetically allowed.

Propagation of final-state products through the nucleus

Once the hard-scatter is simulated, the final-state hadrons are propagated through the nuclear remnant with subsequent possible interactions handled by the NEUT cascade model. These are referred to as "Final-State Interactions" (FSI), in which hadrons can scatter, be produced or be absorbed whilst passing through the nucleus. Charged products may also be affected by the electrostatic Coulomb potential of the nucleus.

Nuclear remnant deexcitation

Finally, for interactions in oxygen, the nuclear remnant may be in an excited state and so deexcitations and energy released in these deexcitations must be simulated to take into account any subsequent photons produced which may be detected in a water-Cherenkov detector such as SK.

A.3 FLUKA and JNUBEAM

The FLUktuierende KAskade (FLUKA) 2011 MC simulation package [102][121] is used to simulate the production of hadrons in the T2K target when bombarded with protons. This is tuned to external data fits of replica-target hadron production data at the NA61/SHINE experiment [101][124], and simulated hadrons exiting the target were then propagated to the magnetic horns, decay volume and beam dump of the target hall using JNUBEAM, which uses the GEANT [116] geometry package to simulate interactions of the hadrons and their decay products. The GEometry ANd Tracking package [116] is a MC method library used to propagate events and particle tracks within a specified detector geometry. GEANT uses the flux profile over the volume of the detector, the geometry of the detector, the properties of the materials, and particle interaction cross sections to simulate where in the detector events occur, and then draws events from the output of NEUT to simulate interaction vertices within the detector. The final-state products are then propagated through the detector geometry and secondary interactions of the final-state products are then simulated. Energy deposition and detector response is simulated in a way that produces an output which is equivalent to that of the data. JNUBEAM and the SK detector simulation use GEANT3, whilst the ND280 detector geometry is simulated in GEANT4.

A.4 Non-generator Monte Carlo methods

MC methods are not exclusively used for event generation and can also be used to solve other mathematical problems. Numerical integration of complex high-dimensional problems may be far easier approached with Markov Chain Monte Carlo (MCMC) methods rather than trying to analytically calculate the integral or step through a grid of points, such as with a generalised trapezoid rule-like method. Another important other use of an MC method by T2K is that of the MCMC fitting method outlined in section 5.5.3 in which parameters of a model are thrown randomly and the parameter space is stepped through according to some likelihood criteria. In this sense MCMC fitters can be thought of as computing a numerical integration of the likelihood to build a posterior probability function by marginalising or integrating over nuisance parameters. Marginalisation and nuisance parameters are discussed in chapter 5.

A.5 Reconstruction and analysis

Once the initial processing of the data collected from ND280 and of the simulated detector response in the MC have been finalised, events must be processed into an analysable format. The simulated detector response, or "reconstructed information" (reco) is treated in the same way as the data, although the MC truth information is available. First the data and MC undergo calibration using the oaCalib[2] packages in which low-level variables are corrected according to collected calibration data from test beam or cosmic-ray muon data. After this, the data and MC information is reconstructed into tracks and vertices by oaRecon. This takes information such as charge deposits in scintillator bars or charge collected in the TPCs and uses clustering algorithms to reconstruct tracks (or showers in the ECal) which represent the path of particles in the detector. Likelihood fits are used to discriminate particle types based upon properties of the track such as charge deposited, curvature, and distance travelled. Once tracks are reconstructed, intersecting tracks, or track start and finish points, are used to reconstruct vertices which indicate the interactions of a particle. These can either be interactions which produce the particle, scatter it, or cause it to convert into other particles. Finally, this is passed to the oaAnalysis package which strips out unnecessary information and packages the events in a format that is more useful for high-level analysis.

A.6 Highland and psyche

Once the data and MC have been processed into an analyser-friendly format, physics analyses are performed using the High Level Analysis ND280 (`Highland`) package. `Highland` uses the MC truth and reco information to assess the purity and efficiency of specific event selections as well as evaluating the size and impact of systematic uncertainties on observables. This can then be used to conduct analyses such as cross-section measurements, or to prepare samples for the ND280 fit in the oscillation analysis which is the focus of this thesis. The details of the event selections in this work are described in chapter 5. `Highland` uses the ROOT Tree structure to place "cuts" on the data, excluding events which fail some criteria, and to create a selection that is a subset of the full set of ND280 events. The optimal cuts to maximise metrics such as efficiency $\times$ purity are found using various techniques such as Multi-Variate Analyses (MVA) or Boosted Decision Trees (BDTs). The output of these selections are stored in a reduced "FlatTree" or "MicroTree" format, which takes only the necessary information for the oscillation analysis from the ROOT Tree. `Highland` FlatTrees files are used to produce inputs to the ND280 fit. Once a selection is finalised

[2]Here the prefix 'oa' denotes off-axis, signifying these packages are for T2K's off-axis near detector, ND280. This is not to be confused with the use of 'OA' in later sections when used to refer to the Oscillation Analysis, in which the PMNS parameters are extracted from fits to the T2K data.

it can be ported to the Propagation of SYstematics and CHaractarisation (`psyche`) package which can be used to read the `Highland` FlatTree information, select the events based on cuts as developed in `Highland`, and apply detector systematic variations to the events, but without much of the overhead of the more detailed low-level analysis information that is accessed in `Highland`. The BANFF and MaCh3 fitters are able to load in events and apply systematic uncertainties through an interface to `psyche`. `psyche` is also used to generate the detector covariance matrix as described in subsection 5.10.1. The BANFF and MaCh3 fitters are described in more detail in chapters 5 and 6. The analysis outlined in this thesis uses `psyche` version v3r47 and `Highland` version v2r43.

A.7 Modifying the MC

Given the computational expense of running an entire production of MC from start to finish, running a production of MC for each combination of model parameters would be impossible. Even running these simulations a handful of times would use months of computing resources and so other methods must be used to generate MC at different values of the underlying model parameters. Two methods are used in T2K analyses to change the MC without regenerating from scratch. The first is "reweighting", which changes the number of events and is used to change the cross-section or probability of an event occurring. The second is the application of an observable "varier", which changes the property of an event which is used to compare it to data, such as the kinematics.

A.7.1 Reweighting

When generating MC, the events are often "weighted" to reflect the amount of data that is expected to be collected. For example, a flat weight is applied to all MC events to scale the simulated POT in the MC to the recorded POT-equivalent data collected at the detectors which pass data quality cuts. Other weights are also applied such as the nominal flux weights which tune the MC to reflect fits to external hadron production data from NA61/SHINE or in situ beam monitor and INGRID measurements. In a fit, however, parameter values need to be varied and they are "reweighted" from the weight calculated at the nominal or prior central value of the parameters. To do this the fitter machinery must interface in some way to reweighting libraries which apply these reweights based on generated MC at alternative model parameter values. Weights typically correspond to a bin of some physical quantity, such as true neutrino energy, E_ν, or four-momentum transfer, Q^2, and templates are made from MC generated at alternative values of the parameter. The weight, $w(p,x)$, which is dependent on physical quantity x, can then be calculated for the alternative value,

p, of the parameter as

$$w(p,x) = \frac{N_x(p)}{N_x(p_{\text{gen}})} \tag{A.5}$$

where $N_x(p)$ is the number of MC events in a given bin of x at parameter value p, and $N_x(p_{\text{gen}})$ is the number of MC events in that same bin of x at the generated value of the parameter p_{gen}. The reweighted number of events N_{RW} at some value of the parameter, p, in a fitted observable bin then becomes

$$N_{RW}(p) = \sum_{i}^{N} w_i(p,x_i) \tag{A.6}$$

where N is the number of unweighted, generated MC events, and w_i is the weight of the ith event.

T2K's reweighting libraries used for the oscillation analysis are NEUTReWeight and NIWGReWeight which are interfaced by the weight file generation executables through the T2KReWeight wrapper. The `genWeightsFromNRooTracker_BANFF_2020` app was used for the 2020 oscillation analysis to produce weights for the near detector fit by both the BANFF group and the MaCh3 group. This app stores all of the relevant information from the processed Highland2 flat trees needed for the ND280 fit. Events are loaded using `psyche` in which nominal flux and total detector weights are stored for each event, then the weights for complex response parameters are calculated at different values of the parameter, which can be "splined" using a third-order polynomial spline to interpolate between the precalculated values. Typically these variations are integer multiples of the standard deviation, σ, of the parameter. This allows for cross-validations of the weights to be calculated exactly without any interpolation effects at these intervals and compared between the two fitters. These weights are then stored as a ROOT TGraph for every parameter that has a "splined" response, for each event, along with the reduced `psyche` FlatTree. The fitter then uses a `psyche` interface to load in the events before matching the parameters to the appropriate events and converting the TGraphs into splines using the ROOT TSpline3 class. The splines ensure that the parameter response varies smoothly continuously. This allows continuous interpolation of the weights to be applied by the fitter at any value along the spline's range.

A.7.2 Observable variers

In addition to reweighting events, another way of adjusting the MC is through the use of an observable varier. These can be constructed in a similar way to reweights, using the underlying MC to estimate a change in value to some underlying physical or measured variable of the event in order to give freedom to cover an uncertainty, or to apply a correction, which does not directly affect the probability of an event by changing the size of the cross section. In a similar manner to the interpolation of weights for events in between precalculated values using splines for cross-section variations, variations of the observables are then calculated from interpolation splines using val-

ues from templates in the desired observables and variables. These are handled by the reweight packages although they are not reweight parameters. This method is used for the treatment of the nucleon removal energy, E_{rmv}, in CCQE interactions. For the E_{rmv} lepton-momentum varier, templates are generated by profiling MC distributions in E_ν, θ_l and p_l into E_ν and θ_l to find the average value of p_l for a bin in E_ν and θ_l at some value of E_{rmv}, and then using the difference between these templates to calculate the shift in momentum Δp_l. Splines can then be generated for each bin in E_ν and θ_l, interpolating between knots corresponding to the values of E_{rmv} at which the template MC was generated.

Appendix B

ND280 sample data

B.1 Efficiency and Purity

Beam	Topology	Target	Eff. (%)	Pur. (%)
FHC ν_μ	0π	FGD1	48.0	71.3
		FGD2	48.0	68.2
	$1\pi^+$	FGD1	29.0	52.4
		FGD2	24.0	51.2
	Other	FGD1	30.0	71.4
		FGD2	30.0	71.2
RHC$\bar{\nu}_\mu$	0π	FGD1	70.0	74.5
		FGD2	69.0	72.7
	$1\pi^-$	FGD1	19.3	45.4
		FGD2	17.2	41.0
	Other	FGD1	26.5	26.3
		FGD2	24.2	26.0
RHC ν_μ	0π	FGD1	60.3	55.9
		FGD2	60.3	52.8
	$1\pi^+$	FGD1	30.3	44.4
		FGD2	26.0	44.8
	Other	FGD1	27.4	68.3
		FGD2	27.1	69.5

Table B.1: Efficiencies and purities for each of the selections at ND280 included in this analysis for their defined target interaction mode.

Beam	Topology	Target	Uncertainty (%)
FHC	0π	FGD1	1.20
		FGD2	1.40
	$1\pi^+$	FGD1	2.65
		FGD2	2.57
	Other	FGD1	2.33
		FGD2	2.19
RHC	0π	FGD1	1.96
		FGD2	2.08
	$1\pi^-$	FGD1	4.04
		FGD2	3.63
	Other	FGD1	3.61
		FGD2	3.23
ν_μ in RHC	0π	FGD1	1.61
		FGD2	1.76
	$1\pi^+$	FGD1	3.00
		FGD2	2.72
	Other	FGD1	2.35
		FGD2	2.35

Table B.2: Uncertainties on the total number of events in the ND280 selections from detector systematics, broken down by selection

Appendix C

Values of the flux and cross-section parameters

The values and uncertainties for each of the Super-K flux (tables C.1 and C.2) and T2K interaction cross-section systematics split into CCQE and 2p2h (table C.3), single pion production (table C.4), Multi-π/DIS processes (table C.5), FSI (table C.7) and other (table C.6) in the T2K 2020 analysis before and after the BANFF fit to the ND280 data. The uncertainty given is taken from the diagonal element of the covariance matrix. Carbon-only interaction systematics are not propagated to the P-Theta fit to Super-K data but are presented here for completeness. These parameters are only valid in the context of the exact model used in this analysis, and the context shouldn't be interpreted as a "global" constraint on the parameters in NEUT or the models in general.

Flux parameter	Prefit	Postfit
SK ν_μ [0.0, 0.4]	1.00 ± 0.07	1.11 ± 0.05
SK ν_μ [0.4, 0.5]	1.00 ± 0.06	1.10 ± 0.04
SK ν_μ [0.5, 0.6]	1.00 ± 0.05	1.08 ± 0.03
SK ν_μ [0.6, 0.7]	1.00 ± 0.05	1.07 ± 0.03
SK ν_μ [0.7, 1.0]	1.00 ± 0.07	1.06 ± 0.04
SK ν_μ [1.0, 1.5]	1.00 ± 0.07	1.03 ± 0.04
SK ν_μ [1.5, 2.5]	1.00 ± 0.06	1.03 ± 0.04
SK ν_μ [2.5, 3.5]	1.00 ± 0.07	1.01 ± 0.04
SK ν_μ [3.5, 5.0]	1.00 ± 0.09	0.98 ± 0.04
SK ν_μ [5.0, 7.0]	1.00 ± 0.10	0.92 ± 0.04
SK ν_μ [7.0, ∞]	1.00 ± 0.12	0.91 ± 0.04
SK $\bar{\nu}_\mu$ [0.0, 0.7]	1.00 ± 0.09	1.06 ± 0.08
SK $\bar{\nu}_\mu$ [0.7, 1.0]	1.00 ± 0.06	1.04 ± 0.05
SK $\bar{\nu}_\mu$ [1.0, 1.5]	1.00 ± 0.07	1.04 ± 0.06
SK $\bar{\nu}_\mu$ [1.5, 2.5]	1.00 ± 0.08	1.05 ± 0.07
SK $\bar{\nu}_\mu$ [2.5, ∞]	1.00 ± 0.08	1.04 ± 0.06
SK ν_e [0.0, 0.5]	1.00 ± 0.06	1.09 ± 0.04
SK ν_e [0.5, 0.7]	1.00 ± 0.05	1.08 ± 0.04
SK ν_e [0.7, 0.8]	1.00 ± 0.05	1.06 ± 0.04
SK ν_e [0.8, 1.5]	1.00 ± 0.06	1.04 ± 0.04
SK ν_e [1.5, 2.5]	1.00 ± 0.08	1.00 ± 0.04
SK ν_e [2.5, 4.0]	1.00 ± 0.09	0.98 ± 0.04
SK ν_e [4.0, ∞]	1.00 ± 0.09	0.98 ± 0.05
SK $\bar{\nu}_e$ [0.0, 2.5]	1.00 ± 0.10	1.02 ± 0.09
SK $\bar{\nu}_e$ [2.5, ∞]	1.00 ± 0.13	1.09 ± 0.11

Table C.1: SK ν-beam (FHC) flux parameters before and after the BANFF fit to the ND80 data, including the uncertainty from the diagonal of the covariance matrix. The values in brackets show the range of E_ν (GeV) for each parameter.

Flux parameter	Prefit	Postfit
SK ν_μ [0.0, 0.7]	1.00 ± 0.09	1.11 ± 0.06
SK ν_μ [0.7, 1.0]	1.00 ± 0.06	1.07 ± 0.05
SK ν_μ [1.0, 1.5]	1.00 ± 0.06	1.07 ± 0.04
SK ν_μ [1.5, 2.5]	1.00 ± 0.07	1.07 ± 0.04
SK ν_μ [2.5, ∞]	1.00 ± 0.07	1.02 ± 0.04
SK $\bar{\nu}_\mu$ [0.0, 0.4]	1.00 ± 0.07	1.09 ± 0.05
SK $\bar{\nu}_\mu$ [0.4, 0.5]	1.00 ± 0.06	1.09 ± 0.04
SK $\bar{\nu}_\mu$ [0.5, 0.6]	1.00 ± 0.06	1.07 ± 0.04
SK $\bar{\nu}_\mu$ [0.6, 0.7]	1.00 ± 0.05	1.06 ± 0.03
SK $\bar{\nu}_\mu$ [0.7, 1.0]	1.00 ± 0.08	1.09 ± 0.04
SK $\bar{\nu}_\mu$ [1.0, 1.5]	1.00 ± 0.08	1.06 ± 0.04
SK $\bar{\nu}_\mu$ [1.5, 2.5]	1.00 ± 0.06	1.01 ± 0.04
SK $\bar{\nu}_\mu$ [2.5, 3.5]	1.00 ± 0.07	1.01 ± 0.05
SK $\bar{\nu}_\mu$ [3.5, 5.0]	1.00 ± 0.09	0.95 ± 0.06
SK $\bar{\nu}_\mu$ [5.0, 7.0]	1.00 ± 0.09	0.95 ± 0.06
SK $\bar{\nu}_\mu$ [7.0, ∞]	1.00 ± 0.12	0.93 ± 0.09
SK ν_e [0.0, 2.5]	1.00 ± 0.09	1.03 ± 0.07
SK ν_e [2.5, ∞]	1.00 ± 0.08	1.03 ± 0.07
SK $\bar{\nu}_e$ [0.0, 0.5]	1.00 ± 0.06	1.08 ± 0.04
SK $\bar{\nu}_e$ [0.5, 0.7]	1.00 ± 0.05	1.07 ± 0.04
SK $\bar{\nu}_e$ [0.7, 0.8]	1.00 ± 0.06	1.06 ± 0.04
SK $\bar{\nu}_e$ [0.8, 1.5]	1.00 ± 0.06	1.04 ± 0.04
SK $\bar{\nu}_e$ [1.5, 2.5]	1.00 ± 0.08	1.01 ± 0.06
SK $\bar{\nu}_e$ [2.5, 4.0]	1.00 ± 0.09	1.01 ± 0.07
SK $\bar{\nu}_e$ [4.0, ∞]	1.00 ± 0.15	1.09 ± 0.13

Table C.2: SK $\bar{\nu}$-beam (RHC) flux parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. The values in brackets show the range of E_ν (GeV) for each parameter.

