Neutrino induced events in the MINOS detectors

Reuben Phillip Litchfield

Keble College, Oxford

Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy
at the University of Oxford

Hilary Term, 2008

Neutrino induced events in the MINOS detectors

Reuben Phillip Litchfield
Keble College, Oxford

Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy
at the University of Oxford.

Hilary Term, 2008

Abstract

The MINOS experiment is designed to study neutrino oscillations. It uses an accelerator generated beam of neutrinos and two detectors, the smaller at a distance of 1 km and the larger at 735 km. By comparing the spectrum and flavour composition of the beam at the two detectors precise determinations of the oscillation parameters are possible.

This thesis concentrates on the analysis of data from the larger Far Detector. By studying the spectrum of neutral current events it is possible to look for evidence of non-interacting 'sterile' neutrinos. The thesis describes how events are selected for this analysis, and a method for discriminating between charged current and neutral current events. The systematic uncertainties resulting from these cuts are evaluated. Several techniques for using Near Detector data to eliminate systematic uncertainties in the predicted Far Detector spectrum are compared.

An oscillation analysis, based on the first year of MINOS data, uses the selected events to make a measurement of f_{s}, the fraction of unseen neutrinos that are sterile. The measured value is $f_{\mathrm{s}} = 0.07^{+0.32}$ at 68%C.L. , and is consistent with the standard three-neutrino picture, which has no sterile neutrino.

Acknowledgements

Thank yous to those who should be thanked

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Of all the constituents of the Standard Model, neutrinos are perhaps the most elusive. Although their history stretches back almost as far as the birth of quantum mechanics, our understanding of their nature has changed radically in the last few decades. We now know that neutrinos have mass, and oscillate between the three known flavours. Neutrino physics is now a thriving field, with many experiments either planned or running.

The Main Injector Neutrino Oscillation Search (MINOS) is one such experiment. It uses two functionally identical detectors, separated by 734 km to characterise a beam of neutrinos. By comparing the energy spectrum and flavour composition in these detectors it is possible to make detailed measurements of the parameters governing the flavour oscillations.

This thesis describes a study of an extended model of neutrino oscillations, in which there are additional 'sterile' neutrinos that do not interact except through oscillations with the three known 'active' flavours.

Chapter 2 of starts with the history of neutrino oscillation, and introduces the major experimental results. A derivation of the oscillation formula is presented, followed by a sketch of how the many experimental results are incorporated in the standard interpretation. Then a simple four-flavour extension is introduced to demonstrate that a sterile neutrino can be accommodated by current data. The chapter ends with a brief summary of future experiments in the field.

The MINOS experiment is introduced in detail in Chapter 3. The NuMI beamline is described first, followed by a description of the detector technology. A brief outline of the analysis software, including the Monte Carlo simulation, is given, followed by a description of the oscillation physics that can be studied with MINOS.

In Chapter 4, a description of how events in the larger 'Far' Detector are selected for the sterile neutrino analysis is given, followed by estimates of the systematic uncertainties relating to this selection. The chapter then details how the selected events are separated into charged current and neutral current samples for analysis.

Chapter 5 takes a slightly different direction, and describes another part of the analysis—how data from the 'Near' Detector is used to improve the prediction of the spectrum of events in the Far Detector energy spectrum. It describes four different methods that have been used in the analysis of charged current events. The resilience of these different approaches to various systematic uncertainties is evaluated, and more general comparisons are made between methods.

Chapter 6 brings together the selection described in Chapter 4 and the techniques introduced in Chapter 5. A study is done to evaluate the systematic uncertainties affecting the sterile oscillation analysis. Finally the first year of MINOS data is analysed, to look for evidence of oscillations to sterile neutrinos.

Chapter 2

Neutrino physics

2.1 A history of the neutrino

The history of the neutrino arguably began on the 4th of December, 1930 with a letter written by Wolfgang Pauli to Lise Meitener and the participants of a radioactivity conference in Tübingen.[1] At the time it appeared that β-decay processes did not conserve angular momentum or even energy, and Pauli proposed the existence of an electrically neutral, spin $\frac{1}{2}$ particle that was emitted along with the electron in β-decays. Pauli was wary of postulating a particle that had somehow escaped all observation, and the idea was next developed by Enrico Fermi in 1934 as part of his theory of β-decay.[2] In this paper Fermi gave Pauli's particle a name, the neutrino, and a way to interact with other particles, shown in Figure 2.1(L). The success of Fermi's theory went a long way towards establishing the neutrino as one of nature's building blocks, but it was another 22 years before it would be directly observed.

[1] Typed copy from CERN's Pauli Letter Collection [1]. Reproduced in English in [2]

[2] Original German [3] and English translation from 1968 [4]

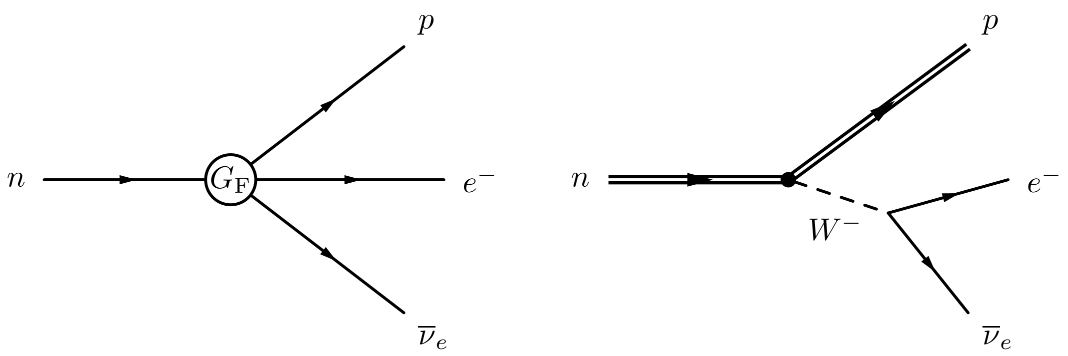

Figure 2.1: (L) Diagram of the β-decay process $n \to p + e^- + \overline{\nu}_e$ in Fermi's theory, and **(R)** the same process in the Standard Model.

In the intervening time there were several developments that would eventually impact neutrino physics. In the mid-forties it was suggested by Sakata and Inoue [5] and independently by Marshak [6] that Yukawa's meson [7] was in fact two distinct particles; the pion and the muon.[3] The suggestions were confirmed by the observations of Lattes *et al.* [8]. It was quickly realized that the decay of pions to muons was a similar process to β-decay and that the muons were accompanied by a neutrino. What was not known was whether this was the same neutrino as that produced in ordinary β-decay.

The field of experimental neutrino physics began in 1953 when Reines and Cowan made the first tentative observation [9] of (anti)neutrinos from the Hanford reactor, via the reaction $\overline{\nu}+p \rightarrow n+e^{+}$. In 1955 Ray Davis demonstrated [10] that no signal was seen in an experiment (also based at Hanford) that was sensitive to the reaction $\nu + {}^{37}\mathrm{Cl} \rightarrow e^{-} + {}^{37}\mathrm{Ar}$. Confirmation of antineutrino observation came in 1956 when Reines *et al.* performed a more sensitive version of their earlier experiment at Savannah River. [11] The next year Goldhaber *et al.* published the result [12] of a widely-admired experiment, which determined the helicity of the neutrino to be negative (left-handed). In 1962, Lederman, Schwartz, Steinberger *et al.* observed neutrinos produced in the decay of pions ($\pi^{\pm} \rightarrow \mu^{\pm} + \nu/\overline{\nu}$). [13] These neutrinos produced muons, not electrons, indicating that they were distinct from the neutrinos involved in β-decay.

Parity violation, charged and neutral currents

At about the same time as Reines and Cowan's first observations, a theoretical paper was published that had startling implications. Lee and Yang pointed out the possibility that parity ($\mathcal{P}$, symmetry under spatial inversion) might not be conserved in the weak interaction, and this would explain some puzzling features of the decay of strange particles (kaons). [14] Their hypothesis was quickly confirmed by Wu *et al.* [15] and subsequently Garwin *et al.* [16]. At the same time, it was realised that these results also imply the violation of charge conjugation symmetry ($\mathcal{C}$). [17] The discovery of parity violation lead to the suggestion by Sudarshan and Marshak [18], and by Feynman and Gell-Mann [19] that weak interactions have have a $V{-}A$ structure. This form results in interactions that maximally violate $\mathcal{P}$ and $\mathcal{C}$ but preserve $\mathcal{CP}$ and time reversal symmetry ($\mathcal{T}$). The fact that $\mathcal{P}$ is maximally violated also reopened the question of whether the neutrino and antineutrino are really distinct.[4] Davis's 1955 result can instead be explained by the fact that the weak interaction does not couple the particles involved.

By the time the $V{-}A$ model was developed it was realised that the point

[3]The paper by Sakata and Inoue is particularly interesting. It distinguishes between between spin 1 'Yukawa particles' and spin $\frac{1}{2}$ 'mesons' and postulates the existence of a neutral counterpart to the meson with negligible mass "equivalent [to] the neutrino"...

[4]I follow common practice and continue to distinguish between the two. It is easier to ignore the 'anti-' prefix than work out where it should have gone.

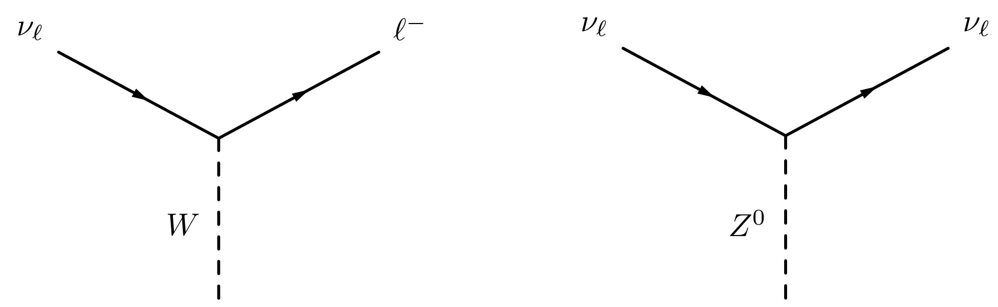

Figure 2.2: Standard Model interactions of the neutrino.

interaction of Fermi's theory must be a low energy approximation, and that the weak interaction was probably due to the exchange of a massive spin 1 boson, as shown in Figure 2.1(R). But simply putting in a mass for the boson lead to divergences in the theory. The solution was developed over the next decade, into what is now known as the Glashow–Weinberg–Salam (GWS) model. [20] As well a providing a mass for the mediating $W^\pm$ bosons, it also unified the weak interaction with electromagnetism into a single electroweak force. The theory also made a prediction: the existence of a massive neutral counterpart to the $W^\pm$, the Z^0 boson. Z^0-mediated *neutral current* interactions would be rather like those of a heavy (therefore short-ranged) photon, but unlike electromagnetic interactions they could involve the neutrino. This made a neutrino beam the obvious place to look for evidence of the neutral current and in 1973 it was seen by the Gargamelle collaboration; both in the reaction $\nu_\mu + e^- \to \nu_\mu + e^-$ [21] and in the reaction $\nu_\mu + N \to \nu_\mu + X$ [22]. Confirmation, and a measurement of the Z^0 mass, came in 1983 from the experiments at CERN's Sp$\bar{\text{p}}$S. [23] Combined with a model of strong interactions, the GWS model forms the basis of what is now known as the standard model. For neutrinos the standard model is unchanged from the GWS model and allows only interactions with the $W^\pm$ and Z^0, bosons, as shown in Figure 2.2.

Flavour change and oscillations

The history of neutrino oscillations proceeded initially quite separately from the development of the theory of weak interactions. It can be said to have begun with a suggestion from Bruno Pontecorvo in 1957. Inspired by papers [24, 25] by Pais, Gell-Mann and Piccioni that predicted mixing[5] between K^0 and $\overline{K}^0$, Pontecorvo wondered if any other system might exhibit particle/antiparticle transitions. He first considered the bound $(\mu^\pm e^\mp)$ system [26], then neutrino/antineutrino transitions [27] and concluded that mixing would be inevitable if not prohibited by some kind of neutrino charge conservation. In 1967, after the discovery of the

[5] The second paper (by Pais and Piccioni) mentions an oscillating term

muon neutrino, he suggested the possibility of flavour-changing $\nu_\mu \leftrightarrow \nu_e$ oscillations. [28]

In 1963 Cabibbo proposed an explanation [29] for the fact that strangeness is not always conserved in weak interactions (e.g. in $K^+ \rightarrow \pi^+ + \pi^0$). Cabibbo suggested that the weak interaction coupled to a mixture of strangeness conserving and strangeness violating hadronic currents.[6] Cabibbo's model is now understood in terms of mixing between quark states, so that the u couples not to the d, but a mixture of d and s.

In 1970 Glashow, Iliopoulos and Maiani showed that by adding a fourth quark to Cabibbo's model, flavour changing neutral currents could be suppressed. [31] Three years later Kobayashi and Maskawa noted [32] that the $\mathcal{CP}$ violation observed in the neutral kaon system by Cronin, Fitch *et al.* [33] could be explained by including a third generation of quarks. The unitary matrix that transforms between the strong and weak quark eigenstates is now known as the CKM matrix (for Cabibbo, Kobayashi and Maskawa). An equivalent matrix for the lepton sector was thought not to be necessary because neutrinos were assumed to be massless so could be defined only by their weak mixing.

The first hints that there was more to the neutrino sector than had been assumed came in the late 1960s from the Homestake mine, when Davis measured the interaction rate of solar neutrinos (produced in nuclear reactions in the sun's core), in a larger version of the ^{37}Cl apparatus he used in 1955. [34] Davis worked closely with theorist John Bahcall, whose predictions of the expected rate were about three times higher than the measured rate. At the time it was generally thought there was 'something wrong with the theory, the experiment, or both'. Further refinements of both theory and experiment improved the precision to which the discrepancy was known, but did not make it go away. This became known as the solar neutrino problem.

In 1989, the Kamioka Nucleon Decay Experiment (Kamiokande) provided independent confirmation [35] of the Homestake result, detecting solar neutrinos by the Čerenkov light emitted by electrons produced in neutrino interactions in a 0.68 kt fiducial volume of water. This was followed in 1991 by results from SAGE [36] and GALLEX (later GNO) [37], two experiments based on the reaction $\nu_e + {}^{71}\mathrm{Ga} \rightarrow e^- + {}^{71}\mathrm{Ge}$. These three experiment all saw about half as many neutrinos as predicted by the standard solar model. This largely ruled out experimental error as a cause of the solar neutrino problem, but it was ten years later (and 30 years after the discrepancy was identified) that Bahcall's predictions were finally confirmed.

At around the same time another 'neutrino problem' was coming to light. Kamiokande [38], and another proton decay experiment, IMB [39], were inves-

[6] A paper the previous year by Maki, Nakagawa and Sakata [30] attempted to explain strangeness non-conservation using the Nagoya model of baryons. In this model baryons are composed of some kind of 'baryonic matter' plus a lepton, so the mixing introduced in the paper was between ν_e and ν_μ.

tigating their backgrounds due to neutrinos produced in the upper atmosphere. They found the rate of interactions of ν_μ was only about 60% of what they expected. By comparing this rate to the rate of ν_e interactions they were able to rule out an error in the flux calculations. In the late 1980s two other experiments, Fréjus [40] and NUSEX [41], reported results consistent with no depletion but their statistical errors were too large to convincingly refute the earlier results. It was speculated that the problem lay with water-Čerenkov experiments. This became known as the atmospheric neutrino problem.

The fact that depletion of specific flavours of neutrinos had been observed from two unrelated sources lent support to the idea flavour change could be responsible for both phenomena. The oscillation mechanism suggested by Pontecorvo was a leading candidate for such a flavour change mechanism. Like kaon oscillations it could be parameterised by an amplitude and a length scale, but compared to oscillations in the quark sector the parameters of oscillation could be related to the mixing of neutrino states in a very direct way. The amplitude of oscillation corresponded to the amount of mixing between mass and flavour eigenstates, while the length scale (for neutrinos of a given energy) could be related to the difference in the mass-squared of the mass eigenstates. The main perceived difficulties with oscillations as explanation of the two 'neutrino problems' were also related to these quantities. Firstly the neutrino was supposed, in the standard model, to be massless and therefore have an infinite oscillation length. Secondly, it was expected that the neutrino mixing matrix would be the same as, or at least similar to, the CKM matrix. The amount of mixing required to solve either neutrino problem would be much larger than observed in the quark sector.

In 2002 the Sudbury Neutrino Observatory (SNO) measured the flux of ^{8}B solar neutrinos using the reaction $\nu + d \rightarrow p + n + \nu$. [42] This was the first measurement of the solar neutrino flux that was equally sensitive to all active flavours[7] of neutrino. This measurement of the flux *did* agree with the solar models, leaving flavour change between active neutrinos as the most plausible explanation of the solar neutrino problem. We now believe that electron neutrinos coming from the sun are depleted by the MSW effect [45, 46], a resonant phenomenon related to neutrino oscillations, that occurs for neutrinos passing through matter.

Another way to study the mixing responsible (assuming $\mathcal{CPT}$) for the solar neutrino problem is to look at electron antineutrinos produced in nuclear fission. Several experiments were sited near ($\sim 1\,\mathrm{km}$ from) nuclear reactors in order to search for neutrino oscillations. None of these experiments saw any depletion of the $\overline{\nu}_e$ flux. However, the limits set by these experiments (in particular CHOOZ [47] and Palo Verde [48]) would turn out to be relevant to the *atmospheric* neutrino problem.

The reason that none of these reactor experiments saw any evidence of mixing was that they were simply too close to their neutrino sources. The Kamioka Liquid

[7] An earlier SNO paper in 2001 [43], in combination with results from the Super-Kamiokande experiment [44], provided evidence that the flux of ^{8}B neutrinos contained a non-ν_e component.

AntiNeutrino Detector (KamLAND) experiment was sensitive to the same transition, but placed much further (~ 180 km) from the reactors that were its source of antineutrinos. In 2005, KamLAND reported not only only seeing depletion of the $\overline{\nu}_e$ flux, but also an energy dependence, shown in Figure 2.3 that is characteristic of the neutrino oscillations. [49]

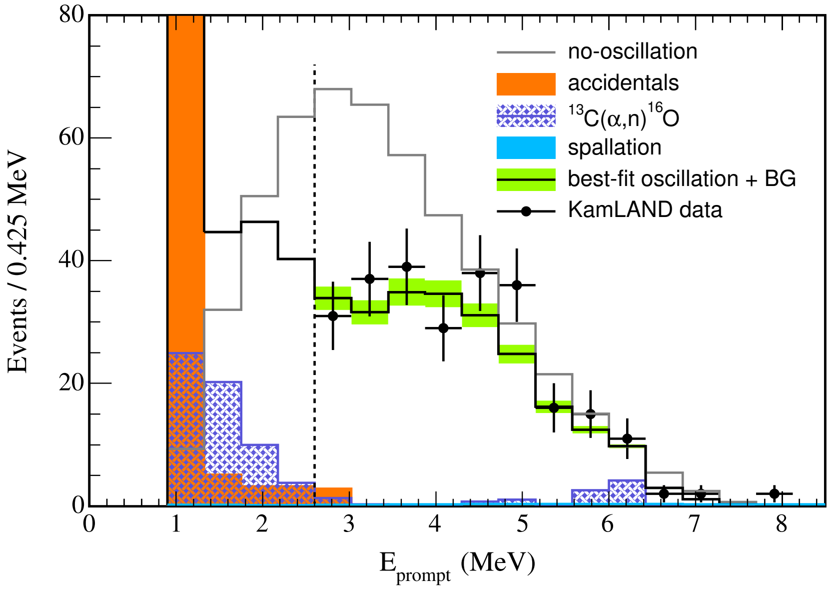

Figure 2.3: KamLAND prompt energy spectrum with best-fit oscillation spectrum shown in solid black. The green bands indicate the systematic error. Figure from [49].

By this time neutrino flavour change was also well accepted as the explanation of the atmospheric neutrino problem. In 1996 the Soudan-2 collaboration published a result [50] that was in agreement with the IMB and Kamiokande experiments. Soudan-2 was an iron tracking calorimeter, so this result ruled out a systematic problem affecting water-Čerenkov experiments. The most convincing evidence came in 1998 from Super-Kamiokande, the gigantic successor to the Kamiokande experiment. Both experiments are very similar in construction, being cylindrical tanks of water surrounded by Photomultiplier Tubes. But at 22.5 kt, Super-Kamiokande has a fiducial volume that is around 30 times larger than its predecessor, providing much higher statistics and the containment of higher energy events. Super-Kamiokande's most recent results [51] are shown in Figure 2.4, which shows interaction rates for electron and muon type neutrinos, separated into two ranges of energy and binned as a function of the zenith angle ϕ_z. Neutrinos coming from above have a $\cos\phi_\mathrm{z} \simeq 1$ and travel around 30 km before passing through the detector. Those coming from below have $\cos\phi_\mathrm{z} \simeq -1$ and travel around 13 000 km before passing through the detector. The observed zenith angle dependence in the ν_μ sample is therefore a dependence on the distance travelled by the neutrino—a hallmark of flavour change. The dependency is well fit by oscillations, although a few other models of flavour change can also provide acceptable agreement with data. These models differ in how they depend on L/E, the distance travelled by a neutrino divided by its energy. A recent analysis by the Super-Kamiokande collaboration [52] uses only events which have a good resolution in L/E. This analysis is shown in Figure 2.5 and favours the oscillation model.

Neutrino oscillation physics

With neutrino mass and oscillation now widely accepted the emphasis has turned to measuring the parameters that govern the oscillations. Figure 2.6 shows a summary of neutrino oscillation experiments up to 2006. Experiments reporting oscillations signals have fallen into three regimes, corresponding to the mass splitting that is observed and (to a lesser extent) the neutrino flavours involved. The first of these is the solar regime, as these oscillations are responsible for the solar neutrino problem. The second is, in an analogous way, called the atmospheric regime. The third is called the the LSND regime, after the only experiment where an effect has been observed. In each regime oscillations can be parameterised by a oscillation amplitude $\sin^2 2\vartheta$, and mass (squared) splitting δm^2, in vacuum. Assuming that only two neutrinos participate in each regime the probability of oscillation from flavour α to flavour β is given by[8]:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2 2\vartheta \sin^2 \left(\frac{1.267(\delta m^2/\,\mathrm{eV}^2)(L/\,\mathrm{km})}{(E/\,\mathrm{GeV})} \right). \tag{2.1}$$

In the solar regime we identify the parameters as ($\Delta m^2_\odot$, $\sin^2 2\theta_\odot$), and their measurement comes from several sources: the original Chlorine and Gallium experiments, Super-Kamiokande, SNO and KamLAND, although it is dominated by the last two. In Figure 2.6 these contributions can be seen (below $\delta m^2 \sim 10^{-3}$). The triangular shape of the Chlorine and Gallium allowed regions is due to the MSW effect, which produces a similar flavour composition emerging from the sun over a wide range of vacuum parameters. Super-Kamiokande and SNO are able to measure the angle and energy of incident neutrinos which allows them to look for day-night effects and distortions in the energy spectrum, which helps to further constrain the mixing parameters [53]. Finally the KamLAND experiment is able to give a very precise measurement of the mass splitting because it can observe the shape of the oscillated energy spectrum directly, without the complication of the MSW effect. The resulting limits can be seen in Figure 2.7. The world average [54] for the solar parameters is:

$$\Delta m^2_\odot = \left(8.0^{+0.6}_{-0.4}\right) \times 10^{-5}\,\mathrm{eV}^2, \qquad \theta_\odot = \left(33.9^{+2.4}_{-2.2}\right)^\circ. \tag{2.2}$$

[8] This formula is derived (excluding numerical factors) in Section 2.3.3

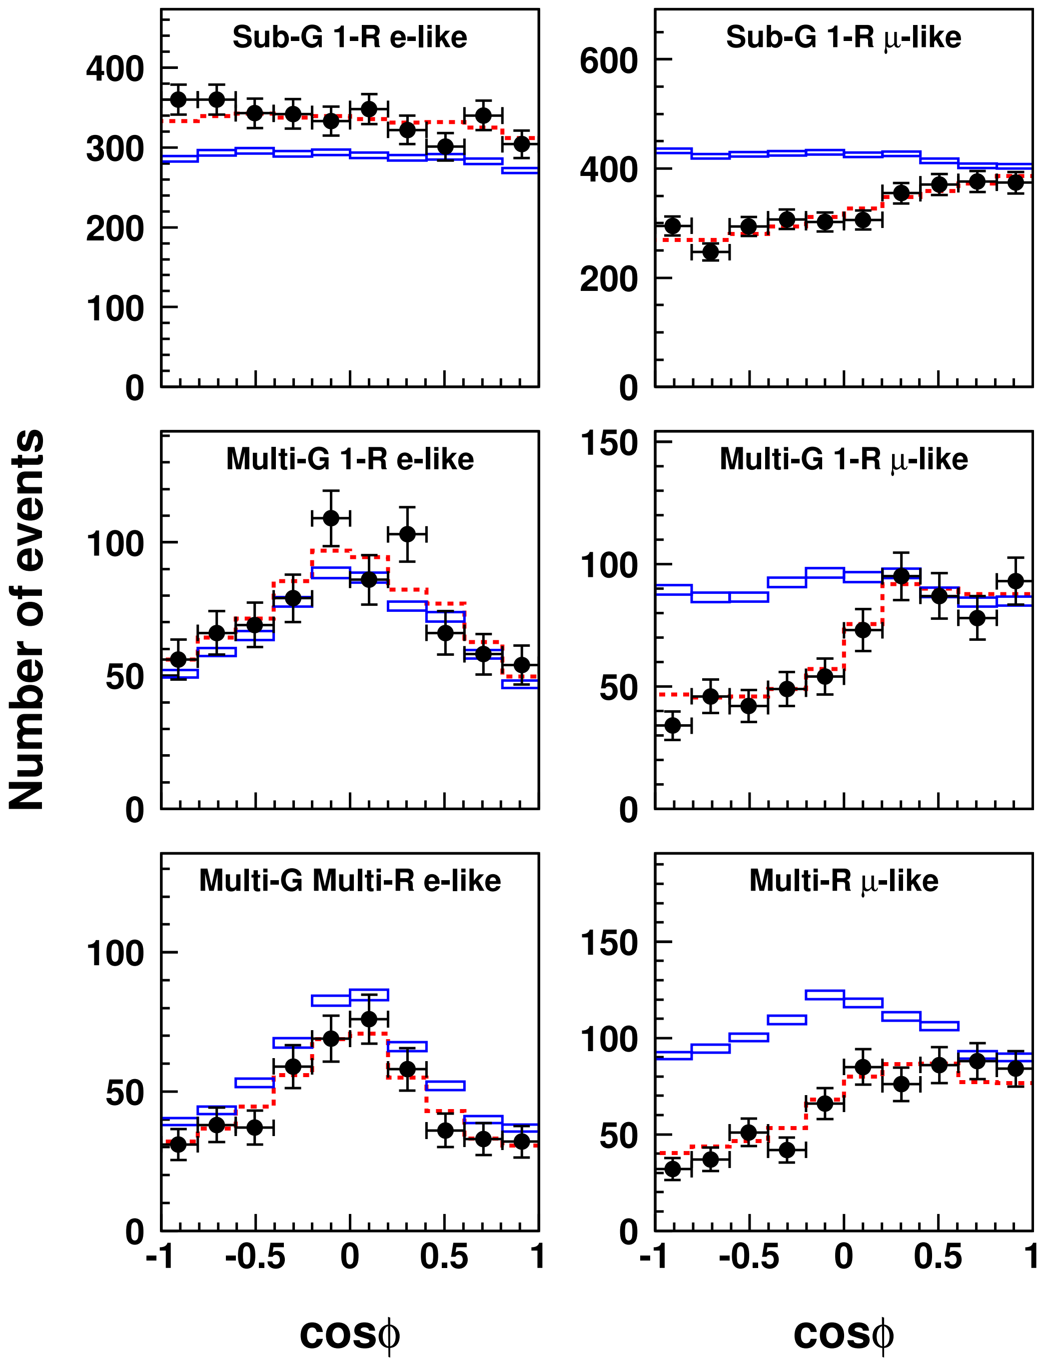

Figure 2.4: Zenith angle distributions of e-like and μ-like neutrino events in Super-Kamiokande. The markers are data (with statistical error bars). The dashes and boxes represent the predicted distributions with and without oscillations respectively. The box height shows the statistical error on the predicted spectrum. Figure adapted from [51].

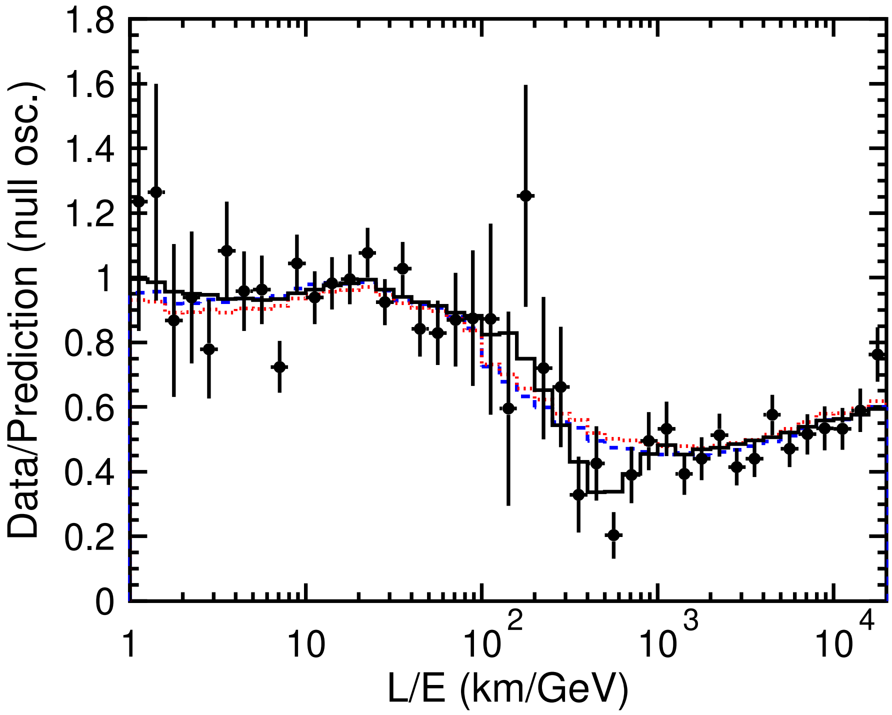

Figure 2.5: Ratio of ν_μ-like events seen in Super-Kamiokande to the number expected without flavour change, as a function of reconstructed L/E. Markers show the observed ratio with statistical errors. Lines show the expectation for neutrino oscillations (black solid), neutrino decay (blue dashed), and neutrino decoherence (red dotted). Decoherence and decay are described in Section 2.4. Figure from [52].

In the atmospheric regime, measurements have been made by the Soudan-2 and Super-Kamiokande experiments mentioned previously, and by the MACRO collaboration [55]. These experiments use atmospheric neutrinos and are primarily sensitive to the mixing angle $\sin^2 2\theta_{\rm atm}$. Another approach that is primarily sensitive to $\Delta m^2_{\rm atm}$ is taken by the K2K [56] and MINOS experiments. Both these experiments look for muon neutrino disappearance use a beam of ν_μ produced by the decay of charged pions. A detector is sited some distance (K2K: 250 km, MINOS: 735 km) away from the beam source so the beam energy is chosen so maximise the oscillation signal at the detector site. Because there are a range of neutrino energies in the beam and the distance travelled is known the L/E dependence of the oscillation probability can be observed, which makes these experiments good at measuring $\Delta m^2_{\rm atm}$. Prior to the MINOS result, the allowed range [54] for the atmospheric parameters were (90% C.L.):

$$1.9 \times 10^{-3} < \left|\Delta m^2_{\rm atm}\right| < 3.0 \times 10^{-3}, \qquad \sin^2 2\theta_{\rm atm} > 0.9. \tag{2.3}$$

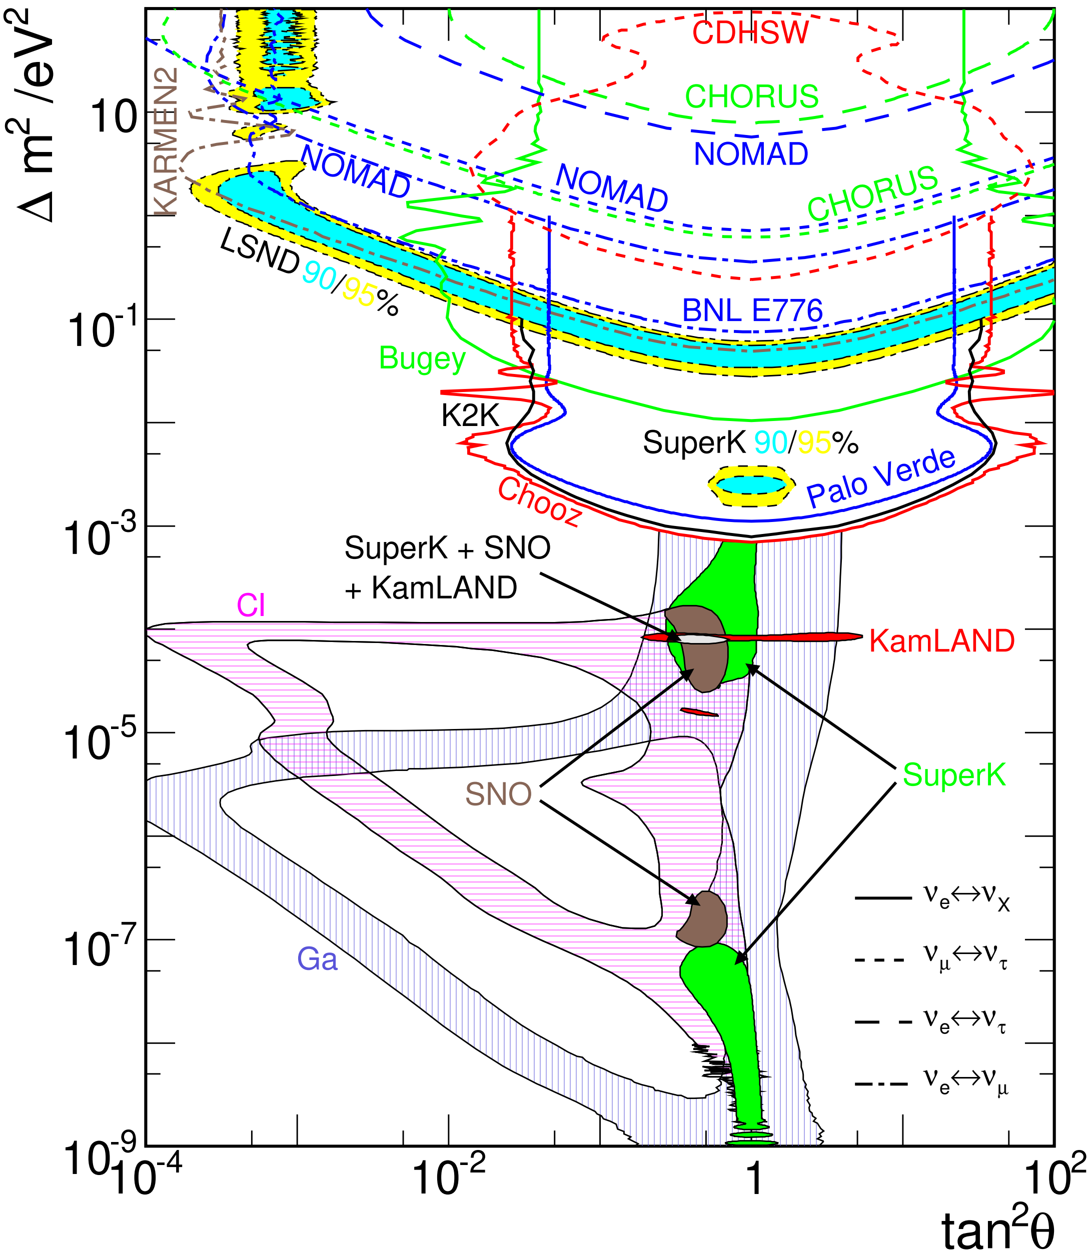

Figure 2.6: Summary of neutrino oscillation experiments up to 2006. The flavour states involved are identified by the line style. Filled regions denote a positive signal at 95% C.L. unless specified, whilst lines denote limits at 90% C.L. Figure adapted from *hitoshi.berkeley.edu/neutrino/*.

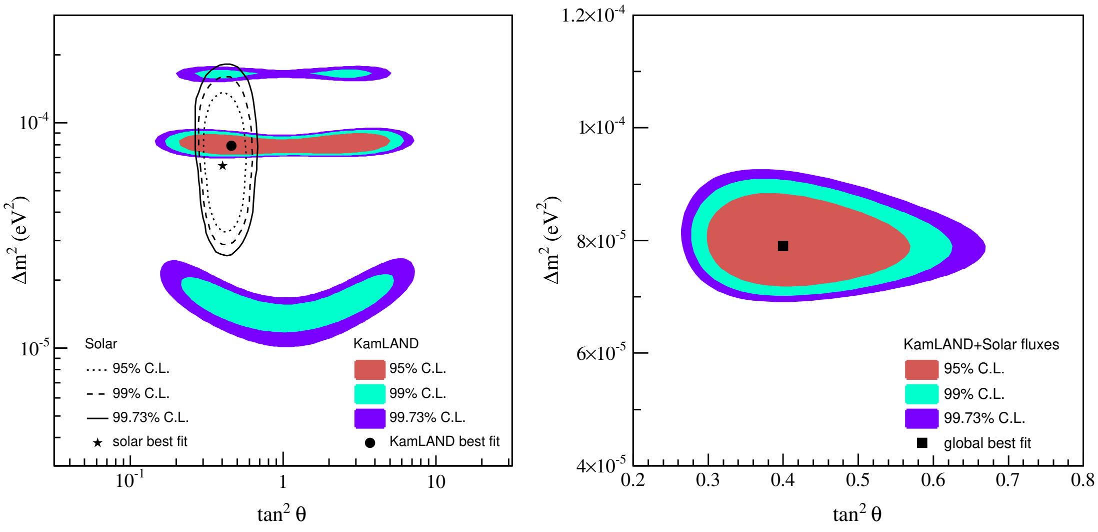

Figure 2.7: Allowed values of $\Delta m^2_\odot$ and $\sin^2 2\theta_\odot$, as determined by SNO, KamLAND, and Super-Kamiokande experiments. Figure from [49].

The first results from MINOS [57], suggest that $|\Delta m^2_{\rm atm}|$ is in the upper part of this range. From MINOS alone, the parameters are (again 90% C.L.):

$$2.31 \times 10^{-3} < \left|\Delta m^2_{\rm atm}\right| < 3.43 \times 10^{-3}, \qquad \sin^2 2\theta_{\rm atm} > 0.78. \tag{2.4}$$

The 2D confidence limits on the atmospheric parameters are shown in Figure 2.8.

In the case of solar neutrinos, we can only observe the disappearance of electron neutrinos, as the muon and tau are too massive to be created in neutrino interactions on a fixed target. However in the atmospheric sector it is possible to observe all three types of lepton. Observation of the tau is difficult because of its short lifetime and the fact that the energy required to make a tau is generally higher than the oscillation maxima unless the baseline is very long. However electrons could be produced, so their non-observation indicates that the ν_e component of one of the mass states involved in atmospheric oscillations is very small or zero.

As indicated previously, the L/E scale of the CHOOZ and Palo Verde reactor experiments makes them sensitive to oscillations at $\Delta m^2_{\rm atm}$, and the non-observation of oscillations in these experiments provides the most sensitive limit on the parameter θ_{13}, which controls the amount of ν_e involved in atmospheric oscillations. A global analysis [58] using this and other data gives (90% C.L.):

$$\sin^2(\theta_{13}) < 0.02 \quad \left(\sin^2(2\theta_{13}) < 0.07\right) \tag{2.5}$$

The third regime in which evidence for oscillations is claimed is characterised by a substantially larger mass-squared splitting, of the order of 1 eV². The signal seen be the Liquid Scintillator Neutrino Detector (LSND) experiment is an *appearance*

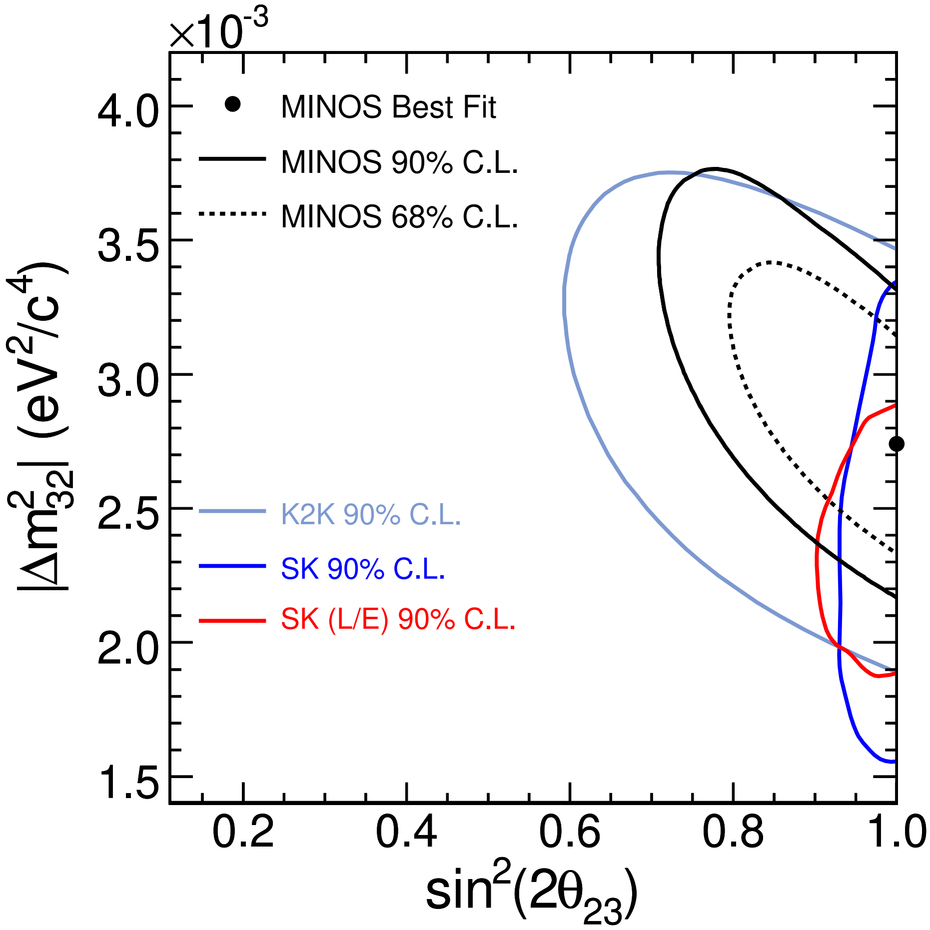

Figure 2.8: Allowed values of $\Delta m^2_{\rm atm}$ and $\sin^2 2\theta_{\rm atm}$, as determined by K2K, MINOS, and Super-Kamiokande experiments.

of $\overline{\nu}_e$ in a beam of $\overline{\nu}_\mu$. [59] LSND saw a significant $(87.9 \pm 22.4 \pm 6.0)$ excess of $\overline{\nu}_e$-like events. The range of parameters allowed by LSND is quite large, as shown on Figure 2.9, but all allowed values of $|\Delta m^2_{\rm LSND}|$ are much larger than $|\Delta m^2_{\rm atm}| + |\Delta m^2_\odot|$, which is the maximum value consistent with a three neutrino framework.

Several other experiments in the same regime, notably the KARMEN2 [60], NOMAD [61], and Bugey[9] [62] experiments, have failed to corroborate the LSND result. A combined analysis of data from LSND and KARMEN2 [63] significantly reduces the size of the allowed region, as shown in Figure 2.9(L). As of April 2007, results [64] from the MiniBooNE experiment refute a simple (2-mass ap-

[9]Bugey is a rector experiment, sensitive to $\overline{\nu}_e \to \overline{\nu}_e$, and the applicability of this result is model dependent.

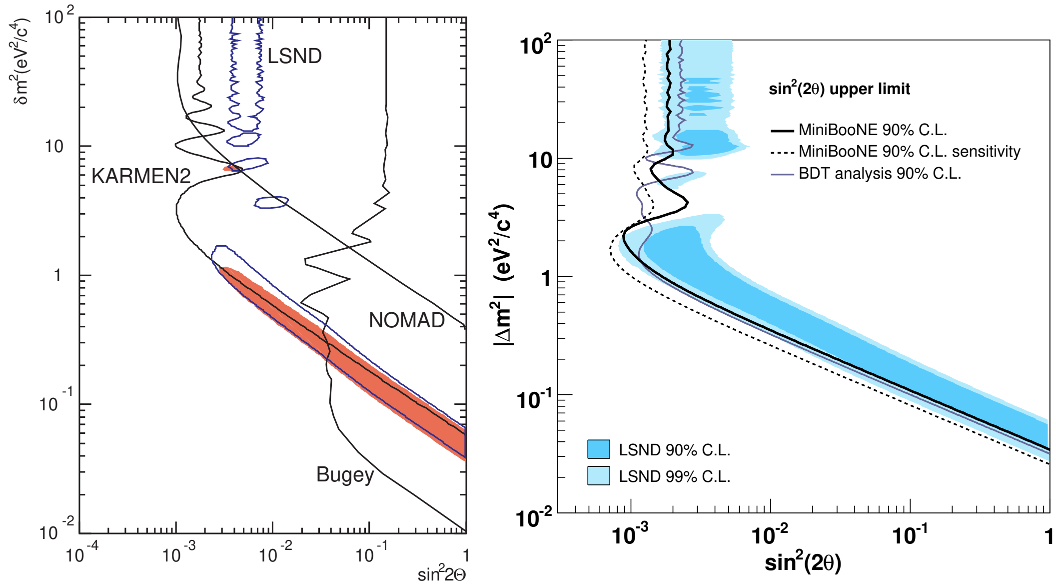

Figure 2.9: **(L)** 90% C.L. regions from the KARMEN2 and LSND experiments as calculated in [63]. The joint 90% region (assuming compatibility) is shown in red. Also shown are the 90% C.L. exclusion regions from NOMAD and Bugey. Figure adapted from [63]. **(R)** MiniBooNE excluded region (and alternative analysis).

proximation[10]) oscillation explanation of LSND result, as shown in Figure 2.9(R).

2.2 Neutrino masses

Although the phenomena of neutrino flavour change does not necessarily imply that neutrinos oscillate, one thing is unavoidable: neutrinos must have masses. This conclusion comes direct from special relativity, since a massless particle travels at the speed of light and therefore cannot measure the passage of time. The fact that neutrinos change as they propagate demonstrates that they *do* 'experience time' and therefore cannot be massless. This is probably the most important change to our understanding of neutrinos since they were discovered.

Massive neutrinos are considerably more interesting than massless ones. The question is immediately raised: why are the masses so small? The limits on neutrino mass coming from the β-decay of the triton $t \to d + p + e^- + \overline{\nu}_e$ is $m_\nu < 2\,\mathrm{eV}$ (90% C.L.) [54]. Even if neutrino masses are only just below this limit,

[10]The MiniBooNE result is based on the non-observation of $\nu_\mu \to \nu_e$ oscillations. Oscillation models with more than two neutrinos allow $\mathcal{CP}$ violation and may be able to accommodate both results.

that is 10^6 time smaller than the mass of the electron. In contrast the top quark mass, commonly regarded as surprising, is only around 40 times bigger than that of the bottom quark.

The smallness of neutrino masses suggests their origin may be somewhat different from that of the other fermions in the Standard Model. One possible explanation is the see-saw mechanism [65], which (in its simplest incarnation) is not to difficult to explain, although it requires a small amount of Quantum Field Theory. In the Standard Model, normal (Dirac) fermions gain masses through the coupling of right and left handed components to the Higgs field, with terms in the Lagrangian of the form:

$$\mathcal{L}_{\mathrm{D}} = -m_{\mathrm{D}} \left(\overline{\psi}_{\mathrm{R}} \psi_{\mathrm{L}} + \overline{\psi}_{\mathrm{L}} \psi_{\mathrm{R}} \right) . \tag{2.6}$$

The Dirac mass m_{D} is actually a coupling to the Higgs field and has gauge quantum numbers (Electroweak charges). The two fields ψ_{L} and ψ_{R} are the left- and right-handed (chirality) fermion fields respectively. If neutrinos were massless, then this can be explained by eliminating the right-handed neutrino field N_{R} from the theory, leaving just the left-handed neutrino field ν_{L}. This would prohibit a mass generating term like Equation (2.6) for neutrinos, and is aesthetically pleasing since the right-handed neutrino has no other interactions in the theory.

But if neutrinos do have a mass then we must introduce a term like in Equation (2.6) and a right-handed neutrino field. Having introduced a right-handed neutrino we find that it carries no gauge quantum numbers and it is also possible to introduce terms such as

$$\mathcal{L}_{\mathrm{M}} = -\frac{1}{2} M \left(\overline{N}_{\mathrm{R}} N_{\mathrm{R}}^{c} + \overline{N}_{\mathrm{R}}^{c} N_{\mathrm{R}} \right) , \tag{2.7}$$

where $N_{\mathrm{R}}^{c} = \mathrm{i}\gamma^0\gamma^2(\overline{N}_{\mathrm{R}})^{\mathrm{T}}$ is the charge conjugate of a right-handed neutrino and M is the Majorana mass. Such terms allow a neutrino to turn into an antineutrino, violating total lepton number by two units, and making neutrinos indistinguishable from antineutrinos.

We can combine the Dirac and Majorana terms into a single term, restricting ourselves to a single flavour for clarity:

$$\mathcal{L}_{\mathrm{D+M}} = -\frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \overline{N}_{\mathrm{R}}^{c} \end{pmatrix} \begin{pmatrix} 0 & m_{\mathrm{D}} \\ m_{\mathrm{D}} & M \end{pmatrix} \begin{pmatrix} \nu_{\mathrm{L}}^{c} \\ N_{\mathrm{R}} \end{pmatrix} \quad (+h.c.) \tag{2.8}$$

In many theories the natural size of the Majorana mass is of the order of 10^{15} GeV, much larger than Dirac masses generated by the Higgs mechanism. In this case the eigenstates (obtained by diagonalising the matrix) are a (nearly[11]) right-handed Majorana neutrino with mass of approximately M, and a (nearly) left-handed Majorana neutrino with a mass of approximately m_{D}^2/M. If m_D is taken to be

[11]The fraction of the opposite handed component is of the order $(m_{\mathrm{D}}/M)^2$

around 100 GeV, the left-handed neutrinos will be given masses in the region of 10^{-2} eV. This is the see-saw mechanism. One of its most important predictions is that neutrinos are Majorana particles, a question that has yet to be settled experimentally. Most of neutrino physics, and neutrino oscillations in particular, is insensitive to the Majorana/Dirac question, but for the rest of this thesis we will assume (unless specifically stated) that neutrinos are Dirac fermions.

If neutrinos are indeed Majorana particles, this allows the existence of phenomena such as neutrinoless double beta decay ${}^{A}_{Z}\mathrm{X} \rightarrow {}_{Z+2}^{A}X' + 2e^{-}$, which, in addition to being a significant discovery in itself, would allow us to measure the absolute scale of neutrino masses. Several experiments [66] are now looking for evidence of this process.

2.3 Neutrino oscillations

When a neutrino is produced (via a $W^{\pm}$ interaction) it is in association with a charged lepton. The neutrino is in a state with a well defined flavour: the same as that of the lepton. But to describe neutrino propagation the appropriate states to use are those with well defined mass. There are therefore two ways to describe a neutrino: what kind of charged lepton it would produce in a charged current interaction (the flavour basis), and how it propagates in free space (the mass basis)

There is no reason why these two bases should be the same, but it should be possible to describe any neutrino in both bases. This requires that the transformation between the flavour and mass bases to be unitary. If we use $|\nu_i\rangle$ for the mass basis, and $|\nu_\alpha\rangle$ for the flavour basis, then

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha i} |\nu_i\rangle \,, \tag{2.9}$$

where $U^\dagger U = \mathbb{I}$. U is named the PMNS matrix (sometimes just MNS) for Pontecorvo, Maki, Nakagawa and Sakata. The oscillation mechanism is then simply an expression of quantum mechanics.

2.3.1 The oscillation probability

First we consider the variation of the neutrino state $|\nu(\mathbf{x}, t)\rangle$ as a function of position $\mathbf{x}$ and time t, taking $|\nu(t = 0)\rangle = |\nu\rangle$. For simplicity's sake we consider neutrino states $|\nu(\mathbf{x}, t)\rangle$ that are infinite plane waves. This is actually an over-simplification, and will lead to some difficulties[12], which we will have to sidestep. Immediately we have a problem, in that a plane wave has equal probability to be at all points in space at all times. By hand we put in that the neutrino is localised in space and propagating at a fixed velocity so that the evolution of the state is

[12] See [67] for a concise summary. More sophisticated treatments, such as the use of wavepackets [68], help to resolve these issues.

a function only of the distance L that the neutrino has travelled. We then have that $|\nu(\mathbf{x},t)\rangle = |\nu(L)\rangle$ and

$$|\nu(L)\rangle = \mathrm{e}^{-\mathrm{i}p\cdot x}\,|\nu(0)\rangle\,, \tag{2.10}$$

using the notation $p\cdot x = Et - \mathbf{p}\cdot\mathbf{x}$. The values of E and $\mathbf{p}$ depend on the mass of the neutrino so we expand $|\nu(0)\rangle$ in the mass basis:

$$|\nu(L)\rangle = \sum_i a_i \mathrm{e}^{-\mathrm{i}p_i\cdot x}\,|\nu_i\rangle\,. \tag{2.11}$$

The a_i are the amplitudes of the mass states $|\nu_i\rangle$ at $L=0$. If we take the neutrino to be in a flavour eigenstate $|\nu_\alpha\rangle$ at $L=0$ then we find

$$a_i = \langle\nu_i|\,\nu_\alpha\rangle = U^*_{\alpha i} \tag{2.12}$$

which we can substitute into Equation (2.11).

If now want to find out the flavour state of the neutrino after travelling a distance L we use the conjugate of Equation 2.9:

$$\langle\nu_\beta| = \sum_j U_{\beta j}\,\langle\nu_j|\,. \tag{2.13}$$

Substituting this into Equation (2.11), and taking the initial state to be a flavour eigenstate so that $a_i = U^*_{\alpha i}$ we get

$$\langle\nu_\beta|\,\nu_\alpha(L)\rangle = \sum_j\sum_i U_{\beta j}\,\langle\nu_j|\,U^*_{\alpha i}\mathrm{e}^{-\mathrm{i}p_i\cdot x}\,|\nu_i\rangle \tag{2.14}$$

$$= \sum_i U_{\beta i}U^*_{\alpha i}\mathrm{e}^{-\mathrm{i}p_i\cdot x}. \tag{2.15}$$

Now we turn our attention to the phase $\phi_i = p_i\cdot x$. First we note another difficulty with the formalism. A plane wave has a well defined energy and momentum. By the relation

$$E_i^2 = m_i^2 + |\mathbf{p}_i|^2 \tag{2.16}$$

a neutrino represented by a plane wave also has a well defined *mass*. But there must be some ambiguity about the measured energy and momentum of the neutrino or we will have made a measurement of the neutrino mass, and eliminated the oscillation. More sophisticated treatments integrate over a range of neutrino momenta, avoiding this problem and localising the neutrino in spacetime, but in the plane wave formalism we must perform a sleight of hand to represent our ignorance of which mass state the neutrino was in. We follow a common choice which is to say that all mass states have the same momentum. We can then write

$$\phi_i = E_i t - \mathbf{p}{\cdot}\mathbf{x}. \tag{2.17}$$

Next we assume that all mass states are highly relativistic, so that $m_i \ll E$ where E is the average of all E_i. We can do a binomial expansion

$$E_i = \sqrt{|\mathbf{p}|^2 + m_i^2} \simeq |\mathbf{p}| \left(1 + \frac{m_i^2}{2\,|\mathbf{p}|^2}\right), \tag{2.18}$$

and substitute this into Equation 2.17. At the same time we use the (assumed) highly relativistic nature of the neutrinos to equate $t = L$, and resolve the vector product $\mathbf{p} \cdot \mathbf{x} = |\mathbf{p}|\, L$, giving

$$\begin{aligned} \phi_i &\simeq |\mathbf{p}| \left(1 + \frac{m_i^2}{2\,|\mathbf{p}|^2}\right) L - |\mathbf{p}|\, L && (2.19)\\ &= \frac{m_i^2 L}{2E}, && (2.20) \end{aligned}$$

where we have used $|\mathrm{p}| \simeq E$ in the last line.

Returning to Equation 2.15 we can now write:

$$\langle \nu_\beta | \, \nu_\alpha(L) \rangle = \sum_i U_{\beta i} U^*_{\alpha i} \mathrm{e}^{-\mathrm{i} m_i^2 L/2E}, \tag{2.21}$$

which gives

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= \left| \, \langle \nu_\beta | \, \nu_\alpha(L) \rangle \, \right|^2 \\ &= \sum_{i\,j} U_{\beta i} U^*_{\alpha i} U^*_{\beta j} U_{\alpha j} \mathrm{e}^{-\mathrm{i}(m_i^2 - m_j^2)L/2E}. && (2.22) \end{aligned}$$

Writing $\Delta m_{ij}^2 = m_i^2 - m_j^2$ we expand this

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= \sum_i U_{\beta i} U^*_{\alpha i} \sum_j U^*_{\beta j} U_{\alpha j} \\ &\quad + \sum_{i\,j} U_{\beta i} U^*_{\alpha i} U^*_{\beta j} U_{\alpha j} \left(\mathrm{e}^{-\mathrm{i}\Delta m_{ij}^2 L/2E} - 1\right). && (2.23) \end{aligned}$$

The first term is $\delta_{\alpha\beta}$ because U is unitary. To make sense of the second term we note that the terms with $j > i$ are complex conjugates of the terms with $j < i$, and that when $j = i$ the phase difference is zero. Hence

$$P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} + 2 \sum_{i>j} \mathfrak{Re}\Big[U_{\beta i} U^*_{\alpha i} U^*_{\beta i} U_{\alpha i} \underbrace{\left(\mathrm{e}^{-\mathrm{i}\Delta m_{ij}^2 L/2E} - 1\right)}_{\star}\Big] \tag{2.24}$$

Finally we split the marked ($\star$) term into real and imaginary parts, and use a little trigonometry to give:

$$\begin{aligned}P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta}& \\ & -4\sum_{i>j} \mathfrak{Re}\left[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}\right]\sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right) \\ & +2\sum_{i>j}\ \mathfrak{Im}\left[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}\right]\sin\left(\frac{\Delta m^2_{ij}L}{2E}\right).\end{aligned} \tag{2.25}$$

Equation (2.25) is the observation probability of a ν_β at a distance L from a source of ν_α. The first term is is the observation probability in the absence of oscillations. The second term is the 'standard' oscillation term which gives rise to oscillations unless either ν_α or ν_β are also mass eigenstates. The third term is more subtle. It disappears if U can be expressed as a real matrix, which must be the case if there are two neutrinos (but can be true for any number). Also, if $\beta = \alpha$

$$\mathfrak{Im}\left[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}\right] = \mathfrak{Im}\left[U_{\alpha i}U^*_{\alpha i}U^*_{\alpha j}U_{\alpha j}\right] = \mathfrak{Im}\left[|U_{\alpha i}|^2\,|U_{\alpha j}|^2\right] = 0 \tag{2.26}$$

so this term disappears if we are considering $P\left(\nu_\alpha \to \nu_\alpha\right)$. This is a common situation in neutrino oscillation experiments, and in this case Equation (2.25) simplifies to:

$$P(\nu_\alpha \to \nu_\alpha) = 1 - 4\sum_{i>j}|U_{\alpha i}|^2\,|U_{\alpha j}|^2\sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right), \tag{2.27}$$

which is referred to as the *survival probability.*

2.3.2 $\mathcal{CP}$ violation

The third term of Equation (2.25) is particularly interesting because, if $\mathcal{CPT}$ symmetry holds then:

$$\begin{aligned}\mathcal{CP}\ P\left(\nu_\alpha \to \nu_\beta\right) &= \mathcal{T}\ P\left(\nu_\alpha \to \nu_\beta\right) \\ P\left(\overline{\nu}_\alpha \to \overline{\nu}_\beta\right) &= P\left(\nu_\beta \to \nu_\alpha\right).\end{aligned} \tag{2.28}$$

This means that the oscillation probability for antineutrinos is equivalent to the probability for neutrinos with the exchange $\alpha \leftrightarrow \beta$. Examination of Equation (2.25) reveals that this is equivalent to the exchange $U \leftrightarrow U^*$, and the only effect of such an exchange is to flip the sign of the third term. An important consequence of this is that the survival probability is invariant under $\mathcal{CP}$.

2.3.3 The 'standard model' of neutrino oscillations

Current data[13] is consistent with the inclusion of three neutrino mass eigenstates. The absolute masses of these are not known but at least two must be massive. Various limits on neutrino masses are given in Table 2.1. Two of the mass states (ν_1, ν_2) are comparatively close in m^2, with the third (ν_3) being more distant, as shown in Figure 2.10. The splitting $\Delta m^2_{\odot}$ which give rise to the oscillations seen in the solar neutrino experiments and KamLAND is identified as Δm^2_{21}, and it is known from the solar MSW effect that $m_2^2 > m_1^2$. The atmospheric neutrino oscillations seen by Super-Kamiokande and the K2K and MINOS accelerator-based experiments are a combination of oscillations at Δm^2_{31} and Δm^2_{32} and we can associate Δm^2_{atm} with a weighted average of these two quantities[14]. It is not yet known if m_3^2 is larger or smaller than the other two masses; the former situation is referred to as the Normal Hierarchy, while the latter is termed the Inverted Hierarchy. Both can be seen in Figure 2.10.

Measurement	Origin of limit	Limit on m_ν
Tritium decay	$m^2_{\nu_e}$	$< 2.0\,\text{eV}$
π^+ decay	$m^2_{\nu_\mu}$	$< 0.19\,\text{MeV}$
τ decay	$m^2_{\nu_\tau}$	$< 18.2\,\text{MeV}$
SN1987A time-of-flight	$m^2_{(\nu_e \to \nu)}$	$< 5.7\,\text{eV}$
Terrestrial time-of-flight [69]	$m^2_{(\nu_\mu \to \nu_\tau)}$	$< 50\quad\text{MeV}$
Cosmology	$\sum_i m_{\nu_i}$	$< 0.7\,\text{eV}$

Table 2.1: Neutrino mass measurements. Limits on neutrino masses come from several sources, some of which are model dependent. In addition, the exact mass involved is often unclear, as production and detection of a neutrino generally involve flavour eigenstates. Conservation of $\mathcal{CPT}$ is assumed. Data from [54], unless otherwise specified.

With three neutrino masses the PMNS matrix is a 3×3 unitary matrix, with nine real parameters. If neutrinos are Dirac fermions, up to five of these can be absorbed as phase differences [32] between the three leptons and three neutrinos, leaving four real parameters. Three of these can be expressed as rotations, and

[13]excluding the LSND result

[14]It is common for the terms Δm^2_{32}, Δm^2_{31}, and Δm^2_{atm} to be used almost interchangeably. The spread is about an order of magnitude smaller than the precision of current experiments, so the distinction is not significant, although this is likely to change with the next generation of experiments.