Parameter	Prefit	Postfit	Comment
M_A^{QE} (GeV/c^2)	1.03 ± 0.06	1.17 ± 0.04	
$Q^2 < 0.05$ GeV^2	0.50 ± 0.17	0.78 ± 0.05	
$0.05 < Q^2 < 0.10$ GeV^2	0.70 ± 0.15	0.89 ± 0.04	
$0.10 < Q^2 < 0.15$ GeV^2	0.78 ± 0.13	1.03 ± 0.05	Norm. on true CCQE events in true Q^2.
$0.15 < Q^2 < 0.20$ GeV^2	0.89 ± 0.15	1.03 ± 0.08	
$0.20 < Q^2 < 0.25$ GeV^2	0.93 ± 0.16	1.09 ± 0.10	
$0.25 < Q^2 < 0.50$ GeV^2	1.00 ± 0.11	1.26 ± 0.06	
$0.50 < Q^2 < 1.00$ GeV^2	1.00 ± 0.18	1.14 ± 0.08	
$Q^2 > 1.00$ GeV^2	1.00 ± 0.40	1.26 ± 0.14	
$\Delta E_{rmv}^{\nu,\mathrm{C}} \nu$ (MeV)	2.00 ± 6.00	-2.38 ± 1.75	
$\Delta E_{rmv}^{\bar{\nu},\mathrm{C}} \bar{\nu}$ (MeV)	0.00 ± 6.00	1.64 ± 1.93	
$\Delta E_{rmv}^{\nu,\mathrm{O}} \nu$ (MeV)	4.00 ± 6.00	2.55 ± 3.08	
$\Delta E_{rmv}^{\bar{\nu},\mathrm{O}} \bar{\nu}$ (MeV)	0.00 ± 6.00	-1.26 ± 3.19	
2p2h norm. ν	1.00 ± 1.00	1.06 ± 0.15	
2p2h norm. $\bar{\nu}$	1.00 ± 1.00	0.72 ± 0.16	
2p2h norm. C$\rightarrow$O	1.00 ± 0.20	1.05 ± 0.15	
2p2h shape C	0.00 ± 3.00	0.97 ± 0.46	-1 is non-Δ-like, 0 is Valencia [38], +1 is Δ-like. +1 is Valencia-like [38], 0 is Martini-like [128]. Not fit at ND.
2p2h shape O	0.00 ± 3.00	0.00 ± 0.17	
2p2h low-E_ν ν	1.00 ± 1.00	1.00 ± 1.00	
2p2h high-E_ν ν	1.00 ± 1.00	1.00 ± 1.00	
2p2h low-E_ν $\bar{\nu}$	1.00 ± 1.00	1.00 ± 1.00	
2p2h high-E_ν $\bar{\nu}$	1.00 ± 1.00	1.00 ± 1.00	

Table C.3: CCQE and 2p2h cross-section parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. The parameters are detailed in section 5.8.

Parameter	Prefit	Postfit	Comment
C_5^A	0.96 ± 0.15	0.98 ± 0.06	
M_A^{RES} (GeV/c^2)	1.07 ± 0.15	0.79 ± 0.05	
$I_{1/2}$ non-res norm. low-p_π $\bar{\nu}_\mu$	0.96 ± 0.96	0.96 ± 0.96	Not fit at ND.
$I_{1/2}$ non-res norm.	0.96 ± 0.40	0.87 ± 0.23	
CC coh. C norm.	1.00 ± 0.30	0.61 ± 0.22	
CC coh. O norm.	1.00 ± 0.30	0.61 ± 0.22	

Table C.4: Single pion production cross-section parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. The parameters are detailed in section 5.8.

Parameter	Prefit	Postfit	Comment
CC Bodek-Yang on/off DIS	0.00 ± 1.00	1.04 ± 0.19	+1 is B-Y supp. off,
CC Bodek-Yang on/off multi-π	0.00 ± 1.00	-0.03 ± 0.18	0 is B-Y supp. on.
CC multiplicity multi-π	0.00 ± 1.00	0.14 ± 0.71	+1 is AGKY-like, 0 is NEUT-like.
CC misc. norm.	1.00 ± 1.00	2.28 ± 0.43	
CC DIS+multi-π norm. ν	1.00 ± 0.04	1.06 ± 0.03	
CC DIS+multi-π norm. $\bar{\nu}$	1.00 ± 0.07	0.94 ± 0.06	

Table C.5: Multi-π and DIS cross-section parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. These control Bodek-Yang corrections, and cover some of the expected differences between the NEUT and AGKY model [133]. The parameters are detailed in section 5.8.

Parameter	Prefit	Postfit	Comment
Coulomb corr. ν	1.00 ± 0.02	1.00 ± 0.02	
Coulomb corr. $\bar{\nu}$	1.00 ± 0.01	1.00 ± 0.01	
ν_e/ν_μ norm.	1.00 ± 0.03	1.00 ± 0.03	No ND selection
$\bar{\nu}_e/\bar{\nu}_\mu$ norm.	1.00 ± 0.03	1.00 ± 0.03	
NC coh. norm.	1.00 ± 0.30	1.02 ± 0.30	No ND selection
NC 1γ norm.	1.00 ± 1.00	1.00 ± 1.00	
NC other ND norm.	1.00 ± 0.30	1.66 ± 0.13	Not propagated to FD.
NC other FD norm.	1.00 ± 0.30	1.00 ± 0.30	Not fit at ND.

Table C.6: NC and miscellaneous cross-section parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. The parameters are detailed in section 5.8.

Parameter	Prefit	Postfit	Comment
Pion FSI Quasi-Elastic	1.00 ± 0.29	0.83 ± 0.09	Scaling of pion scattering probabilities relative to NEUT.
Pion FSI Quasi-Elastic $p_\pi > 500$ MeV/c	1.00 ± 0.47	0.75 ± 0.16	
Pion FSI Inelastic	1.00 ± 1.10	1.71 ± 0.31	
Pion FSI Absorption	1.00 ± 0.31	1.19 ± 0.12	
Pion FSI Charge Exchange	1.00 ± 0.44	0.78 ± 0.34	

Table C.7: Pion Final State Interaction (FSI) parameters before and after the BANFF fit to the ND280 data, including the uncertainty from the diagonal of the covariance matrix. The parameters are detailed in section 5.8.

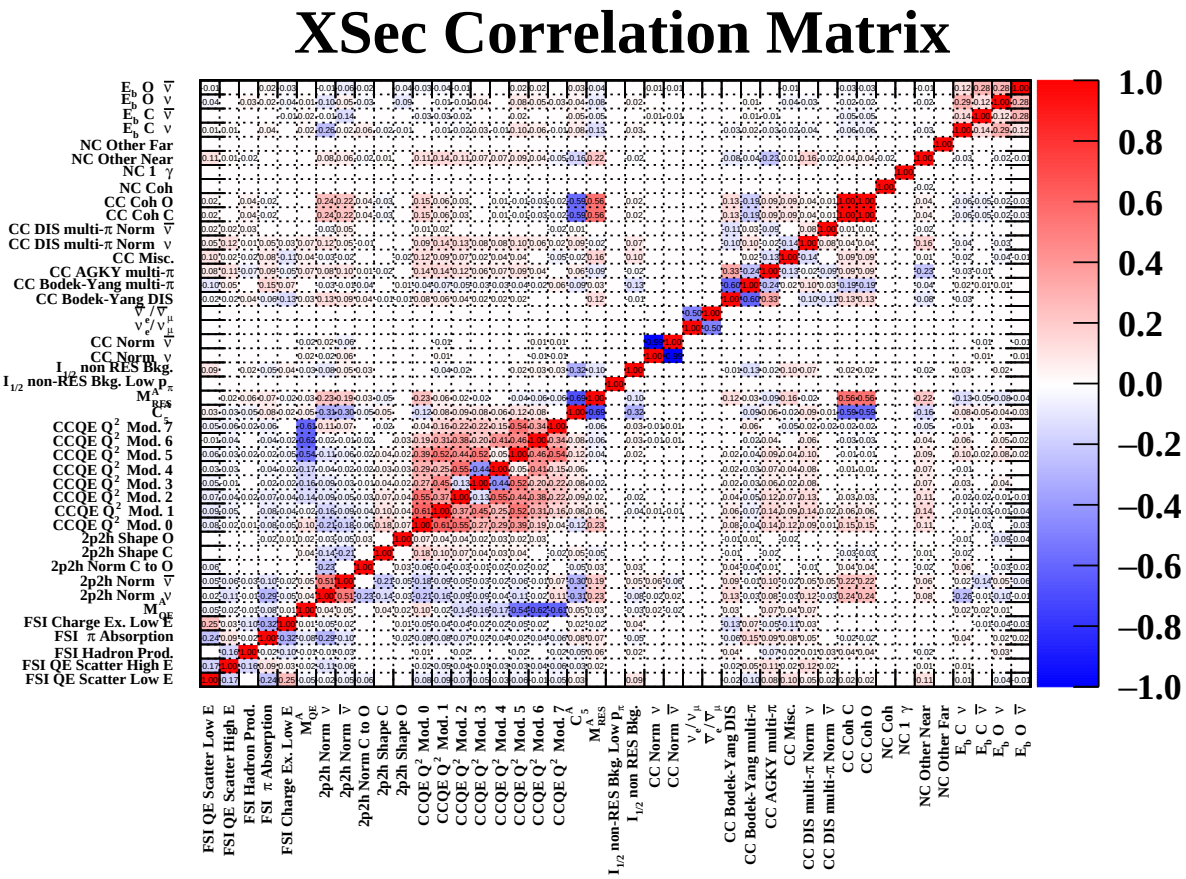

Figure C.1: Postfit correlation matrix for the cross-section parameters. The most significant features of the CCQE model are the correlations introduced between the previously uncorrelated Q^2 normalisation parameters. There is an increasingly strong anti-correlation between these and M_A^{QE}. This, along with the significant correlations of these parameters with the flux parameters, shows how M_A^{QE} is able to move several standard deviations away from its prior central value through the large number of associate degrees of freedom in both E_ν and Q^2. A clear indication of the (anti)correlations which can be introduced between signal and background processes can be seen with the resonant and coherent single-pion production model parameters. Similarly, the Q^2 parameters have some correlations with the Multi-π and DIS as the fit suppresses the CCQE contribution to the high E_ν region of the samples. Whilst the nucleon removal energy parameters have strongly correlated priors, the postfit correlations are weaker suggesting more independence, though this may be expected due to the separation of target and neutrino beam mode in the different samples. Correlations are introduced between the removal energy parameters and the 2p2h shape and pion production models as the are able to change the shape of the CCQE contribution beyond that of the Q^2 freedoms.

Appendix D

Event rates for 1σ variations of the ΔE_{rmv} parameters

This appendix shows the full 1σ variations for the cross-section parameters for both BANFF and MaCh3 frameworks in terms of the effect on the total number of events in each of the fitted samples.

The differences found are sub-percent, in agreement with previous analyses. In the past, these have been attributed to MaCh3 using a GPU (with 32-bit precision) to calculate the cross-section weights, whereas BANFF uses a CPU (with 64-bit precision), and the methods used to multiply multiple weights together.

Parameter	1σ *value*	*Sample*	*BANFF* −1σ	MaCh3 −1σ	*BANFF nominal*	*MaCh3 nominal*	*BANFF* +1σ	*MaCh3* +1σ	(*BANFF - MaCh3 / BANFF*) −1σ
$\Delta E_{\mathrm{rmv}}^{\nu,\mathrm{C}}$	***6.00***	FGD1 ν_μ $CC0\pi$	27953.72	27951.10	27951.80	27951.10	27950.78	27951.10	0.000094
		FGD1 ν_μ $CC1\pi$	8359.41	8358.97	8358.62	8358.97	8358.97	8358.97	0.000053
		FGD1 ν_μ CCother	7032.80	7031.47	7031.49	7031.47	7032.40	7031.47	0.000189
		FGD2 ν_μ $CC0\pi$	27558.07	27556.20	27558.87	27556.20	27562.82	27556.20	0.000068
		FGD2 ν_μ $CC1\pi$	6724.25	6723.98	6723.79	6723.98	6723.73	6723.98	0.000041
		FGD2 ν_μ CCOther	6454.78	6454.68	6454.74	6454.68	6455.72	6454.68	0.000015
		FGD1 anti-ν_μ $CC0\pi$	7270.61	7270.56	7270.33	7270.56	7269.91	7270.56	0.000007
		FGD1 anti-ν_μ $CC1\pi$	694.32	694.32	694.32	694.32	694.33	694.32	0.000005
		FGD1 anti-ν_μ CCOther	1286.73	1286.78	1286.79	1286.78	1286.78	1286.78	0.000036
		FGD2 anti-ν_μ $CC0\pi$	7036.29	7036.71	7036.50	7036.71	7036.14	7036.71	0.000060
		FGD2 anti-ν_μ $CC1\pi$	624.88	624.76	624.69	624.76	624.72	624.76	0.000195
		FGD2 anti-ν_μ CCOther	1176.47	1176.62	1176.53	1176.62	1176.53	1176.62	0.000129
		FGD1 ν_μBkg $CC0\pi$ in AntiNu Mode	3035.99	3035.85	3035.59	3035.85	3035.16	3035.85	0.000045
		FGD1 ν_μBkg $CC1\pi$ in AntiNu Mode	1159.04	1159.02	1159.01	1159.02	1159.06	1159.02	0.000016
		FGD1 ν_μBkg CCOther in AntiNu Mode	1073.66	1073.16	1073.13	1073.16	1073.26	1073.16	0.000466
		FGD2 ν_μBkg $CC0\pi$ in AntiNu Mode	3012.63	3013.01	3012.40	3013.01	3012.37	3013.01	0.000127
		FGD2 ν_μBkg $CC1\pi$ in AntiNu Mode	930.65	930.64	930.64	930.64	930.59	930.64	0.000010
		FGD2 ν_μBkg CCOther in AntiNu Mode	1000.37	1000.03	1000.39	1000.03	1000.41	1000.03	0.000343

Table D.1: Event rate broken by sample for one σ variation of the carbon ν ΔE_{rmv} parameter. MaCh3 and BANFF are presented for comparison.

Parameter	1σ *value*	*Sample*	*BANFF* −1σ	MaCh3 −1σ	*BANFF nominal*	*MaCh3 nominal*	*BANFF* +1σ	*MaCh3* +1σ	(*BANFF - MaCh3 / BANFF*) −1σ
$\Delta E_{\mathrm{rmv}}^{\bar{\nu},\mathrm{C}}$	***6.00***	FGD1 ν_μ $CC0\pi$	27951.77	27951.10	27951.80	27951.10	27951.71	27951.10	0.000024
		FGD1 ν_μ $CC1\pi$	8358.62	8358.97	8358.62	8358.97	8358.62	8358.97	0.000042
		FGD1 ν_μ CCother	7031.49	7031.47	7031.49	7031.47	7031.49	7031.47	0.000003
		FGD2 ν_μ $CC0\pi$	27558.76	27556.20	27558.87	27556.20	27558.85	27556.20	0.000093
		FGD2 ν_μ $CC1\pi$	6723.79	6723.98	6723.79	6723.98	6723.79	6723.98	0.000028
		FGD2 ν_μ CCOther	6454.74	6454.68	6454.74	6454.68	6454.74	6454.68	0.000010
		FGD1 anti-ν_μ $CC0\pi$	7271.18	7270.56	7270.33	7270.56	7270.02	7270.56	0.000086
		FGD1 anti-ν_μ $CC1\pi$	694.27	694.32	694.32	694.32	694.32	694.32	0.000070
		FGD1 anti-ν_μ CCOther	1286.98	1286.78	1286.79	1286.78	1286.79	1286.78	0.000158
		FGD2 anti-ν_μ $CC0\pi$	7037.13	7036.71	7036.50	7036.71	7036.99	7036.71	0.000059
		FGD2 anti-ν_μ $CC1\pi$	624.69	624.76	624.69	624.76	624.73	624.76	0.000103
		FGD2 anti-ν_μ CCOther	1177.04	1176.62	1176.53	1176.62	1176.53	1176.62	0.000355
		FGD1 ν_μBkg $CC0\pi$ in AntiNu Mode	3035.59	3035.85	3035.59	3035.85	3035.45	3035.85	0.000087
		FGD1 ν_μBkg $CC1\pi$ in AntiNu Mode	1159.01	1159.02	1159.01	1159.02	1159.01	1159.02	0.000005
		FGD1 ν_μBkg CCOther in AntiNu Mode	1073.13	1073.16	1073.13	1073.16	1073.13	1073.16	0.000024
		FGD2 ν_μBkg $CC0\pi$ in AntiNu Mode	3012.40	3013.01	3012.40	3013.01	3012.31	3013.01	0.000202
		FGD2 ν_μBkg $CC1\pi$ in AntiNu Mode	930.64	930.64	930.64	930.64	930.64	930.64	0.000005
		FGD2 ν_μBkg CCOther in AntiNu Mode	1000.39	1000.03	1000.39	1000.03	1000.39	1000.03	0.000358

Table D.2: Event rate broken by sample for one σ variation of the carbon $\bar{\nu}$ ΔE_{rmv} parameter. MaCh3 and BANFF are presented for comparison.

Parameter	1σ *value*	*Sample*	*BANFF* −1σ	MaCh3 −1σ	*BANFF nominal*	*MaCh3 nominal*	*BANFF* +1σ	*MaCh3* +1σ	(*BANFF - MaCh3 / BANFF*) −1σ
$\Delta E_{\mathrm{rmv}}^{\nu,\mathrm{O}}$	***6.00***	FGD1 ν_μ $CC0\pi$	27951.72	27951.10	27951.80	27951.10	27953.71	27951.10	0.000022
		FGD1 ν_μ $CC1\pi$	8358.86	8358.97	8358.62	8358.97	8358.95	8358.97	0.000013
		FGD1 ν_μ CCother	7031.43	7031.47	7031.49	7031.47	7031.50	7031.47	0.000006
		FGD2 ν_μ $CC0\pi$	27557.96	27556.20	27558.87	27556.20	27560.69	27556.20	0.000064
		FGD2 ν_μ $CC1\pi$	6724.10	6723.98	6723.79	6723.98	6723.81	6723.98	0.000018
		FGD2 ν_μ CCOther	6455.34	6454.68	6454.74	6454.68	6455.92	6454.68	0.000103
		FGD1 anti-ν_μ $CC0\pi$	7270.33	7270.56	7270.33	7270.56	7270.33	7270.56	0.000031
		FGD1 anti-ν_μ $CC1\pi$	694.32	694.32	694.32	694.32	694.32	694.32	0.000000
		FGD1 anti-ν_μ CCOther	1286.79	1286.78	1286.79	1286.78	1286.79	1286.78	0.000006
		FGD2 anti-ν_μ $CC0\pi$	7036.58	7036.71	7036.50	7036.71	7036.48	7036.71	0.000019
		FGD2 anti-ν_μ $CC1\pi$	624.75	624.76	624.69	624.76	624.68	624.76	0.000019
		FGD2 anti-ν_μ CCOther	1176.53	1176.62	1176.53	1176.62	1176.53	1176.62	0.000073
		FGD1 ν_μBkg $CC0\pi$ in AntiNu Mode	3035.74	3035.85	3035.59	3035.85	3035.60	3035.85	0.000035
		FGD1 ν_μBkg $CC1\pi$ in AntiNu Mode	1159.01	1159.02	1159.01	1159.02	1159.01	1159.02	0.000005
		FGD1 ν_μBkg CCOther in AntiNu Mode	1073.13	1073.16	1073.13	1073.16	1073.13	1073.16	0.000024
		FGD2 ν_μBkg $CC0\pi$ in AntiNu Mode	3013.18	3013.01	3012.40	3013.01	3012.49	3013.01	0.000055
		FGD2 ν_μBkg $CC1\pi$ in AntiNu Mode	930.63	930.64	930.64	930.64	930.64	930.64	0.000013
		FGD2 ν_μBkg CCOther in AntiNu Mode	1000.01	1000.03	1000.39	1000.03	1000.60	1000.03	0.000021

Table D.3: Event rate broken by sample for one σ variation of the oxygen ν ΔE_{rmv} parameter. MaCh3 and BANFF are presented for comparison.