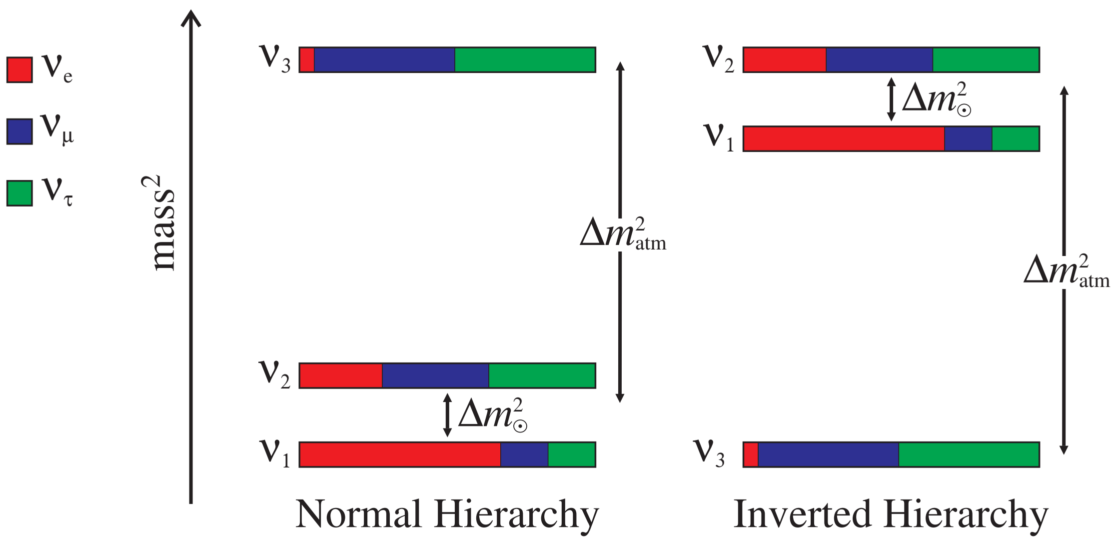

Figure 2.10: Normal and inverted Neutrino mass hierarchies. The flavour content of each mass eigenstate is shown, approximately reflecting current data, with the addition that U_{e3} is set to $0.05\mathrm{e}^{\mathrm{i}\pi/2}$.

the fourth must be a complex phase. The matrix is commonly parameterised [70]:

$$U = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{R_{23}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}\mathrm{e}^{-\mathrm{i}\delta} \\ 0 & 1 & 0 \\ -s_{13}\mathrm{e}^{\mathrm{i}\delta} & 0 & c_{13} \end{pmatrix}}_{U_{13}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{R_{12}} \tag{2.29}$$

with $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. Using this parameterisation

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}\mathrm{e}^{-\mathrm{i}\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}\mathrm{e}^{\mathrm{i}\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}\mathrm{e}^{\mathrm{i}\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}\mathrm{e}^{\mathrm{i}\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}\mathrm{e}^{\mathrm{i}\delta} & c_{23}c_{13} \end{pmatrix}. \tag{2.30}$$

If neutrinos are Majorana fermions then only three of the complex phases can be absorbed, and the PMNS matrix must be extended: $U \to UK$, where $K = \mathrm{diag}(\mathrm{e}^{\mathrm{i}\alpha_1/2}, \mathrm{e}^{\mathrm{i}\alpha_2/2}, 1)$. This has no impact on the oscillation probabilities as the additional phases cancel out in Equation (2.15).

It is typical, however, for experimental results to be analysed as if there were only two types of neutrino. In that case there is only one mass splitting δm^2, and the mixing matrix has only one real parameter:

$$U_2 = \begin{pmatrix} \cos\vartheta & \sin\vartheta \\ -\sin\vartheta & \cos\vartheta \end{pmatrix}. \tag{2.31}$$

Putting this into Equation (2.25) gives

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2(2\vartheta)\sin^2\left(\frac{\delta m^2 L}{4E}\right) \tag{2.32}$$

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\vartheta)\sin^2\left(\frac{\delta m^2 L}{4E}\right). \tag{2.33}$$

If we do the same thing with three neutrinos and examine the resulting expression it is somewhat surprising that the two flavour approximation works at all. Turning first to the CHOOZ/Palo Verde results, where we are looking for $\overline{\nu}_e$ disappearance. Putting Equation (2.30) into Equation (2.27) we obtain

$$P(\overline{\nu}_e \to \overline{\nu}_e) = 1 - 4\left[\begin{array}{r} |U^*_{e1}|^2\,|U^*_{e2}|^2 \sin^2\left(\frac{\Delta m^2_{12} L}{4E}\right) \\ + |U^*_{e1}|^2\,|U^*_{e3}|^2 \sin^2\left(\frac{\Delta m^2_{13} L}{4E}\right) \\ + |U^*_{e2}|^2\,|U^*_{e3}|^2 \sin^2\left(\frac{\Delta m^2_{23} L}{4E}\right) \end{array}\right]. \tag{2.34}$$

The $\overline{\nu}_e$ disappearance is seen in a regime where $\Delta m^2_{23} L/4E \sim 1$, which implies $\Delta m^2_{12} L/4E \ll 1$, so the first term is negligible.[15] This also allows us to approximate $\Delta m^2_{13} \simeq \Delta m^2_{23} \simeq \Delta m^2_{\mathrm{atm}}$, and hence

$$P(\overline{\nu}_e \to \overline{\nu}_e) \simeq 1 - 4\left(|U^*_{e1}|^2 + |U^*_{e2}|^2\right)|U^*_{e3}|^2 \sin^2\left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right) \tag{2.35}$$

$$\simeq 1 - 4\left(1 - |U^*_{e3}|^2\right)|U^*_{e3}|^2 \sin^2\left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right) \tag{2.36}$$

$$\simeq 1 - \sin^2(2\theta_{13})\sin^2\left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right), \tag{2.37}$$

again making use of the unitary of U. This is equivalent to Equation (2.32) with $(\vartheta, \delta m^2) \to (\theta_{13}, \Delta m^2_{\mathrm{atm}})$. Using $\Delta m^2_{\mathrm{atm}} \sim 2.5 \times 10^{-3}$ from the atmospheric experiments we get the 'CHOOZ limit': $\sin^2 2\theta_{13} < 0.15$. This means that $|U_{e3}|^2$ must either be very large ($\nu_e \sim \nu_3$) or very small (ν_e is a mixture of ν_1 and ν_2). Unitarity requires that if $|U_{e3}|^2$ is large then $|U_{e1}|^2$ and $|U_{e2}|^2$ must be small, which prohibits the existence of large-amplitude oscillations driven by Δm^2_{12}. This is incompatible with the oscillations seen by KamLAND so $|U_{e3}|^2$ must be small, and in many cases we can neglect it, taking the U_{13} sub-matrix to be the unit matrix.

Next we turn to the KamLAND experiment. This also measures the $\overline{\nu}_e$ survival

[15] This approximation is good enough here. If $\sin^2(2\theta_{13})$ is very small ($\sim 10^{-3}$) we cannot neglect the Δm^2_{12} term.

probability but over a much longer baseline, so that $\Delta m_{12}^2 L/4E \sim 1$. From the limit on $\sin^2 2\theta_{13}$, we know that the amplitude of atmospheric scale oscillations is small, so the second and third terms of Equation (2.34) can be neglected, and so

$$P(\overline{\nu}_e \to \overline{\nu}_e) \simeq 1 - 4\left|U_{e1}^*\right|^2 \left|U_{e2}^*\right|^2 \sin^2\left(\frac{\Delta m_{12}^2 L}{E}\right) \tag{2.38}$$

$$\simeq 1 - 4\sin^2(2\theta_{12})\sin^2\left(\frac{\Delta m_{12}^2 L}{E}\right), \tag{2.39}$$

again recovering Equation (2.32), this time with $(\vartheta, \delta m^2) \to (\theta_{12}, \Delta m_{12}^2)$.

Now for the atmospheric sector. The experiments here measure the ν_μ survival probability:

$$P(\nu_\mu \to \nu_\mu) = 1 - 4\left[\begin{array}{r} \left|U_{\mu 1}\right|^2 \left|U_{\mu 2}\right|^2 \sin^2\left(\frac{\Delta m_{12}^2 L}{4E}\right) \\ + \left|U_{\mu 1}\right|^2 \left|U_{\mu 3}\right|^2 \sin^2\left(\frac{\Delta m_{13}^2 L}{4E}\right) \\ + \left|U_{\mu 2}\right|^2 \left|U_{\mu 3}\right|^2 \sin^2\left(\frac{\Delta m_{23}^2 L}{4E}\right) \end{array}\right]. \tag{2.40}$$

This time we take both $\Delta m_{12}^2 L/E$ and θ_{13} to be negligible and get

$$P(\nu_\mu \to \nu_\mu) \simeq 1 - 4\left(1 - \left|U_{\mu 3}\right|^2\right)\left|U_{\mu 3}\right|^2 \sin^2\left(\frac{\Delta m_{\mathrm{atm}}^2 L}{4E}\right) \tag{2.41}$$

$$\simeq 1 - \sin^2(2\theta_{23})\sin^2\left(\frac{\Delta m_{\mathrm{atm}}^2 L}{4E}\right), \tag{2.42}$$

again obtaining the 2 flavour approximation, with $(\vartheta, \delta m^2) \to (\theta_{23}, \Delta m_{\mathrm{atm}}^2)$.

2.3.4 Matter effects

Analysis of solar neutrino data is more complicated. As mentioned earlier the depletion is attributed to the MSW effect, not simple oscillations. This effect is named for Mikheev and Smirnov, who suggested it [46], and for Wolfenstein, who first considered the behaviour of neutrinos passing through matter [45]. Restricting ourselves to a two neutrino (ν_e, ν_α) model for simplicity, we first write down the Hamilton for the propagation of ultra-relativistic neutrinos in vacuum:

$$\mathcal{H}_0 \begin{pmatrix} \nu_e \\ \nu_\alpha \end{pmatrix} = \mathrm{i}\frac{\partial}{\partial t}\begin{pmatrix} \nu_e \\ \nu_\alpha \end{pmatrix} = \begin{pmatrix} A_{ee} & T_{e\alpha} \\ T_{e\alpha}^* & A_{\alpha\alpha} \end{pmatrix}\begin{pmatrix} \nu_e \\ \nu_\alpha \end{pmatrix}. \tag{2.43}$$

If ν_e and ν_α were identical to the mass eigenstates ν_1 and ν_2 the matrix $\mathcal{H}_0$ would be diagonal with elements $m_i^2/2E$. Instead there is a rotation between mass and flavour eigenstates, as in Equation (2.31):

$$\begin{aligned}\mathcal{H}_0 = U_2^\dagger \begin{pmatrix} \frac{m_1^2}{2E} & 0 \\ 0 & \frac{m_2^2}{2E} \end{pmatrix} U_2 &= U_2^\dagger \left[\frac{m_1^2}{2E}\mathbb{I} + \begin{pmatrix} 0 & 0 \\ 0 & \frac{\delta m^2}{2E} \end{pmatrix} \right] U_2 \\ &= \frac{m_1^2}{2E} + U_2^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{\delta m_2^2}{2E} \end{pmatrix} U_2. \end{aligned} \tag{2.44}$$

In this form the oscillation parameters can be easily identified. We also see the insensitivity of oscillations to the absolute mass scale (m_1^2). We are not interested in this diagonal component, but will retain a term k to remind us that it is non-zero.

Now we add a term for neutrino interactions in matter. In normal matter this comes from coherent forward scattering of neutrinos on protons, neutrons and electrons. All active neutrinos scatter equally of matter via neutral current interactions, but there is additional charged current scattering for ν_e. For an electron (anti)neutrino a matter potential is introduced by charged current interactions with electrons [71]:

$$V_{\mathrm{m}}^e = \pm\sqrt{2}G_{\mathrm{F}}n_e \tag{2.45}$$

where n_e is the number density of electrons and the lower sign applies for antineutrinos. A similar potential applies for all active flavours due to neutral current interactions:

$$V_{\mathrm{m}}^{\mathrm{NC}} = \mp\frac{G_{\mathrm{F}}n_n}{\sqrt{2}} \tag{2.46}$$

where n_n is the number density of neutrons. It is convenient to use two related quantities:

$$A_{\mathrm{m}} = \frac{2E}{\delta m^2}V_{\mathrm{m}}^{\mathrm{NC}} \qquad \text{and} \qquad \epsilon = \frac{2E}{\delta m^2}V_{\mathrm{m}}^e, \tag{2.47}$$

with which the Hamiltonian for propagation in matter can be written:

$$\mathcal{H}_{\mathrm{m}} = k + U_2^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{\delta m^2}{2E} \end{pmatrix} U_2 + \begin{pmatrix} \frac{\delta m^2}{2E}\left(A_{\mathrm{m}} + \epsilon\right) & 0 \\ 0 & \frac{\delta m^2}{2E}A_{\mathrm{m}} \end{pmatrix} \tag{2.48}$$

Multiplying out the second term and removing a multiple of the unit matrix from the third term gives:

$$\mathcal{H}_{\mathrm{m}} = k + \frac{\delta m^2}{2E} \begin{pmatrix} \sin^2\vartheta + \epsilon & -\sin\vartheta\cos\vartheta \\ -\sin\vartheta\cos\vartheta & \cos^2\vartheta \end{pmatrix}. \tag{2.49}$$

To obtain oscillation parameters in matter we want to arrange this expression in the form

$$\mathcal{H}_{\mathrm{m}} = k + U_{\mathrm{m}}^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{\delta m_{\mathrm{m}}^2}{2E} \end{pmatrix} U_{\mathrm{m}}, \tag{2.50}$$

which defines $\delta m^2_{\rm m}$ and $\vartheta_{\rm m}$, the effective mixing parameters in matter. To get the required form we diagonalise the dimensionless matrix in Equation (2.49). Solving the quadratic equation for the eigenvalues we get $\lambda_\pm = \lambda_0 \pm \delta\lambda/2$, where:

$$\lambda_0 = \frac{1+\epsilon}{2}\,;$$
$$\delta\lambda = \sqrt{(1+\epsilon)^2 - 4\epsilon\cos^2\vartheta} = \sqrt{[\epsilon - \cos(2\vartheta)]^2 + \sin^2(2\vartheta)}. \tag{2.51}$$

As in the vacuum case, the effective mass splitting in matter is given by the difference of the eigenvalues, so:

$$\frac{\delta m^2_{\rm m}}{2E} = \frac{\delta m^2}{2E}\delta\lambda \tag{2.52}$$

Multiplying the matrices in Equation (2.50) we get:

$$\mathcal{H}_{\rm m} = k + \frac{\delta m^2_{\rm m}}{2E}\begin{pmatrix} \sin^2\vartheta_{\rm m} & -\sin\vartheta_{\rm m}\cos\vartheta_{\rm m} \\ -\sin\vartheta_{\rm m}\cos\vartheta_{\rm m} & \cos^2\vartheta_{\rm m}\end{pmatrix}. \tag{2.53}$$

Then we can equate the off-diagonal terms with those in Equation (2.49) to get an expression for $\vartheta_{\rm m}$:

$$-\frac{\delta m^2_{\rm m}}{2E}\sin\vartheta_{\rm m}\cos\vartheta_{\rm m} = -\frac{\delta m^2}{2E}\sin\vartheta\cos\vartheta \tag{2.54}$$
$$\delta\lambda\,\sin(2\vartheta_{\rm m}) = \sin(2\vartheta). \tag{2.55}$$

Putting in the expression for $\delta\lambda$ we get:

$$\delta m^2_{\rm m} = \delta m^2\sqrt{[\epsilon - \cos(2\vartheta)]^2 + \sin^2(2\vartheta)}\,; \tag{2.56}$$

$$\sin^2(2\vartheta_{\rm m}) = \frac{\sin^2(2\vartheta)}{[\epsilon - \cos(2\vartheta)]^2 + \sin^2(2\vartheta)}, \tag{2.57}$$

with $\epsilon = \pm 2\sqrt{2}G_{\rm F}n_e E/\delta m^2$ and the minus sign applies for antineutrinos. These expressions have several interesting consequences for neutrinos propagating through matter:

- The effective mixing parameters depend on the energy of the neutrino and the matter density. The L/E dependence of the oscillation will not be (exactly) sinusoidal, and may be extremely complicated if the matter density changes over the baseline of the experiment.

- The effect is different for neutrinos and antineutrinos, and depends on the sign of δm^2. This makes it possible to determine which of the two mass eigenstates is heavier, something that oscillations in vacuum are inherently insensitive to.

- If $\sin^2(2\vartheta)$ is non-zero, there can be a resonance where (for a particular neutrino energy) the effective mixing in matter is maximal, i.e. $\sin^2(2\vartheta_\mathrm{m}) = 1$. Whether this resonance occurs for neutrinos or antineutrinos depends on the sign of δm^2.

- For sufficiently high matter densities, $\sin^2(2\vartheta_\mathrm{m}) \simeq 0$, implying the in-matter mass eigenstates are nearly coincident with the flavour eigenstates.

This is enough to sketch how the MSW effect suppresses the solar ν_e flux. By convention, the heavier of vacuum mass eigenstates is labelled ν_2, and the lighter eigenstate ν_1. The equivalent states in matter are identified as $\nu_{\mathrm{m}2}$ and $\nu_{\mathrm{m}1}$ respectively. In the core of the Sun, where the ν_e are produced, the density of electrons is very high and the flavour eigenstate ν_e corresponds to a mass state that is almost pure $\nu_{\mathrm{m}2}$. As the ν_e propagate outward, the electron density falls gradually. With no sudden changes in density, the neutrinos remain in the same $\nu_{\mathrm{m}2}$ state as it gradually returns to the ν_2 vacuum state when the neutrinos emerge from the Sun. The neutrinos then propagate through the vacuum as (nearly) pure ν_2, and therefore do not oscillate. For a wide range of energies the solar neutrinos detected on Earth are approximately ν_2, and therefore a mixture of ν_e, ν_μ and ν_τ. So roughly, the suppression of the ν_e flux gives a direct measurement of $|U_{e2}|^2$.

2.4 Extended models

There are several alternative models of neutrino flavour change. Many appeared as alternative explanations of the solar neutrino problem, and most have been ruled out over time. Some require neutrinos to interact with the sun's magnetic field and cannot easily explain any other observations of neutrino flavour change. Two of the most interesting are neutrino decay and neutrino decoherence. Another alternative, described in Section 2.4.3 it to extend the oscillation mechanism to include additional, sterile, neutrinos.

2.4.1 Neutrino decay

One possible alternative explanation of neutrino disappearance is that neutrinos decay. With mixing between mass and flavour eigenstates, the survival probability is not simply an exponential, and can result in a finite flux of the initial neutrino flavour even after travelling arbitrarily long distances from the source. As the observed deficits of neutrinos cannot be accounted for by decay in the Standard Model, decay mechanisms must invoke new physics in the neutrino sector. As a result there is no single decay model, but for illustrative purposes a simple version

[72] is outlined below. In this model the initial flavour state is a mixture of two (or more) mass eigenstates:

$$|\nu_\alpha\rangle = \sin\vartheta\,|\nu_1\rangle + \cos\vartheta\,|\nu_2\rangle\,.$$

If one of the mass states (ν_2) decays with lifetime m_2/κ to a state that does not mix with the initial flavour state, then

$$|\nu_\alpha(L)\rangle = \sin\vartheta\,|\nu_1\rangle + \cos\vartheta\,\mathrm{e}^{-\kappa L/2E}\,|\nu_2\rangle\,, \tag{2.58}$$

and the survival probability is simply

$$P(\nu_\alpha \to \nu_\alpha) = \left(\sin^2\vartheta + \cos^2\vartheta\,\mathrm{e}^{-\kappa L/2E}\right)^2. \tag{2.59}$$

Note that it is not necessary to reject the existence of oscillations in this model, it may simply be that the relevant mass splitting is too small for oscillations to develop before the ν_2 decays.

2.4.2 Neutrino decoherence

Another possible flavour change mechanism is decoherence. One possible form of decoherence arises naturally in wavepacket treatments of the oscillation probability. If the neutrinos (initially in flavour state ν_α) propagate far enough, eventually the wave packets for each mass eigenstate will separate and there will be distinct wavepackets containing pure mass eigenstates. The relative amplitudes for each mass eigenstate are $A_i = \langle\nu_i|\nu_\alpha\rangle = U^*_{\alpha i}$. The probability to observe a neutrino as a flavour state ν_β is then simply the weighted sum of the ν_β component of each wavepacket, i.e.

$$P\left(\nu_\alpha \to \nu_\beta\right) = \sum_i \left|A_i\,\langle\nu_\beta|\nu_i\rangle\right|^2 = \sum_i |U_{\alpha i}|^2\,|U_{\beta i}|^2. \tag{2.60}$$

The separation of wavepackets naturally increases as the neutrinos propagate further, and the L dependence of the decoherence can be parameterised as $\exp\left[-(\Gamma L)^2\right]$ [73]. The parameter $1/\Gamma$ is the coherence length and, being related to the velocity difference between mass states, increases for higher neutrino energies and smaller mass splittings. It is also dependent on the width of the initial wave packet, and for practically achievable neutrino sources the coherence length is longer than terrestrial baselines.

However, other forms of decoherence have been suggested. If there is a way for neutrinos to interact with their environment in a way that distinguishes between flavours this can cause a decoherence effect. The energy dependence of the decoherence effect depends on the assumptions made about the neutrino's interaction with the environment, but in the case of pure decoherence (no oscillations) with two neutrinos the survival probability can be expressed as [74]:

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \tfrac{1}{2}\sin^2 2\vartheta \left(1 + e^{-\gamma(E)L}\right), \tag{2.61}$$

where γ is an (unknown[16]) function of energy.

Analyses that have looked for distortion of the neutrino spectrum [52] [49] have found that the data agree better with neutrino oscillations than with pure decay or pure decoherence. As a result, interest in these models is moving towards mixed 'oscillations plus ...' scenarios. [73]

2.4.3 Sterile Neutrinos

An obvious extension to to the standard oscillation model is to ask if we can add another mass eigenstate. One attraction of this idea is that it potentially provides an explanation of the LSND result. But to add an additional mass eigenstate we must also add an additional flavour eigenstate and this leads to a problem. The Z^0 width results from LEP [75] rule out the existence of more than three light neutrinos with normal couplings to the Z^0, so the neutrino we add must have no weak couplings. This additional neutrino is referred to as a sterile neutrino, ν_s; the original three neutrinos are termed 'active', in contrast.

We are not limited to adding just one sterile neutrino. Indeed, the theoretical problem is that having added one sterile neutrino we have no reason to stop. Discussions of sterile neutrinos tend to stick to adding as few as possible for convenience sake, as we cannot observe mixing between the sterile states.

If we take the simplest possible case of one sterile neutrino the first thing we must do is expand the mixing matrix U to a 4×4 matrix, U_4. This matrix will have 16 real parameters, of which seven (assuming Dirac neutrinos) can be absorbed into the lepton phases, six become mixing angles, and three become $\mathcal{CP}$ violating phases. Compared to the three neutrino model there are three new mixing angles and two more $\mathcal{CP}$ violating phases, so the phenomenology is much more complicated in general.[17]

Another new parameter of the extended model is the mass of the new mass eigenstate. If the model is to explain the LSND anomaly then this mass must be chosen such that the new $\Delta M^2 \sim 1\,\mathrm{eV}^2$. Even with this assumption there are a range of possibilities. Compared to the two possible mass hierarchies shown in Figure 2.10 there are now six possible hierarchies.[18] These divide into two quite distinct schemes, shown schematically in Figure 2.11. One, the 3+1 scheme, has the large ΔM^2 between the new mass state (ν_4) and the three 'normal' mass

[16] In general, agreement with data is poor if $\gamma(E)$ is large at high energy, so comparisons to decoherence models commonly assume $\gamma = \gamma_0/E$.

[17] The number of free parameters is $(n-1)^2$ where n is the number of neutrinos—another reason to restrict the discussion to as few sterile neutrinos as possible.

[18] Assuming $\Delta M^2 \gg \Delta m^2_{\mathrm{atm}}$. If the new mass scale is similar to the existing scales the phenomenology is even more complicated.

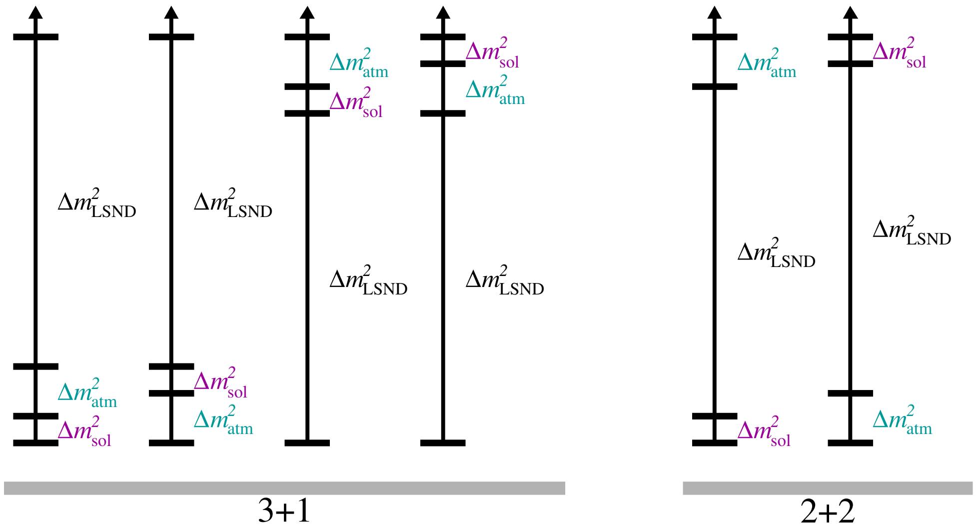

Figure 2.11: 3+1 and 2+2 neutrino hierarchies

states. In the limit that $\nu_s \rightarrow \nu_4$ the 3+1 scheme reduces to the standard 3 neutrino scheme. In the alternative 2+2 scheme, the large ΔM^2 separates the smaller $\Delta m^2_{\rm atm}$ and $\Delta m^2_{\odot}$.

In the 2+2 scheme, two of the mass states are involved in atmospheric oscillations, and the other two are involved in solar oscillations. Therefore, at least one of these regimes must involve a large amplitude for oscillating to the sterile neutrino. However, the neutral current data from SNO disfavours sterile neutrino involvement in the solar regime, whilst in the atmospheric regime the absence of a matter effect as the neutrinos pass through the earth's core [76] disfavours sterile neutrino involvement here also. Taken together these results make 2+2 models unlikely [77], even if we choose to believe the LSND result.

The 3+1 scheme has its own problems. To a good approximation the atmospheric and solar mass splittings are negligible when considering the LSND result. This suppresses $\mathcal{CP}$ violating effects in the LSND regime (it approximates a 2-neutrino case) and makes it hard to accommodate LSND with the negative results from MiniBooNE and NOMAD. Adding a second sterile neutrino, and another large mass splitting, makes a large $\mathcal{CP}$ violating effect possible and improves things somewhat. [78]

Because there is no clearly favoured model for including a fourth neutrino, a search for sterile neutrinos in MINOS is most sensibly done with an empirical parameterisation. In contrast to LSND, the approach taken in MINOS is to look for depletion of neutral current (NC) events. Oscillations between active flavours will not affect the NC spectrum because of lepton universality, and the fact we do not observe the outgoing neutrino. Oscillations to a sterile neutrino, however, will deplete the neutral current spectrum, by definition. The neutral current spectrum is therefore sensitive to the active neutrino survival probability:

$$P(\nu_\mu \nrightarrow \nu_s) = P(\nu_\mu \to \nu_\mu) + P(\nu_\mu \to \nu_\tau) + P(\nu_\mu \to \nu_e) \tag{2.62}$$

If oscillations in the atmospheric sector were $\nu_\mu \to \nu_s$ this probability would show the standard oscillation signature in neutrino energy. But results from Super-Kamiokande and MACRO favour oscillations to ν_τ so we consider the more general case that the of oscillations to a mixed final state:

$$|\nu_\mu\rangle \longrightarrow \sqrt{f_s}\,|\nu_s\rangle + \left(1 - \sqrt{f_s}\right)|\nu_\tau\rangle\,. \tag{2.63}$$

The introduces a new parameter f_s, the sterile fraction, and gives an active neutrino survival probability

$$P(\nu_\mu \nrightarrow \nu_s) = 1 - f_s \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m^2_{\rm atm} L}{4E}\right), \tag{2.64}$$

which is different from the ν_μ survival probability only by the inclusion of the sterile fraction.

So the signal for sterile neutrinos in MINOS is rather different to LSND, instead of evidence of an extra mass eigenstate, we look for the extra flavour eigenstate.

A toy model

Although the oscillation probability used is empirical, it is interesting to ask whether it is physically reasonable; does it lead to a contradiction with other MINOS searches for example? We would like a model that

- Is consistent with MINOS dissaperence result: $P\left(\nu_\mu \to \nu_\mu\right)$ should be unaffected.
- Mixes ν_τ and ν_s: $P\left(\nu_\mu \to \nu_s\right)$ as given by Equation (2.64).
- Is consistent with other oscillation analyses: $P\left(\nu_\mu \to \nu_e\right)$ should be unaffected.

We start with a 3+1 scheme, as it can be reduced to the standard three neutrino model in the limit that $f_s \to 0$. We first extend the mixing matrix to 4 neutrinos, but with the 4$^{\rm th}$ element diagonal. For clarity of presentation we take $\sin\theta_{13} = 0$, but this is not necessary. The matrix U_4^* that transforms mass eigenstates to flavour eigenstates is then:

$$U_4^* = \begin{pmatrix} U_3^* & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} c_{12} & s_{12} & 0 & 0 \\ -s_{12}c_{23} & c_{12}c_{23} & s_{23} & 0 \\ s_{12}s_{23} & -c_{12}s_{23} & c_{23} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{2.65}$$

At this point $\nu_s \equiv \nu_4$, and we have a choice of rotations to introduce. Using the usual convention these are R_{14}, R_{24}, and R_{34}, and we can potentially introduce two new complex phases. We can place these additional rotation matrices either on the right or left sides of the existing matrix. But what we are trying to achieve is a mixture between ν_τ and ν_s so that:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1-f_s & f_s \\ 0 & 0 & f_s & 1-f_s \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix}. \tag{2.66}$$

This is precisely the effect of left-multiplying by R_{34}.[19] so we have

$$U_4^* = R_{34} R_{23} \left(U_{13}^*\right) R_{12} \tag{2.67}$$

$$= \begin{pmatrix} c_{12} & s_{12} & 0 & 0 \\ -s_{12}c_{23} & c_{12}c_{23} & s_{23} & 0 \\ c_{34}s_{12}s_{23} & -c_{34}c_{12}s_{23} & c_{34}c_{23} & s_{34} \\ -s_{34}s_{12}s_{23} & s_{34}c_{12}s_{23} & -s_{34}c_{23} & c_{34} \end{pmatrix}. \tag{2.68}$$

Substituting this matrix into Equation (2.25), we find that the probabilities $P(\nu_\mu \rightarrow \nu_\mu)$ and $P(\nu_\mu \rightarrow \nu_e)$ are unaffected because the $U_{\mu i}$ and U_{ei} elements are unchanged. In particular, there are no oscillations involving the new mass state because $U_{\mu 4} = 0$.

Using the same approximations ($\Delta m_{12}^2 L/E$ and θ_{13} are negligible) as before, the active neutrino survival probability is:

$$P(\nu_\mu \nrightarrow \nu_s) = 1 - \sin^2(\theta_{34}) \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m_{\mathrm{atm}}^2 L}{4E}\right), \tag{2.69}$$

with the expected relation $f_s = \sin^2(\theta_{34})$.

The empirical form given by Equation (2.64) is therefore physically justifiable, but there is no reason nature should choose this model. In particular, any model that attempts to explain the result from LSND must necessarily be more complicated. Such a model would in any case have many parameters and require several experiments to pin down. For this reason, we stick to the empirical f_s parametrisation to analyse MINOS data.

2.5 Future experiments

Stepping back and assuming the three-flavour model of neutrino oscillations is correct, there are still unresolved questions about the oscillation model. It is not known if the parameter $\sin(\theta_{13})$ is identically zero or just very small. The phase δ

[19] In general any matrix appended from the left acts in the existing flavour basis.

is even harder to measure as it always appears in conjunction with $\sin(\theta_{13})$, and requires sensitivity to effects on the scale of $\Delta m_{12}^2/\Delta m_{13}^2$. Finally, the hierarchy of neutrino masses is still unknown, as is the absolute scale of neutrino mass.

The easiest of these parameters to measure in an oscillation experiment is the mixing angle θ_{13}. This can be done by building more sensitive reactor-based $\overline{\nu}_e$ disappearance experiments, the first of which that is likely to come online is Double Chooz [79]. Another experiment is currently under construction at Daya Bay in China. [80] Both Double Chooz and Daya Bay will use detectors near the reactors to constrain systematic errors on the neutrino flux, and will be sensitive to values of $\sin^2 2\theta_{13}$ of $\mathcal{O}(0.01)$.

If θ_{13} is non-zero then ν_μ can turn into ν_e in experiments sensitive to the atmospheric oscillation scale. The probability is given by

$$P(\nu_\mu \to \nu_e) = 4 \begin{bmatrix} U_{e1}U_{\mu 1}^* U_{e2}^* U_{\mu 2} \sin^2\left(\frac{\Delta m_{12}^2 L}{4E}\right) \\ +U_{e1}U_{\mu 1}^* U_{e3}^* U_{\mu 3} \sin^2\left(\frac{\Delta m_{13}^2 L}{4E}\right) \\ +U_{e2}U_{\mu 2}^* U_{e3}^* U_{\mu 3} \sin^2\left(\frac{\Delta m_{23}^2 L}{4E}\right) \end{bmatrix}. \tag{2.70}$$

Assuming Δm_{12} is negligible, the first term disappears so we can use the unitarity condition $\sum_i U_{ei}U_{\mu i}^* = 0$ to reduce this to:

$$\begin{aligned} P(\nu_\mu \to \nu_e) &= 4\,|U_{e3}|^2\,|U_{\mu 3}|^2 \sin^2\left(\frac{\Delta m_{\mathrm{atm}}^2 L}{4E}\right) \\ &= \sin^2(\theta_{23})\sin^2(2\theta_{13})\sin^2\left(\frac{\Delta m_{\mathrm{atm}}^2 L}{4E}\right) \end{aligned} \tag{2.71}$$

to lowest order. However, a more careful expansion in terms of the the small parameters $\Delta m_{12}^2/\Delta m_{13}^2$ and $\sin^2(2\theta_{13})$ [81] shows that the discarded terms are not negligible in this case, and if $\sin^2(2\theta_{13}) \lesssim 0.01$ they can be quite significant. In addition, matter effects are also important over long enough baselines. With the inclusion of these additional terms, the appearance probability $P(\nu_\mu \to \nu_e)$ depends not only on θ_{13}, $|\Delta m_{atm}^2|$ and θ_{23}, but also on the $\mathcal{CP}$ violating phase δ, and whether the the neutrino mass spectrum follows the normal or inverted hierarchy.

Two experiments are in development that hope to measure $P(\nu_\mu \to \nu_e)$, NOνA [82] and T2K [83]. Like MINOS, both beams of muon neutrinos originating from pion decay. However the detectors of both experiments are placed at a small angles to the beam axis. Because of the relatavistic kinematics, sampling the neutrino beam 'off-axis' makes it possible to achieve a much narrower energy spectrum, which significantly reduces the backgrounds from neutral current interactions of higher energy neutrinos. It also lowers the average energy of the neutrinos, so that oscillations can be observed over shorter baselines (and therefore with higher fluxes).

T2K will use a beam originating from the J-PARC facility in Tokai, Japan, and directed toward Super-Kamiokande, over a baseline of 295 km. At Kamioka, the peak energy of the neutrino spectrum will be around 700 MeV, although the exact value will depend on the off-axis angle. If θ_{13} is found to be at the upper end of the allowed range, the experiment will be able to look for $\mathcal{CP}$ violation by running alternately with neutrino and antineutrino beams.

NOνA will use the same NuMI beamline as MINOS, and a slightly longer baseline of 810 km. NuMI will be run in a higher energy (on-axis) mode so as to achieve a similar peak energy off-axis. Because of the longer baseline NOνA will have greater sensitivity to matter effects which may enable it to resolve the mass hierarchy.

On their own, measurements of $P(\nu_\mu \to \nu_e)$ from these experiments cannot uniquely determine the value of θ_{13} or any of the other parameters. This is because there are degenerate combinations of parameters that produce the same oscillation probability in a given experiment. One potential degeneracy can be seen in Equation (2.71). For a non-maximal value of $\sin^2(2\theta_{23})$ there are two possible values of $\sin^2(\theta_{23})$, and hence for a measured value of $P(\nu_\mu \to \nu_e)$ there are (usually) two possible values of $\sin^2(2\theta_{13})$. Two-fold degeneracies also arise due to the unknown mass hierarchy (sign of Δm^2_{31}) and between different combinations of θ_{13} and δ. Therefore any one experiment in isolation may have an eightfold degeneracy when trying to extract parameter values. The combination of multiple experiments (T2K & NOνA & Reactor) is very important as it can help to eliminate these degeneracies.

Chapter 3

The MINOS experiment

The MINOS experiment is designed to observe the phenomena of neutrino disappearance, confirm that the mechanism for disappearance is neutrino oscillations, and measure the parameters that characterise the oscillation. The primary source of neutrinos for MINOS is the NuMI beam at Fermilab, and the experiment uses two detectors, one (the Near Detector) on-site at Fermilab, and the other (the Far Detector) at the Soudan Underground Laboratory.

MINOS is termed a long-baseline experiment; the distance between the Near and Far Detectors is 734 km. The energy spectrum of the NuMI beam is typically peaked around 5 GeV. By Equation (2.25), MINOS is therefore sensitive to oscillations arising from mass-squared splittings of around $10^{-3} \rightarrow 10^{-2}\,\mathrm{eV}^2$. This is the 'atmospheric sector' of neutrino oscillations. Since the NuMI beam is primarily composed of ν_μ, the dominant physics is expected to be oscillation to ν_τ. Most neutrino interactions seen by MINOS are charge current interactions of the form $\nu_\ell + X \rightarrow \ell^- + Y$. Below 5 GeV, the rate of τ^- production from ν_τ charged current interactions is negligible, so MINOS instead observes a deficit of μ^- events. As well as this primary measurement, MINOS will also look for sub-dominant oscillations to ν_e and ν_s. Both MINOS detectors are magnetised, which also allows ν_μ and $\overline{\nu}_\mu$ to be studied separately, putting constraints on exotic models of neutrino oscillation.

3.1 Overview of the experiment

To accurately determine the neutrino oscillation parameters, MINOS needs to measure the probability that a ν_μ produced at Fermilab is observed as a ν_μ at Soudan, and how this probability depends on the neutrino energy. To achieve this it is necessary to observe the energy spectrum of ν_μ at the Far Detector, and compare it to the spectrum predicted in the absence of oscillations. This prediction can be made by simulation, but uncertainties in the inputs will translate directly into errors on the measured oscillation parameters.

The dependence on a correct simulation be substantially reduced by measuring the neutrino spectrum close to the point of production, where oscillations are negligible, and using the simulation to extrapolate this measurement to the Far Detector. This is the one purpose of the Near Detector. The other benefit of the Near Detector is that the event rate is much higher, providing plenty of data on neutrino cross sections and detector performance with which to tune the the simulation. For this reason the Near Detector is designed to be as similar as possible to the the Far Detector, so measurements made at the Near Detector can be used in analysis of Far Detector data with a minimum of error.

3.2 The NuMI beam

The NuMI (Neutrinos at the Main Injector) beam uses protons accelerated to an energy of 120 GeV, by the Main Injector at Fermilab. In normal operation, the Main Injector can accelerate six bunches of protons at a time, of which five are extracted to NuMI, in a 8.7 µs spill[1]. Under favourable conditions, extraction can be up to every 1.87 s, although times of 2.4 s and longer are more typical.

Extracted protons are directed along the NuMI beamline, a cartoon of which is shown in Figure 3.1. The beamline points downward, at −3.3° to the horizontal, towards Soudan. The protons are incident on a 94 cm long target comprised of 47 fins of graphite. The target is mounted inside the barrel of a magnetic horn, and a second horn lies about 10 m downstream. Secondary particles (mainly pions and kaons) are produced by protons interacting in the target, and positive secondaries are focused down a 675 m long evacuated pipe by the horns. Neutrinos are produced by decay-in-flight of the secondaries (primarily $\pi^+ \rightarrow \mu^+ + \nu_\mu$). At the end of the decay pipe the remaining protons and secondaries are stopped by a water cooled absorber, while the muons stop in the 240 m of dolomite downstream of the beam.

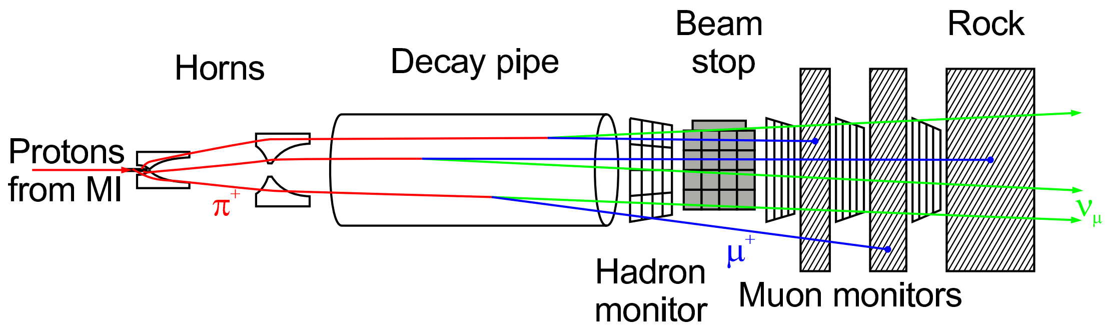

Figure 3.1: Cartoon of the NuMI beamline. In addition to the elements described in the text, ionisation chambers are used to monitor the hadrons and muons in the beam.

[1]The final batch is used to produce antiprotons. When antiprotons are not being produced six batches can go to NuMI, with a correspondingly longer spill time

Focusing by the horns is achromatic, and the momenta at which secondaries are focused can be chosen by adjusting the relative position of the target and horns, as depicted in Figure 3.2 The standard mode of operation is referred to as the the low energy or LE10 (denoting the target position) mode, where positive secondaries are focused such that the resulting spectrum of ν_μ is peaked at about 4 GeV. Other modes of operation sample secondaries with different ranges of momentum, and can be used to improve the beamline simulation and hadron production models. Production of neutrinos at higher energies can be optimised by moving both the target and the second horn, but for purpose of constraining the beam simulation only the target is moved.

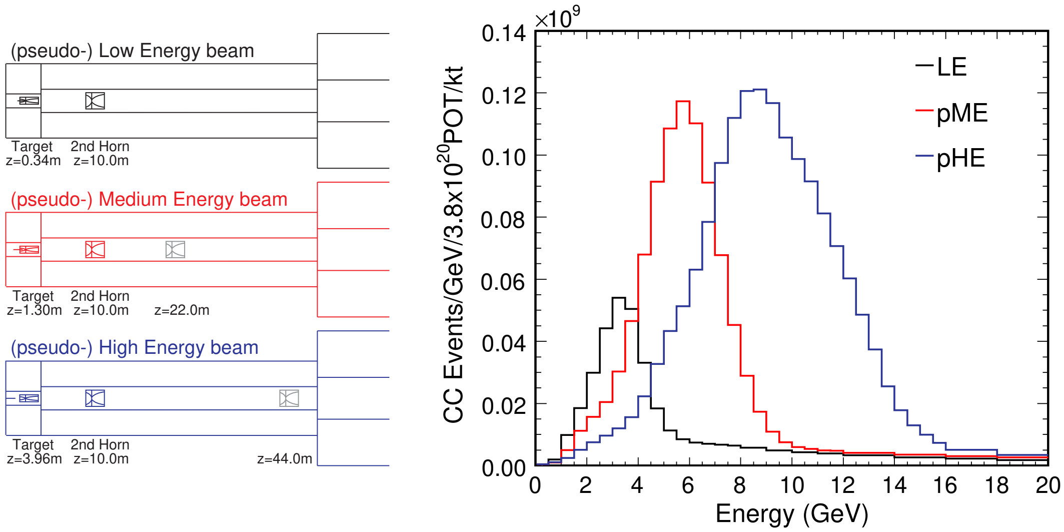

Figure 3.2: Beam spectra for different target positions. The horn 2 positions for the 'full' ME and HE beams are also shown in grey

In low energy mode, the tertiary neutrino beam is 98.7% ν_μ. Around 1.8% is $\overline{\nu}_\mu$, arising mainly from π^- that pass though the field-free centres of the horns, but also from μ^+ decays. Most of the remainder (roughly 0.5%) are ν_e, either from μ^+ decay or from K^+ produced at the target, with a small ($\sim 0.3\%$) contribution from $\overline{\nu}_e$.

3.3 The MINOS detectors

As implied above, the Near and Far Detectors share a similar design. Both detectors are sampling calorimeters [84], consisting of alternating planes of scintillator (1 cm thick) and steel (2.54 cm thick). The scintillator planes are divided into modules, encased in aluminium, and mounted on the steel planes. The combined steel/scintillator planes rest on a steel support structure at a 5.95 cm intervals, with the scintillator side facing the beam. Large coils passing through the planes provide a toroidal magnetic field in the steel. The exact construction of each detector is different, reflecting the much lower neutrino flux at Soudan.

3.3.1 Detector Volume

The number of events observed in the Far Detector is expected to be the most significant factor limiting the sensitivity of MINOS.[2] Because of this the Far Detector is as large as possible, with a total mass of 5.4 kt. Far Detector planes are regular octagons, eight meters from edge to edge. There are 484 planes of scintillator, grouped into two supermodules, of 248 and 236 scintillator planes each. Both supermodules are 'bookended' by an additional steel plane, giving a total of 486 planes of steel. Each supermodule has its own coil, which pass through the centres of the planes, and provide a toroidal magnetic field of 1.4 T (in the steel) at a radius of 2 m.

The Near Detector differs from the Far Detector primarily in the fact that at 0.98 kt, it is much smaller. The overall plane count is much lower at 282 steel planes, and the detector is not divided in two. The steel planes are flattened octagons measuring 4.8 m across and 3.8 m vertically. The coil is mounted off-centre, and passes through the right (looking in the beam direction) hand side of the steel planes. The fiducial volume of the detector is in the larger region to the left of coil, while the smaller region to the right of the coil primarily provides a flux return to shape the magnetic field in the fiducial volume. In addition, not all of the detector is instrumented with scintillator. Every fifth plane is a 'full' plane with scintillator covering the area to the left of the coil and some of the area to the right. Planes 1–120 are refereed to as the calorimeter, and the remaining planes in this region are 'partial' planes, covering just the area to the left hand side of the coil. From plane 121 onwards is the spectrometer, which is intended to be used only for measuring the momentum of muons emerging from the downstream end of the calorimeter. To keep costs down, the spectrometer does not have partial planes of scintillator, so only 33 of the 161 spectrometer planes are instrumented. Figure 3.3 shows the outline of full and partial planes.

3.3.2 Magnetic fields

For charged particles travelling in the general direction of beam axis, the toroidal magnetic field has a (de)focusing effect, depending on the particle's charge. Normally the fields in each detector are orientated to focus negatively-charged particles and defocus positively-charged particles, but special runs have been taken with the fields reversed. Because negatively-charged particles are focused inward, more of the μ^- produced in ν_μ charged current interactions are contained within the detector volume, which allows their momentum to be measured from their pene-

[2] The proposed lifetime of the experiment is around five years.

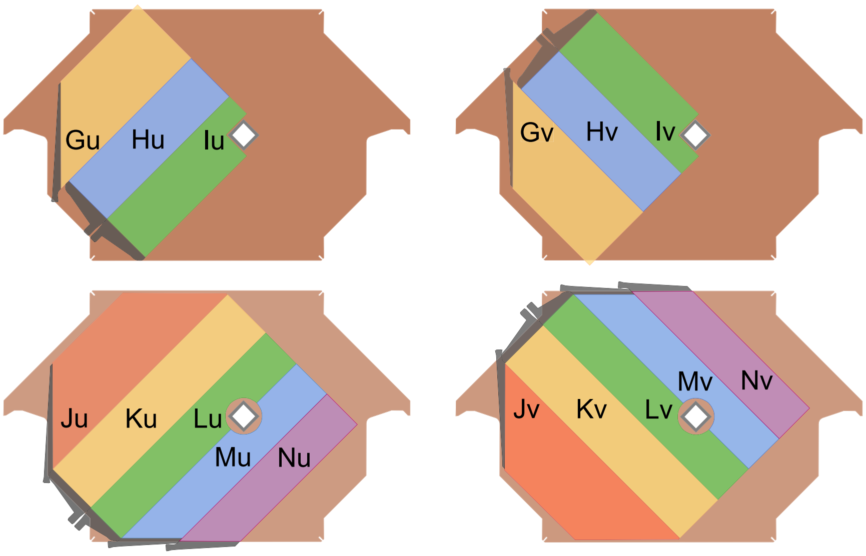

Figure 3.3: Outline of the Near Detector scintillator planes, showing 'Full' and 'Partial' planes as they appear when facing downstream. The coloured regions identify individual scintillator modules.

tration range. In addition, if a muon exits the detector, its momentum can still be determined from how it curves in the magnetic field.

3.3.3 Light detection

In both detectors, the scintillator planes are composed of rectangular strips of solid polystyrene doped with PPO and POPOP scintillators. The strips have a cross section of 4.1 cm $\times$ 1 cm and have a reflective TiO_2 coating to maximise light yield. Running down the centre of one side of the strip is a groove into which is glued a 1.2 mm Kuraray wavelength shifting (WLS) fibre doped with Y11 fluor. Strips are grouped into modules, and each module is optically sealed in aluminium. At the ends of the modules the wavelength shifting fibre leads to connectors where a small enclosure allows the fibres to be flashed with light as part of the calibration system. In the Far Detector there is readout (and therefore connectors) on both ends, while in the Near Detector readout is from one end only, with mirrors on the other end of each strip. The Far Detector strips are around twice as long as those in the Near Detector, so the detectors see a similar overall light yield at the longitudinal centre of a strip. Far Detector planes have 192 strips, Near Detector full planes have 96 strips, and Near Detector partial planes have 64 strips.

From the connectors, bundles of clear optical fibres connect the strips to Photo-Multiplier Tubes (PMT) housed in light-tight enclosures in racks around the Detector. As well as preventing light leakage, the steel enclosures significantly reduce the magnetic field strength to which the PMTs are exposed. The Near Detector

PMTs are housed individually, while at the Far Detector three PMTs sit in a single enclosure.

The PMTs used in the Near and Far detectors are Hamamatsu R5900-00-M64 and R5900-00-M16 multi-anode PMTs (commonly just M64s and M16s). The faces of the M64s are about 2.8 cm square, and they have a depth of 2 cm[85] There is a single photocathode maintained at a potential of around −800 V with respect to the anode.[3] Behind this are twelve dynodes and the anode which is segmented into a grid of 8×8 pixels, each of which can be read out independently. The M16s have a similar construction, but the anodes are segmented into a 4×4 grid of pixels. The larger size of the M16 pixels allows eight fibres to be read out by a single pixel. This optical multiplexing reduces the number of electronics channels required, but introduces an eightfold ambiguity as to which Far Detector strip was really hit. The mapping of strips to pixels is different at each end however, and this helps to resolve the ambiguity. This information is used by a dedicated *demultiplexing* algorithm during reconstruction.

Compared to regular PMTs, the multi-anode PMTs have several advantages. Asides from cost, their small size means they are more resilient to magnetic fields [86] and have comparatively low dark noise [87]. Unfortunately there are also disadvantages. Firstly, it is possible for photoelectrons to move between pixels, causing cross-talk between strips. This complicates reconstruction, but is fairly well understood. Secondly, the small size of of the pixels makes them susceptible to space-charge effects, causing the response at high light levels to be a non-linear function of the number of photons incident on the photocathode.

3.3.4 Digitisation

The second major difference between Near and Far detectors is also due to event rate. The Far Detector sees neutrino events on the order of once per day, and the overall event rate is dominated by cosmic ray muons which occur with a frequency of about 1 Hz. In contrast, there are several tens of neutrino interactions in the Near Detector during a single spill, giving an instantaneous event rate of a few MHz. The net result is that, while the readout of the Near Detector must be continuous, the provision of similar electronics at the Far Detector would be inappropriate and expensive.

Near Detector front-end electronics

The readout for the Near Detector is based on the Fermilab QIE (*Charge Integrator and Encoder*) ASIC, and digitises continually in time buckets of 18.8 ns[4]. The QIE uses an 8 bit ADC but divides charge across a sequence of capacitors in

[3]The high voltage supplied is tuned to keep the average gain on all pixels at 0.8×10^6 electrons per photoelectron.

[4]This is to match the 53.1 MHz RF used by the Main Injector

order to achieve a dynamic range close to that of a 16 bit ADC. A full description of the QIE ASIC can be found in [88]. Each PMT anode pixel is read out by a single QIE, which sits with its own ADC and FIFO buffer on a small circuit board called a MENU (*MINOS Electronics for Neutrinos*). Sixteen MENUs sit together on another board, the MINDER (*MINOS Near DEtector Readout*) board, which provides timestamps for the digits, and charge injection circuitry to calibrate the QIE-ADC. Control and read-out of the MINDERS is via a MASTER (*MINOS Acquisition, Sparsifier, and Time-stamper for Event Records*) board which is housed separately, up to eight together, in a VME crate. Each MASTER uses look-up tables to calibrate the MENU digits and discards digits below a threshold corresponding to around 1/3 of a single-photoelectron signal or *pe*.

Readout of the Near Detector is triggered in one of two ways. When protons are sent to NuMI a signal from the accelerator complex is use to initiate an acquisition period of up to 19 µs during which every time bucket is digitised and stored in the FIFO. After the acquisition period has ended the FIFO buffers are readout by the the MASTER crates. Between these 'spill gate' triggers an alternative trigger, derived from the last dynode of each PMT, initiates a 150 ns readout period for the PMT and is used to record cosmic ray muons.

Far Detector front-end electronics

In the Far Detector, the PMT anode signals are amplified by the multi-channel VA32_HDR11 (usually just VA), a modified version of the *Viking* chip manufactured by IDE AS of Norway [89]. Three VA chips are mounted on a VFB (*VA Front-end Board*) which is mounted on the side of the PMT enclosure. Each VA has 32 channels, and reads out a single M16 PMT, with one of the unused channels used for calibration and four used for common-mode noise subtraction . Each channel consists of a pre-amplifier and shaper, which produce a unipolar output that continuously tracks the PMT anode signals, and a sample and hold circuit.

In parallel with the VA, a signal is taken from the last dynode, which is common to all 16 pixels of the PMT. This dynode signal is fed to discriminator on the VFB (the *ASDLite*[5]), which provides an output signal if the input exceeds a threshold of around 1/3 of a *pe*. This threshold is tuned to reject sub-*pe* noise as described in [91]. The signal from the ASDLite is sent to a remote VME board called the VARC (*VA Readout Controller*), which can handle twelve VFBs or 36 PMTs. The trigger signal is used to generate a timestamp for the digits, with a precision of roughly 1.5 ns, and a signal that is sent to the sample and hold circuit of the VA chip. There is a tuned delay of around 500 ns so that the sample and hold occurs when the output of the VA peaks after a charge signal on the PMT anode. The held anode signals are sent to an ADC which sits on a mezzanine card[6] on the VARC, and serves two VFBs, or six PMTs. Following a single dynode trigger, digitisation

[5] A prototype of the octal ASD used in ATLAS muon chambers. [90]

[6] The VMM or *VARC Mezzanine Module*. Each VARC has six VMMs.

by the ADC takes 5 µs, during which the VA is dead to further signals. As the ADC serves six VA chips, and reads each out in a round-robin fashion, a single VA can be dead for up to 30 µs. This should not be a problem because of the low event rate at the Far Detector, but during commissioning an unexpectedly high trigger rate of around 5 kHz per VA was observed, and eventually traced to single photon noise in the scintillator/WLS fibre system [92]. This level of noise would result in an unacceptable level of dead time, so the digitisation sequence is only started if two dynode triggers are seen on the same VARC within 400 ns of each other. This is referred to as the 2:36 trigger as it correspond to a hit seen in 2 of the 36 (sometimes 30) PMTs read-out by a single VARC. Three PMTs are needed to read out a single plane, and the mapping of channels is such that the 36(30) PMTs correspond to a contiguous set of 24(20) planes. The trigger therefore has little impact on real events, and is in any case modelled in the simulation.

3.3.5 Data acquisition and software triggers

The electronics of both detectors have a VME interface to dedicated read-out-processors that communicate that transfer the data via PVIC [93] buses to small trigger farms of PCs.[7] The trigger farms are responsible for identifying data of physics interest, which are recored for offline analysis. The output of a single software trigger can contain multiple events (particularly in the Near Detector) and is termed a *snarl*. In addition to a number of special triggers for calibration and diagnostic purposes there are several physics triggers:

Spill triggers are generated when it is known that a NuMI beam spill was passing through the detector. At the Near Detector data from the spill gate readout is identified by the front-end electronics and grouped as a single snarl by the trigger PCs.

At the Far Detector such direct triggering is unavailable. Instead the time of the spill gate triggers are recorded by a GPS system that also provides global timing information for the Near Detector. These times are transmitted to the Far Detector over the Internet, and can be used to estimate the arrival time of the neutrino beam. Although these signals arrive after the neutrino beam, there is plenty of buffering capacity, and ample time to wait for the arrival of spill information. A window of ±50 µs is recorded around the estimated arrival time (along with a 30 µs pre-trigger) during which all hits passing the 2:36 trigger are written to file. This trigger is properly referred to as the remote spill trigger, to distinguish it from the direct spill trigger in the Near Detector, but when talking about the Far Detector alone it is commonly to refer to it as simply the spill trigger.

As well as real (remote) spill trigger, triggers are generated at random in

[7]The data acquisition system is detailed in [94].

the periods between spills. These 'fake spill triggers' are treated exactly like real spill triggers, and are used to estimate detector backgrounds.

Non-spill triggers are use to record events based on the presence of a signal in the detector. The utility of these triggers is twofold: to collect cosmic-ray and atmospheric neutrino data; and, in the case of the Far Detector, to collect beam neutrinos in the event of a loss of communication with Fermilab. Before it was realised that it was possible to implement the remote spill trigger, these were to be the primary triggers for Far Detector beam data. There are three such triggers: The *plane trigger*, which requires that four out of five contiguous planes contain hits. The *energy or E4 trigger*, which requires a total raw pulse height of 1500 ADC counts in four consecutive planes, and at least six hits in two or more consecutive planes. And the *activity trigger* which requires hits in any 20 planes of the detector.

3.3.6 Detector calibration

Shower energies in MINOS are determined by summing the energy seen in individual strips of the detector. To do this it is necessary to account for variations in strip response and this requires calibration. Many of the required calibrations make use of the cosmic ray muons that pass through (and occasionally stop in) the detector. There is also a dedicated Light Injection system that illuminates the WLS fibres.

The Light Injection system

In order to monitor the long term stability of each fibre it is useful to have a consistent source of illumination. This is provided by a LED-based Light Injection (LI) system. [95]. The system consists of a number of *Pulser Boxes*, each of which contain 20 blue LEDs, plus control circuitry. When flashed each LED illuminates 71 optical fibres, 64 of which are routed to illuminate the WLS fibres (the rest are either used for monitoring or are spare). Each one of 64 fibres is routed to an enclosure near the edge of the scintillator modules where it illuminates either eight or ten WLS fibres simultaneously. By flashing all the LEDs in all pulser boxes in turn every optical channel in the detector can be tested. At the same time one of the seven additional fibres is routed direct to a PMT (the trigger PMT, or tPMT) to provide clear indication that the Light Injection has taken place.

The LI system is used in two modes. By far the most common mode is the *drift point* LI which illuminates each strip regularly with a fixed amount of light. This makes it possible to monitor the Gain of the PMTs on hourly timescales. By comparing mean of many pulses to their RMS withs, it is possible to calculate the number of *pe* seen by the PMT pixel, and therefore the gain (in ADC/*pe*) of that pixel. This information is useful for refining the simulation, and various reconstruction tasks, such as deciding which hits are genuine and which are detector noise.

The other mode in which the LI system is used is to provide a mapping of the PMT/electronics gain as a function of light level, which is not linear at high light levels. A sequence of this *gain curve* LI requires flashing each strip at several different intensities, and takes much longer than a single sequence of drift point LI. Gain curves are taken on a monthly basis, however, and therfore comprise a much smaller fraction of all LI flashes.

The calibration chain

The raw pulse height Q_{raw} of a strip is corrected be several multiplicative factors, dependent on time t, the hit strip s, the position of the hit along the strip l, and differing between detectors:

$$Q_{\text{calib}} = Q_{\text{raw}} \times M \times D(t) \times L(s, Q_{\text{raw}}) \times S(s, t) \times A(s, l) \tag{3.1}$$

Each of the factors M, D, L, S and A represent a step in the calibration and each is described below.

Perhaps the most important calibration is the inter-detector calibration M, often referred to as the stopping muon calibration. This compares the overall response of the detectors, using muons that stop in the detector volume as a standard candle (hence the name). A detailed description of this calibration can found in [96].

Apart from M, which corrects for differences between the detectors, there are four intra-detector calibrations. Each of these four factors is scaled such that the average correction is close to unity, so that uncalibrated values can be used to approximate the correctly calibrated value.

$D(t)$ is the drift calibration, which compensates for the change of readout components with time, particularly ageing of the PMTs and scintillator. It uses through-going cosmic ray muons as a standard candle, and tracks the response of the entire detector on a daily basis.

$L(s, Q_{\text{raw}})$ is the linearity calibration, which corrects for the nonlinear response of the PMTs, and (in the Far Detector) the front-end electronics. This makes use of the light injection calibration system described above. The corrected ADC value at this point is given the unit “SigLin” to distinguish it from a raw ADC value.

The next calibration is $S(s, t)$, the strip-to-strip calibration, which also uses through-going muons to correct for variation between scintillator strips, PMT pixels, and electronics channels. This requires that the scintillator strip (rather than just the read-out channel) can be identified. At the Far Detector this means information about other hits must be used to break the ambiguity introduced by the multiplexing. The corrected ADC value after this step is given the unit “SigCor”

The final intra-detector calibration is the attenuation correction, $A(s, l)$. This corrects for light attenuation along the scintillator strip, using measurements from a radioactive source scanned along every strip before installation.[8] Because the calibration requires knowledge of the position of the hit along the length of the strip, information from the two detector views must be combined. This can be reliably done for muon tracks, but it is harder for showers. As a result the calibrated ADC quantity ("SigMap") is not available for showers.

Even after all these calibrations have been performed, all that is known is the ADC response to a scintillator hit. We must additionally know the scintillator response to particles of a given energy, and the particle composition (and energies) of neutrino-induced showers.

The single-particle response was measured directly, using a scaled-down version of the MINOS detectors. This detector, called the Calibration Detector, was functionally identical to the two main detectors, similarly consisting of interleaved steel[9] and scintillator planes exposed to test beams at CERN. The response of the MINOS detectors to single hadrons and to single electrons was measured in CalDet and is described in [97] and [98] respectively. By instrumenting CalDet with both Near and Far detector electronics, it was also possible to confirm that the differences in detector read-out were well understood and correctly modelled. [99]

Results from CalDet can be utilised by including them in the simulation of particles propagating through the detector, and informed the decision to use the GCALOR [100] hadronic interaction simulation code for the propagation of shower particles. There is, however, one final factor affecting shower response. This is the types and multiplicities of particles produced in a hadronic shower from neutrino-iron interactions, and the way energy is divided among them. It is extremely difficult to measure this with MINOS, because the steel planes are relatively thick, and the scintillator bars ar too wide to resolve individual particles. As a result the composition of showers is accounted for only by the Monte Carlo simulation, which is constrained by previous experiments.[10] The resulting uncertainties (particularly with regard to the amount of energy absorbed in excitations of the Iron nucleus) are one of the most significant systematic uncertainties remaining for MINOS analyses. Careful analysis of Near Detector data may eventually reduce these uncertainties, as may results from future experiments such as MINERνA, [102] which will also use the NuMI beam line.