Parameter	1σ *value*	*Sample*	*BANFF* −1σ	MaCh3 −1σ	*BANFF nominal*	*MaCh3 nominal*	*BANFF* +1σ	*MaCh3* +1σ	(*BANFF - MaCh3 / BANFF*) −1σ
$\Delta E_{\mathrm{rmv}}^{\bar{\nu},\mathrm{O}}$	***6.00***	FGD1 ν_μ $CC0\pi$	27951.79	27951.10	27951.80	27951.10	27951.74	27951.10	0.000025
		FGD1 ν_μ $CC1\pi$	8358.62	8358.97	8358.62	8358.97	8358.62	8358.97	0.000042
		FGD1 ν_μ CCother	7031.49	7031.47	7031.49	7031.47	7031.49	7031.47	0.000003
		FGD2 ν_μ $CC0\pi$	27558.87	27556.20	27558.87	27556.20	27558.80	27556.20	0.000097
		FGD2 ν_μ $CC1\pi$	6723.79	6723.98	6723.79	6723.98	6723.75	6723.98	0.000029
		FGD2 ν_μ CCOther	6454.74	6454.68	6454.74	6454.68	6454.74	6454.68	0.000010
		FGD1 anti-ν_μ $CC0\pi$	7271.51	7270.56	7270.33	7270.56	7271.41	7270.56	0.000131
		FGD1 anti-ν_μ $CC1\pi$	694.32	694.32	694.32	694.32	694.32	694.32	0.000001
		FGD1 anti-ν_μ CCOther	1286.60	1286.78	1286.79	1286.78	1286.78	1286.78	0.000138
		FGD2 anti-ν_μ $CC0\pi$	7036.77	7036.71	7036.50	7036.71	7037.22	7036.71	0.000008
		FGD2 anti-ν_μ $CC1\pi$	624.69	624.76	624.69	624.76	624.69	624.76	0.000108
		FGD2 anti-ν_μ CCOther	1176.39	1176.62	1176.53	1176.62	1176.58	1176.62	0.000192
		FGD1 ν_μBkg $CC0\pi$ in AntiNu Mode	3035.59	3035.85	3035.59	3035.85	3035.59	3035.85	0.000086
		FGD1 ν_μBkg $CC1\pi$ in AntiNu Mode	1159.01	1159.02	1159.01	1159.02	1159.01	1159.02	0.000005
		FGD1 ν_μBkg CCOther in AntiNu Mode	1073.13	1073.16	1073.13	1073.16	1073.13	1073.16	0.000024
		FGD2 ν_μBkg $CC0\pi$ in AntiNu Mode	3012.36	3013.01	3012.40	3013.01	3012.49	3013.01	0.000217
		FGD2 ν_μBkg $CC1\pi$ in AntiNu Mode	930.64	930.64	930.64	930.64	930.64	930.64	0.000005
		FGD2 ν_μBkg CCOther in AntiNu Mode	1000.39	1000.03	1000.39	1000.03	1000.39	1000.03	0.000358

Table D.4: Event rate broken by sample for one σ variation of the oxygen $\bar{\nu}$ ΔE_{rmv} parameter. MaCh3 and BANFF are presented for comparison.

Appendix E

Prefit and postfit event rates at ND280

This appendix shows the prefit (left) and postfit (right) predictions, broken down by interaction mode, for each of the 18 ND280 samples, as well as the ratio of the data to the MC prediction before and after the BANFF ND280 fit. Samples included in the fit are binned in the reconstructed momentum and cosine of the reconstructed angle between the muon and the detector axis. The prefit underestimates the ν_μ CC0π rate in both FGDs for both horn currents, which drives the most significant changes between the prefit and postfit models. There is also an underestimation of the ν_μ CC-Other samples, though these have less statistical power to impact the fit. The ν_μ CC1π samples are slightly over estimated in comparison. The RHC $\bar{\nu}_\mu$ samples are, with the exception of CC-Other, closer to the data than the RHC ν_μ samples.