[8] In the Near Detector this has been superseded by a correction using cosmic ray muons.

[9] The steel used for CalDet was slightly thinner, at 2.5 cm, a difference that is taken into account in analysis of CalDet data.

[10] For this part of the simulation, MINOS uses the NEUGEN [101] neutrino event generator.

3.4 MINOS Monte Carlo simulation

Although MINOS is designed to be a data-to-data comparison between the Near and Far Detectors, in practice it is still necessary to have a large quantity of simulated data (often referred to as MC as it uses the Monte Carlo method). There several reasons for this:

- Simulated data allows us to correct the for differences between neutrino event spectra in the Near and Far Detectors that are not due to oscillations. There are two main sources for these differences, the detector acceptance and the spectrum of neutrinos in the beam passing through each site. Chapter 5 discusses ways the simulation is used to account for these differences.

- The number of neutrino events at the Far Detector is limited, which makes it difficult to design and test reconstruction algorithms. In the Near Detector the opposite problem is encountered; there are multiple events in normal beam snarls and it it is useful to have a sample of single event snarls for initial tuning. It is also useful to use the detailed 'truth' information available about the origins of each hit in the detector to tune the reconstruction of normal snarls.

- The simulation provides a way to evaluate different analysis approaches, so that less sensitive methods are not pursued. This includes different hardware conditions, so that they can (where possible) be optimised without wasting running time of the experiment.

- Analyses can be tested on simulated data for a range of oscillation hypotheses to check that they do not have a systematic bias. The simulation also provides a way to evaluate the systematic errors induced by various uncertainties that might be impossible or impractical to study with data alone.

There is one (highly pervasive) downside to using simulated data; inaccuracies in the model may lead to systematic errors in the analysis. This is outweighed by the benefits, and in any case the Near Detector allows MINOS to re-tune the simulation or simply avoid making use of aspects that are inaccurate.

The simulation itself can be roughly divided into three parts:

GNUMI is the simulation of the beamline. It is based on GEANT3 [103], and uses FLUKA [104] to simulate the production of hadrons from the target. The primary output of GNUMI are fluxes of neutrinos (ν_μ, ν_e $\overline{\nu}_\mu$, and $\overline{\nu}_e$) at each detector site.

GMINOS is the simulation of neutrino events in the detectors and surrounding rock. It selects neutrinos to interact based on provided fluxes and models the interaction using NEUGEN [101]. The long-lived final-state particles are tracked through the detector using GEANT-3, with shower development handled by GCALOR [100]. Energy depositions in the scintillator are recorded

PhotonTransport and DetSim are C++ based simulations of the scintillator and electronics systems respectively. PhotonTransport generates photons in the scintillator from the GMINOS truth hits, handles absorption, re-emission and transport in the WLS fibres and provides a list of photoelectrons generated at each PMT photocathode. DetSim simulates amplification by PMTs, digitisation of PMT anode signals and the hardware triggering systems. The final output of DetSim are digitised hits, similar to what is produced by the real detectors. These hits can then be fed into the same reconstruction and analysis code as real data.

3.5 Reconstructing events

Both simulated and real data are passed through the same reconstruction chain so they can be analysed. The reconstruction is characterised by the formation of increasingly sophisticated objects from hits in the detector. A (necessarily brief) summary is given here.

The individual hits in the detector are referred to as *digits*; each digit corresponds to an single channel readout, with a time and pulse height. In the Far Detector each digit corresponds to eight physical scintillator strips, and one of the first steps in the reconstruction is to resolve this ambiguity. This is the demultiplexing, or demuxing step.

Before demultiplexing can take place however, an algorithm called the chopper is run. The chopper is configured to find 150 ns gaps in detector activity, and break the snarl into 'chops' at these gaps. It then passes the largest (by summed pulse height) chop to the demultiplexer. This prevents out-of-time noise from confusing the demultiplexing algorithm, and improves the agreement between spill data (which has snarls of 100 µs duration) and simulation (where noise is generated in a window of only 80 ns).

The next level of object is the *strip*. Strips are amalgamations of digits (but can correspond to single digits) and are intended to correspond to a single energy deposition. In the Near Detector, where digits are made continually, strips are made up from concurrent (in time) digits on a single readout channel. In the Far Detector, digits corresponding to opposite ends of a single physical strip are combined into reconstructed strips.

In the Near Detector (where there are many neutrino interactions in each snarl) the strips that are close in both space and time are grouped together to form *slices*. A single slice is supposed to correspond to a single neutrino interaction, although this is not always the case. In the Far Detector this stage is not necessary. From the strips in each slice the reconstruction attempts to form *tracks* and *showers*.

Tracks are formed in two steps. Firstly the track finder looks for linear clusters of hits in each view and combines adjacent clusters together to form the tracks. Then the track fitter runs over these tracks using a Kalman filter to obtain measurements of the momentum and charge of the particle that made the track, as well as identifying the start (vertex) and end points. Normally tracks correspond to muons, but may be pions or even protons.

Showers are more generic clusters of strips in a particular area of the detector, which can correspond to the hadronic part of a neutrino-nucleus interaction, to a electron shower from a ν_e charged current interaction, or to a delta ray from a muon. The shower finding code estimates the energy of the shower, and the location of the vertex from which it was produced.

Finally, showers and tracks that are close in space and time are combined to form *events*. Events are intended to represent complete single neutrino interactions. The canonical example is a muon track being associated with a hadronic shower at the neutrino interaction point, but other combinations are possible. In particular, most neutral current events have no track and are reconstructed as an isolated shower.

3.6 Blind Analysis

In common with many other experiments, MINOS pursues a policy of blind analysis. This is to reduce the possibility of experimenter bias, which presents an a potential source of systematic error that is extremely difficult to quantify. There are many strategies for blinding an analysis [105], but for MINOS oscillation analyses this is achieved by distorting the Far Detector data spectrum.

The blinding scheme is applied to the Raw Data at the very start of reconstruction. As the early stages of reconstruction are not performed by users, so this effective centralises the application of blinding and reduces the possibility of an accidental unblinding. To scheme involves placing Far Detector events into a hidden dataset according to an unknown probability, with the remainder being open to all analyses (although they are free to apply additional blinding schemes as desired). The probability to place an event in the open set is bounded between 0.25 and 1, and is an oscillating function of the number of planes containing hits, and the total (uncalibrated) pulse height in the detector. The oscillating function has four unknown parameters that are obfuscated using a pseudo-random number generator. To unblind, the open and hidden data sets are simply recombined. Permission to unblind is granted to a particular analysis once it is finalised, this requires the selection criteria, Near-Far extrapolation, and data sets to be defined, although bugs found after unblinding can be fixed at the approval of the collaboration. Analysers are expected to avoid looking at distributions that are important for other analyses. This thesis abides by the blinding policy: the data used were unblinded twice, once for the analysis of ν_μ charged current events, and once for the neutral current events.

3.7 Neutrino physics with MINOS

After 5 years of running, the NuMI beam will have delivered somewhere in the region of 10^{21} protons on target (POT). In the absence of oscillations this should result in around 2 000 identifiable ν_μ charged current events. This should be sufficient to determine $\Delta m^2_{\rm atm}$ and $\sin^2 2\theta$ to a precision of 10% and 0.1 (at 90% C.L.) respectively. The sensitivity of MINOS comes from the fact that it is able to observe the spectral distortion caused by oscillations, as shown in Figure 3.4. Neutrino energies can be reconstructed by summing the energy from the track and from the shower. Track energies can be obtained with a typical resolution of 13% by measuring the curvature in the magnetic field, or 6% from range if the track stops in the detector. [57] The calorimetric response to single hadrons gives a resolution of around $55\%/\sqrt{E/\,\mathrm{GeV}}$ [97] and the overall resolution of hadronic shower energies is similar. In the sample of events selected as ν_μ charged current the muon usually caries most of the neutrino energy, and so the resolution of the neutrino energy is of the order of 10%.

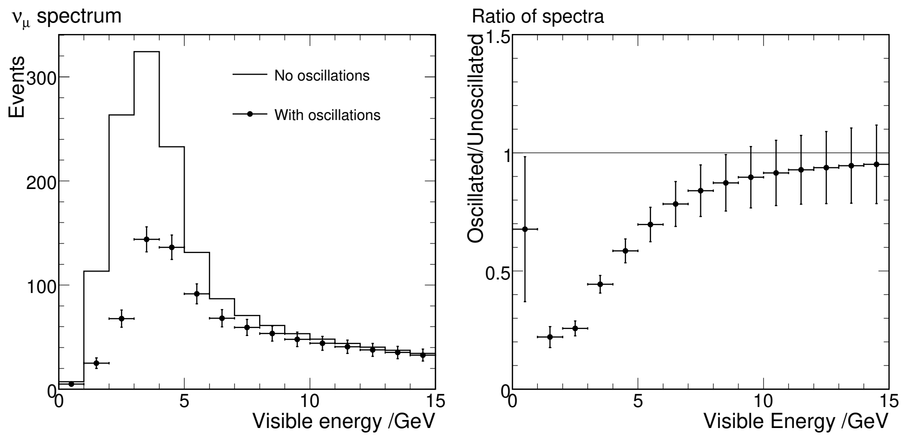

Figure 3.4: Cartoon of the ν_μ disappearance analysis, showing the dip signature of oscillations. The position of the dip gives information on $|\Delta m^2_{\rm atm}|$, while its depth gives information about $\sin^2 2\theta_{23}$.

The ability of MINOS to observe the neutrino spectrum allows direct comparison between different models of neutrino disappearance. It should be possible to convincingly distinguish between oscillations and alternatives such as the decay or decoherence models described in Section 2.4. Figure 3.5 shows how the reconstructed energy spectrum varies between models of neutrino disappearance.

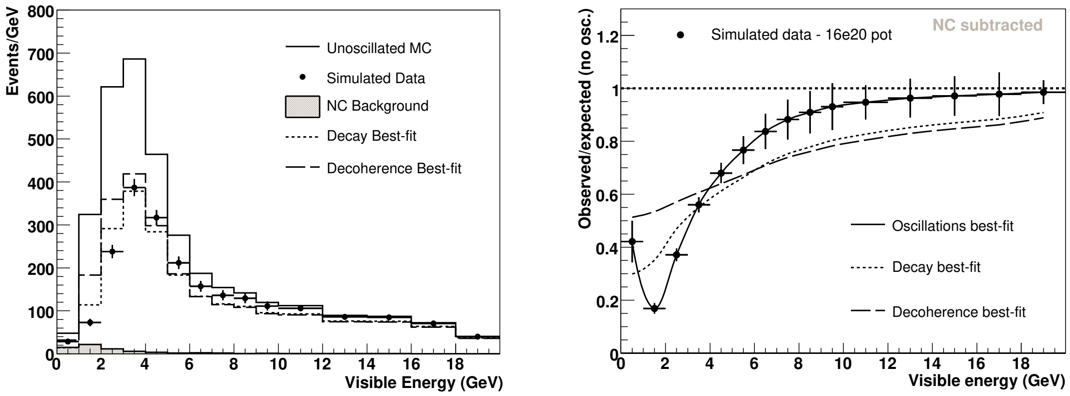

Figure 3.5: Simulated oscillation signal in MINOS, showing the spectrum and ratio to the 'no disappearance' case. Best fits for decay and decoherence models are also shown, demonstrating MINOS's ability to discriminate between models.

Appearance analyses

The search for oscillations in the sub-dominant channel $\nu_\mu \to \nu_e$ entails looking for electromagnetic showers from ν_e interactions, above the level expected from the beam contamination. This analysis is quite challenging, requiring first that the electromagnetic (π^0) component of hadronic showers is well modelled, but also a good understanding of the small ν_e component in the beam. Nevertheless, MINOS hopes to be able to make a more precise measurement (limit) of θ_{13} than currently set by the CHOOZ experiment. Figure 3.6 shows the sensitivity to $\sin^2(2\theta_{13})$, which depends on the value of the $\mathcal{CP}$ violationg parameter, δ.

Another possible sub-dominant channel for oscillations is the $\nu_\mu \to \nu_s$ channel, as described in Section 2.4.3. It is possible to look for oscillations into sterile neutrinos by looking for a depletion in the rate of neutral current events,[11] and (as with the ν_e appearance analysis) this requires a good understanding of showering events. A specific difficulty with the analysis of neutral current events is the fact that the neutrinos only deposit part of their energy in the detector, making the reconstructed energy spectrum substantially less useful than in other analyses. Chapter 6 of this thesis describes a $\nu_\mu \to \nu_s$ analysis based on the first year of MINOS data.

Antineutrino analyses

By making use of the magnetic field to determine the charge of final state muons, it is possible to measure the survival probabilites seperately for ν_μ and $\overline{\nu}_\mu$. Dif-

[11] This is still an appearence channel, in the sense that it looks for a flavour of neutrino that is not present in the original beam, and $\mathcal{CP}$ violating effects are observable, at least in principle.

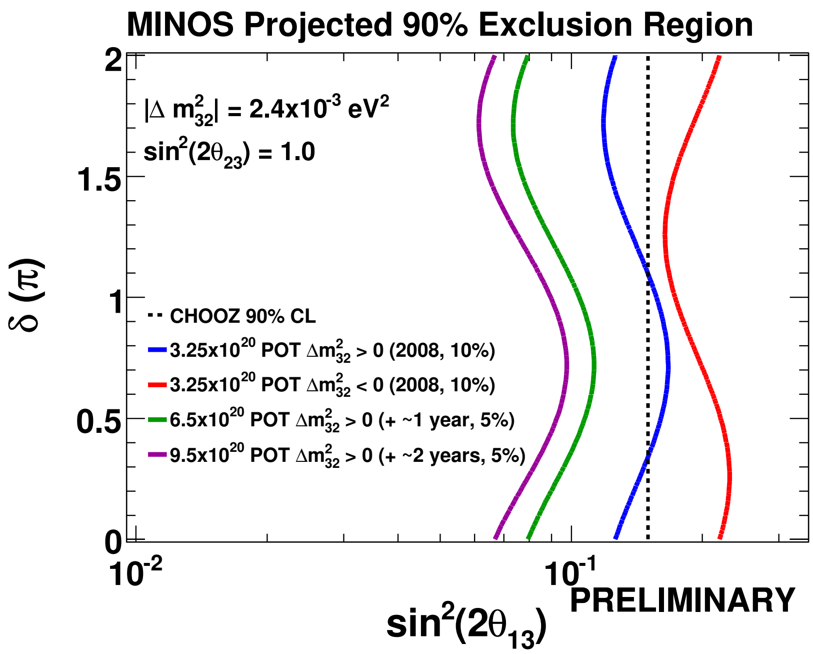

Figure 3.6: MINOS sensitivity to ν_e appearence, for various numbers of POT. For 3.25×10^{20} POT the effect of the mass hierachy is also shown.

ferences between the survival probabilites is a violation of $\mathcal{CPT}$, so would be a signature of new physics. The much lower flux of $\overline{\nu}_\mu$ makes this analysis much more difficult than ν_μ dissapearence analysis.

Another μ^+-based analysis that instead benefits from the low $\overline{\nu}_\mu$ flux is a search for $\nu_\mu \to \overline{\nu}_\mu$ transitions. This would require neutrinos to be Majorana particles, and also that there is a spin-flip transiton, so the $\overline{\nu}_\mu$ can interact weakly. Needless to say, observation of $\nu_\mu \to \overline{\nu}_\mu$ would also be evidence of new physics.

Chapter 4

Identifying neutral current events at the Far Detector

Neutral current (NC) events are the second most common form of beam neutrino event in MINOS. The general form for a neutral current neutrino interaction is

$$\nu + X \rightarrow \nu + X'.$$

If X is one of the fundamental fermions then it will not change flavor as a result of the interaction (i.e. $X' = X$) but compound hadrons such as protons and nuclei can fragment as a result of the energy transfered. The dominant process for the hadronic part of the interaction depends on the energy transfered to the hadronic system. The energy range of interest in MINOS goes from a few hundred MeV to a few GeV. At the low end of this range a nucleon can scatter elastically, whilst for higher energies deep inelastic scattering (DIS) dominates. Between these extremes both processes can occur, along with formation of short-lived resonant states.

The visible result of such an interaction is typically a hadronic shower. Although other final state topologies are possible, such as elastic scattering resulting in short proton tracks, these states are difficult to observe with the resolution of the MINOS detectors. Because the escaping neutrino still carries an unknown fraction of its original energy, it is not possible to know what the energy of the incident neutrino was, only that it must have been higher than the observed shower energy.

These factors make NC events much harder to reconstruct and analyse than the ν_μ charged current interactions. Firstly, the hadronic showers are much less well defined than muon tracks, and this makes it harder to separate individual events, especially in the high-rate environment of the Near Detector. This is compounded by the fact that NC showers always have lower energy than the incident neutrino, so many are close to the lower limit of what MINOS is able to reconstruct. In the Far Detector this is the more significant issue, especially since detector noise is not well understood. Secondly, because NC events 'feed down' from higher energies, a event observed with a few GeV of energy could have come from an interaction at any higher energy, up to the maximum beam energy (around 120 GeV). It is therefore possible for uncertainties affecting a small part of the neutrino energy spectrum to affect a broad range in the reconstructed energy spectrum, making it difficult to disentangle various sources of uncertainty. Finally, a sterile oscillation signal affecting neutrinos with energies in the $1 \sim 2\,\text{GeV}$ range will cause the greatest depletion of neutral current events in the lowest energy range, making it important to keep reconstruction and selection thresholds low.

Outline

This chapter describes how neutral current events are selected in the Far Detector. This problem is conveniently split into two steps. In the first step, neutrino events are separated from the non-neutrino background. In the second step step the neutrino events are separated into neutral current like, and (ν_μ) charged current like categories. The two steps are rather different, firstly because we wish to retain both the CC-like and NC-like neutrino events. The second and more important difference is that there is no good simulation of non-neutrino backgrounds, so they must be estimated from data. Fortunately, most of these backgrounds can be estimated from data taken out-of-spill (i.e. when it is know that there in no beam from NuMI.

The chapter is divided into three sections. The first details the the non-neutrino backgrounds, and the cuts that are used to remove them. The next section presents an evaluation of the systematic uncertainties resulting from the backgrounds and their removal. The final section details the the method used to discriminate between NC-like and CC-like neutrino events.

When considering cuts we have two (usually) competing aims: to retain as much of the signal as possible, while reducing the fraction of background events. To quantify this; if S and s are the number of signal events before and after a cut and b is the number of background events after the cut then we can characterise the efficiency of the cut $\mathcal{E} = s/S$, and the purity of the resulting selection $\mathbb{P} = s/(s+b)$. If several selection cuts are applied in sequence the overall efficiency of the full selection is simply the product of the individual efficiencies, while the purity is just that of the final set of events.

4.1 Background rejection

In the MINOS Far Detector, there are 3 main sources of non-neutrino background: light injection, cosmic ray muons, and fibre noise.

Light injection (LI) snarls occur when the light injection calibration system flashes during a beam spill. Because the beam originates a long way from the Far Detector site it is not practical to avoid flashing during beam spills, and the LI flashes around 10^5 times each hour [84]. As a result a great many snarls contain LI activity.

A more worrying form of detector noise is that which originates from the the readout system itself. This could be noise in the electronics, dark noise originating in the PMTs, or spontaneous photon emission from the WLS fibres and scintillator bars. The M16s have a very low rate of 1 *pe* dark noise, as mentioned in Section 3.3.3. Sub-*pe* noise from PMTs and electronics is eliminated by a having a threshold on the dynode trigger, dynamically set at about $\frac{1}{3}$ *pe* [91].

The dominant source of 1 *pe* noise comes from the scintillator/WLS fibre system, [92] and is hence referred to as fibre noise. The dead time incurred by digitizing this noise would be unacceptable, so the 2:36 trigger described in Section 3.3.4 is applied. This dramatically reduces the amount of noise that makes it into the data stream. But since it is clustered noise that mimics events, the 2:36 trigger has a much less significant impact on the rate of events reconstructed from the noise, which occur at about 50 times the rate of real neutrino interactions.[1]

The final common background is muons coming from cosmic rays. The rate of cosmic ray muons entering the Far Detector is around 0.5 Hz. [106] Since spill-triggered snarls last around 100 µs we should expect to see a cosmic event in one of every 20 000 snarls—about the same rate as beam neutrino events.

4.1.1 Removal of light injection

LI typically illuminates several channels on three VARCs, corresponding to up to 64 planes, simultaneously. Cross-talk in the PMTs means that almost all channels on the illuminated planes will be hit. Following a LI flash, these channels will be dead for up to 30 µs. For this reason, and because of the low prior probability of getting a beam neutrino event coincident with a LI flash, we chose to discard the *snarl* if there is evidence of LI activity, as opposed to discarding just an event. The difference is subtle; because we ignore all the data when LI is seen, the live time of the detector should in principle be adjusted accordingly. In practice, this correction is of order 10^{-5}, so can be ignored.

The primary means to identify LI is via a dedicated piece of hardware, the trigger PMT (tPMT). This PMT is illuminated directly by the LEDs of the LI system, so receives plenty of light. As a result the sparsification threshold can be turned up very high, eliminating accidental triggers and dead-time inefficiencies. A cut on the existence of a tPMT hit is therefore a very effective (inefficiency $\lesssim$ 0.01%) way of identifying LI snarls. The tPMT selected LI therefore also provides a good sample to test subsequent LI cuts

The next approach to removing LI is a set of conditions, together referred to as the LISieve, *all* of which must be be failed by an snarl for it to be identified as LI. These are:

[1]In low energy configuration at typical intensities of 20×10^{12} protons per spill.

Hits per plane During LI there are either eight or ten strips illuminated simultaneously on a single plane. After cross talk occurs at the PMT most of the channels corresponding to an illuminated plane will be hit. The reconstructed number of strips hit on a plane should therefore be close to the limit set by the demultiplexing algorithm, which assumes all hit strips are within 1 m of each other (giving a maximum of 24 reconstructed strips on a plane).

In gain curve LI, where the illumination level can be lower, it is possibly for some strips to remain un-hit. This is particularly true for low efficiency channels, but the mean is still high compared to neutrino events, which typically hit fewer than five strips per plane. The cut is shown graphically in Figure 4.1

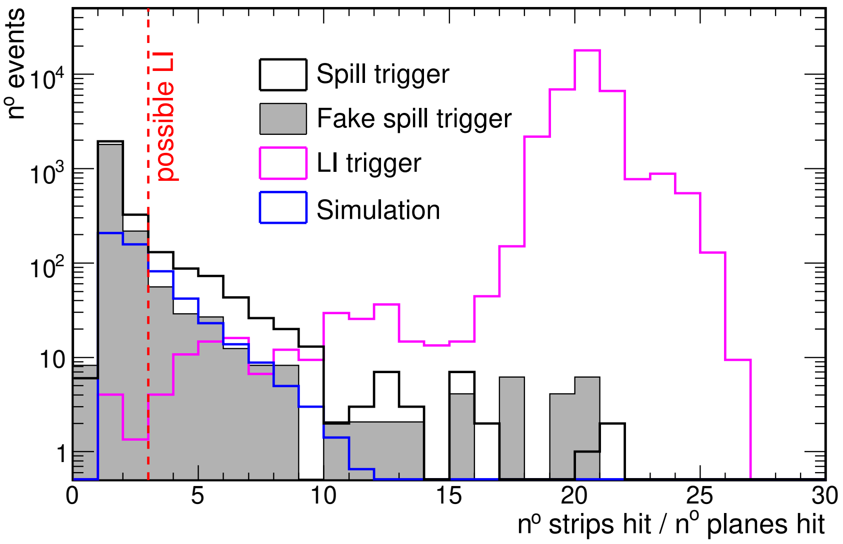

Figure 4.1: Mean number of hits per plane for events in spill and fake spill-trigger snarls (black), tPMT tagged snarls (magenta) and simulated beam neutrino events (blue). More than 3 hits per plane is considered indicative of LI.

West-East asymmetry The Pulser boxes where the LEDs are housed sit along the sides of the detector, and the fibre fan-outs illuminate the WLS fibres emerging from one end of each scintillator strip. The re-emitted light in the WLS fibres is isotropic, but because of attenuation in the fibres the light level will be around five times higher at the nearby readout than at the opposite end. [92] For snarls containing real events, the light in the fibres originates more centrally within the scintillator volume, so the asymmetry will be lower. We can define an asymmetry variable:

$$A_{\mathrm{WE}} = \frac{\sum PH_{\mathrm{W}} - \sum PH_{\mathrm{E}}}{\sum PH_{\mathrm{W}} + \sum PH_{\mathrm{E}}}, \tag{4.1}$$

and take high values of $|A_{\mathrm{WE}}|$ to be indicative of LI.

There is a complication here however. Gain curve LI, used to calibrate the nonlinear response of the PMTs and VA electronics, necessarily flashes a long way into the the nonlinear regime. In these cases the near end readout saturates and the measured asymmetry is reduced as the opposite end readout 'catches up'. This effect is partially removed by the linearity calibration, but at sufficiently high pulse heights the VA is saturated and no calibration is possible. The solution is that the LI sieve does not require asymmetry in snarls where the *smaller* of $(\sum PH_{\mathrm{W}}, \sum PH_{\mathrm{E}})$ is just below the saturating region (1.7×10^6 sigcor), as this indicates the larger (near end read out) will be saturated. Figure 4.2 shows the distribution of $|A_{\mathrm{WE}}|$ for non-saturating snarls.

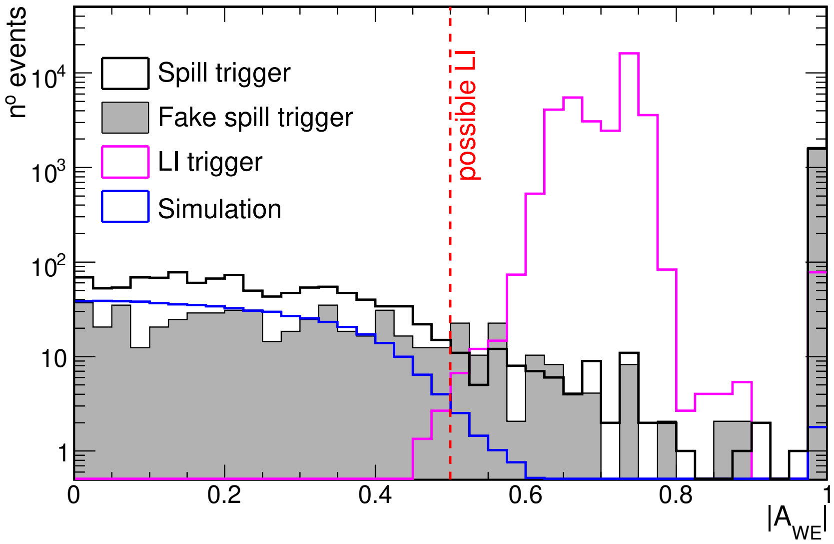

Figure 4.2: West-East asymmetry for events in spill and fake spill-trigger snarls (black), tPMT tagged snarls (magenta) and simulated neutrino events (blue). Values of $|A_{\mathrm{WE}}|$ above 0.5 are considered indicative of LI.

Highest pulser box fraction Because the Pulser boxes flash a contiguous range of planes in the detector, one way of identify LI would be to look for events that start and finish on particular pairs of planes. The problem with this approach is that it can fail if there is fibre noise on the immediately adjacent planes, or for low pulse height LI during gain curves. An alternative is to look for the number of hit planes in regions corresponding to each pulser box, as a fraction of the total number of planes illuminated by the pulser box. If the largest of these hit fractions is above 0.85 then this is considered indicative of LI.

Pulser box ratio Simply hitting a lot of the planes associated with a single pulser box is not enough to determine that an event is LI, since this will also include any physics event with a long muon track. LI events will also have a much lower fraction of planes hit in any other pulser box region, something which is not generally true of large physics events. We therefore also look at the double ratio of the second largest to largest fraction of hit planes. If this ratio is below 0.05 this is considered indicative of LI. The combination of pulser box fraction conditions are shown in Figure 4.3

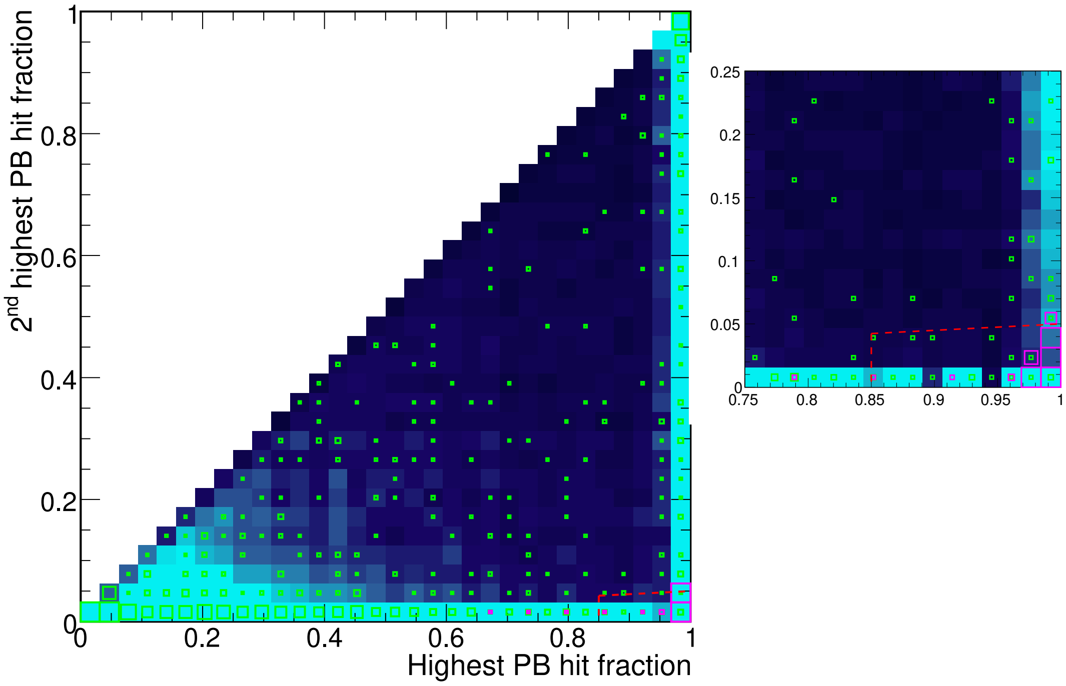

Figure 4.3: Highest and second highest fraction of planes hit in regions corresponding to a single pulser box. The neutrino beam simulation events are the cyan surface (lighter shades indicate more events), spill triggered snarls are green, and tPMT flagged LI snarls are magenta. Inset shows a zoom of the lower right corner.

Each condition individually would generate a lot of false positives, but the probability for a non-LI snarl containing to fail all conditions is very low.As a result, the effectiveness of the LISieve is very good, comparable to that of the tPMT cut. Since the inefficiencies should be uncorrelated with the those for the the tPMT, the LI snarls identified by each method should provide an unbiased sample for the evaluation of the other.

4.1.2 Event quality

Two cuts (a fiducial volume cut, and the 'main event' cut) are made to avoid systematic errors resulting from mis-modelling of events in the simulation. As they also have an impact on detector backgrounds they are considered here. Subsequent cuts are shown with these cuts already made, although the fiducial volume was usually enlarged when choosing cut positions, in order to obtain a larger sample of events.

Fiducial cut

Since oscillations are an energy dependent phenomenon it is important that we can accurately model the unoscillated spectrum. For neutral current events this means correctly modelling the energy deposited in the detector by the shower. As the exact distribution of energy deposition in the hadronic shower is not as well modelled as the total energy distribution data/simulation agreement will be adversely affected by the inclusion of showers that spill into uninstrumented regions of the detector. In addition to this, calibration is hardest in the outside regions of the detector where the response varies most strongly with position. To avoid these problems a fiducial volume cut is made on the basis of the proximity of the event vertex to the edges of the detector. This also has the benefit of improving CC/NC discrimination (since a track is more likely to be seen).