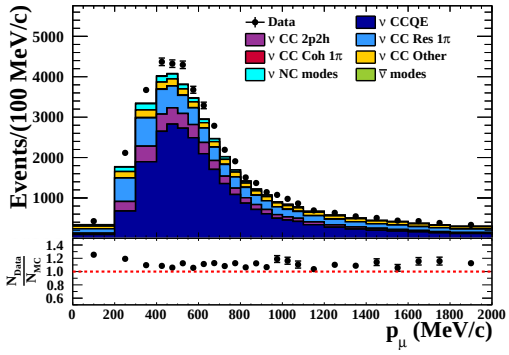

(a1) Prefit

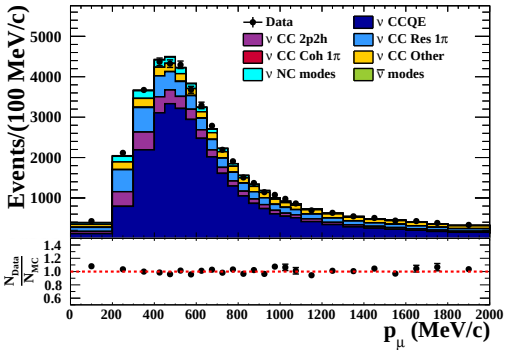

(a2) Postfit

(a) Reconstructed muon momentum

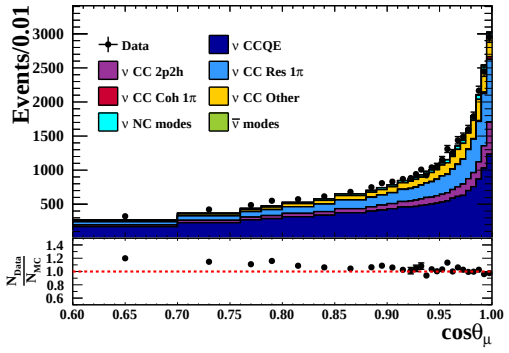

(b1) Prefit

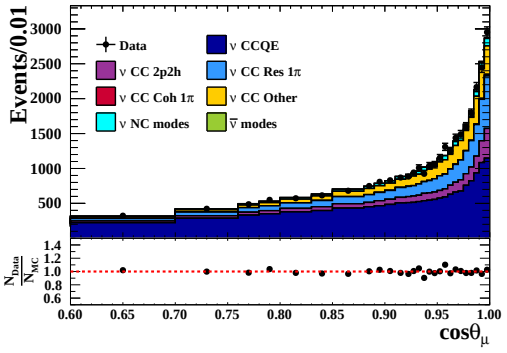

(b2) Postfit

(b) Cosine of reconstructed muon angle

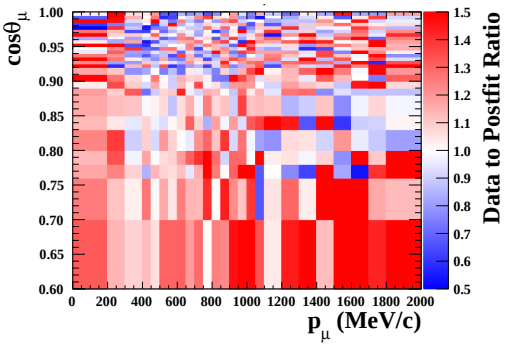

(c1) Prefit

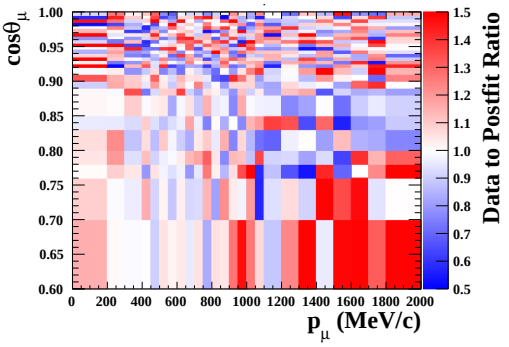

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.1: FHC FGD1 ν_μ CC0π

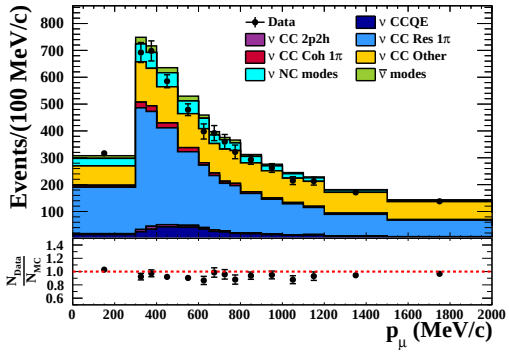

(a1) Prefit

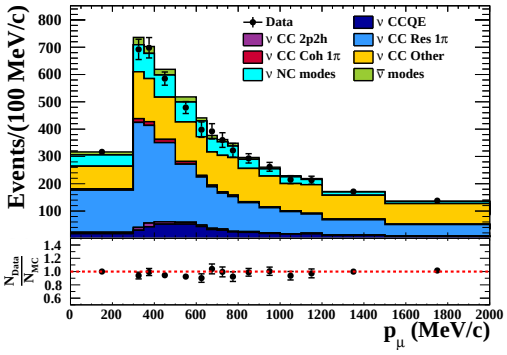

(a2) Postfit

(a) Reconstructed muon momentum

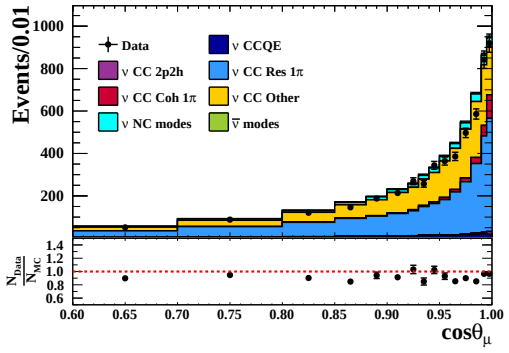

(b1) Prefit

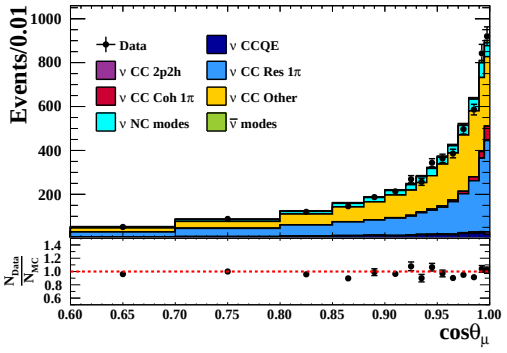

(b2) Postfit

(b) Cosine of reconstructed muon angle

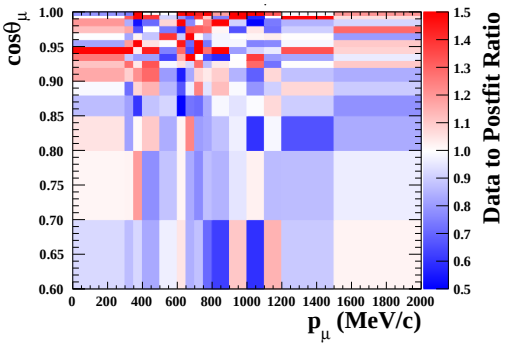

(c1) Prefit

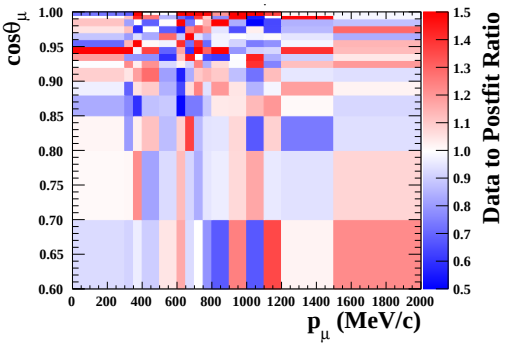

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.2: FHC FGD1 ν_μ CC1π

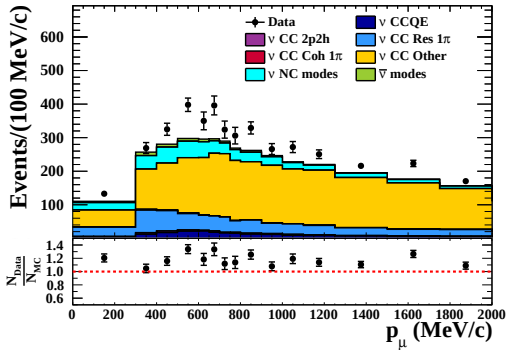

(a1) Prefit

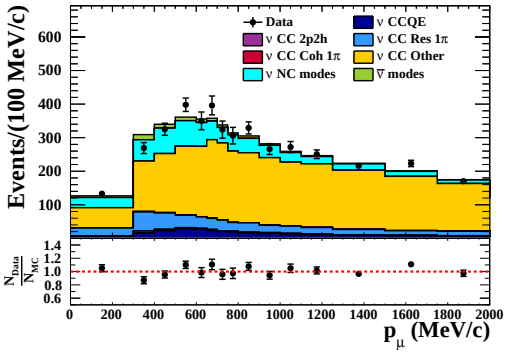

(a2) Postfit

(a) Reconstructed muon momentum

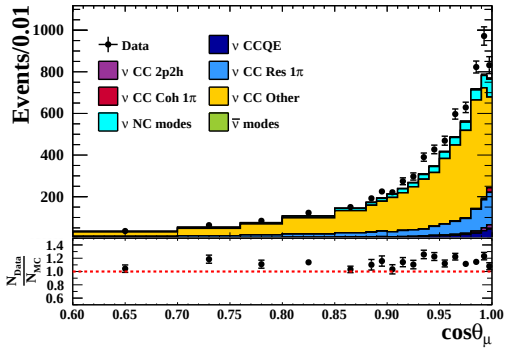

(b1) Prefit

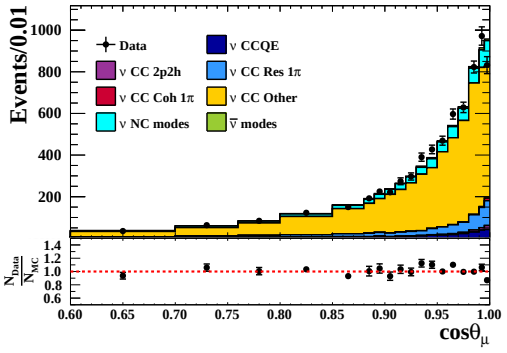

(b2) Postfit

(b) Cosine of reconstructed muon angle

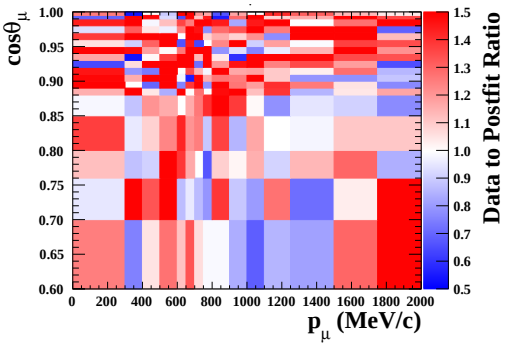

(c1) Prefit

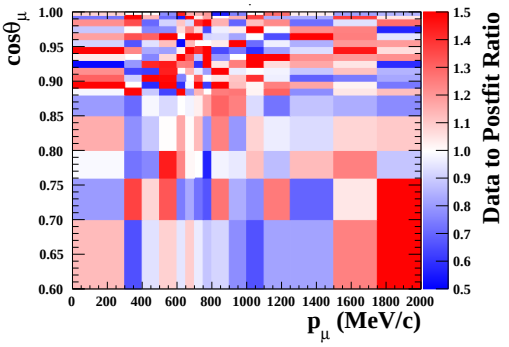

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.3: FHC FGD1 ν_μ CC-Other

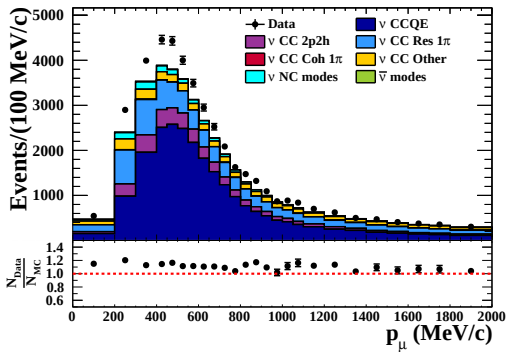

(a1) Prefit

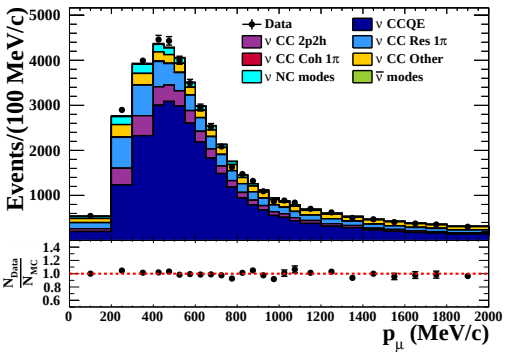

(a2) Postfit

(a) Reconstructed muon momentum

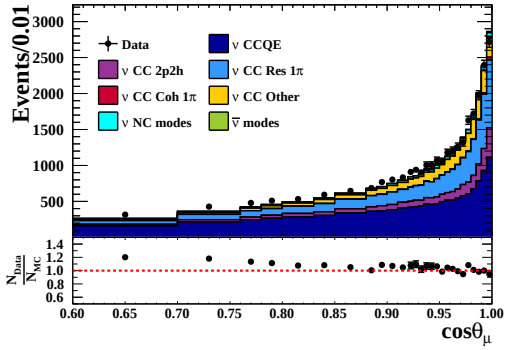

(b1) Prefit

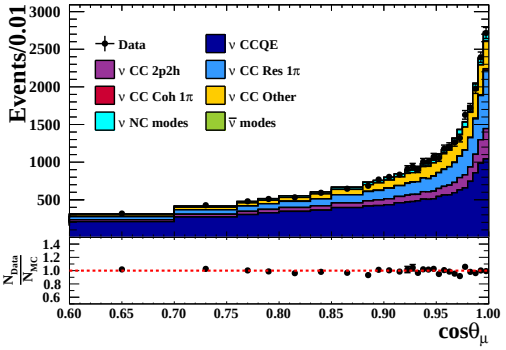

(b2) Postfit

(b) Cosine of reconstructed muon angle

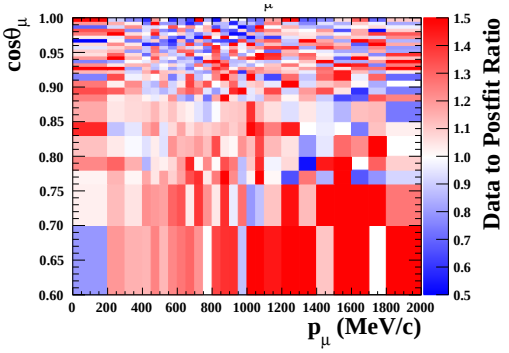

(c1) Prefit

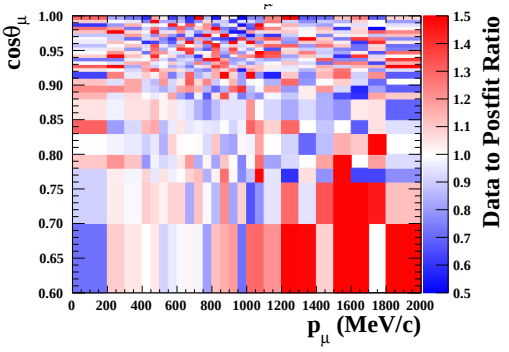

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.4: FHC FGD2 ν_μ CC0π

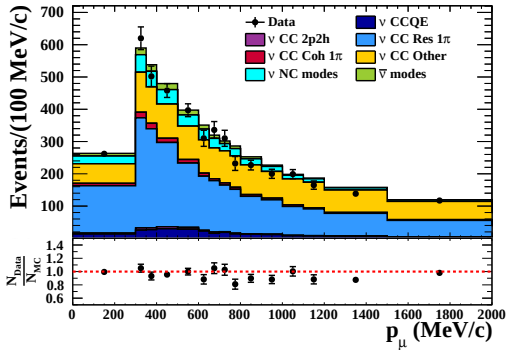

(a1) Prefit

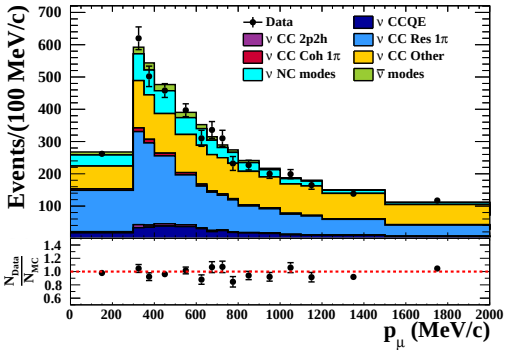

(a2) Postfit

(a) Reconstructed muon momentum

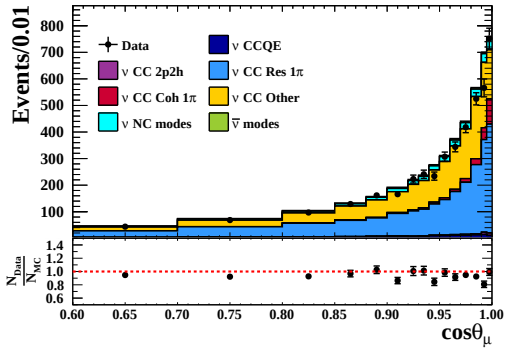

(b1) Prefit

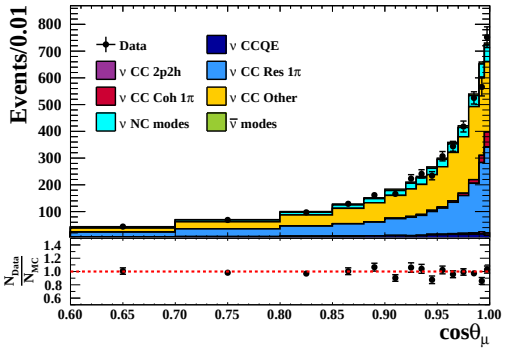

(b2) Postfit

(b) Cosine of reconstructed muon angle

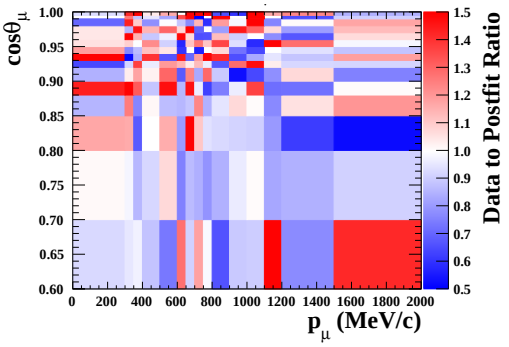

(c1) Prefit

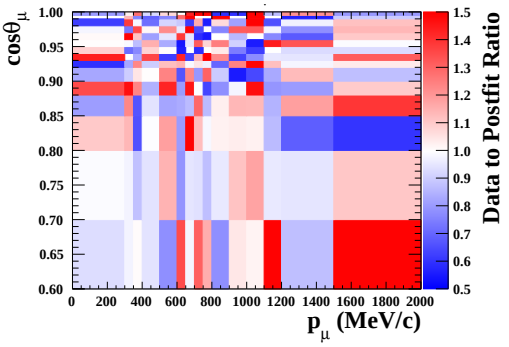

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.5: FHC FGD2 ν_μ CC1π

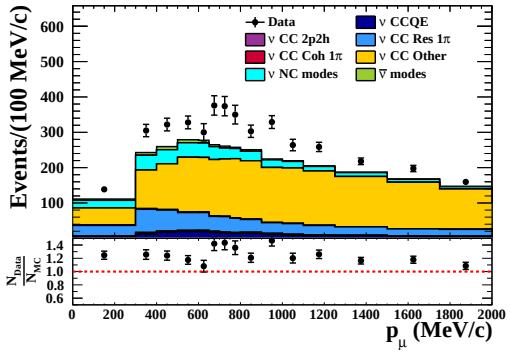

(a1) Prefit

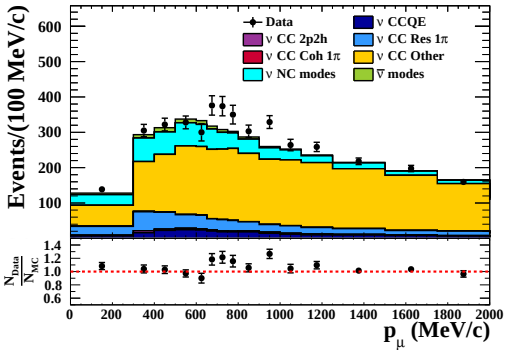

(a2) Postfit

(a) Reconstructed muon momentum

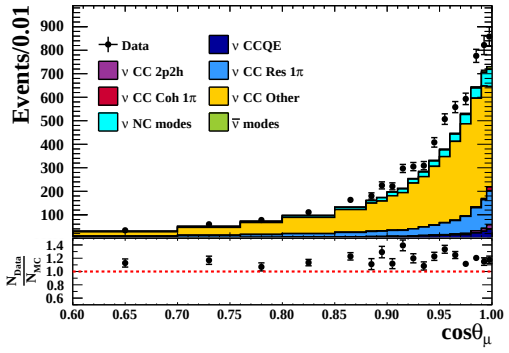

(b1) Prefit

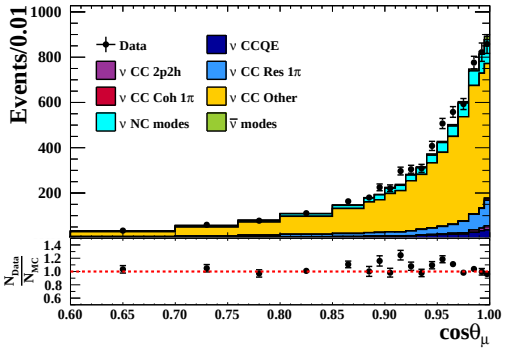

(b2) Postfit

(b) Cosine of reconstructed muon angle

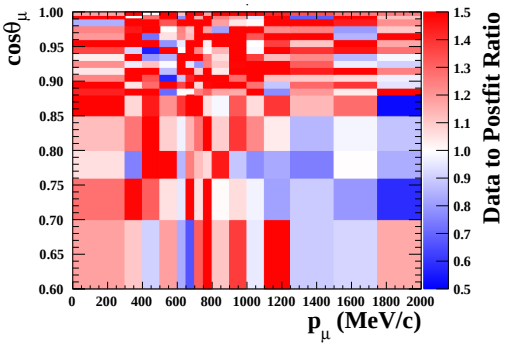

(c1) Prefit

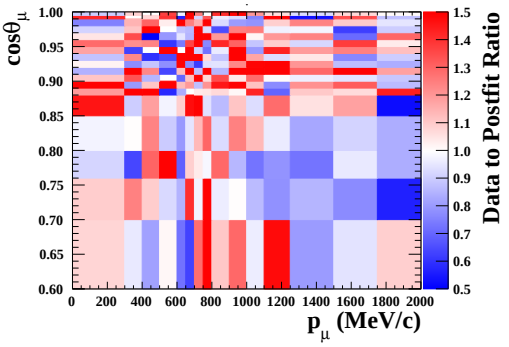

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.6: FHC FGD2 ν_μ CC-Other

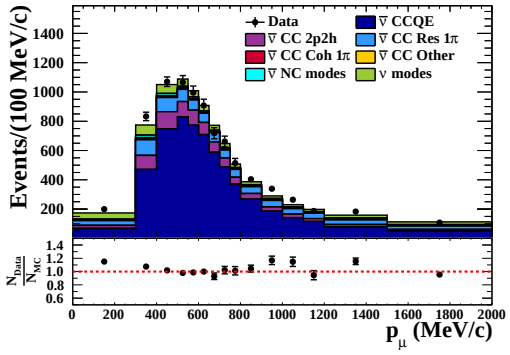

(a1) Prefit

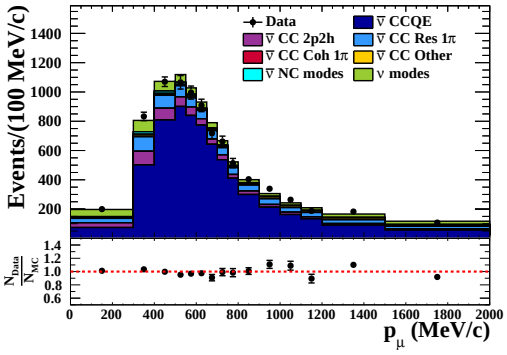

(a2) Postfit

(a) Reconstructed muon momentum

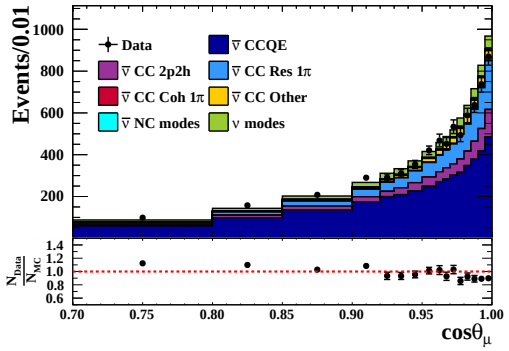

(b1) Prefit

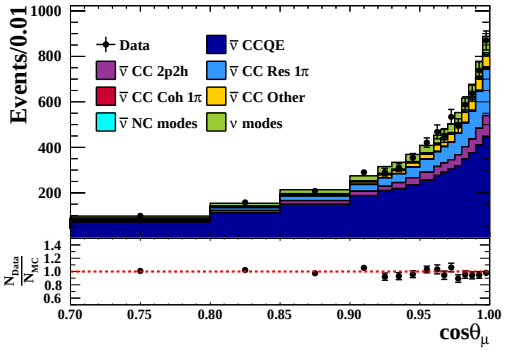

(b2) Postfit

(b) Cosine of reconstructed muon angle

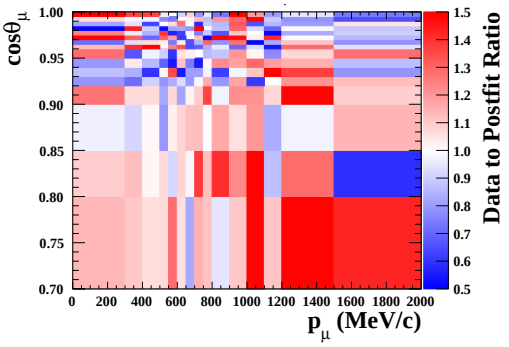

(c1) Prefit

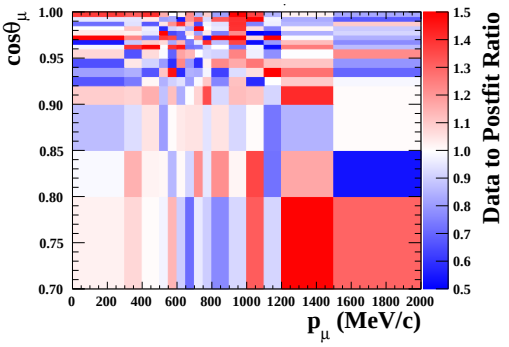

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.7: RHC FGD1 $\bar{\nu}_\mu$ CC0π

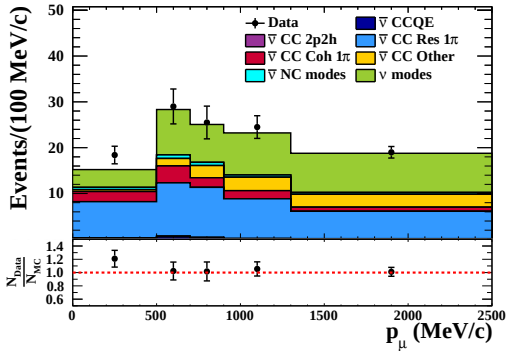

(a1) Prefit

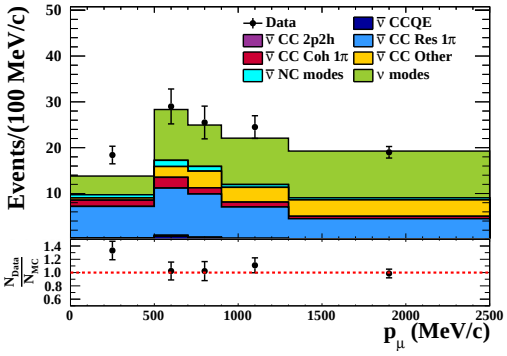

(a2) Postfit

(a) Reconstructed muon momentum

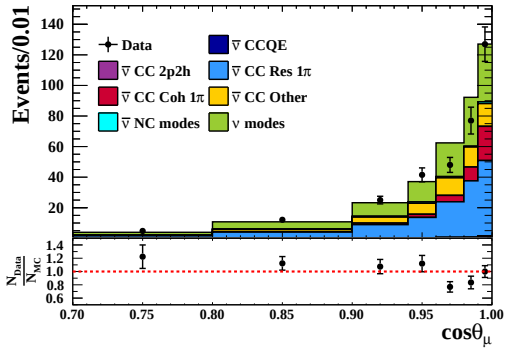

(b1) Prefit

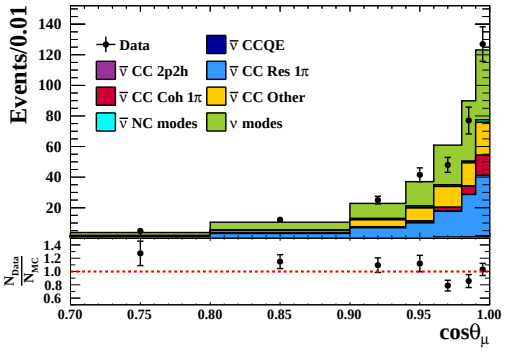

(b2) Postfit

(b) Cosine of reconstructed muon angle

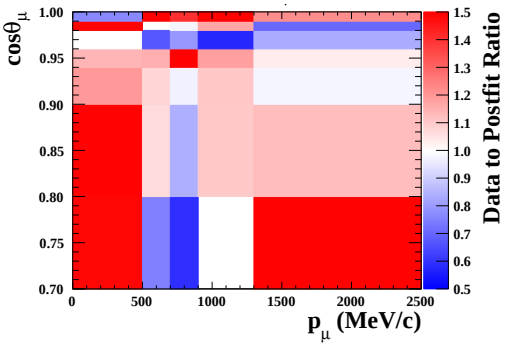

(c1) Prefit

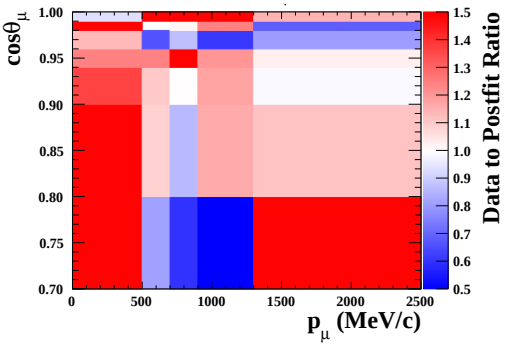

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.8: RHC FGD1 $\bar{\nu}_\mu$ CC1π

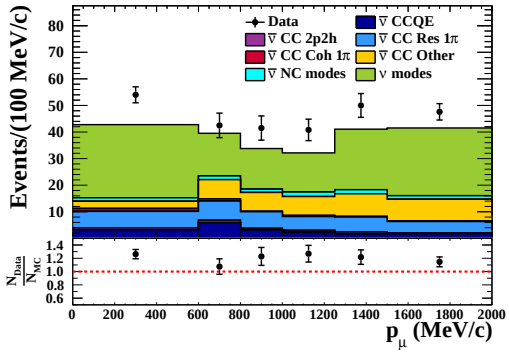

(a1) Prefit

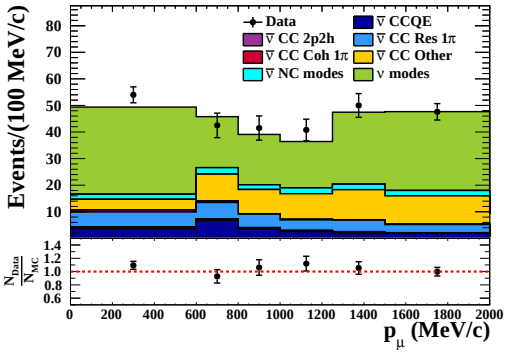

(a2) Postfit

(a) Reconstructed muon momentum

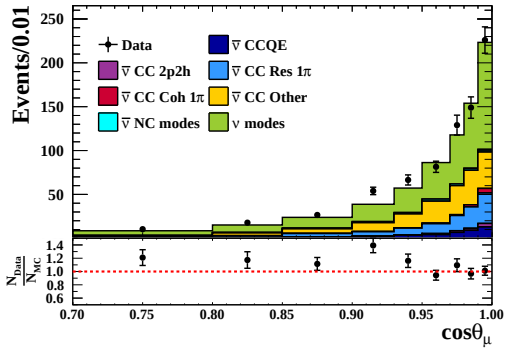

(b1) Prefit

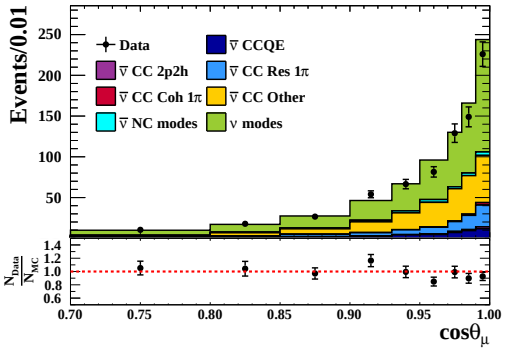

(b2) Postfit

(b) Cosine of reconstructed muon angle

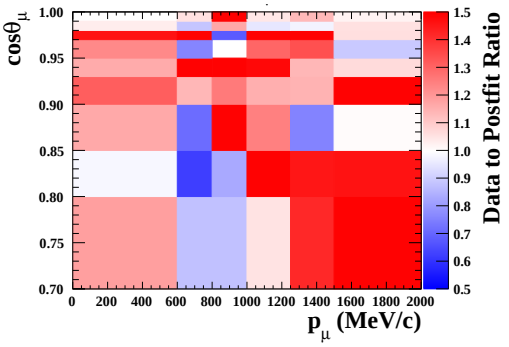

(c1) Prefit

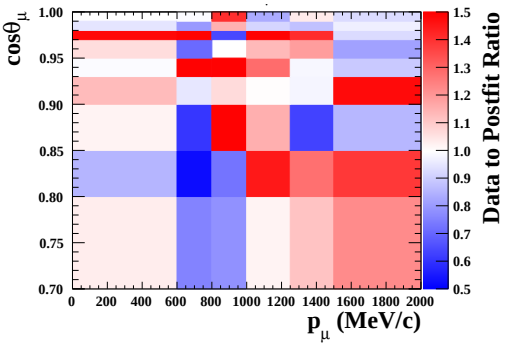

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.9: RHC FGD1 $\bar{\nu}_\mu$ CC-Other

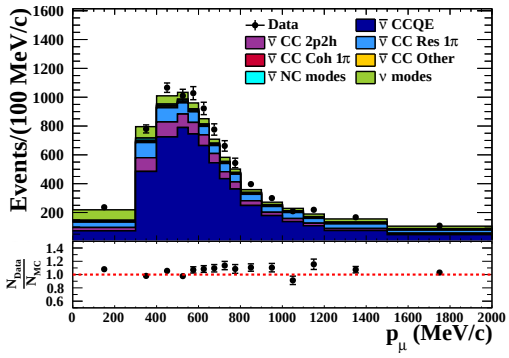

(a1) Prefit

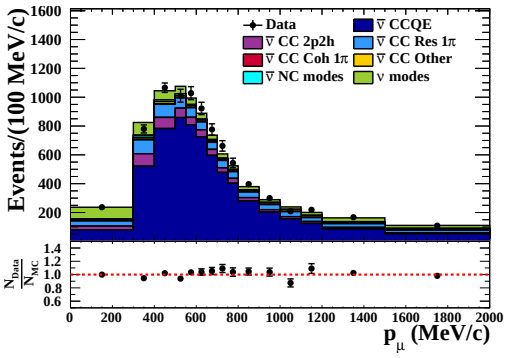

(a2) Postfit

(a) Reconstructed muon momentum

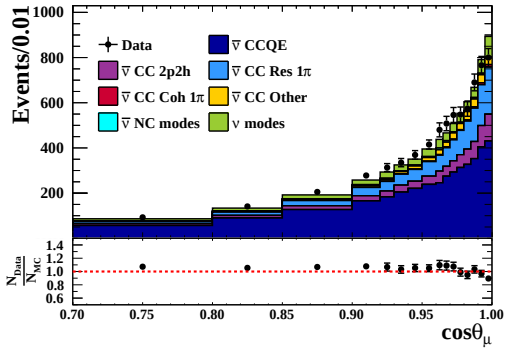

(b1) Prefit

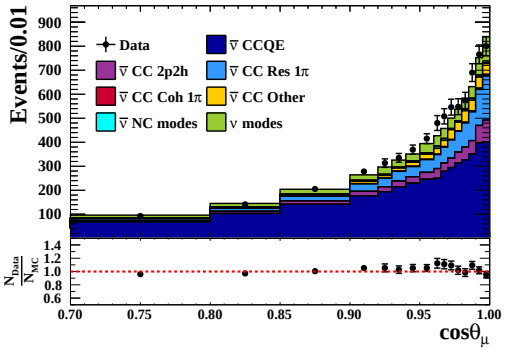

(b2) Postfit

(b) Cosine of reconstructed muon angle

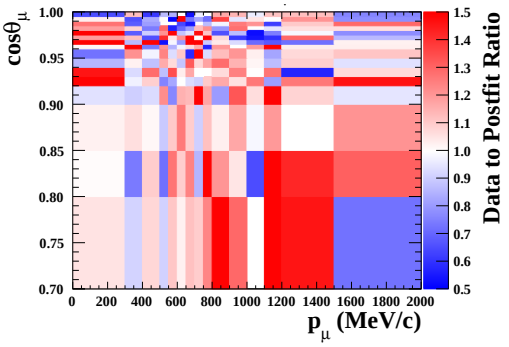

(c1) Prefit

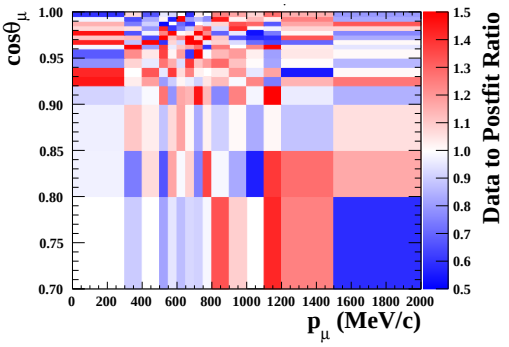

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.10: RHC FGD2 $\bar{\nu}_\mu$ CC0π

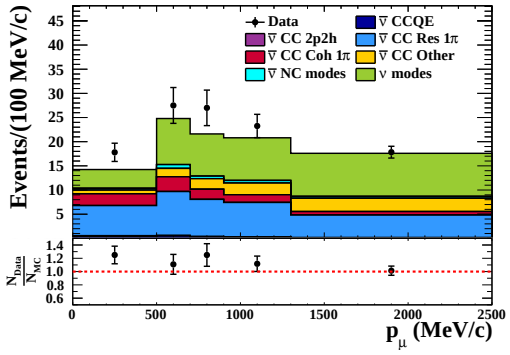

(a1) Prefit

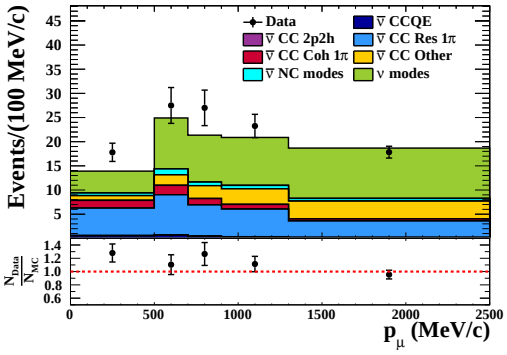

(a2) Postfit

(a) Reconstructed muon momentum

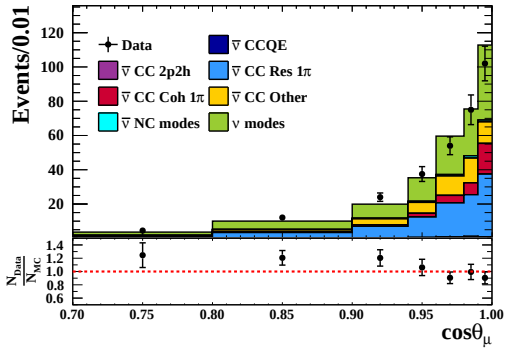

(b1) Prefit

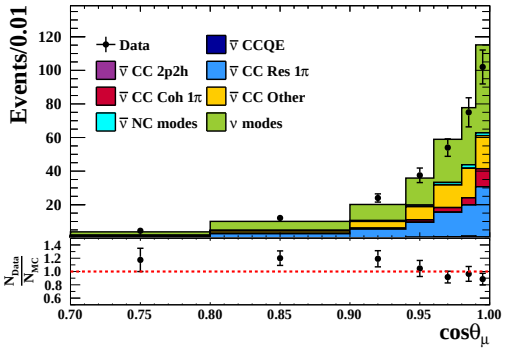

(b2) Postfit

(b) Cosine of reconstructed muon angle

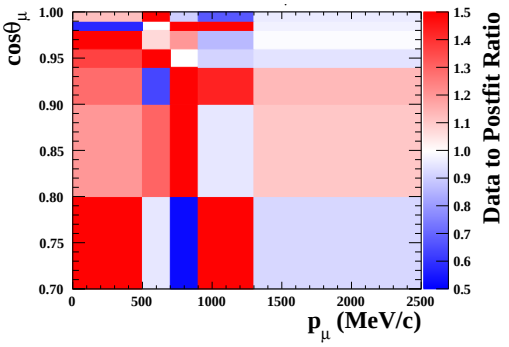

(c1) Prefit

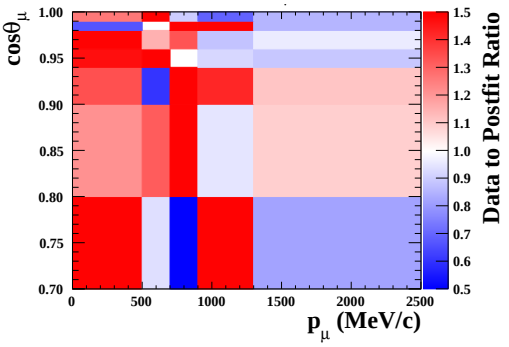

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.11: RHC FGD2 $\bar{\nu}_\mu$ CC1π

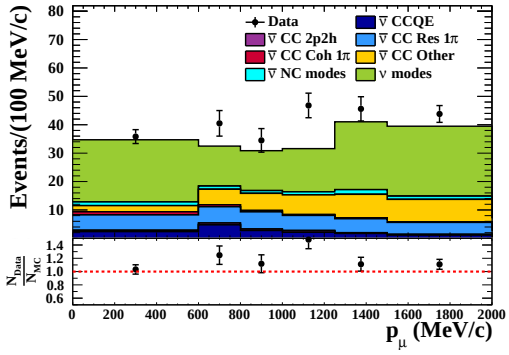

(a1) Prefit

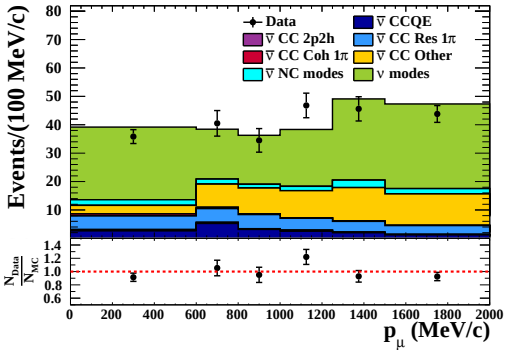

(a2) Postfit

(a) Reconstructed muon momentum

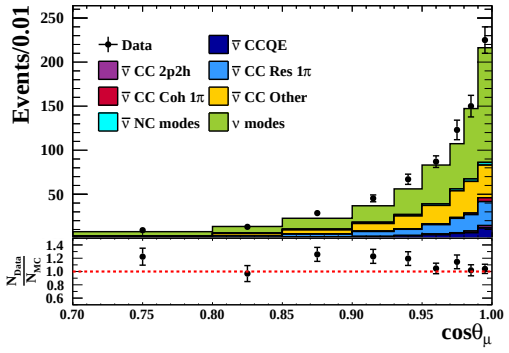

(b1) Prefit

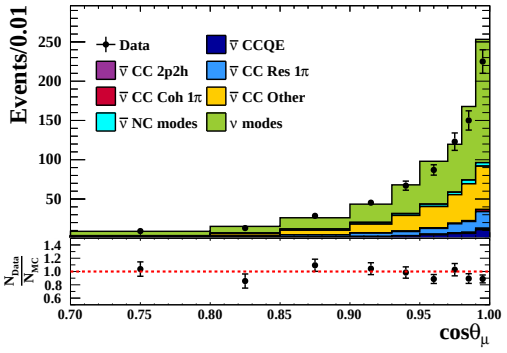

(b2) Postfit

(b) Cosine of reconstructed muon angle

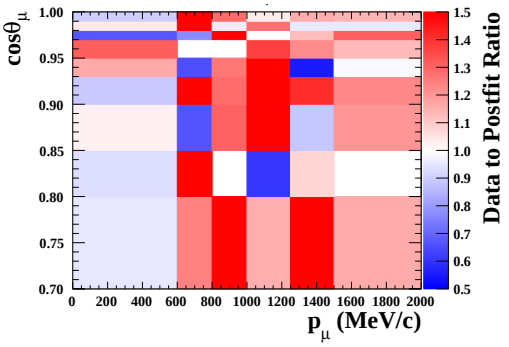

(c1) Prefit

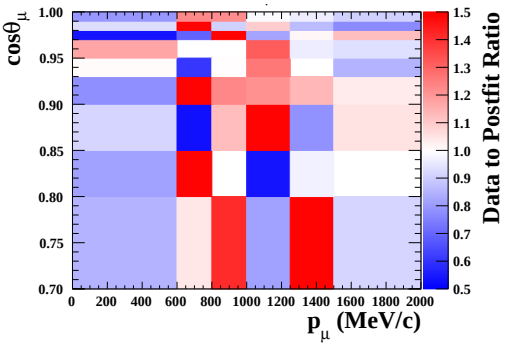

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.12: RHC FGD2 $\bar{\nu}_\mu$ CC-Other