Another important effect of the fiducial cut is that it significantly reduces the non-neutrino backgrounds. The cosmic ray background is of course suppressed as the muons enter from the edges of the detector, but there is also a substantial effect on the fibre noise background. This is because many fibre noise events are too small to be demultiplexed unambiguously, and the algorithm tends to put these at the edges of the detector.

The event vertex is our best estimate of the position of the neutrino interaction point. If a clear track is present this is generally well estimated by using the start point of the track as reported by the Kalman filter. If no clear track is present in a showering event, the track-finding algorithm may still find a track inside the shower, by picking arbitrary lines of hits in the U and V views. Such "drawn-on" tracks can have reconstructed vertices quite far from the true interaction point, so the event vertex is determined directly from the pattern of hits in the detector. To decide which vertex to use when making a fiducial cut, a quantity *track extension* is defined:

$$\delta_{\text{ts}} = \text{ planes in track} - \text{planes in shower.}$$

If the event lacks a track or shower (it cannot have neither) the relevant number of planes is set to zero. The fiducial cut is then made on the track vertex if $\delta_{\text{ts}} > 0$ planes, and on the directly determined event vertex otherwise.

Having defined the vertex to be used the cut can be made. In the transverse (x–y) plane it requires events to be further than 50 cm from the nearest edge of a scintillator plane, and more than 45 cm from the centre of the coil hole. Figure 4.4, shows the effect of the fiducial cut on reconstructed energy, and Figure 4.5 shows the effect on backgrounds.

Longitudinal cuts treat both supermodules identically, owing to the difficulty in modelling and reconstructing a shower that crosses the gap between supermodules. A minimum depth of four scintillator planes (0.2328 m) are required from the front of a supermodule, to veto events entering from the front face. From the rear of each supermodule, a depth of 1 m is required, to provide sufficient containment for the majority of showers. Figure 4.6 shows how the reconstructed shower energy is dependent on the vertex position near the end of the detector.

Snarl dominance ('Main event') cut

Another quantity that is potentially sensitive to the way hadronic showers are modelled is the number of events in a snarl. The beam neutrino event rate in the far detector is low enough that the probability of seeing two neutrino interactions in the same snarl is essentially nil. Therefore, if multiple events are reconstructed this must either signify a reconstruction failure or that at least one of the events is from non-neutrino background. Simply removing snarls that have multiple events reconstructed is likely to introduce discrepancies between data and simulation, since reconstruction failures often occur in 'unusual' events, that are the hardest to model. To mitigate against this events in two-event snarls are not automatically discarded. Rather an event can be kept if it dominates the snarl it is in, having greater than 75% of the total pulse height in the snarl. Figure 4.7 shows this cut graphically.

4.1.3 Removal of fibre noise

For fibre noise to be made up into an event the reconstruction must match hits in both U and V views, so at least two noise hits are required. Random clusters of single fibre noise become exponentially less probable with increasing numbers of strips in the cluster, so the noise events are typically rather small. To remove this type of noise we make use of two related characteristics: the low number of strips, and the low pulse height. The number of strips and pulse height are highly correlated over the full energy range of real neutrino events, but in the narrow range where noise events are a background they are effectively independent.

Events are accepted if they have 9 or more strips, or if they have a total pulse height of at least 5000 sigcor. Events are also accepted if they have 5 or more strips and a total pulse height of at least 2500 sigcor The overall cut can be seen in Figure 4.8. Using pulse height and number of strips removes less neutrino events than either alone and therefore makes the selection less susceptible to uncertainties in the exact pattern of energy deposition in the shower.

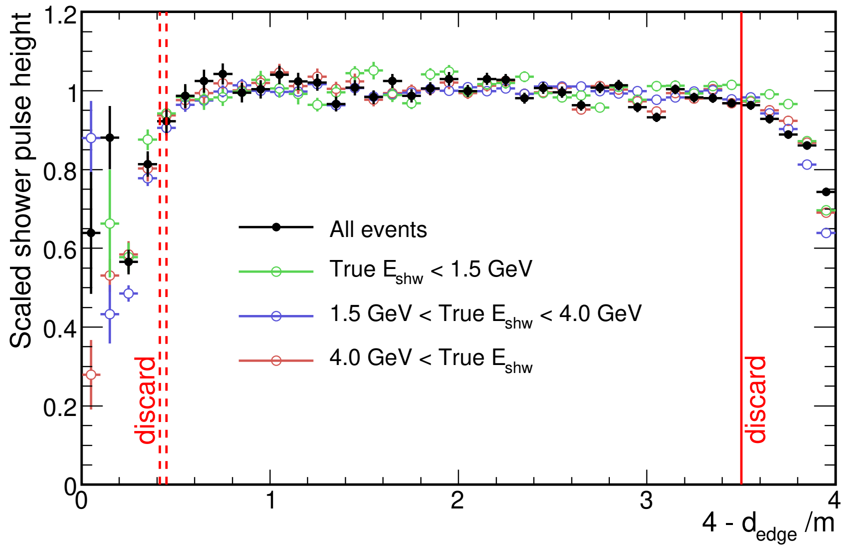

Figure 4.4: Mean pulse height of showers as a function of the reconstructed vertex position, for various ranges of true shower energy. Each is scaled so that the average pulse height for $1\,\mathrm{m} < d_{\mathrm{edge}} < 3\,\mathrm{m}$ is one. Also shown are the fiducial cuts in the x–y plane. Longitudinal fiducial cuts are already made.

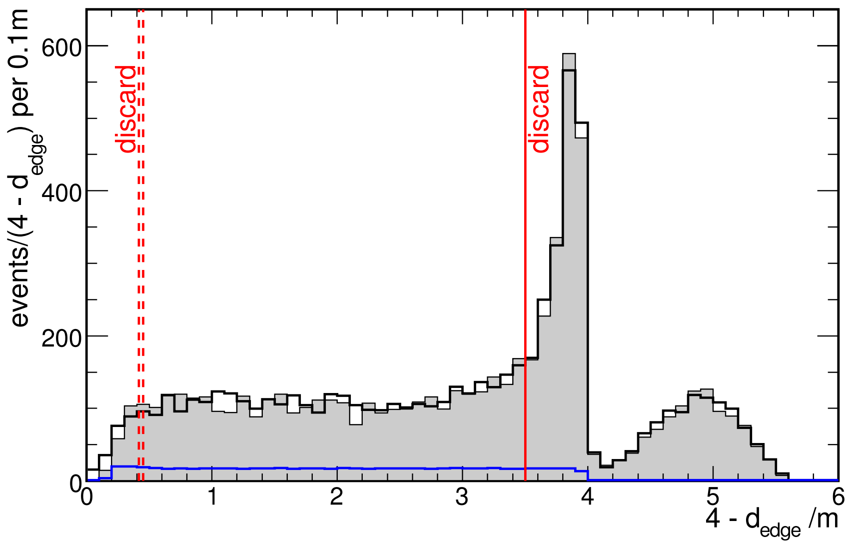

Figure 4.5: Comparison of event vertex distributions in spill triggered data (black), fake spill-trigger data (grey fill) and beam neutrino simulation (blue). Fiducial cuts in the x–y plane are also shown.

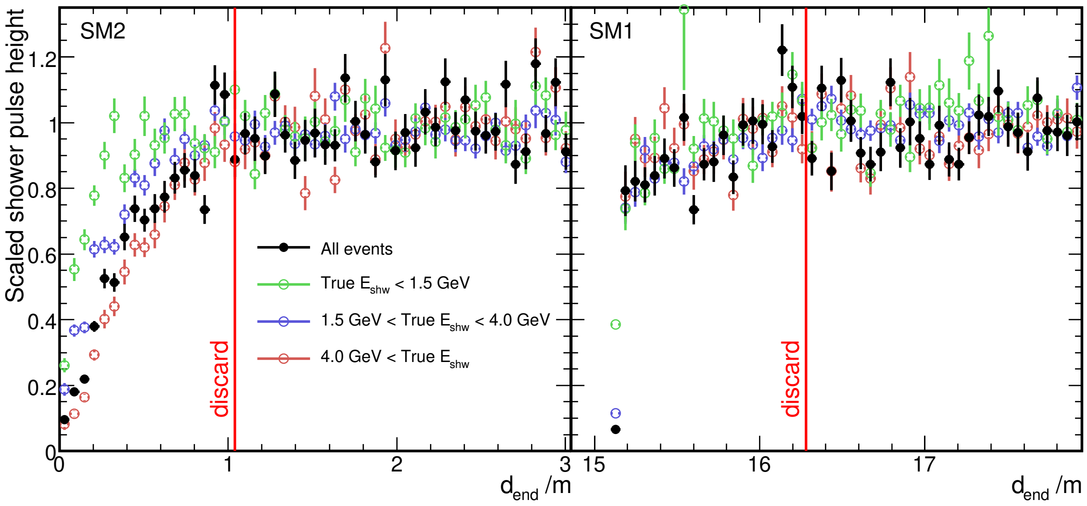

Figure 4.6: Average pulse height of showers as a function of the reconstructed vertex position, for various ranges of true shower energy. Each is scaled so that the average pulse height for $3\,\mathrm{m} < d_{\mathrm{end}} < 12\,\mathrm{m}$ is one. The fiducial cut shown is derived from consideration of SM2 (left panel), but applied identically to SM1 (right panel). Transverse fiducial cuts are already made.

4.1.4 Removal of cosmic ray muons

The reconstruction of events in time with a beam spill[2] assumes that events in the detector come from beam neutrinos, for example by expecting muons to be travelling close to the beam direction. This improves the reconstruction of beam neutrino events but tends to cause failures when reconstructing cosmic ray muons. Even so most cosmic rays are reasonably well reconstructed, and most are removed by a fiducial cut. The ones that remain tend to be the less well reconstructed ones. These seem to fall into a few categories, dependent on the angle the muon makes with the z axis:

If the cosmic track is less than about 70° to the z axis it will normally be well reconstructed, and most are removed by a fiducial cut. Those that remain can be removed by a cut on track angle. The variable used is the absolute value of the z direction cosine, $|p_z|/E$, of the muon (calculated at the vertex). If this is less than 0.4 then the event is rejected, as shown in Figure 4.9.

Cosmics that are at high angles to the z axis (referred to as steep because the are usually close to the vertical) are sometimes overlooked by the track finding algorithms of the reconstruction, which is expecting tracks to be close to the z axis. These events are then formed into showers. Subsequent reconstruction of

[2]there is a separate reconstruction for atmospheric events

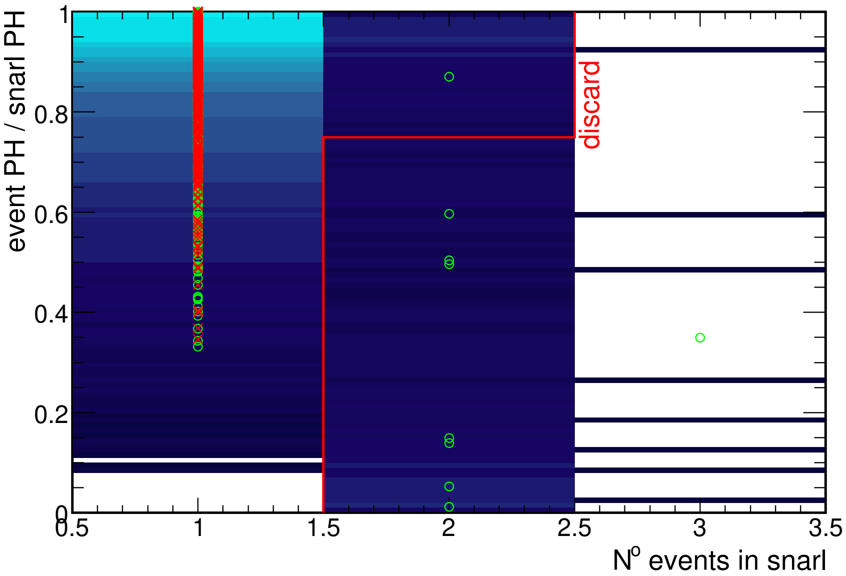

Figure 4.7: Fraction of snarl pulse height included in an event against number of events in the snarl. Beam neutrino simulation is the cyan surface, spill trigger data is the green circles, and fake spill-triggers are red crosses.

such events is often poor because the unusual shower shape, with the result that many of these events are not removed by a fiducial cut. These showers tend to be very wide in comparison to their length, which provides a means to remove them. A variable *transverse RMS* is defined by calculating the RMS of strip positions separately in U and V planes, then adding the two values in quadrature. This is plotted as a function of the number of planes in the shower. It is found that for beam neutrino induced showers the transverse RMS rises approximately linearly with the logarithm of the length. A linear fit is made to the neutrino showers to obtain the gradient of the cut (0.1901 m),and the cut is made parallel to this, with an offset (of 0.3 m) determined from the one dimensional distribution of transverse RMS $-$ 0.1901 $\log_{10}$[planes in shower]. Figure 4.10 shows the cut.

If the cosmic ray muon passes even more steeply through the detector, at only a few degrees to the x–y plane, it will hit very few planes. Because the strips are multiplexed this can have the side effect of suppressing the transverse width of the showers. But these events are highly distinctive, since they have very few planes (only one or two in each view) but many strips in each view. An alternative width:length ratio cut is to use strips per plane in the event as a width-like measure, and number of planes in the event as a measure of length. The ratio of the two is strips / planes^2 and the cut is shown in Figure 4.11

Finally, there is a specific reconstruction failure associated with cosmic ray muons that stop in the detector. The reconstruction assumes particles are heading downstream, so if the muon enters the detector heading upstream the track is

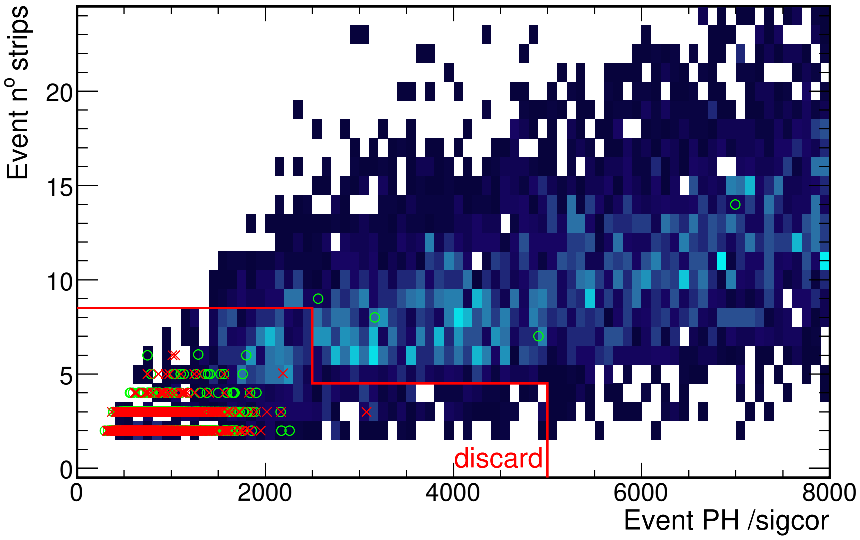

Figure 4.8: Number of strips in an event against reconstructed event pulse height in sigcor. Events from the beam neutrino simulation are the cyan surface (lighter shades are higher), while spill and fake spill-trigger events are green and red points respectively.

reconstructed backwards, and the end of the stopping muon track is recorded as the vertex. These events end up looking like charged current events with small vertex showers. To identify these events it is necessary to spot that the track is reconstructed backwards, which is achieved by making a linear fit of the time of strips on the track as a function of their z position. If a track is suspected of being reconstructed backwards, and it exits the top, sides or downstream end of the detector then it is cut out if its (corrected) y direction cosine is less than -0.04, implying that it is a downward going cosmic ray.

4.1.5 Event timing

One of the most effective cuts that can be made is to only accept events that occur in coincidence with the beam. Without the ability to synchronise both detectors to a global time provided by their respective GPS timing systems this would be impossible since the only indication that neutrinos are passing through the detector would be the events themselves. During reconstruction the recorded observation times of neutrino events can be related to the time of the signals from the Main Injector that identify when beam has been sent to NuMI. There is an offset, firstly because it takes around 2.5 ms for neutrinos to travel between Fermilab and Soudan, and secondly because of the finite time it takes for the various electronics systems to operate, and for events to be recorded. The offset is systematic, however, and has bee calculated in [69]. Without even knowing the

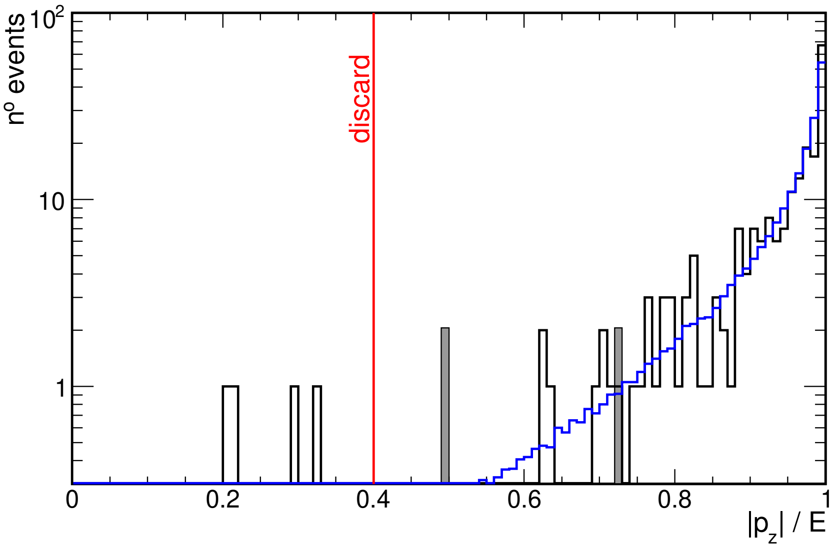

Figure 4.9: Absolute value of z direction cosine of tracks, for spill (black) and fake spill-trigger (grey fill) snarls and neutrino simulation (blue). All other preselection cuts have been applied, and the simulation is normalised to have the same number of events as the spill data after all preselection cuts.

size of this offset, it is possible to look for a peak in the time difference between when an event is seen, and when the neutrino beam was thought to be passing through the detector. All real events should be in this peak, so we can discard events that lie outside of it to gain a significant reduction in background rates. A plot showing the timing distribution is shown in Figure 4.12. The predicted arrival time of the beam and its actual arrival time are in good agreement at the microsecond level. The width observed is due to the duration of the beam pulse, and cuts are made at $-2\,\mu s$ and $+12\,\mu s$.

4.2 Evaluation of systematic errors

The cuts to remove backgrounds will unfortunately remove some fraction of real beam neutrino events. Table 4.1 shows the efficiencies for each cut, together with the cumulative efficiency for all cuts in order. The only cut that removes a significant fraction of simulated events is the fiducial cut. As part of the rational of imposing a fiducial volume is to remove badly reconstructed events this is inevitable. What is more problematic is the systematic error associated with these cuts. There are two origins for these systematic errors: backgrounds events remaining in the data sample, and uncertainty on how many beam neutrino events are removed.

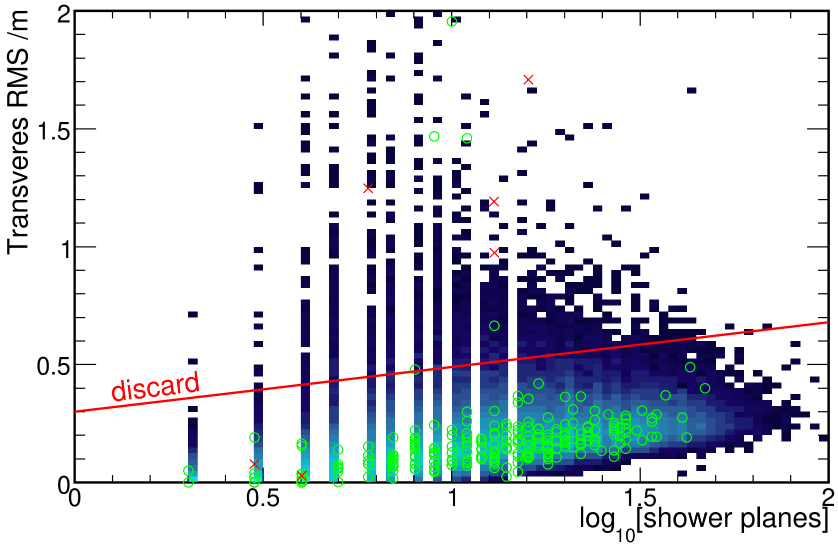

Figure 4.10: A two dimensional plot of the shower transverse RMS against $\log_{10}$[planes in shower]. Beam neutrino events are the cyan surface, spill triggered snarls are green, and fake spill-trigger snarls are red.

4.2.1 Uncertainties from LI removal

The light injection cuts are a special case as they remove snarls, not events. The two cuts used are the tPMT cut and the LISieve. As they are independent, and both are very good, each cut provides a sample to gauge the effectiveness of the other. The tPMT fails to identify around one in 10^4 LI snarls, while the LISieve is not quite as good, missing around five in 10^4 snarls. Applying both will therefore miss something in the region of one in 10^7 LI snarls, corresponding to around a thousand years of data taking, with the current frequency of LI calibration. Unless there is a unnoticed hardware problem, the background from LI should be

Cut	Events remaining	Efficiency this cut	total
After LI removal	181473	—	100.0%
Fiducial volume	124262	68.5%	68.5%
Main event	122805	98.8%	67.7%
Noise	122088	99.4%	67.3%
Cosmic ray	118215	96.8%	65.1%

Table 4.1: Efficiency (fraction of beam neutrino events retained, as calculated from simulated interactions in the detector volume) of background removal cuts when applied in sequence.

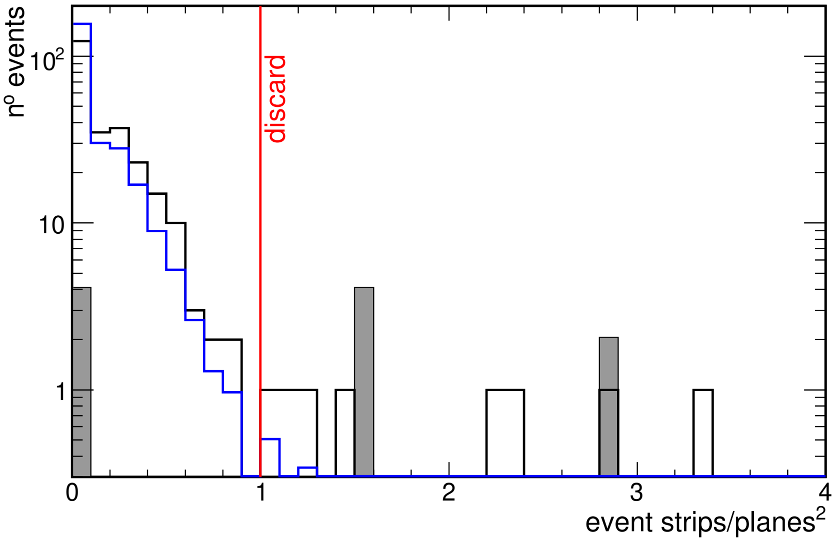

Figure 4.11: Number of strips per planes-squared, for spill (black) and fake spill-trigger snarls (grey fill) and beam neutrino simulation (blue). Events at 1 or above are removed. All other preselection cuts have been applied, and the simulation is normalised to the same number of events as the spill data after all preselection cuts.

negligible.

In terms of real events removed by these cuts there is no error induced by the tPMT cut as it simply acts as a veto, and the livetime of the detector can be corrected for this. The LISieve can potentially affect the spectrum of events, but when applied to a sample of 190135 simulated snarls, none were identified as LI, so this error is also small.

4.2.2 Uncertainty on the detector backgrounds

To estimate how many background events are left in the data sample, we can use the events recorded by the fake spill-trigger. As the event timing is found to be reliable, we can also use the out-of-time events to improve this estimate.[3] The background estimate is then the total number of background events in these samples scaled by the total exposure. Table 4.2 shows the calculation of the expected contamination from background events, which is 0.1%. The inclusion of background events is therefore not a significant source of systematic error.

[3]In principle we must correct for the fact the chopper selects only a single (the largest) chop from the data. In practice, the fraction of spill triggers containing a beam neutrino event is around 10^{-4}, so the correction is negligible.

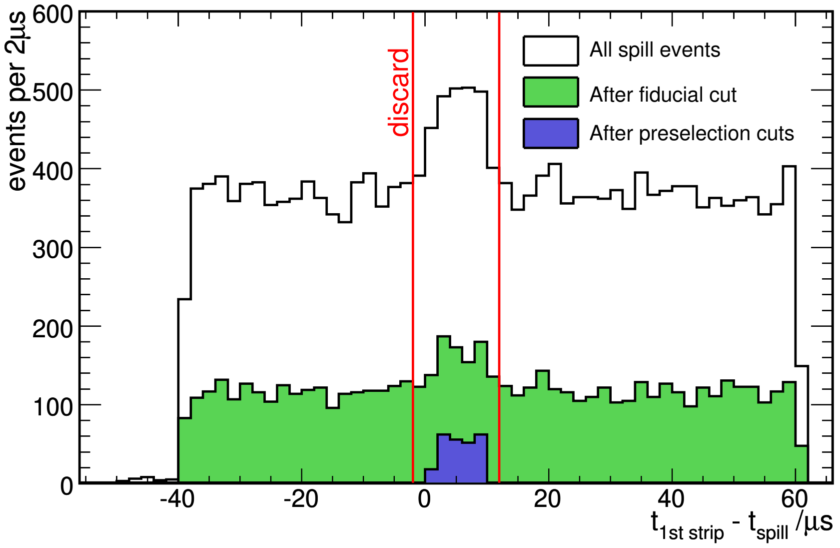

Figure 4.12: Difference between spill arrival times and event times. After the fiducial cut, almost all of the remaining background is removed by the noise cut. The events that remain all lie between 0 and 10 µs

Beam induced backgrounds

One form of 'background' cannot be estimated using out of spill data. These are beam induced backgrounds, namely rock events that leak into the detector. As the simulated events used for comparison are all contained vertex events this can potentially lead to a systematic underestimate of the number of events seen in data. Most such events will be removed by the fiducial cut, but neutrons from rock interactions will still get through. As neutrons will often be accompanied by a rock muon or other shower particles they can be vetoed by the main event cut. In addition, they carry only part of the energy of the shower that created them so

Trigger condition	Number of triggers	Window / µs	Exposure / s	Number of events
Fake spill	1464005	100	146.4	2
Spill (out of time)	3017209	86	259.5	0
Total Background			405.9	2
Spill (in time)	3017209	14	42.2	246
Estimated background				0.208

Table 4.2: Calculation of the expected number of background events after cleaning cuts. Also shown is the number of events in the signal sample. Numbers are taken from the open dataset.

they will tend to result in low energy events.

To get an estimate of the rate of unaccompanied rock neutrons we can look at neutrino interactions in the supermodule one that result in events in the fiducial volume of supermodule two. There are around 100 such events in the set of events used in Table 4.1. Of these, most are accompanied by another reconstructed 'event' originating from the same interaction and entering the front face of supermodule two, leaving around 30 unaccompanied neutrons, almost all of which are in the first metre of the fiducial volume. This number can be used as an estimate of the number of unaccompanied rock neutrons entering the front of supermodule one, and should be compared to the number of contained vertex events passing the fiducial, main event and noise cuts (122088). The rock face in front of the detector has a larger area than the back of supermodule one however, and neutrons can also enter from the detector sides. A rough guess is a factor of two for the area of the rock face and a further factor of two for including side-entering neutrons. This suggests something in the region of one unaccompanied rock neutron for every 1000 contained events, with large errors. Differences between data and simulation may raise or lower this number, but with the data set used it is just about small enough to be ignored.

4.2.3 Uncertainty on cut efficiencies

A more substantial problem is the uncertainty on the number of events removed. In general this arises because the simulation of events in the detector is not perfect. The main concern relates to the modelling of hadronic activity, which is difficult for a number of reasons. Firstly, the underlying neutrino-nucleon interaction straddles the transition where elastic scattering, resonance and DIS processes are all important, making the underlying physics complicated. Secondly, the nucleon in the target is (usually) part of an Iron nucleus, and secondary interactions can absorb or redistribute the energy of the hadronic system. [107] Finally, hadronic shower development is intrinsically difficult to model correctly because of the many processes that can take place, and the non-trivial dependence on the calorimeter material. [108] Many hadronic shower models focus on slightly higher energies and on more coarsely grained readout than used in MINOS, [97] so the detailed distribution of energy in the shower is less certain than the overall scale. This can lead to mis-modelling of quantities such as the dimensions and number of strips hit in a shower. Although the Near Detector can be used to check this modelling, differences between detectors such as the multiplexing and different PMT cross-talk patterns still leaves some room for error. Shower mis-modelling also contributes to other reconstruction differences between data and simulation, such as the probability that a track is found if no clear muon is present

Aside from the uncertainty associated with the modelling of neutrino interactions, there are the known discrepancies in the noise that is superimposed on events. This problem is substantially reduced by the use of the chopper in reconstruction but is still present at some level.

Any systematic errors arising because of these background removal cuts may be particularly problematic for analyses because the cuts are only applied at the Far Detector. Most systematic uncertainties affecting MINOS analyses can be substantially reduced by using simulated events and data from both detectors to achieve (partial) cancellation of errors. Since many background removal cuts are designed to remove pathological side-effects of the detector read out, the systematic errors induced would be different between detectors even if identical cuts were applied in both, and in some cases this is not even possible.

Quantifying the level of systematic uncertainty is also difficult. In the Near Detector there is plenty of data, and it is possible to make detailed comparisons of the level of agreement between data and simulation. In the Far Detector there are only a few hundred events, enough to compare the total number of events to the level of a few percent. Making detailed 'shape' comparisons of measured quantities requires substantially more events even for this level of accuracy. Another complication is the presence of blinding and oscillations (of the ν_μ charge current background) in the Far Detector data. Blinding presents a chicken-and-egg problem. It can of course be removed once the analysis is finalised, but to be finalised, the errors must first be estimated. Oscillations, on the other hand, could be applied in the simulation, but to do so would require knowledge of the true values. This hints at the philosophical problem with the whole approach of using data/simulation agreement to quantify systematic error: any systematic error that could affect the result must by definition be (partially) attributable to the effect one wishes to study. For all of the above reasons, the Far Detector data is of limited use for estimating systematic errors, so other approaches must be found.