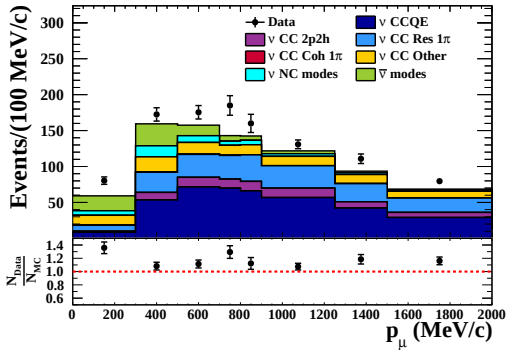

(a1) Prefit

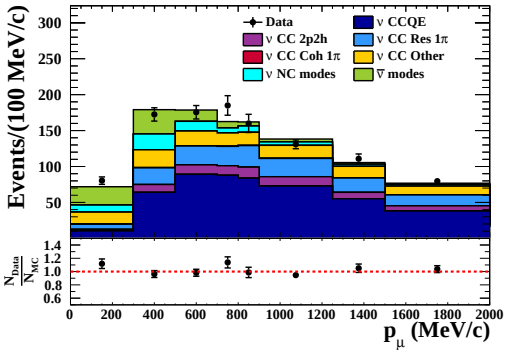

(a2) Postfit

(a) Reconstructed muon momentum

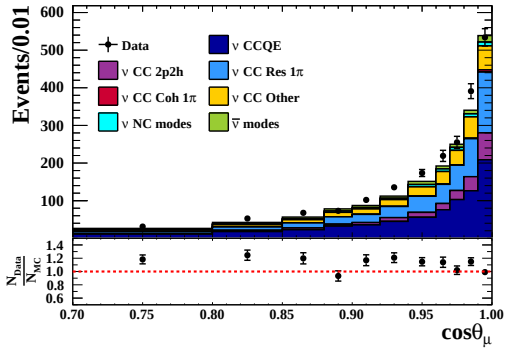

(b1) Prefit

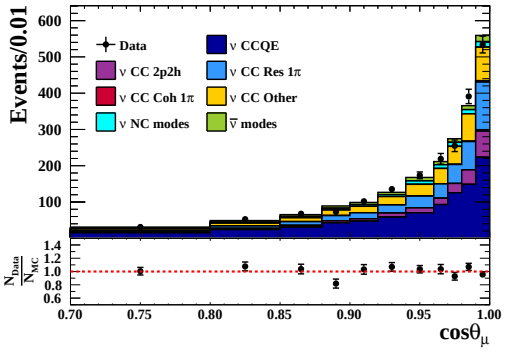

(b2) Postfit

(b) Cosine of reconstructed muon angle

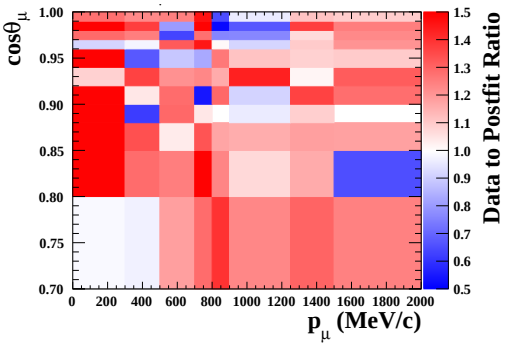

(c1) Prefit

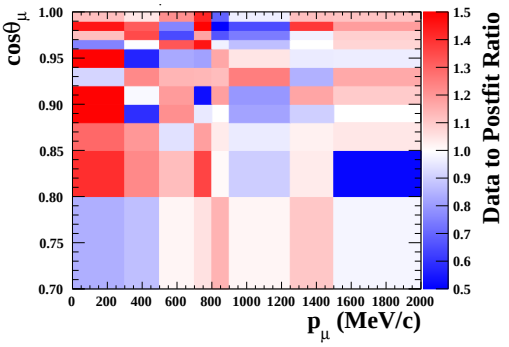

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.13: RHC FGD1 ν_μ-bkg CC0π

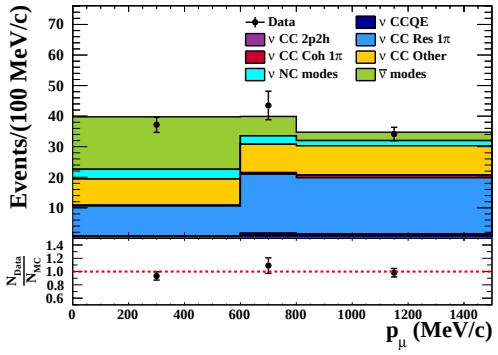

(a1) Prefit

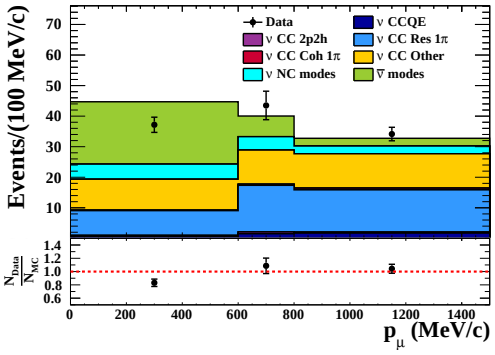

(a2) Postfit

(a) Reconstructed muon momentum

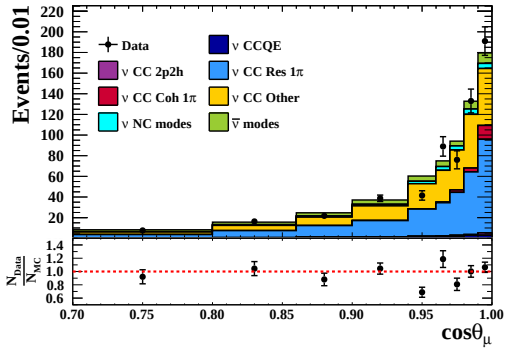

(b1) Prefit

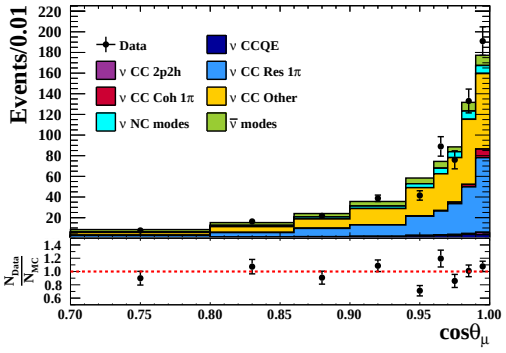

(b2) Postfit

(b) Cosine of reconstructed muon angle

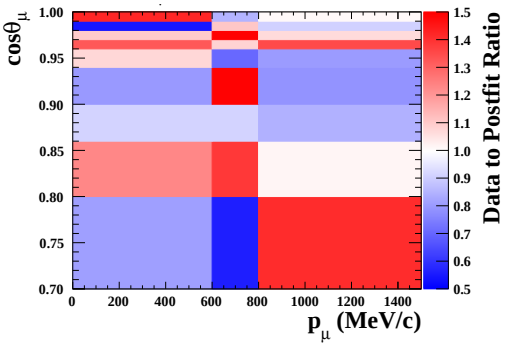

(c1) Prefit

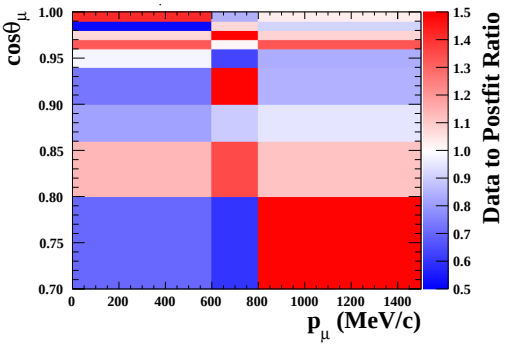

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.14: RHC FGD1 ν_μ-bkg CC1π

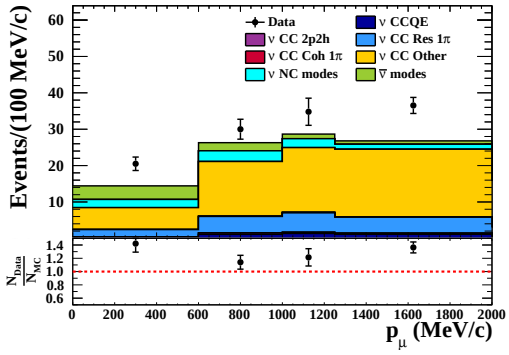

(a1) Prefit

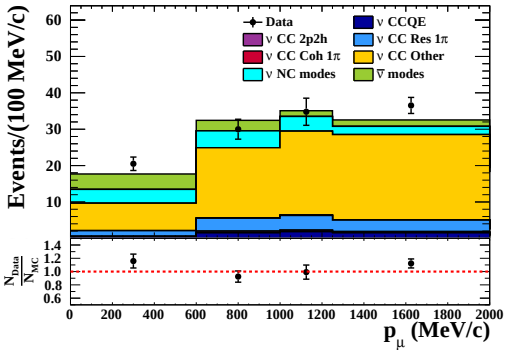

(a2) Postfit

(a) Reconstructed muon momentum

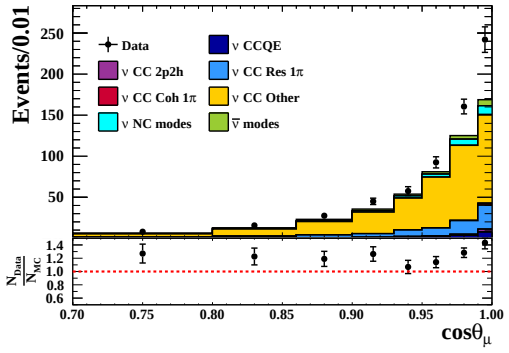

(b1) Prefit

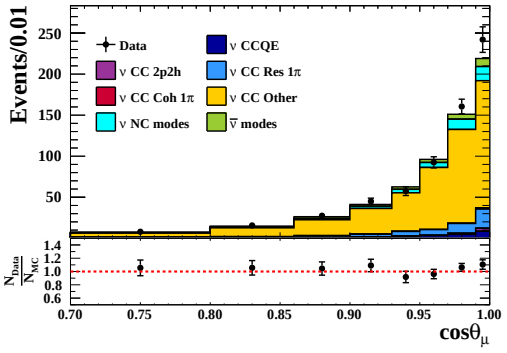

(b2) Postfit

(b) Cosine of reconstructed muon angle

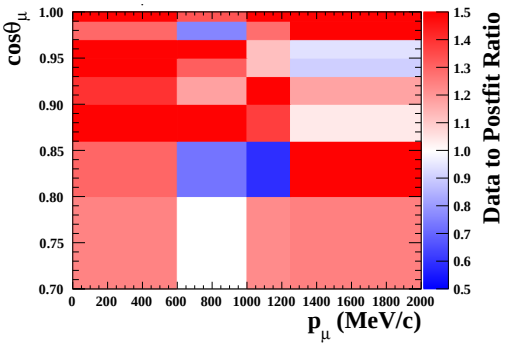

(c1) Prefit

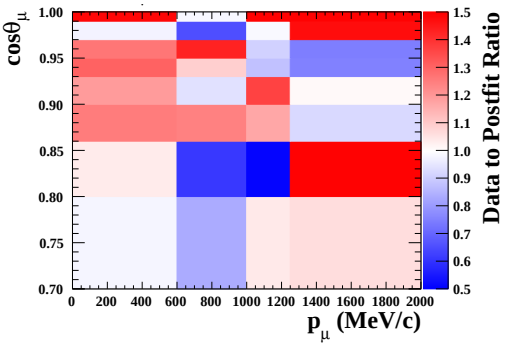

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.15: RHC FGD1 ν_μ-bkg CC-Other

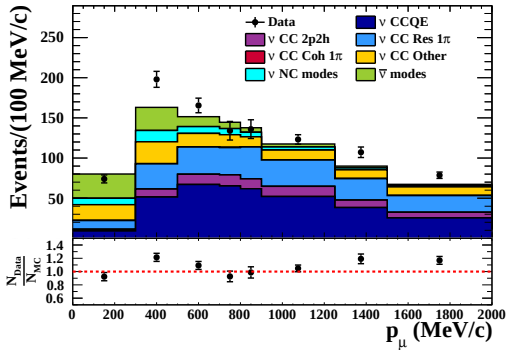

(a1) Prefit

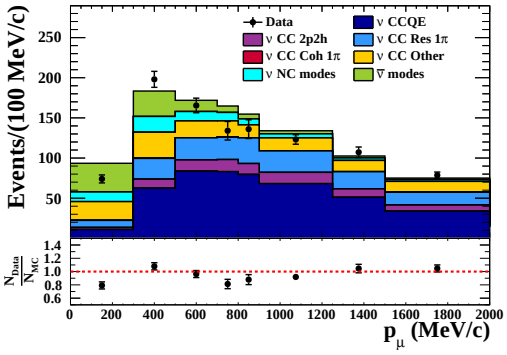

(a2) Postfit

(a) Reconstructed muon momentum

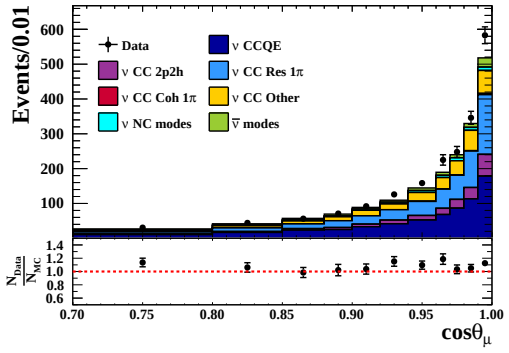

(b1) Prefit

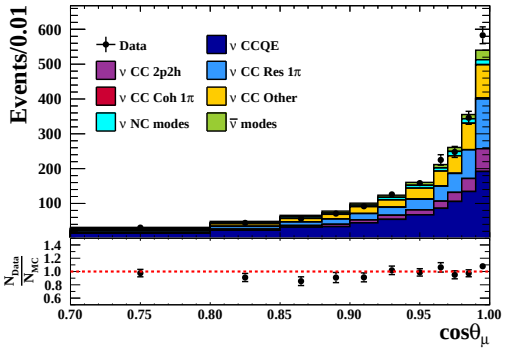

(b2) Postfit

(b) Cosine of reconstructed muon angle

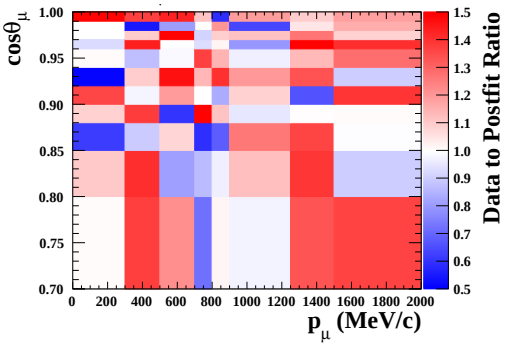

(c1) Prefit

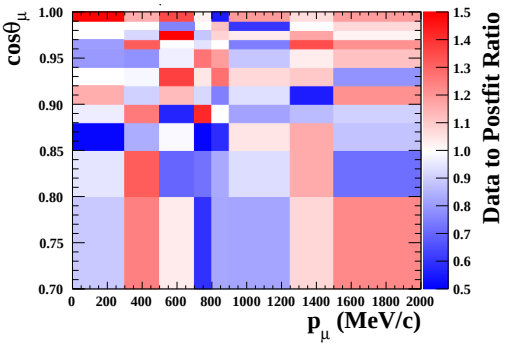

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.16: RHC FGD2 ν_μ-bkg CC0π

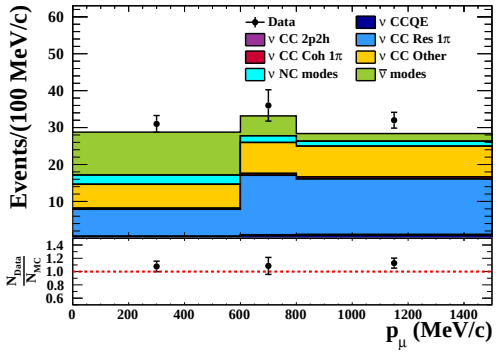

(a1) Prefit

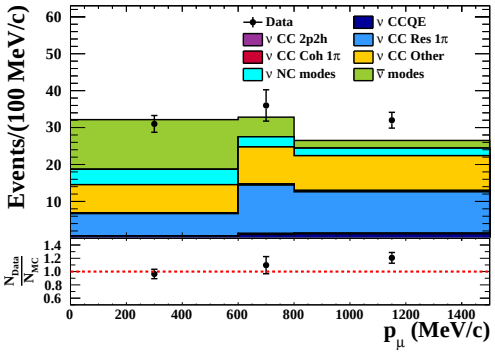

(a2) Postfit

(a) Reconstructed muon momentum

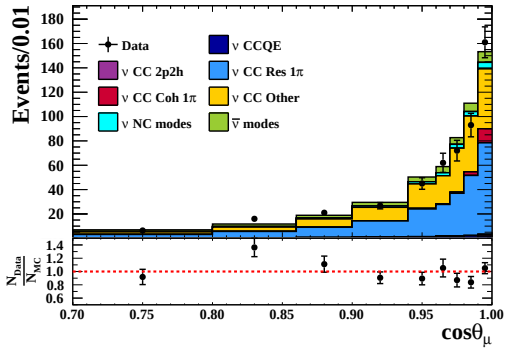

(b1) Prefit

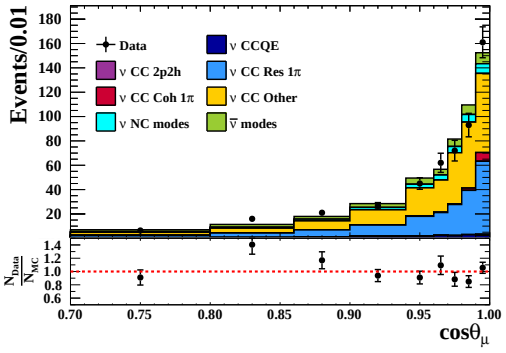

(b2) Postfit

(b) Cosine of reconstructed muon angle

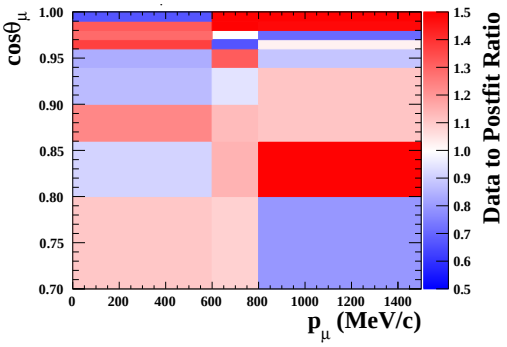

(c1) Prefit

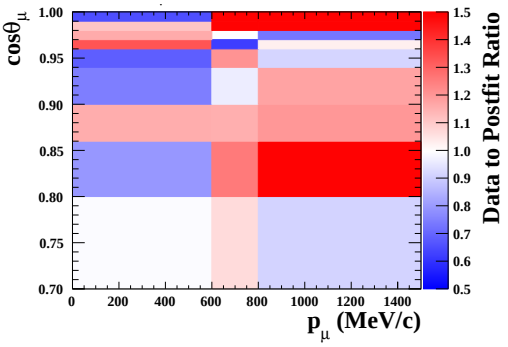

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.17: RHC FGD2 ν_μ-bkg CC1π

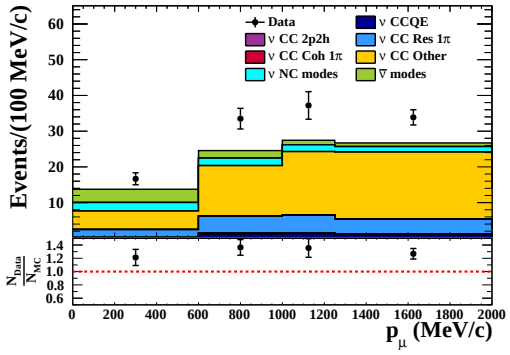

(a1) Prefit

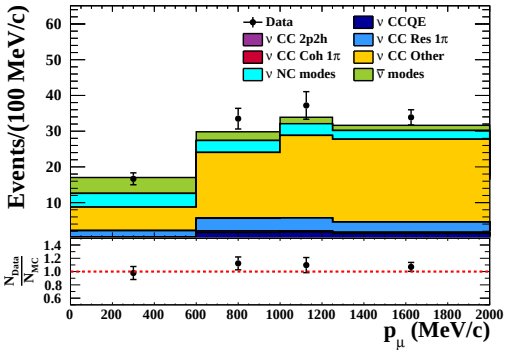

(a2) Postfit

(a) Reconstructed muon momentum

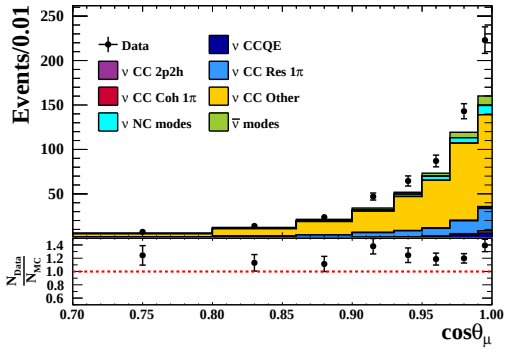

(b1) Prefit

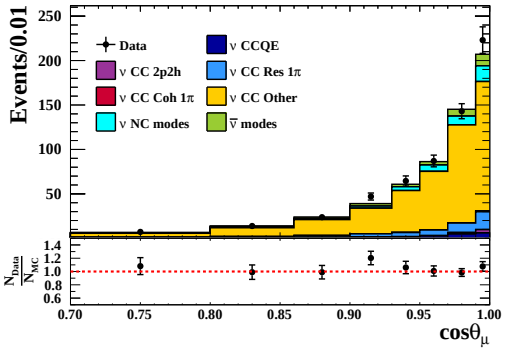

(b2) Postfit

(b) Cosine of reconstructed muon angle

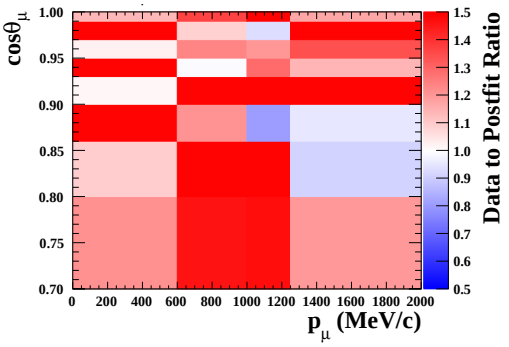

(c1) Prefit

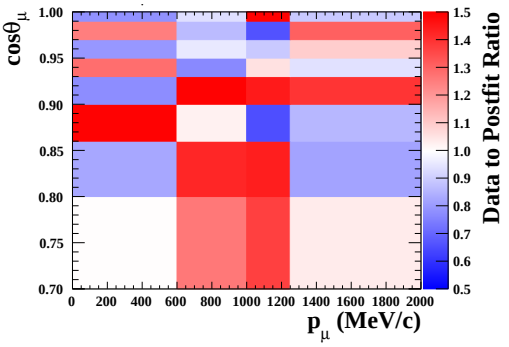

(c2) Postfit

(c) Ratio of data to MC prediction in reconstructed muon momentum and angle

Figure E.18: RHC FGD2 ν_μ-bkg CC-Other

The Bibliography

[1] Sheldon L. Glashow. "Partial-symmetries of weak interactions". In: *Nuclear Physics* 22.4 (1961), pp. 579–588. ISSN: 0029-5582. DOI: https://doi.org/10.1016/0029-5582(61)90469-2. URL: https://www.sciencedirect.com/science/article/pii/0029558261904692.

[2] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264. URL: https://link.aps.org/doi/10.1103/PhysRevLett.19.1264.