Fiducial cut

Not using a fiducial cut can introduce significant systematic uncertainties, as described above, but it introduces others. The most problematic error results from the interaction between the demultiplexing algorithm and the fiducial cut. As mentioned previously, if the event is too small to be unambiguously demultiplexed there is a tendency for the algorithm to place hits at the edge of the detector, outside of the fiducial volume. This can result in a systematic loss of events at low energy. Figure 4.13 shows the effect, after all cuts have been applied apart from the x–y cut. The plot shows (using simulated data) the the fraction f^{out} of neutral current events that lie less than 0.5 m from the the detector edge, as a function of the 'true visible' energy of the shower.[4] This should in principle be determined entirely by the geometry of the detector, and the prediction of the geometric calcu-

[4]The reconstructed energy cannot be used, as it is systematically lower for events near the edge of the detector. The true visible energy is an estimate of the energy visible to the detector, obtained by summing the energy of final state particles, and subtracting the rest masses of any baryons.

lation is shown by the red dotted line. If the true interaction point of the neutrino is used, the result is the blue points. These points are consistent with the geometrical prediction for most event energies, but are systematically low for low energy events. This is attributable to the fact that small showers that 'lose' energy into the uninstrumented part of the detector may end up below the threshold for reconstruction. However, if we look at the distance of the *reconstructed* vertex of the event to the edge (the black points) the dip is not as marked, indicating an excess of low energy events reconstructed in the outer regions of the detector. To demonstrate that demultiplexing really has this effect, the same fraction can be plotted using only neutrino interactions that occurred more than 1.5 m from the detector edges (green points). Errors on the vertex position cannot be larger than the maximum width of the shower, so should be much less than one meter for all except the largest showers. On the other hand, demultiplexing errors should typically move events at least a meter. Demultiplexing errors should also be more common for small showers, as there is less information for the algorithm to make use of. The fact that this fraction rises furthest from zero at the lowest true shower energies demonstrates that this is a demultiplexer effect. There

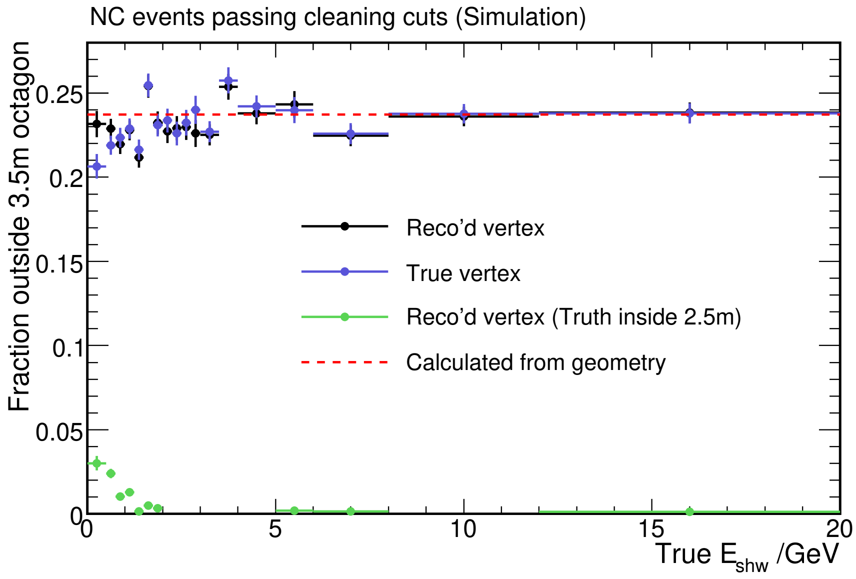

Figure 4.13: Fraction of NC events outside transverse fiducial region, as a function of hadronic energy.

are noticeable differences in the fraction f^{out} (based on using the true interaction point and reconstructed vertex) for the two lowest energy bins. Below 0.5 GeV the difference is 0.025, and between 0.5 and 0.75 GeV the difference is 0.010. From these differences we can construct the quantity

$$\left(\frac{\delta n}{n}\right)^{\text{fid}} = \frac{-\left(f^{\text{out}}_{\text{true}} - f^{\text{out}}_{\text{reco}}\right)}{1 - f^{\text{out}}_{\text{reco}}},$$

which is the fractional increase in the number of events that would pass the fiducial cut if the true interaction points were used.

The quantity $(\delta n/n)^{\mathrm{fid}}$ gives a measure of how much the demuxer / fiducial volume interaction affects the number of events selected, but if the data is well described by the simulation there is no problem. Table 4.3 shows the calculated values of $(\delta n/n)^{\mathrm{fid}}$, for the energy ranges where the effect is significant, for the current analysis and for the previous version of the simulation which used a different hadron shower model. The difference between these numbers is quite significant, and it is difficult to believe the agreement with data would be better. As there are additional effects that may degrade data / simulation agreement further, the full size of the effect (in the newer Simulation) is used as a systematic uncertainty, as shown in Table 4.3. It should be noted that the systematic errors are applied to bins of reconstructed shower energy, while the errors are calculated in terms of the true visible energy. Although the correspondence between true visible energy and reconstructed shower energy is not one-to-one, it is difficult to estimate as there is a strong bias toward losing events with low reconstructed energy. For the purposes of estimating the systematic uncertainty we therefore assume a direct correspondence between reconstructed shower energy and true visible energy.

Energy range	$(\delta n/n)^{\mathrm{fid}}$	$(\delta n/n)^{\mathrm{fid}}$(old)	$(\sigma\mathcal{E}/\mathcal{E})^{\mathrm{fid}}$
$0. \sim 0.5\,\mathrm{GeV}$	0.033	0.016	$\pm 3.3\%$
$0.5 \sim 0.75\,\mathrm{GeV}$	0.013	0.011	$\pm 1.3\%$

Table 4.3: Systematic uncertainty on the number of events selected, resulting from imposing a fiducial volume cut after demultiplexing.

Main Event cut

Although it is suspected that the fraction of neutrino interactions that are split into multiple events is not perfectly described in data, the fraction of events lost in the beam neutrino simulation is only 1.2%. This corresponds to a loss of around 0.6% of the neutrinos interacting in the fiducial volume. Low energy events are less likely to be split, but apart from this the events removed by the Main Event cut show no strong energy bias. Even assuming a 100% error on the number of neutrino interactions removed by the Main Event cut, the impact on both the rate and spectrum of events is negligible.

Noise cut

Although only a small fraction of events are removed by the noise cut they are likely to be concentrated in the lowest energy ranges, so a systematic uncertainty on the cut efficiency will cause distortions of the energy spectrum. To estimate the systematic uncertainty on the number of events removed, the cut was adjusted, corresponding to adding or removing a single strip of 355 sigcor. This is the mean energy of a strip in a noise event, and coincidentally, roughly the pulse height deposited in a strip by a minimum ionising particle. Figure 4.14 shows the fraction of events surviving the noise cut in the nominal and adjusted cases. A study of a similar cut used the Near Detector [91] suggested that the difference between data and simulation was about half of the effect of adding a single strip, and it is assumed that this is appropriate in the Far Detector also. The estimated systematic uncertainties (as a fraction of the number of events retained) are ±7% below 0.5 GeV and ±1% between 0.5 and 0.75 GeV.

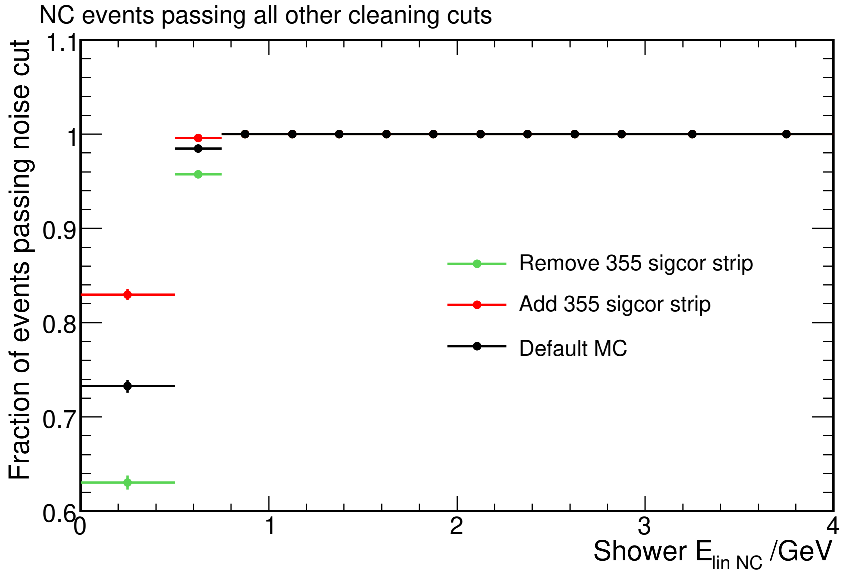

Figure 4.14: Fraction of simulated events passing the the noise cut as a function of reconstructed energy. The nominal cut is shown in black, with cuts equivalent to adding (removing) a single strip shown in green (red).

Cosmic ray cuts

The most significant loss of events after the fiducial cut comes from the four cuts used to remove cosmic rays. In combination, these cuts remove 3.2% of simulated events. The effect, as a function of reconstructed energy is shown in Figure 4.15. A fit is performed with the simple parametrization

$$\mathcal{E}^{\text{CR}}(E) = 1 - A\mathrm{e}^{-E/Q}, \tag{4.2}$$

and the values obtained for the parameters are $A = 0.0514$, $Q = 6.96\,\text{GeV}$. When the same cuts are applied to the previous version of the simulated data set, the fraction of events lost overall was 2.7%, a reduction of around one sixth. The

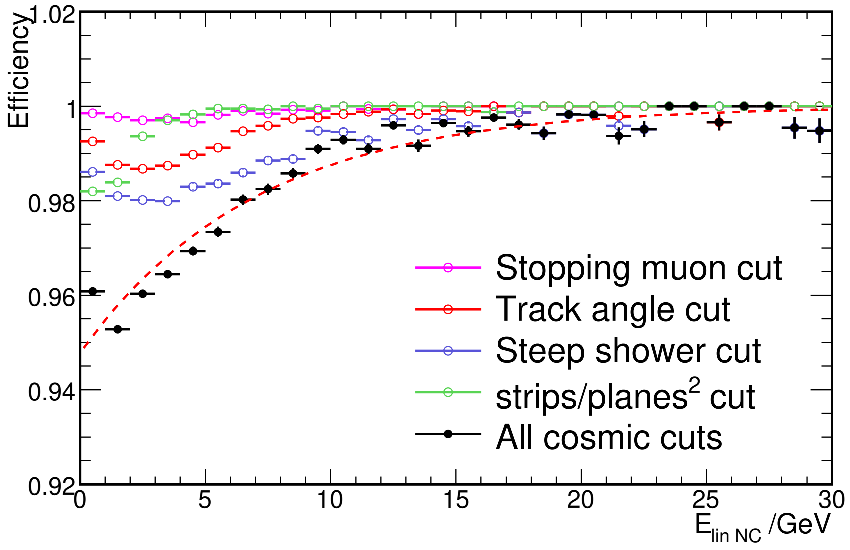

Figure 4.15: Fraction of simulated events passing all cosmic ray cuts as a function of reconstructed energy (black) and each cut individually (colours). Note that an event may fail more than one cut. The red dotted line shows an exponential fit to the total fraction lost.

difference between simulation and data is estimated as twice the difference between the two sets of simulated data, and the the estimated uncertainty on the cut efficiency is $\pm 0.0161\mathrm{e}^{-E/6.96}$.

4.2.4 Timing cut

Unless there is a problem with the timing system, the timestamping of events is accurate to 200 ns. [69] The existence of a peak in the number of events between 0 and 10 µs in Figure 4.12 that persists through all the other cleaning cuts is also good evidence that there is no problem with the timing system. The event time is taken from the earliest strip time, so one possible way for an event to fail the timing cut is by the inclusion of noise hits from before the event. However, gaps between successive strips in the event can be no longer 150 ns, and the cut is placed 2 µs before the steep rise of the beam peak so at least 14 noise strips would be required to cause the event to fail the timing cut. The systematic uncertainty from using the timing cut is therefore negligible.

4.2.5 Summary of systematic uncertainties

There are three non-negligible systematic uncertainties arising from the cleaning cuts: from the demuxer / fiducial volume interaction, from the noise cut, and from the cosmic ray cuts. The first two uncertainties primarily affect small events and are grouped together as a systematic uncertainty on the efficiency to reconstruct events that (if reconstructed) have a reconstructed energy below 0.75 GeV:

$$\left(\frac{\sigma\mathcal{E}}{\mathcal{E}}\right)^{0.75} = \begin{cases} \pm 7.8\% & E_{\mathrm{shw}} < 0.5\,\mathrm{GeV} \\ \pm 1.5\% & 0.5 \leq E_{\mathrm{shw}} < 0.75\,\mathrm{GeV} \end{cases}. \tag{4.3}$$

The cosmic ray cuts have a less significant impact at low energies but affect the high energy events more so are treated as a separate systematic uncertainty on the efficiency:

$$\left(\frac{\sigma\mathcal{E}}{\mathcal{E}}\right)^{\mathrm{CR}} = \pm(0.0161\mathrm{e}^{-E_{\mathrm{shw}}/6.96}). \tag{4.4}$$

4.3 NC/CC discrimination

With a selection of neutrino-induced events the next step is to identify those events that are likely to originate from neutral current interactions. Although there are events from ν_e and ν_τ charged-current interactions in the data sample, they are not a significant fraction. For the purposes of a neutral current analysis it is only necessary to discriminate between ν_μ charged current and neutral current events.

The main discriminating characteristic is the existence of the muon track in ν_μ charged current interactions. But the presence of a reconstructed track is not sufficient to identify the event as charged current like (CC-like) as tracks are sometimes found that correspond to particles (usually pions) from the hadronic system. In addition to 'genuine' hadronic tracks, the track finding algorithm can also find tracks in dense showers by making linear associations of hits that don't correspond to a real particle.

Conversely, there are plenty of events without any track that are ν_μ charged current events. If the neutrino transfers nearly all of its energy to the hadronic system the muon will not travel far enough to protrude out of the shower, and leaves no obvious evidence of its existence.

A useful first step is to divide the events into three topological classes, as shown in Table 4.4. Two of the classes, track only and shower only events, can be approximately identified as CC-like and NC-like events, respectively. It may be possible to improve over these simple identifications, but clearly it is more important to discriminate between NC and CC events in the track and shower sample.

For the subset of events that have both tracks and showers, a powerful variable has already been introduced. This is δ_{ts}, which was introduced in Section 4.1.2. Figure 4.16 shows how δ_{ts} can be used to discriminate between CC and NC events. Taking events with $\delta_{\mathrm{ts}} \lesssim 10$ planes, or no reconstructed track would provide a sample of NC-like events.

Event topology	Fraction of NC events	Sample purity	with ν_μ disappearance
Track only	0.4%	0.04	0.07
Shower only	48.5%	0.73	0.82
Track & shower	51.1%	0.13	0.17

Table 4.4: Topological classes of events, showing what fraction of preselected NC events are in each classification, and the purity (fraction of events that are from NC interactions) of each classification. The purity is also shown with CC events weighted by $P(\nu_\mu \to \nu_\mu)$ with $\Delta m^2_{\mathrm{atm}} = 2.5 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2 2\theta_{23} = 1$.

4.3.1 Improving the NC selection

Although the track extension is a very good variable for identifying neutral current events it is worth looking for other variables to refine the selection. The inclusion of more variables will in general provide better performance at the cost of increasing the complexity of the selection algorithm and making it harder to judge how well the simulation describes the data, potentially increasing the systematic error. As it turned out, very few variables were found that had any additional discriminating power.

As well as simply making cuts on more variables, some improvement can be made by using multivariate methods (such as artificial neural networks [91]) that take into account correlations between different variables. Several people within MINOS are working on such approaches, so they are not discussed here.

This analysis instead makes use of two simple properties of the track, containment and curvature direction. Track containment simply refers to whether a track appears to exit, or stop inside the detector volume. Track curvature direction is whether the track appears to curve outward (defocusing) or inward (focusing) in the magnetic field of the detector, and therefore whether the tracked particle appears to have positive or negative charge.

The use of track containment is a straightforward refinement of the track extension variable. If the muon exits the detector then the reconstructed track can be much shorter than it would be otherwise. Events where the track exits the detector are therefore more likely to be classed as NC-like. But since the CC interaction rate is higher than the NC interaction rate, most of these events are actually CC interactions, as illustrated in Figure 4.17. It therefore makes sense to regard events with a track exiting the detector as CC-like. 'Exiting the detector' is defined as having the track endpoint within 17 cm of the sides or downstream end of the detector.

The more interesting additional variable is the track curvature direction. The reason this variable is interesting is that the ν_μ CC events that form the majority of the background produce negatively curving μ^-, while positively curving μ^+ come

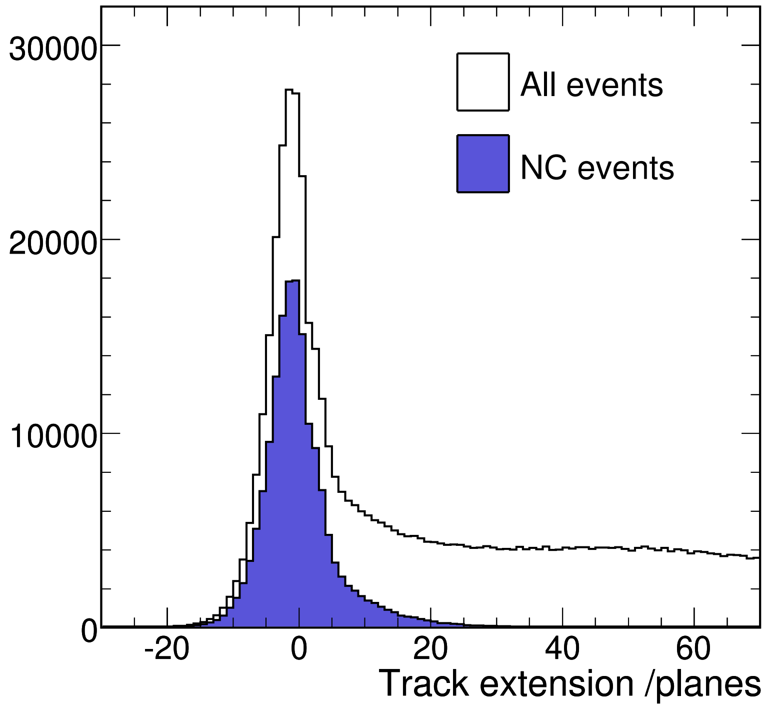

Figure 4.16: Track extension (δ_{ts}) of events with tracks and showers in the Far Detector. Events from neutral current interactions are shown in blue, and tend to have smaller values of δ_{ts}. CC events are weighted by the ν_μ survival probability, with $\Delta m^2_{\mathrm{atm}} = 2.5 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2 2\theta_{23} = 1$.

only from the small $\overline{\nu}_\mu$ component. Tracks in NC events, on the other hand, should be distributed evenly between positive and negative curvature. It is therefore possible to cut more loosely on the positively curving sample to obtain more NC events with a smaller increase in the background than if the cut was applied symmetrically to the positively and negatively curving samples. Figure 4.18 shows the track extension separately for positively and negatively curving tracks, along with the cut positions. Events with positively curving track are accepted as NC-like if $\delta_{\mathrm{ts}} < 20\,$planes. Events with negatively curving tracks are accepted as NC-like if they have $\delta_{\mathrm{ts}} < 8\,$planes.

The exact value of the cuts can be tuned to maximise the sensitivity to the disappearance of neutral current events. A possible cut tuning procedure is described in Appendix A. This was not done here as it requires the beam flux to be reweighted. The general principle of beam reweighting will be described in Section 5.2, but the correct beam weights for the neutral current analysis were

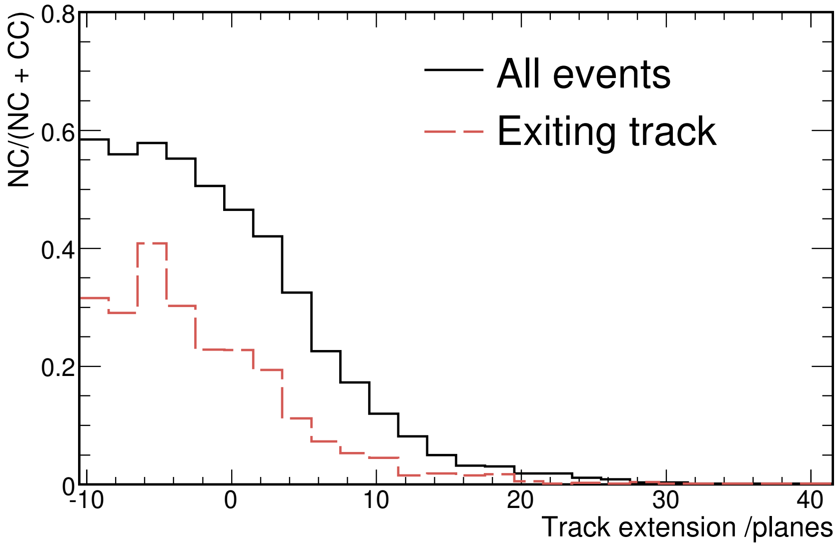

Figure 4.17: Fraction of events (for a given value of δ_{ts}) that are from NC interactions. CC events are weighted by $P(\nu_\mu \to \nu_\mu)$. The black solid line uses all events in the fiducial volume, and the red dashed line uses only events where the track appears to exit the detector.

not available[5] when the cuts were decided, so reasonable values were chosen to allow work to progress.

Systematic errors

It is important to check that the simulation describes the cut variables reasonably well, otherwise the number of signal and background events identified as NC-like in data and simulation will be different. To check this not only requires the correct beam weight, but also for the effect of oscillations to be taken into account. The parameters governing the ν_μ survival probability are reasonably well constrained, so input values of $\sin^2 2\theta_{23} = 1$, and $\Delta m^2_{\mathrm{atm}} = 2.5 \times 10^{-3}\,\mathrm{eV}^2$ can be taken as representative. Measuring the value of f_{s} is the aim of this analysis so anything in the range $0 \leq f_{\mathrm{s}} \leq 1$ should be considered. The agreement between simulation and data in the Far Detector is shown in Figure 4.19. The agreement is reasonably good and there is no cause for concern.

Another approach is to consider ways in which the simulation may not describe the real detector. A known deficiency is that the detector has small misalignments between planes which are accounted for in reconstruction whereas the simulation

[5]Because of the feed down from higher neutrino energies, weights must be calculated for neutrino energies up to 120 GeV, while the analysis in Section 5.2 ignores neutrinos with energies over 30 GeV. Plots sensitive to the reweighting were remade after it became available, so comparisons between data and simulation are valid.

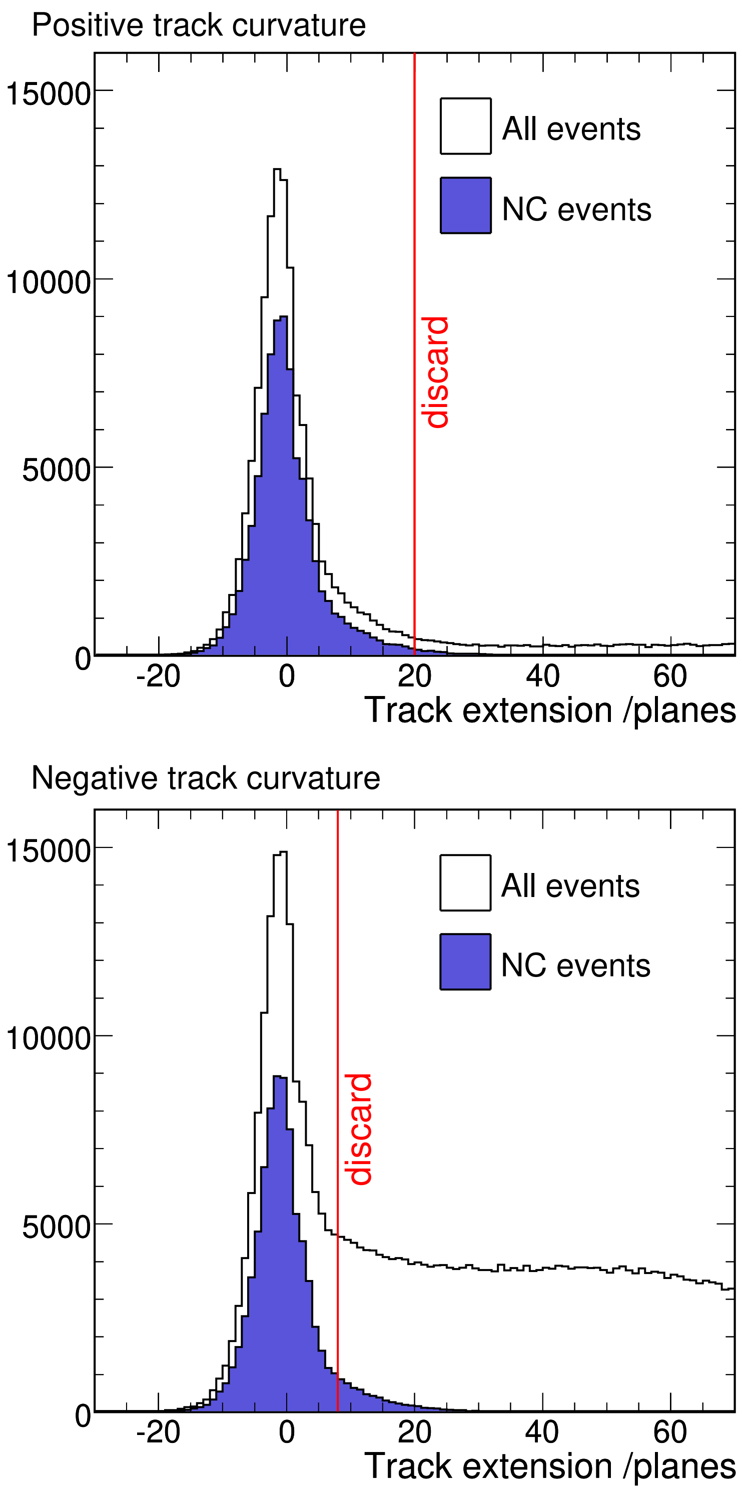

Figure 4.18: Track extension for positive and negative curvature tracks. CC events are weighted by $P(\nu_\mu \to \nu_\mu)$. If $\delta_{\rm ts}$ is less than the indicated values the event is regarded as NC-like.

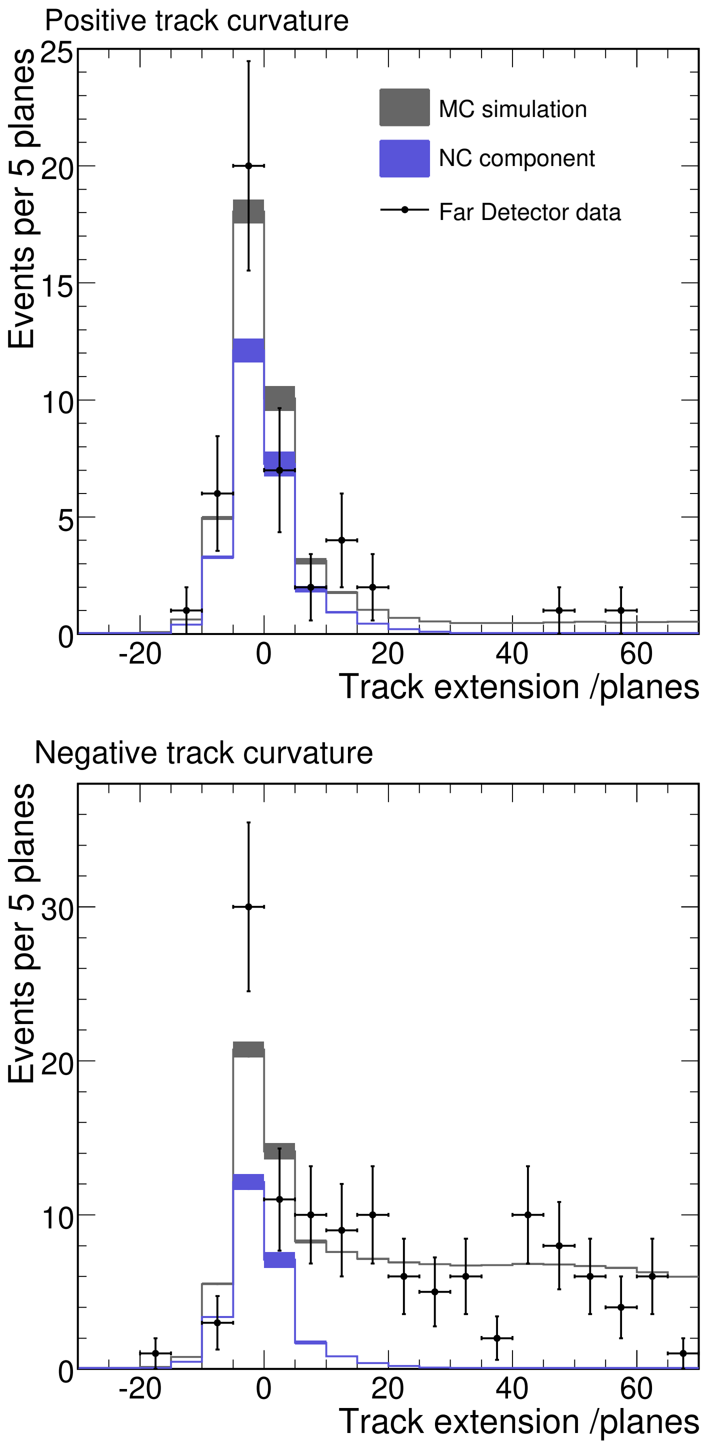

Figure 4.19: Comparison of track extension in data and simulation. Oscillation weights are applied as described in the text, with the allowed range of f_s represented by the width of the simulation bands.

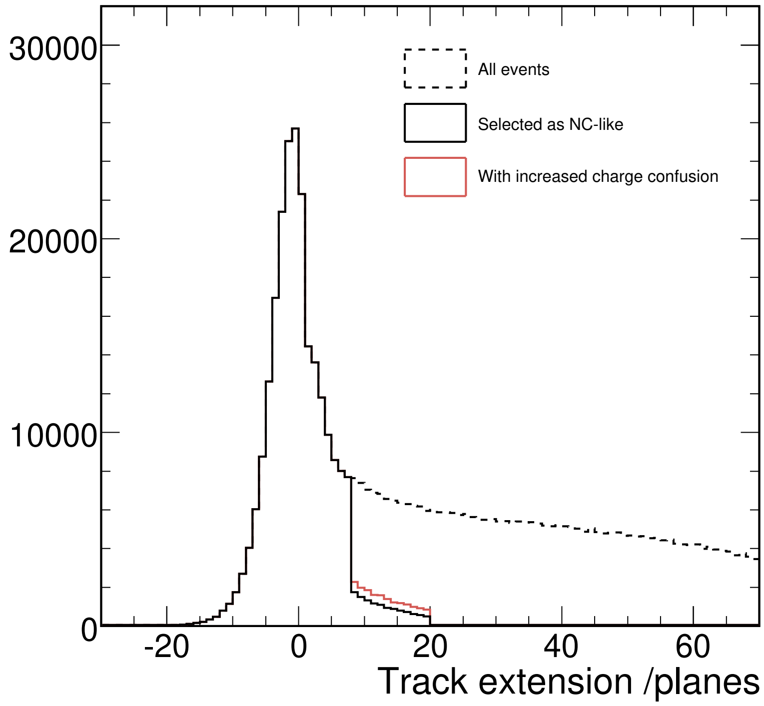

Figure 4.20: Effect (in simulation) on the NC-like selection of increasing the rate of negatively curving tracks identified as positive and vice versa.

has perfect alignment. This can lead to differences in the accuracy of the track fitter and therefore a systematic difference in the rate of miss-identified charge.

The fitter is sensitive to q/p, where q is the electric charge on the tracked particle, and $p = |\mathbf{p}|$ is the magnitude of its momentum. It reports an error $\sigma_{q/p}$ on the measured value of q/p. If $\left|\sigma_{q/p}/(q/p)\right| = 1$ then the error is as big as the measured value and the chance of having an incorrect direction of curvature is not small. To get a handle on the size of the effect, the simulated events are tweaked so their reconstructed curvature direction is flipped if $\left|\sigma_{q/p}/(q/p)\right| > 0.2$. As shown in Figure 4.20, this affects events with $8 \leq \delta_{\mathrm{ts}} < 20$, which are only selected if positively curving, and approximately doubles the charge current background in this range. This is only a small fraction of the total number of events selected as NC-like however, and corresponds to an increase in the ν_μ charged current background of $1 \sim 2\%$.

Chapter 5

Predicting the Far Detector Spectrum

A central part of the MINOS experiment is the use of Near Detector data to predict the Far Detector spectrum in the absence of neutrino oscillations. This chapter compares four different methods developed by the collaboration for doing this, in the context of the ν_μ CC analysis described in [57]. The aim of this analysis is to measure the mixing parameters $|\Delta m^2_{\mathrm{atm}}|$ and $\sin^2 2\theta_{23}$, via the ν_μ survival probability given in Equation (2.42). Neutral current events are the primary background for this analysis, and an approach based on probability density functions (PDFs) is used to select CC-like events. The method is covered briefly here for convenience.

5.1 Selection of CC events

The initial selection uses cuts to remove both the detector backgrounds and the more obviously NC-like events. At least one well-reconstructed track is required, which must pass basic quality cuts and have a vertex within the detector fiducial volume. In the Near Detector this is a cylinder of radius 1 m and length 4 m, centred on the beam and starting 1 m from the from the front face of the detector. In the Far Detector the volume is also cylindrical, with a radius of 3.7 m. The track vertex must be at least 50 cm from the ends of the supermodules, and at least 2 m from the downstream end of the second supermodule (to allow the momentum to be determined).