[3] Abdus Salam. "Weak and Electromagnetic Interactions". In: *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034.

[4] G. Arnison et al. "Experimental observation of isolated large transverse energy electrons with associated missing energy at $s = 540$ GeV". In: *Physics Letters B* 122.1 (1983), pp. 103–116. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(83)91177-2. URL: https://www.sciencedirect.com/science/article/pii/0370269383911772.

[5] G. Arnison et al. "Experimental observation of lepton pairs of invariant mass around 95GeV/c^2 at the CERN SPS collider". In: *Physics Letters B* 126.5 (1983), pp. 398–410. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(83)90188-0. URL: https://www.sciencedirect.com/science/article/pii/0370269383901880.

[6] G. Aad et Al. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Physics Letters B* 716.1 (2012), pp. 1–29. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.08.020. URL: https://www.sciencedirect.com/science/article/pii/S037026931200857X.

[7] Roel Aaij et al. "Test of lepton universality in beauty-quark decays". In: (2021). arXiv: 2103.11769 [hep-ex].

[8] B. Abi et al. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm”. In: *Phys. Rev. Lett.* 126 (14 Apr. 2021), p. 141801. DOI: 10.1103/PhysRevLett.126.141801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.126.141801.

[9] Y. et Al. Fukuda. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. URL: https://link.aps.org/doi/10.1103/PhysRevLett.81.1562.

[10] W. Pauli. “Dear Radioactive Ladies and Gentlemen”. In: (1930).

[11] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Phys. Rev.* 105 (4 Feb. 1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413. URL: https://link.aps.org/doi/10.1103/PhysRev.105.1413.

[12] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. In: *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.138.

[13] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Soviet Journal of Experimental and Theoretical Physics* 26 (May 1968), p. 984.

[14] Maki, Z. Nakagawa, M. and Sakata, S. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. ISSN: 0033-068X. DOI: 10.1143/PTP.28.870. eprint: https://academic.oup.com/ptp/article-pdf/28/5/870/5258750/28-5-870.pdf. URL: https://doi.org/10.1143/PTP.28.870.

[15] K. Abe et al. “The T2K experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (2011), pp. 106–135. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2011.06.067. URL: https://www.sciencedirect.com/science/article/pii/S0168900211011910.

[16] Patrick Dunne. *Latest Neutrino Oscillation Results from T2K*. July 2020. DOI: 10.5281/zenodo.4154355.

[17] M. H. Ahn et al. “Measurement of neutrino oscillation by the K2K experiment”. In: *Phys. Rev. D* 74 (7 Oct. 2006), p. 072003. DOI: 10.1103/PhysRevD.74.072003. URL: https://link.aps.org/doi/10.1103/PhysRevD.74.072003.

[18] The T2K Collaboration. *The T2K experiment*. URL: https://t2k-experiment.org/.

[19] A. D. Sakharov. "Violation of CP in variance, C asymmetry, and baryon asymmetry of the universe". In: *JETP Lett. (USSR)* 126 (5 1967), p. 24.

[20] K. et al. Abe. "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations". In: *Nature* 580.7803 (Apr. 2020), pp. 339–344. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2177-0. URL: http://dx.doi.org/10.1038/s41586-020-2177-0.

[21] S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry". In: *Phys. Rev. D* 2 (7 Oct. 1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285. URL: https://link.aps.org/doi/10.1103/PhysRevD.2.1285.

[22] M. Banner et al. "Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider". In: *Physics Letters B* 122.5 (1983), pp. 476–485. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(83)91605-2. URL: https://www.sciencedirect.com/science/article/pii/0370269383916052.

[23] P. Bagnaia et al. "Evidence for $Z^0 \to e^+e^-$ at the CERN pp collider". In: *Physics Letters B* 129.1 (1983), pp. 130–140. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(83)90744-X. URL: https://www.sciencedirect.com/science/article/pii/037026938390744X.

[24] S. Chatrchyan et Al. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Physics Letters B* 716.1 (2012), pp. 30–61. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.08.021. URL: https://www.sciencedirect.com/science/article/pii/S0370269312008581.

[25] M. Lubej. *A diagram of the Standard Model*. URL: http://www-f9.ijs.si/~lubej/.

[26] R. et Al. Aaij. "Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \to J/\psi K^- p$ Decays". In: *Phys. Rev. Lett.* 115 (7 Aug. 2015), p. 072001. DOI: 10.1103/PhysRevLett.115.072001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.115.072001.

[27] R. et Al. Aaij. "Observation of a Narrow Pentaquark State, $P_c(4312)^+$, and of the Two-Peak Structure of the $P_c(4450)^+$". In: *Phys. Rev. Lett.* 122 (22 June 2019), p. 222001. DOI: 10.1103/PhysRevLett.122.222001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.122.222001.

[28] M. Goldhaber, L. Grodzins, and A. W. Sunyar. "Helicity of Neutrinos". In: *Phys. Rev.* 109 (3 Feb. 1958), pp. 1015–1017. DOI: 10.1103/PhysRev.109.1015. URL: https://link.aps.org/doi/10.1103/PhysRev.109.1015.

[29] Q. R. et Al. Ahmad. "Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. URL: https://link.aps.org/doi/10.1103/PhysRevLett.87.071301.

[30] M Gell-Mann and A Pais. "Behavior of Neutral Particles under Charge Conjugation". eng. In: *Physical review* 97.5 (1955), pp. 1387–1389. ISSN: 0031-899X.

[31] S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry". In: *Phys. Rev. D* 2 (7 Oct. 1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285. URL: https://link.aps.org/doi/10.1103/PhysRevD.2.1285.

[32] B Pontecorvo. "MESONIUM AND ANTIMESONIUM". In: *Journal of Experimental and Theoretical Physics* 33 (Aug. 1957). URL: https://www.osti.gov/biblio/4343073.

[33] Nicola Cabibbo. "Unitary Symmetry and Leptonic Decays". In: *Phys. Rev. Lett.* 10 (12 June 1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531. URL: https://link.aps.org/doi/10.1103/PhysRevLett.10.531.

[34] M. Tanabashi et al. "Review of Particle Physics". In: *Phys. Rev. D* 98.3 (2018). [Updated 2019 tables used for $\sin^2\theta_{13}$ constraint], p. 030001. DOI: 10.1103/PhysRevD.98.030001.

[35] JUNO Collaboration. *Neutrino mass ordering*. 2017. URL: http://www.staff.uni-mainz.de/wurmm/juno.html.

[36] C.H. Llewellyn Smith. "Neutrino reactions at accelerator energies". In: *Physics Reports* 3.5 (1972), pp. 261–379. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(72)90010-5. URL: https://www.sciencedirect.com/science/article/pii/0370157372900105.

[37] R. Bradford et al. "A New parameterization of the nucleon elastic form-factors". In: *Nucl. Phys. B Proc. Suppl.* 159 (2006). Ed. by F. Cavanna, J. G. Morfin, and T. Nakaya, pp. 127–132. DOI: 10.1016/j.nuclphysbps.2006.08.028. arXiv: hep-ex/0602017.

[38] R. Gran et al. "Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV". In: *Phys. Rev. D* 88 (11 Dec. 2013), p. 113007. DOI: 10.1103/PhysRevD.88.113007. URL: https://link.aps.org/doi/10.1103/PhysRevD.88.113007.

[39] Dieter Rein and Lalit M Sehgal. "Neutrino-excitation of baryon resonances and single pion production". In: *Annals of Physics* 133.1 (1981), pp. 79–153. ISSN: 0003-4916. DOI: https://doi.org/10.1016/0003-4916(81)90242-6. URL: https://www.sciencedirect.com/science/article/pii/0003491681902426.

[40] Ch. Berger and L. M. Sehgal. "Partially conserved axial vector current and coherent pion production by low energy neutrinos". In: *Phys. Rev. D* 79 (5 Mar. 2009), p. 053003. DOI: 10.1103/PhysRevD.79.053003. URL: https://link.aps.org/doi/10.1103/PhysRevD.79.053003.

[41] D. Rein. "Diffractive Pion Production in Neutrino Reactions". In: *Nucl. Phys. B* 278 (1986), pp. 61–77. DOI: 10.1016/0550-3213(86)90106-9.

[42] Glück, M., Reya, E. & Vogt, A. "Dynamical parton distributions revisited." In: *Eur. Phys. J. C.* 5 (1998), pp. 461–470. DOI: 10.1007/s100529800978.

[43] Francesca Di Lodovico and. "The Hyper-Kamiokande Experiment". In: 888 (Sept. 2017), p. 012020. DOI: 10.1088/1742-6596/888/1/012020. URL: https://doi.org/10.1088/1742-6596/888/1/012020.

[44] R. et Al. Acciarri. "Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 1: The LBNF and DUNE Projects". In: *arXiv e-prints*, arXiv:1601.05471 (Jan. 2016), arXiv:1601.05471. arXiv: 1601.05471 [physics.ins-det].

[45] Claude Amsler. "The nuclear shell model". In: *Nuclear and Particle Physics*. 2053-2563. IOP Publishing, 2015, 6-1 to 6–12. ISBN: 978-0-7503-1140-3. DOI: 10.1088/978-0-7503-1140-3ch6. URL: https://dx.doi.org/10.1088/978-0-7503-1140-3ch6.

[46] E. J. Moniz et al. "Nuclear Fermi Momenta from Quasielastic Electron Scattering". In: *Phys. Rev. Lett.* 26 (8 Feb. 1971), pp. 445–448. DOI: 10.1103/PhysRevLett.26.445. URL: https://link.aps.org/doi/10.1103/PhysRevLett.26.445.

[47] J. A. Caballero. "General study of superscaling in quasielastic (e,e') and (ν,μ) reactions using the relativistic impulse approximation". In: *Phys. Rev. C* 74 (1 July 2006), p. 015502. DOI: 10.1103/PhysRevC.74.015502. URL: https://link.aps.org/doi/10.1103/PhysRevC.74.015502.

[48] Omar Benhar, Donal Day, and Ingo Sick. "Inclusive quasielastic electron-nucleus scattering". In: *Rev. Mod. Phys.* 80 (1 Jan. 2008), pp. 189–224. DOI: 10.1103/RevModPhys.80.189. URL: https://link.aps.org/doi/10.1103/RevModPhys.80.189.

[49] O. Benhar et al. “Spectral function of finite nuclei and scattering of GeV electrons”. In: *Nucl. Phys. A* 579 (1994), pp. 493–517. DOI: 10.1016/0375-9474(94)90920-2.

[50] Arie Bodek and Tejin Cai. “Removal energies and final state interaction in lepton nucleus scattering”. In: *The European Physical Journal C* 79.4 (Apr. 2019). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-6750-3. URL: http://dx.doi.org/10.1140/epjc/s10052-019-6750-3.

[51] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino cross sections across energy scales”. In: *Rev. Mod. Phys.* 84 (3 Sept. 2012), pp. 1307–1341. DOI: 10.1103/RevModPhys.84.1307. URL: https://link.aps.org/doi/10.1103/RevModPhys.84.1307.

[52] Brent Follin et al. “First Detection of the Acoustic Oscillation Phase Shift Expected from the Cosmic Neutrino Background”. In: *Physical Review Letters* 115.9 (Aug. 2015). DOI: 10.1103/physrevlett.115.091301.

[53] C. L. Cowan et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103. eprint: https://www.science.org/doi/pdf/10.1126/science.124.3212.103. URL: https://www.science.org/doi/abs/10.1126/science.124.3212.103.

[54] F. Reines and C. L. Cowan Jr. “The Neutrino”. In: *Nature* 178.4531 (Sept. 1956), pp. 446–449. ISSN: 1476-4687. DOI: 10.1038/178446a0. URL: https://doi.org/10.1038/178446a0.

[55] S. et al. Abe. “Precision Measurement of Neutrino Oscillation Parameters with KamLAND”. In: *Physical Review Letters* 100.22 (June 2008). ISSN: 1079-7114. DOI: 10.1103/physrevlett.100.221803. URL: http://dx.doi.org/10.1103/PhysRevLett.100.221803.

[56] M. et al. Apollonio. “Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station”. In: *The European Physical Journal C* 27.3 (Apr. 2003), pp. 331–374. ISSN: 1434-6052. DOI: 10.1140/epjc/s2002-01127-9. URL: http://dx.doi.org/10.1140/epjc/s2002-01127-9.

[57] J. K. et al. Ahn. “Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment”. In: *Physical Review Letters* 108.19 (May 2012). ISSN: 1079-7114. DOI: 10.1103/physrevlett.108.191802. URL: http://dx.doi.org/10.1103/PhysRevLett.108.191802.

[58] G. et al. Bak. “Measurement of Reactor Antineutrino Oscillation Amplitude and Frequency at RENO”. In: *Physical Review Letters* 121.20 (Nov. 2018). ISSN: 1079-7114. DOI: 10.1103/physrevlett.121.201801. URL: http://dx.doi.org/10.1103/PhysRevLett.121.201801.

[59] Y. Abe et Al. “Background-independent measurement of $theta_{13}$ in Double Chooz”. In: *Physics Letters B* 735 (2014), pp. 51–56. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2014.04.045. URL: https://www.sciencedirect.com/science/article/pii/S0370269314002858.

[60] D. et al. Adey. “Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay”. In: *Phys. Rev. Lett.* 121 (24 Dec. 2018), p. 241805. DOI: 10.1103/PhysRevLett.121.241805. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.241805.

[61] “JUNO physics and detector”. In: *Progress in Particle and Nuclear Physics* 123 (2022), p. 103927. ISSN: 0146-6410. DOI: https://doi.org/10.1016/j.ppnp.2021.103927.

[62] Stanislav P. Mikheev and Alexei Yu. Smirnov. “Neutrino oscillations in a variable-density medium and ν bursts due to the gravitational collapse of stars”. In: *arXiv: High Energy Physics - Phenomenology* (1986).

[63] L. Wolfenstein. “Neutrino oscillations in matter”. In: *Phys. Rev. D* 17 (9 May 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369. URL: https://link.aps.org/doi/10.1103/PhysRevD.17.2369.

[64] John N. Bahcall and Raymond Davis. “Solar Neutrinos: A Scientific Puzzle”. In: *Science* 191.4224 (1976), pp. 264–267. DOI: 10.1126/science.191.4224.264. eprint: https://www.science.org/doi/pdf/10.1126/science.191.4224.264. URL: https://www.science.org/doi/abs/10.1126/science.191.4224.264.

[65] Vladimir N Gavrin. “The Russian-American gallium experiment SAGE”. In: *Physics-Uspekhi* 54.9 (Sept. 2011), pp. 941–949. DOI: 10.3367/ufne.0181.201109g.0975. URL: https://doi.org/10.3367/ufne.0181.201109g.0975.

[66] D. Vignaud. “The GALLEX solar neutrino experiment”. In: *Nuclear Physics B - Proceedings Supplements* 60.3 (1998), pp. 20–29. ISSN: 0920-5632. DOI: https://doi.org/10.1016/S0920-5632(97)00498-2. URL: https://www.sciencedirect.com/science/article/pii/S0920563297004982.

[67] K. S. Hirata and T. et al. Kajita. "Observation of ^{8}B solar neutrinos in the Kamiokande-II detector". In: *Phys. Rev. Lett.* 63 (1 July 1989), pp. 16–19. DOI: 10.1103/PhysRevLett.63.16. URL: https://link.aps.org/doi/10.1103/PhysRevLett.63.16.

[68] G. et al. Danby. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36. URL: https://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[69] M. Schwartz. "Feasibility of Using High-Energy Neutrinos to Study the Weak Interactions". In: *Phys. Rev. Lett.* 4 (6 Mar. 1960), pp. 306–307. DOI: 10.1103/PhysRevLett.4.306. URL: https://link.aps.org/doi/10.1103/PhysRevLett.4.306.

[70] Ubaldo Dore, Pier Loverre, and Lucio Ludovici. "History of accelerator neutrino beams". In: *The European Physical Journal H* 44.4 (Nov. 2019), pp. 271–305. ISSN: 2102-6467. DOI: 10.1140/epjh/e2019-90032-x. URL: https://doi.org/10.1140/epjh/e2019-90032-x.

[71] H. Deden et al. "Experimental study of structure functions and sum rules in charge-changing interactions of neutrinos and antineutrinos on nucleons". In: *Nuclear Physics B* 85.2 (1975), pp. 269–288. ISSN: 0550-3213. DOI: https://doi.org/10.1016/0550-3213(75)90008-5. URL: https://www.sciencedirect.com/science/article/pii/0550321375900085.

[72] L. Aliaga et al. "Design, calibration, and performance of the MINERvA detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 743 (2014), pp. 130–159. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2013.12.053. URL: https://www.sciencedirect.com/science/article/pii/S0168900214000035.

[73] K. Kodama et al. "Observation of tau neutrino interactions". In: *Physics Letters B* 504.3 (2001), pp. 218–224. ISSN: 0370-2693. DOI: https://doi.org/10.1016/S0370-2693(01)00307-0. URL: https://www.sciencedirect.com/science/article/pii/S0370269301003070.

[74] P. et al. Adamson. "Combined Analysis of ν_μ Disappearance and $\nu_\mu \to \nu_e$ Appearance in MINOS Using Accelerator and Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 112 (19 May 2014), p. 191801. DOI: 10.1103/PhysRevLett.112.191801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.112.191801.

[75] P. et al. Adamson. “Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset”. In: *Phys. Rev. Lett.* 125 (13 Sept. 2020), p. 131802. DOI: 10.1103/PhysRevLett.125.131802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.131802.

[76] P. et al. Adamson. “Search for Sterile Neutrinos in MINOS and MINOS+ Using a Two-Detector Fit”. In: *Phys. Rev. Lett.* 122 (9 Mar. 2019), p. 091803. DOI: 10.1103/PhysRevLett.122.091803. URL: https://link.aps.org/doi/10.1103/PhysRevLett.122.091803.

[77] M. A. Acero et al. “First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA”. In: *Phys. Rev. Lett.* 123.15 (2019), p. 151803. DOI: 10.1103/PhysRevLett.123.151803. arXiv: 1906.04907 [hep-ex].

[78] Hans-Thomas Janka. “Neutrino Emission from Supernovae”. In: *Handbook of Supernovae*. Ed. by Athem W. Alsabti and Paul Murdin. Cham: Springer International Publishing, 2017, pp. 1575–1604. ISBN: 978-3-319-21846-5. DOI: 10.1007/978-3-319-21846-5_4. URL: https://doi.org/10.1007/978-3-319-21846-5_4.

[79] K. Hirata et al. “Observation of a neutrino burst from the supernova SN1987A”. In: *Phys. Rev. Lett.* 58 (14 Apr. 1987), pp. 1490–1493. DOI: 10.1103/PhysRevLett.58.1490. URL: https://link.aps.org/doi/10.1103/PhysRevLett.58.1490.

[80] Pietro Antonioli et al. “SNEWS: the SuperNova Early Warning System”. In: *New Journal of Physics* 6 (Sept. 2004), pp. 114–114. DOI: 10.1088/1367-2630/6/1/114. URL: https://doi.org/10.1088/1367-2630/6/1/114.

[81] Kohta Murase and Floyd W. Stecker. *High-Energy Neutrinos from Active Galactic Nuclei*. 2022. DOI: 10.48550/ARXIV.2202.03381. URL: https://arxiv.org/abs/2202.03381.

[82] R. Aloisio et al. “Cosmogenic neutrinos and ultra-high energy cosmic ray models”. In: *Journal of Cosmology and Astroparticle Physics* 2015.10 (Oct. 2015), pp. 006–006. DOI: 10.1088/1475-7516/2015/10/006. URL: https://doi.org/10.1088%2F1475-7516%2F2015%2F10%2F006.

[83] Francis Halzen and Spencer R. Klein. “Invited Review Article: IceCube: An instrument for neutrino astronomy”. In: *Review of Scientific Instruments* 81.8 (Aug. 2010), p. 081101. DOI: 10.1063/1.3480478. URL: https://doi.org/10.1063%2F1.3480478.

[84] S. Adriá et al. "SOURCES FOR POINT-LIKE AND EXTENDED NEUTRINO SOURCES CLOSE TO THE GALACTIC CENTER USING THE ANTARES NEUTRINO TELESCOPE". In: *The Astrophysical Journal* 786.1 (Apr. 2014), p. L5. DOI: 10.1088/2041-8205/786/1/l5. URL: https://doi.org/10.1088%2F2041-8205%2F786%2F1%2Fl5.

[85] "Precision electroweak measurements on the Z resonance". In: *Physics Reports* 427.5 (2006), pp. 257–454. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2005.12.006. URL: https://www.sciencedirect.com/science/article/pii/S0370157305005119.

[86] A. Aguilar et al. "Evidence for neutrino oscillations from the observation of $\overline{\nu}_e$ appearance in a $\overline{\nu}_\mu$ beam". In: *Phys. Rev. D* 64 (11 Nov. 2001), p. 112007. DOI: 10.1103/PhysRevD.64.112007. URL: https://link.aps.org/doi/10.1103/PhysRevD.64.112007.

[87] A. A. et al. Aguilar-Arevalo. "Updated MiniBooNE neutrino oscillation results with increased data and new background studies". In: *Phys. Rev. D* 103 (5 Mar. 2021), p. 052002. DOI: 10.1103/PhysRevD.103.052002. URL: https://link.aps.org/doi/10.1103/PhysRevD.103.052002.

[88] Basudeb Dasgupta and Joachim Kopp. "Sterile neutrinos". In: *Physics Reports* 928 (2021). Sterile neutrinos, pp. 1–63. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2021.06.002. URL: https://www.sciencedirect.com/science/article/pii/S0370157321002696.

[89] MicroBooNE Collaboration. *Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final State Topologies*. 2021. DOI: 10.48550/ARXIV.2110.14054. URL: https://arxiv.org/abs/2110.14054.

[90] MicroBooNE Collaboration. *Search for Neutrino-Induced Neutral Current Δ Radiative Decay in MicroBooNE and a First Test of the MiniBooNE Low Energy Excess Under a Single-Photon Hypothesis*. 2021. DOI: 10.48550/ARXIV.2110.00409. URL: https://arxiv.org/abs/2110.00409.

[91] Th. A. et al. Mueller. "Improved predictions of reactor antineutrino spectra". In: *Phys. Rev. C* 83 (5 May 2011), p. 054615. DOI: 10.1103/PhysRevC.83.054615. URL: https://link.aps.org/doi/10.1103/PhysRevC.83.054615.

[92] J. Schechter and J. W. F. Valle. "Neutrinoless double-β decay in $SU(2) \times U(1)$ theories". In: *Phys. Rev. D* 25 (11 June 1982), pp. 2951–2954. DOI: 10.1103/PhysRevD.25.2951. URL: https://link.aps.org/doi/10.1103/PhysRevD.25.2951.

[93] P.A. Zyla et al. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[94] M. et al. Agostini. "Final Results of GERDA on the Search for Neutrinoless Double-β Decay". In: *Phys. Rev. Lett.* 125 (25 Dec. 2020), p. 252502. DOI: 10.1103/PhysRevLett.125.252502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.252502.

[95] J Hartnell and. "Neutrinoless Double Beta Decay with SNO+". In: *Journal of Physics: Conference Series* 375.4 (July 2012), p. 042015. DOI: 10.1088/1742-6596/375/1/042015. URL: https://doi.org/10.1088/1742-6596/375/1/042015.

[96] Christiane S. Lorenz et al. "Reconstruction of the neutrino mass as a function of redshift". In: *Phys. Rev. D* 104 (12 Dec. 2021), p. 123518. DOI: 10.1103/PhysRevD.104.123518. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.123518.

[97] M. Aker et al. "Analysis methods for the first KATRIN neutrino-mass measurement". In: *Phys. Rev. D* 104 (1 July 2021), p. 012005. DOI: 10.1103/PhysRevD.104.012005. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.012005.

[98] M. et al. Aker and The KATRIN Collaboration. "Direct neutrino-mass measurement with sub-electronvolt sensitivity". In: *Nature Physics* 18.2 (Feb. 2022), pp. 160–166. ISSN: 1745-2481. DOI: 10.1038/s41567-021-01463-1. URL: https://doi.org/10.1038/s41567-021-01463-1.

[99] Ivan Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2020.9 (Sept. 2020), p. 178. ISSN: 1029-8479. DOI: 10.1007/JHEP09(2020)178. URL: https://doi.org/10.1007/JHEP09(2020)178.

[100] NuFit. *Global fit to experimental neutrino oscillation data*. 2022.03.26. 2022. URL: http://www.nu-fit.org/.

[101] Laura Zambelli, A Fiorentini, and T Vladisavljevic. "Towards T2K neutrino flux predictions using the replica target measurements by NA61/SHINE". In: *J. Phys.: Conf. Ser.* 888 (2017), 012067. 3 p. DOI: 10.1088/1742-6596/888/1/012067. URL: https://cds.cern.ch/record/2684641.

[102] T. T. Böhlen et al. "The FLUKA Code: Developments and Challenges for High Energy and Medical Applications". In: *Nucl. Data Sheets* 120 (2014), pp. 211–214. DOI: 10.1016/j.nds.2014.07.049.

[103] Saba Parsa. *Baby MIND detector first physics run*. 2020. arXiv: 2004.05245 [physics.ins-det].

[104] Y. Fukuda et al. “The Super-Kamiokande detector”. In: *Nucl. Instrum. Meth. A* 501 (2003). Ed. by V. A. Ilyin, V. V. Korenkov, and D. Perret-Gallix, pp. 418–462. DOI: 10.1016/S0168-9002(03)00425-X.

[105] A. Goldsack. “Reactor Neutrinos in Super-Kamiokande”. PhD Thesis. University of Oxford, 2022.

[106] *Multiphysics Software for Optimizing Designs*. URL: https://www.comsol.com/.

[107] S. Assylbekov et al. “The T2K ND280 off-axis pi–zero detector”. In: *Nucl. Instrum. Meth. A* 686 (2012), pp. 48–63. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.05.028. URL: http://www.sciencedirect.com/science/article/pii/S0168900212005153.

[108] P.A. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[109] P.-A. Amaudruz et al. “The T2K fine-grained detectors”. In: *Nucl. Instrum. Meth. A* 696 (2012), pp. 1–31. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.08.020. URL: http://www.sciencedirect.com/science/article/pii/S0168900212008789.

[110] N. Abgrall et al. “Time projection chambers for the T2K near detectors”. In: *Nucl. Instrum. Meth. A* 637.1 (2011), pp. 25–46. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2011.02.036. URL: http://www.sciencedirect.com/science/article/pii/S0168900211003421.

[111] D Allan et al. “The electromagnetic calorimeter for the T2K near detector ND280”. In: *Journal of Instrumentation* 8.10 (Oct. 2013), P10019–P10019. DOI: 10.1088/1748-0221/8/10/p10019. URL: https://doi.org/10.1088/1748-0221/8/10/p10019.

[112] S. Aoki et al. “The T2K Side Muon Range Detector (SMRD)”. In: *Nucl. Instrum. Meth. A* 698 (2013), pp. 135–146. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.10.001. URL: http://www.sciencedirect.com/science/article/pii/S0168900212011242.

[113] Antonin Vacheret et al. “The front end readout system for the T2K-ND280 detectors”. In: *2007 IEEE Nuclear Science Symposium Conference Record* 3 (2007), pp. 1984–1991.

[114] P. Armaudruz S. Ritt and K. Olchanski. *MIDAS (Maximum Integration Data Acquisition System)*. URL: https://midas.psi.ch.

[115] Yoshinari Hayato and Luke Pickering. “The NEUT Neutrino Interaction Simulation”. In: (June 2021). arXiv: 2106.15809 [hep-ph].

[116] René Brun et al. *GEANT: Detector Description and Simulation Tool; Oct 1994*. CERN Program Library. Long Writeup W5013. Geneva: CERN, 1993. DOI: 10.17181/CERN.MUHF.DMJ1. URL: https://cds.cern.ch/record/1082634.

[117] Melissa Clegg and NOvA collaboration. *NOvA 2020 official data release*. Aug. 2021. URL: https://publicdocs.fnal.gov/cgi-bin/ShowDocument?docid=17.

[118] K. Abe et al. "Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target". In: *Phys. Rev. D* 91.7 (2015), p. 072010. DOI: 10.1103/PhysRevD.91.072010. arXiv: 1502.01550 [hep-ex].

[119] Fred James and Matthias Winkler. "MINUIT User's Guide". In: (June 2004).

[120] W. K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57.1 (Apr. 1970), pp. 97–109. ISSN: 0006-3444. DOI: 10.1093/biomet/57.1.97. eprint: https://academic.oup.com/biomet/article-pdf/57/1/97/23940249/57-1-97.pdf. URL: https://doi.org/10.1093/biomet/57.1.97.

[121] Alfredo Ferrari et al. "FLUKA: A multi-particle transport code (Program version 2005)". In: (Oct. 2005). DOI: 10.2172/877507.

[122] N. Abgrall et al. "Measurements of $\pi^{\pm}$, $K^{\pm}$, K^0_S , Λ and proton production in proton–carbon interactions at 31 GeV/c with the NA61/SHINE spectrometer at the CERN SPS". In: *Eur. Phys. J. C* 76.2 (2016), p. 84. DOI: 10.1140/epjc/s10052-016-3898-y. arXiv: 1510.02703 [hep-ex].

[123] N. Abgrall et al. "Measurements of $\pi^{\pm}$ differential yields from the surface of the T2K replica target for incoming 31 GeV/c protons with the NA61/SHINE spectrometer at the CERN SPS". In: *Eur. Phys. J. C* 76.11 (2016), p. 617. DOI: 10.1140/epjc/s10052-016-4440-y. arXiv: 1603.06774 [hep-ex].

[124] N. Abgrall et al. "Measurements of $\pi^{\pm}$, $K^{\pm}$ and proton double differential yields from the surface of the T2K replica target for incoming 31 GeV/c protons with the NA61/SHINE spectrometer at the CERN SPS". In: *Eur. Phys. J.* C79.2 (2019), p. 100. DOI: 10.1140/epjc/s10052-019-6583-0. arXiv: 1808.04927 [hep-ex].

[125] S. Bhadra et al. "Optical transition radiation monitor for the T2K experiment". In: *Nucl. Instrum. Meth. A* 703 (2013), pp. 45–58. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.11.044. URL: https://www.sciencedirect.com/science/article/pii/S0168900212013812.

[126] K. Abe et al. “Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector”. In: *Nucl. Instrum. Meth. A* 694 (2012), pp. 211–223. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.03.023. URL: http://www.sciencedirect.com/science/article/pii/S0168900212002987.

[127] Bhubanjyoti Bhattacharya, Richard J. Hill, and Gil Paz. “Model independent determination of the axial mass parameter in quasielastic neutrino-nucleon scattering”. In: *Phys. Rev. D* 84 (2011), p. 073006. DOI: 10.1103/PhysRevD.84.073006. arXiv: 1108.0423 [hep-ph].

[128] M. Martini et al. “Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei”. In: *Phys. Rev. C* 80 (6 Dec. 2009), p. 065501. DOI: 10.1103/PhysRevC.80.065501. URL: https://link.aps.org/doi/10.1103/PhysRevC.80.065501.

[129] Torbjorn Sjostrand. “High-energy-physics event generation with PYTHIA 5.7 and JETSET 7.4”. In: *Computer Physics Communications* 82.1 (1994), pp. 74–89. ISSN: 0010-4655. DOI: http://dx.doi.org/10.1016/0010-4655(94)90132-5. URL: http://www.sciencedirect.com/science/article/pii/0010465594901325.

[130] A. Bodek and U. K. Yang. “Modeling neutrino and electron scattering cross-sections in the few GeV region with effective LO PDFs”. In: *AIP Conf. Proc.* 670.1 (2003). Ed. by U. Cotti, M. Mondragon, and G. Tavares-Velasco, pp. 110–117. DOI: 10.1063/1.1594324. arXiv: hep-ex/0301036.

[131] Arie Bodek and U. K. Yang. “A Unified model for inelastic e - N and nu - N cross-sections at all Q**2”. In: *AIP Conf. Proc.* 792.1 (2005). Ed. by Wesley H. Smith and Sridhara R. Dasu, pp. 257–260. DOI: 10.1063/1.2122031. arXiv: hep-ph/0508007.

[132] Bo Andersson. *The Lund Model*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 1998. DOI: 10.1017/CBO9780511524363.

[133] T. Yang et al. “A hadronization model for few-GeV neutrino interactions”. In: *The European Physical Journal C* 63.1 (Aug. 2009), pp. 1–10. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-009-1094-z. URL: http://dx.doi.org/10.1140/epjc/s10052-009-1094-z.

[134] C. Andreopoulos et al., GENIE Collaboration. “The GENIE Neutrino Monte Carlo Generator”. In: *Nucl. Instrum. Meth. A* A614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph].

[135] K. Abe et al. "Improved constraints on neutrino mixing from the T2K experiment with $\mathbf{3.13 \times 10^{21}}$ protons on target". In: *Phys. Rev. D* 103.11 (2021), p. 112008. DOI: 10.1103/PhysRevD.103.112008. arXiv: 2101.03779 [hep-ex].

[136] Roger Barlow and Christine Beeston. "Fitting using finite Monte Carlo samples". In: *Computer Physics Communications* 77.2 (1993), pp. 219–228. ISSN: 0010-4655. DOI: https://doi.org/10.1016/0010-4655(93)90005-W. URL: http://www.sciencedirect.com/science/article/pii/001046559390005W.

[137] J.S. Conway. "Incorporating Nuisance Parameters in Likelihoods for Multisource Spectra". In: arXiv:1103.0354 (2011). Comments: Presented at PHYSTAT 2011, CERN, Geneva, Switzerland, January 2011, to be published in a CERN Yellow Report, 115–120. 6 p. DOI: 10.5170/CERN-2011-006.115. URL: http://cds.cern.ch/record/1333496.

[138] Rene Brun and Fons Rademakers. "ROOT - An Object Oriented Data Analysis Framework, Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996". In: vol. 389. 1. See also "ROOT" [software], Release v5.34/34, 2015-10-02. 1997, pp. 81–86.

[139] Veronique Bernard, Latifa Elouadrhiri, and Ulf-G. Meissner. "Axial structure of the nucleon: Topical Review". In: *J. Phys. G* 28 (2002), R1–R35. DOI: 10.1088/0954-3899/28/1/201. arXiv: hep-ph/0107088.

[140] K. L. Miller et al. "STUDY OF THE REACTION $\nu_\mu d \to \mu^- pp_s$". In: *Phys. Rev. D* 26 (1982), pp. 537–542. DOI: 10.1103/PhysRevD.26.537.

[141] S. J. Barish et al. "Study of Neutrino Interactions in Hydrogen and Deuterium. 1. Description of the Experiment and Study of the Reaction Neutrino $\nu_\mu d \to \mu^- pp_s$". In: *Phys. Rev. D* 16 (1977), p. 3103. DOI: 10.1103/PhysRevD.16.3103.

[142] N. J. Baker et al. "Quasielastic Neutrino Scattering: A Measurement of the Weak Nucleon Axial Vector Form-Factor". In: *Phys. Rev. D* 23 (1981), pp. 2499–2505. DOI: 10.1103/PhysRevD.23.2499.

[143] D. Allasia et al. "Investigation of exclusive channels in neutrino / anti-neutrino deuteron charged current interactions". In: *Nucl. Phys. B* 343 (1990), pp. 285–309. DOI: 10.1016/0550-3213(90)90472-P.

[144] T. Kitagaki et al. "High-Energy Quasielastic Muon-neutrino $n \to \mu^- p$ Scattering in Deuterium". In: *Phys. Rev. D* 28 (1983), pp. 436–442. DOI: 10.1103/PhysRevD.28.436.

[145] P. Stowell et al. "Tuning the genie pion production model with MINERvA data". In: *Phys. Rev. D* 100 (7 Oct. 2019), p. 072005. DOI: 10.1103/PhysRevD.100.072005. URL: https://link.aps.org/doi/10.1103/PhysRevD.100.072005.

[146] R. Gonzalez-Jimenez et al. “Extensions of Superscaling from Relativistic Mean Field Theory: the SuSAv2 Model”. In: *Phys. Rev.* C90.3 (2014), p. 035501. DOI: 10.1103/PhysRevC.90.035501. arXiv: 1407.8346 [nucl-th].

[147] R. Gonzaléz-Jiménez et al. “Extensions of Superscaling from Relativistic Mean Field Theory: the SuSAv2 Model”. In: *Phys. Rev. C* 90.3 (2014), p. 035501. DOI: 10.1103/PhysRevC.90.035501. arXiv: 1407.8346 [nucl-th].

[148] et Al. K. Abe. *T2K ND280 Upgrade – Technical Design Report*. 2020. arXiv: 1901.03750 [physics.ins-det].

[149] K. Abe et al. *First Gadolinium Loading to Super-Kamiokande*. 2021. arXiv: 2109.00360 [physics.ins-det].

[150] The Super-K Collaboration. *Gadolinium neutron capture diagram*. Accessed: 2021-10-12. URL: http://www-sk.icrr.u-tokyo.ac.jp/sk/news/2018/06/skopen02-en.html.

[151] Dervis Karaboga. “An idea based on honey bee swarm for numerical optimization”. In: 2005.

[152] Omar Benhar et al. “Electron- and neutrino-nucleus scattering in the impulse approximation regime”. In: *Phys. Rev. D* 72 (5 Sept. 2005), p. 053005. DOI: 10.1103/PhysRevD.72.053005. URL: https://link.aps.org/doi/10.1103/PhysRevD.72.053005.