Removal of non-neutrino backgrounds is much easier than for the neutral current analysis. In the Near Detector there is no significant background that mimics CC-like events. In the Far Detector the relative rate of neutrino events to cosmic rays is lower and two additional cuts are used: a cut on track angle, and a loose (50 µs) timing cut, to remove cosmic ray muons. The fact that a good track is required effectively eliminates backgrounds from detector noise.

Because the flux of parent π^- and $K^\pm$ are not as well understood as π^+, the analysis uses only negatively curving tracks (to eliminate antineutrinos from π^- decays), and also ignores events with a reconstructed neutrino energy of greater than 30 GeV (which are more likely to come from $K^\pm$)

The resulting sample still contains many neutral current events. A PDF-based likelihood is formed to remove these. Three variables are used:

Event length in planes. Long tracks are more likely to be muons.

Track pulse height fraction. The ratio of the track pulse height to the event pulse height (in SigCor). This tends to be small for neutral current events.

Average pulse height per plane of the track. This tends to be larger for tracks in neutral current events.

The distributions of these variables are plotted separately for simulated CC and NC events. Each distribution $p_i(x_i \mid \mathrm{CC})$ and $p_i(x_i \mid \mathrm{NC})$ is normalised to unit area, and gives the probability that a charged (neutral) current event would have value x_i for variable i. A combined probability is formed (ignoring correlations) from the product:

$$P\left(\mathbf{x} \mid \mathrm{CC}, \mathrm{NC}\right) = \prod_{i=1}^{3} p_i\left(x_i \mid \mathrm{CC}, \mathrm{NC}\right). \tag{5.1}$$

These two probabilities are used to make a single ID parameter:

$$S_{\mathrm{ID}} = \sqrt{-\log P\left(\mathbf{x} \mid \mathrm{NC}\right)} - \sqrt{-\log P\left(\mathbf{x} \mid \mathrm{CC}\right)}, \tag{5.2}$$

which is larger for CC-like events and smaller (more negative) for NC-like events. The distribution of S_{ID}, and the variables used in the PDFs, are shown in Figure 5.1

To select a CC-like sample, events were retained if they had $S_{\mathrm{ID}} > -0.1$ in the Near Detector, or $S_{\mathrm{ID}} > -0.2$ in the Far Detector. Events classified as NC-like were not used in this analysis.

5.2 Tuning the simulation

It is to be noticed in Figure 5.1 that the Monte Carlo simulation used is referred to as “Tuned MC”. This refers to the tuning of the beam simulation. The reason for doing this can be seen by looking reconstructed energy spectrum of CC-like events in the standard beam configuration (LE10, 185 kA, the top right pane of Figure 5.2). The blue line is the default simulation, and from the ratio it can be seen clearly that this differs from the data by up to 40%. Without the Near Detector this would be a serious problem for a neutrino oscillation measurement, but with the Near Detector this difference can be extrapolated to the Far Detector prediction.

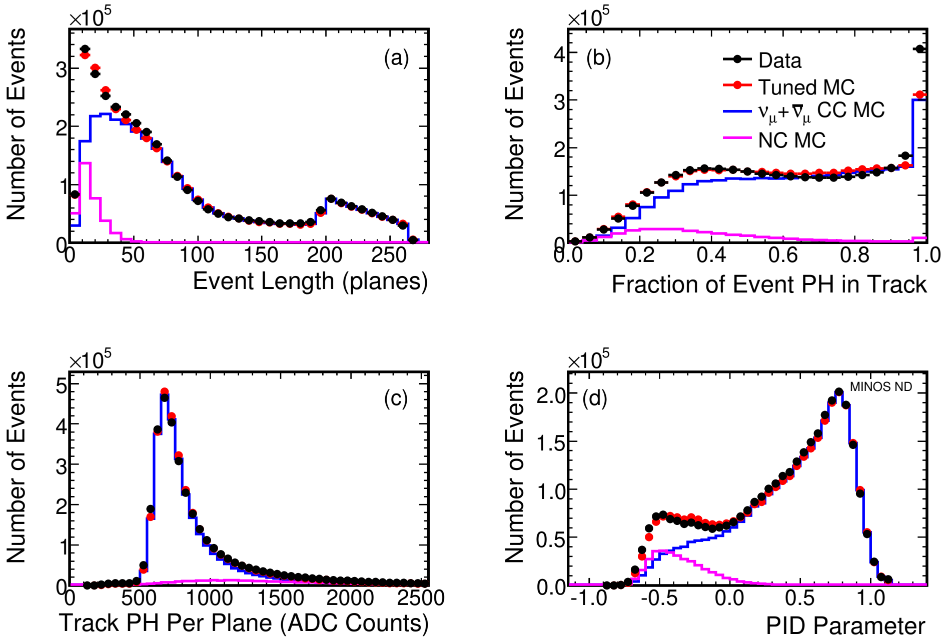

Figure 5.1: Distribution of the variables used for identifying CC-like events, and the ID parameter derived from them. Data from the Near Detector are shown, along with results from simulation showing charged current and neutral current contributions.

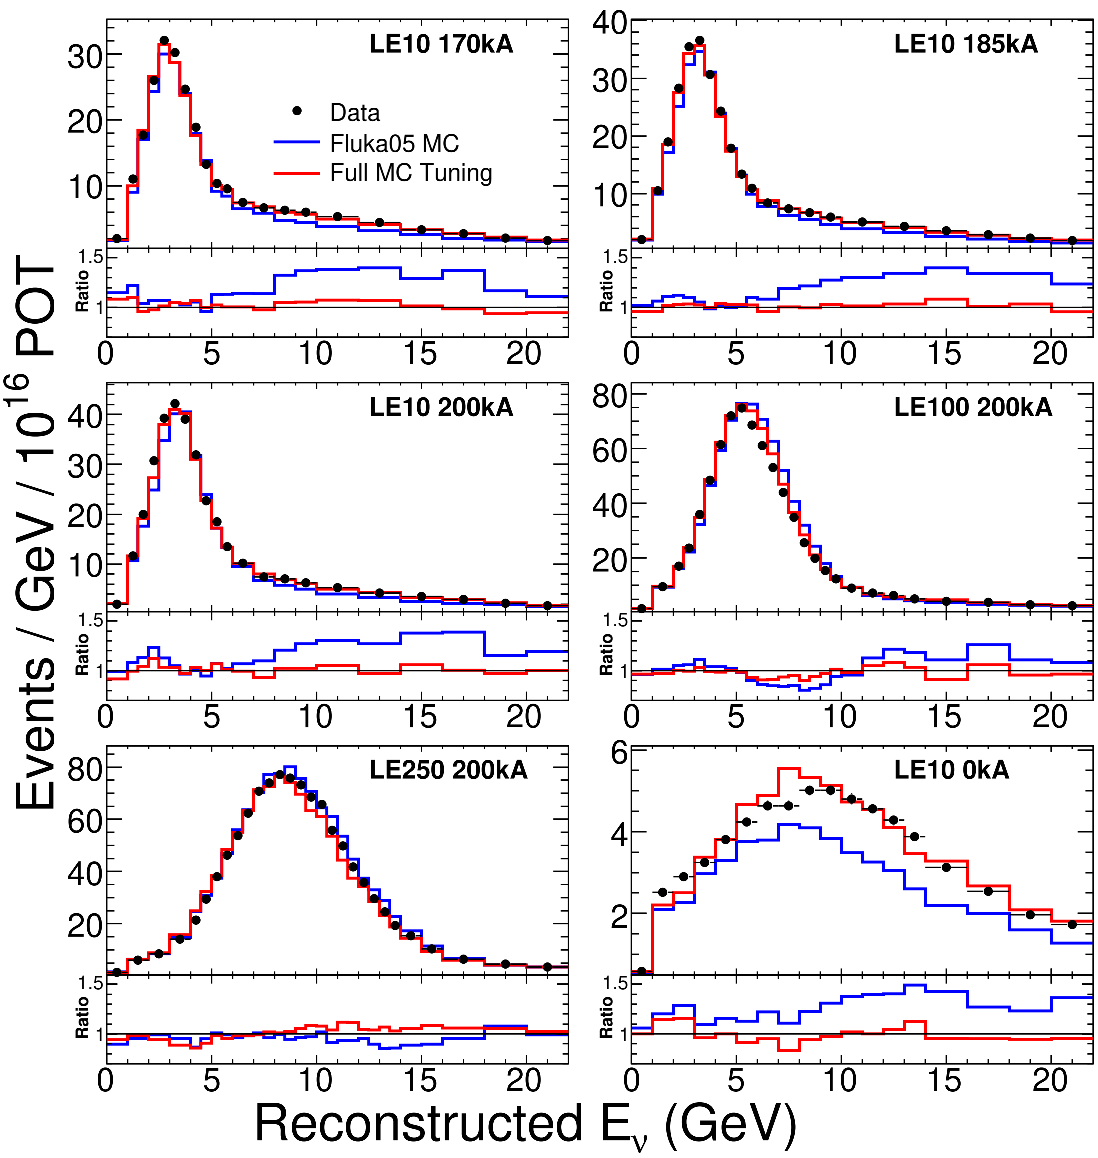

Figure 5.2: Energy spectra of CC-like events in six different beam configurations. Each configuration is identified by the target position *LE[...]* and horn current in kiloamps. The blue line is the default simulation, while the red line is the result of tuning the beam flux to fit the data (black points).

Extrapolation works well, but the discrepancy could be caused by a number of things. Roughly, these can be categorised as uncertainties on the neutrino flux, cross sections, or detector acceptance. Each of these things can be extrapolated, but the exact method to do so is different in each case. It is preferable to identify the source of the discrepancy and MINOS has a way to do so.

The position of the NuMI target and the current in the horns can be altered to adjust the energy spectrum of neutrinos reaching the detectors. The six panes of Figure 5.2 show six of these 'beam configurations' in which data was taken. The fact that the discrepancy between data and simulation moves with the position of the spectrum peak is a clear indication that it is in large part due to the beam model. Errors in the neutrino cross section or detector acceptance would affect the same ranges of neutrino energy regardless of the beam spectrum.

The beam model itself can be regarded as having two separate aspects. One can loosely be called the geometry of the beamline, by which we mean macroscopic parameters like the current in the horns. The other aspect is the hadron production of the thick carbon target. Using the axial symmetry of the beam the yield of hadrons (primarily pions) $\mathrm{d}^2N/\mathrm{d}p_\mathrm{z}\mathrm{d}p_\mathrm{T}$, is represented as a function of the secondary hadrons momentum in the direction of, and perpendicular to, the proton beam (p_z and p_T, respectively). The function is based on the form given in [109], and has seven parameters. The different beam configurations sample different regions of p_z and p_T, so it is possible to identify the momentum regions where the default Fluka'05 hadron production model differs from data.

The beam tuning adjusts the seven hadron production parameters, and five parameters relating to the geometry of the beamline to give the best agreement with all six beam spectra. Three nuisance parameters are also used to take into account uncertainties about the neutrino interactions in the Near Detector. Events are then reweighted according to the twelve beam parameters, resulting in the red curves on Figure 5.2. More information can be found in [110].

A similar tuning is used for the neutral current analysis described in Chapter 6. That tuning is a more recent update that provides weights for secondary beam hadrons up to 120 GeV, and deals more carefully with the subdominant contributions from π^+ and $K^\pm$.

5.3 The extrapolation methods

MINOS takes advantage of the fact that both detectors sit in the same LE10 beam, have similar construction and observe the same interactions. Because of the similarity, it is not necessary to use simulation to predict what will be seen in the Far Detector. Instead, it is only necessary to predict how this is related to what is seen in the Near Detector. Getting this right is a somewhat easier task, and MINOS analyses refer to this step as extrapolation.

Extrapolation, although easier than predicting the Far Detector spectrum by dead reckoning, is not trivial. The different in size of the two detectors leads to differences in acceptance. The beam seen by each detector is also different. For small decay angles θ the energy of the neutrino E_ν is related to the parent pion energy $E_\pi(\gg m_\pi)$ by:

$$E_\nu = \frac{0.43}{1+\gamma^2\theta^2}E_\pi, \tag{5.3}$$

where γ is the Lorentz boost of the pion. The detectors subtend different solid angles at the pion decay points and therefore see a slightly different neutrino spectrum. It is therefore necessary to have complete descriptions of the detector and beam to properly extrapolate the Near Detector data. There are several ways to incorporate the simulations, and the following sections briefly describe how four different extrapolation methods predict the Far Detector spectrum. These methods are also described in [111].

5.3.1 "NDfit" method

This method is closest to the way data is incorporated from independent experiments. Model-uncertainty parameters (two cross-section parameters, track and shower energy scales, and an overall normalisation) are used to represent uncertainties in the simulation common to both detectors. They are allowed to vary with some Gaussian penalty factor, then adjusted so as to minimise a χ^2 fit to Near Detector data. The Far Detector simulation is then updated according to the best fit parameters. Oscillations can easily be applied to the updated simulation by weighting events according to the oscillation parameters and the energy of the incident neutrino.

5.3.2 "2Dfit" method

This method is similar to the NDfit method in that it adjusts several model parameters in order to improve the agreement between Near Detector data and simulation. The method differs in the variables it uses, but mainly because it attempts to fit both the energy spectrum and the kinematic y, defined by

$$y_{\mathrm{kin}} = \frac{E_{\mathrm{shw}}}{E_{\mathrm{shw}}+E_{\mathrm{trk}}}. \tag{5.4}$$

It does this in an empirical way by fitting the simulation to data in a two-dimensional distribution of reconstructed energy versus reconstructed y_{kin}, and having a set of weights applied to the simulation in 'pads' of a two dimensional distribution of true energy versus true y_{kin}. Both the model-uncertainty parameters and the pads have penalty terms to constrain their variation. There are many more bins than pads so the fit is fully constrained. The weights found for the pads are interpreted as differences in the neutrino cross-section so are directly applied to the events with the same energy and y_{kin} in the Far Detector simulation. As with the NDfit, oscillations are then easy to apply to the updated simulation.

5.3.3 "Far over Near" method

This method contrasts quite strongly with the two 'fitting' methods above. The number of events F_i in the ith bin of the Far Detector spectrum is predicted by:

$$F_i = \frac{N_i f_i}{n_i} O_i, \tag{5.5}$$

where N_i is the Near Detector data and n_i and f_i are the Near and Far Detector simulations respectively. O_i is an oscillation weight for the bin and is described below. This method can be thought of as providing a transfer function f_i/n_i to extrapolate the Near Detector data. Alternatively it can be thought of as providing a weighting factor N_i/n_i for the Far Detector simulation. Up to this point the method has only used reconstructed energy spectra, but to apply oscillations the method needs to relate the reconstructed energy of an event to the neutrino energy. This is done with a two-dimensional distribution of reconstructed event energy vs neutrino energy, taken from simulation. Correctly normalised, this can be used as a matrix to convert oscillation weights in bins of neutrino energy into weights for bins of reconstructed energy.

5.3.4 "Beam Matrix" method

This is perhaps the most complicated method to describe. Unlike the other methods the simulation is never used directly to produce spectra. Instead it is used to construct matrices that transform the Near Detector data. In this sense it is the most obviously close to the idea of a Near to Far transfer function.

The method use the simulation to construct matrices that convert spectra of neutrino CC interaction rates (i.e. flux $\times$ CC cross-section), binned in neutrino energy (Greek subscripts), into spectra of CC-selected events, binned in reconstructed energy (Roman subscripts). These matrices account for the overall efficiency to select an interaction as a CC-like event and the purity of the sample, as well as reconstruction bias and the smearing between neutrino energy and reconstructed event energy. The matrix thus constructed for the Far Detector $R^{\mathrm{F}}_{\alpha i}$ is effectively a binned summary of the Far Detector simulation. For the Near Detector, rows and columns of the equivalent true→reco matrix $R^{\mathrm{N}}_{\beta j}$ are scaled to provide an approximation $\overline{R}^{\mathrm{N}}_{\beta j}$ of the inverse reco→true conversion. A third matrix, $B_{\beta\gamma}$, converts interaction rates in the Near Detector to interaction rates in the Far Detector. $B_{\beta\gamma}$ is constructed from the NuMI beam simulation, incorporating our knowledge about the geometry of the beamline and focusing elements, the decay angles of pions and kaons, the solid angle subtended by each detector, and neutrino cross-sections.

The CC-like spectrum at the Far Detector, F_i, is then predicted from the Near Detector data, N_j by

$$F_i = R^{\mathrm{F}}_{i\alpha} O_{\alpha\beta} B_{\beta\gamma} \overline{R}^{\mathrm{N}}_{\gamma j} N_j \ , \tag{5.6}$$

where $O_{\alpha\beta}$ is a diagonal matrix of oscillation weights, and summation over repeated indices is implied.

5.4 Comments and comparisons

By design, the methods are expected to be robust against different kinds of error. The NDfit explicitly specifies the sort of errors that can be handled correctly. The 2Dfit, by virtue of its many pads, should cope well with differences in cross-section. The Far/Near method naturally handles errors in reconstructed energy (such as detector acceptance). And the Beam Matrix is designed to handle the hadron production errors well. But reconstructed event energy, neutrino energy and the hadron yield from the target are all closely related, so all methods cope fairly well with the different types of error.

There is another obvious distinction between the four methods. The NDfit and 2Dfit methods both use a parametric fit to bring the simulation closer to the Near Detector data. The Far/Near and Beam Matrix methods are more naturally thought about as providing transfer functions that operate on the Near Detector event spectrum. It is possible to make too much of this distinction. The Far/Near method can be viewed as a parametric fit in the Near Detector with a scale parameter for each bin of the reconstructed energy spectrum—the scale parameters are then applied to the same bin in the Far Detector spectrum. Similarly it is possible to create variants of the Beam Matrix method that explicitly fit scale factors for the flux of neutrinos in the Far Detector [112]. But there is still an important distinction to be made here. When viewed as a fit, the two 'direct' extrapolation methods adjust the simulation to reproduce the Near Detector data exactly, whilst the two 'fitting' methods can only do so to the extent that their model-uncertainty parameters allow.

The relative lack of flexibility can be both a strength and a weakness for the fitting methods. If the simulation cannot be made to agree with the Near Detector data, then the extrapolation will not completely adjust the Far Detector prediction in line with the Near Detector data. This is not desirable if the the discrepancy is common to both detectors, but might be if the discrepancy was something peculiar to the Near Detector data. In general it is thought that the most important uncertainties for the ν_μ CC analysis are common to both detectors and should be extrapolated, so this is generally regarded as a disadvantage of the fit methods. This issue is clearly more marked for the NDFit method, which has five parameters, than for the 2Dfit, which has around 50 and occupies a more intermediate position.

It is also interesting to compare the Far/Near and the Beam Matrix methods. The matrices used in the Beam Matrix method are derived from the same simulation as used by the Far/Near method. It is therefore possible to describe the Far/Near in the same notation as Equation (5.6). The Near Detector simulation can then be expressed as $n_i = R^{\mathrm{N}}_{\beta i}\mathcal{F}^{\mathrm{N}}_{\alpha}$ using a vector $\mathcal{F}^{\mathrm{N}}_{\alpha}$ of interaction rates (i.e. flux × cross-section) at the Near Detector. A similar equation holds for the Far Detector where, by the definition of $B_{\alpha\beta}$, we have $\mathcal{F}^{\mathrm{F}}_{\beta} = B_{\alpha\beta}\mathcal{F}^{\mathrm{N}}_{\alpha}$ and hence

$$f_i = R^{\mathrm{F}}_{\delta i}O_{\beta\delta}\mathcal{F}^{\mathrm{F}}_{\beta} = R^{\mathrm{F}}_{\delta i}O_{\beta\delta}B_{\alpha\beta}\mathcal{F}^{\mathrm{N}}_{\alpha}. \tag{5.7}$$

Equation (5.5) can then be written:

$$F_i = f_i\, n_i^{-1}\, N_i = R^{\mathrm{F}}_{\delta i}O_{\beta\delta}B_{\alpha\beta}\mathcal{F}^{\mathrm{N}}_{\alpha}\left(R^{\mathrm{N}}_{\gamma i}\mathcal{F}^{\mathrm{N}}_{\gamma}\right)^{-1} N_i \qquad \text{(no sum on } i\text{)}. \tag{5.8}$$

This draws attention to an interesting similarity between the two methods. Although it uses matrices, much of the Beam Matrix method simply reproduces the simulation. The difference arises with the problem of how to make use of the Near Detector data, which ideally requires an inverse of the matrix $R^{\mathrm{N}}_{\alpha i}$. We don't have such a matrix, so the methods must get round this. The Far/Near method is effectively assuming a form for the interaction rates at the Near Detector to create a vector-like object that can be inverted. The Beam Matrix method instead constructs an approximate inverse. Constructing this matrix also requires an input $\mathcal{F}^{N}_{\alpha}$, but the resulting matrix should work reasonably well for similar interaction rate vectors.

5.5 Systematic error evaluation

The statistical sensitivity of each method is very similar as they all use the same Far Detector events. It is in their sensitivity to the various systematic uncertainties that each extrapolation method will differ. A detailed comparison of the effect of systematic uncertainties was undertaken by the Oxford group. Three of the methods (2Dfit, Matrix, NDFit) were evaluated by other members of the group but are presented here for comparison.

To compare each extrapolation method, sets of 'fake data' were prepared, comprising Near and Far Detector Low Energy spectra, and (for the 2Dfit method) two dimensional distributions of reconstructed energy and reconstructed y_{kin}. The Far Detector spectra had $\nu_\mu \to \nu_\tau$ oscillations applied corresponding to $|\Delta m^2_{\mathrm{atm}}| = 3.0 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2 2\theta_{23} = 0.9$. One set of fake data used the nominal (without beam reweighting) simulation, and another 18 'shifted' sets were modified corresponding to reasonable variations in parameters that represented potentially significant sources of systematic error. Each set of fake data was then analysed by performing the extrapolation using the nominal simulation.

A beam fit was performed to each the Near Detector spectra in order to mimic the problem of miss-attributing systematic errors to the beam model. The extrapolation methods were then used to predict the Far Detector spectra based on the shifted Near Detector spectra. A binned likelihood fit assuming Poisson variation in each bin [113] was performed to extract the two-flavour oscillation probability, Equation (2.42). It was agreed that the overall normalisation N would be allowed to vary with a (Gaussian, $\sigma_N = 0.04$) constraint. The Far Detector fit minimises the parameter

$$\begin{aligned}\chi^2(\boldsymbol{\zeta}, N) &= -2\ln\mathcal{L}(\boldsymbol{\zeta}, N)\\ &= 2\sum_{i=1}^{n_{\text{bins}}} N y_i(\boldsymbol{\zeta}) - n_i + n_i \ln\frac{n_i}{N y_i(\boldsymbol{\zeta})} + \frac{(N-1)^2}{\sigma_N^2}.\end{aligned} \tag{5.9}$$

where the observed number of events in the i^{th} energy bin is n_i and the predicted number of events y_i is a function of the oscillation parameters $\boldsymbol{\zeta} = (|\Delta m^2_{\text{atm}}|\,,\ \sin^2 2\theta_{23})$ The number of bins used for fitting was decided by the proponents of each extrapolation method, but the overall energy range was fixed between 0 and 30 GeV. The Far Detector fake data spectra were scaled to an exposure of 1.40×10^{20} POT from an equivalent of 1.13×10^{22} POT, because the relative importance of the penalty term depends on the exposure. The Near Detector fake data spectra were equivalent at a much lower exposure, 1.2552×10^{19} POT, but this has a negligible impact on the statistical error because of the much higher Near Detector event rate.

5.5.1 Systematic shifts

Where possible the systematic shifts correspond to the 1σ Gaussian error on measured input parameters, otherwise they are the best estimate of an appropriate parameter error. The 18 shifted spectra correspond to positive and negative variation of nine parameters. The systematic parameters are summarised in Table 5.1, and described below. Some of these uncertainties are described further in Section 6.4.

Neutrino interaction uncertainties

Uncertainties on the interaction cross sections are modelled by adjusting parameters in NEUGEN. The CCMA uncertainty varies the axial vector mass parameter used in the quasi-elastic and resonance process form factors. The KNO r_{1jk} parameters adjust the fraction of interactions with particular hadron multiplicities. The shower energy models the uncertainty on the amount of energy absorbed by the struck nucleus. Its effect is implemented in the same way as the absolute hadronic energy scale—by multiplying the reconstructed energy of the shower.

Systematic	Shift	Notes
CCMA	±10%	m_a(QE) & m_a(res)
kno_r112122	±10%	KNO r_{112} & r_{122} parameters
kno_r113123	±20%	KNO r_{113} & r_{123} parameters
Shower energy scale	±10%	Energy absorbed by nucleus
NC background	±50%	Neutral current contamination
Absolute calibration	±3%	Energy calibration
ID cut	±0.02	Vary S_{ID} cut applied to data
Track energy	±2%	Muon momentum scale
Beam modelling	±50%	Half of beam tuning

Table 5.1: Systematic uncertainties considered, and their magnitudes.

Reconstruction uncertainties

Two uncertainties are intended to cover mistakes in the selection of CC-like events. The first is the ID cut, which is intended to account for differences between data and simulation in the distributions used to form the ID parameter. This is done by varying the cut on S_{ID} by 0.1 in the fake data sample only.

The second selection uncertainty is the neutral current background, which changes the weight of neutral current events by 50%. This is intended to model the uncertainty on the number of neutral current events *in the CC-like sample*, so is not a cross-section error. For instance, uncertainties on the frequency that tracks are found in neutral current events are represented by this error.

Two 'calibration' errors are included. The absolute hadronic energy scale reflects the uncertainty in the calibration the calorimetric response, and (like the shower energy uncertainty) is implemented as a scale factor for shower energies. The track energy uncertainty reflects uncertainties in the measurement of muon momenta. This comes from two sources (predominantly range at low energies, curvature for higher energies) with different uncertainties, but a flat 2% variation is used for the purposes of studying the systematic.

Beam uncertainties

To allow for the possibility that the the beam fit is incorrect a (conservative) uncertainty was studied, corresponding to 50% of the correction used in the real beam tuning. A beam fit to the fake data was not done in this case as it is supposed to represent the error after the beam tuning.

Relative uncertainties

Two other uncertainties are also potentially important for the final measurement, but not for the comparison of extrapolation methods. The first is the relative normalisation uncertainty, which was estimated at 4%. This is the uncertainty on the ratio of the number of events in each detector. Absolute uncertainties, such as on the total number of POT do not directly contribute to this uncertainty as they affect both detectors equally. Instead the normalisation uncertainty covers how accurately we can measure the target masses of the detectors and their (POT-weighted) livetimes. As the oscillation fit allows a variable normalisation, the effect of this uncertainty is comparative small. The second uncertainty not covered here is the relative calibration uncertainty which was estimated as a 2% variation in the shower energy at one detector. The effect of this error is almost independent of the extrapolation method, so does not affect the comparison. This can be seen easily if it is treated as a 2% shower energy shift applied to the Far Detector (fake) data; the predicted spectrum will be unaffected. In addition, it does not affect the overall error significantly.

5.5.2 Calculating errors

For each systematic uncertainty j, 'shifts' are obtained on $|\Delta m^2_{\mathrm{atm}}|$ and $\sin^2 2\theta_{23}$ by taking the difference between the best fit value with an adjusted uncertainty parameter and the best fit value using the nominal fake data spectra. Two systematic shifts are obtained, η_j^+ and η_j^-, corresponding to positive and negative changes to the uncertainty parameter. In addition, the approximations used in some of the methods can result in an incorrect prediction of the Far Detector spectrum even if the simulation is perfect. This is reflected in the difference between the nominal best fit and input oscillation parameters, and is termed the bias β.

In a few cases the magnitudes of the positive and negative shifts arising from a given uncertainty were significantly different. Drawing any conclusions from this difference is difficult for a number of reasons:

- In most cases the input uncertainty is a best guess and is only symmetric because it is assumed to be so. A symmetric variation using a different parametrisation could give different results.
- The narrow peak of the Low Energy beam causes uncertainties that affect the reconstructed energy of events to give rise to quite complicated distortions of the spectrum. Because of this, the shift in the best fit point is not always a smooth function of the uncertainty parameter. The relative sizes of positive and negative shifts may therefore be sensitive to the exact magnitude of the input uncertainty.
- The shifts in the best fit parameter can also be affected by the input oscillation parameters, as the oscillation maximum lies on or near the peak and the number of affected events can depend sensitively (and again, nonlinearly) on the input value of $|\Delta m^2_{\mathrm{atm}}|$.

For these reasons it is probably unwise to pay attention to the difference between the positive and negative shifts, and to average over points if possible. The systematic sensitivity for each uncertainty is therefore symmetrised:

$$\eta_j = \frac{1}{2}\Big(\left|\eta_j^+\right| + \left|\eta_j^-\right| \Big). \tag{5.10}$$

An overall systematic sensitivity η is then constructed from the sum in quadrature of the individual sensitivities, and the bias:

$$\eta = \sqrt{\beta^2 + \sum_j \eta_j^2}. \tag{5.11}$$

Finally, to obtain a overall figure of merit for comparison of the effect of systematics on different extrapolation methods the systematic sensitivities on $\sin^2 2\theta_{23}$ and $|\Delta m^2_{\mathrm{atm}}|$ are combined:

$$t = \sqrt{1 + \frac{\eta^2\big(|\Delta m^2_{\mathrm{atm}}|\big)}{\overline{\sigma}^2\big(|\Delta m^2_{\mathrm{atm}}|\big)}}\sqrt{1 + \frac{\eta^2\big(\sin^2 2\theta_{23}\big)}{\overline{\sigma}^2\big(\sin^2 2\theta_{23}\big)}} - 1 \tag{5.12}$$

where $\overline{\sigma}$ is the average of the statistical errors from each extrapolation. The averaging ensures t does not improve for methods with larger statistical errors. The statistical sensitivity of each extrapolation method should however be very similar, as they all make use of the same Far Detector data.

5.5.3 Results of the systematic studies

Some examples of the effect of systematic uncertainties are shown in Figure 5.3. The black points show the ratio

$$\frac{\text{FD spectrum with oscillations and systematic shift}}{\text{FD spectrum with oscillations}}$$

The lines show the predicted spectra from each extrapolation method at the input oscillation parameters, also divided by the nominal FD spectrum with oscillations. The fact that the lines generally follow the points shows that the extrapolation methods are working. The residual difference between the points and the lines are what biases the oscillation fit and gives rise to the shifted best fit parameters.

The resulting shifts in the best fit values of $\sin^2 2\theta_{23}$ and $|\Delta m^2_{\mathrm{atm}}|$ are tabulated in Table 5.2 and Table 5.3 respectively. A summary of sensitivities is given in Table 5.4

5.5.4 Comments on systematic studies

In general, these results show that the exact extrapolation method used is not particularly important for this analysis. Statistical uncertainties dominate over

Method		2Dfit	Matrix	Far/ Near	NDfit
Grid spacing		0.001	0.005	N/A	0.001
Systematic:					
CCMA	+	0.002	−0.005	−0.0003	−0.001
	−	−0.003	0.0	0.0017	0.001
kno_r112122	+	0.002	0.0	0.0048	0.004
	−	−0.002	−0.005	−0.0033	−0.002
kno_r113123	+	0.0	0.0	0.0008	0.002
	−	0.0	0.0	0.0007	0.001
Shower energy scale	+	0.001	0.010	0.0022	0.001
	−	−0.001	−0.015	0.0008	−0.001
NC background	+	−0.007	−0.035	−0.0295	−0.032
	−	0.005	0.040	0.0349	0.040
Absolute calibration	+	0.001	0.005	0.0018	0.002
	−	0.0	−0.005	−0.0009	0.0
ID cut	+	0.002	0.005	0.0036	0.008
	−	0.002	−0.005	−0.0011	−0.002
Track energy	+	0.0	0.0	−0.0043	0.001
	−	0.001	−0.005	0.0045	0.003
Beam modelling	+	0.008	0.005	0.0030	0.010
	−	−0.007	−0.005	−0.0011	−0.005
Bias		0.006	0.010	0.0149	0.002
Sum in quadrature		0.010	0.042	0.0334	0.038
Statistical sensitivity		0.120	0.115	0.116	0.115

Table 5.2: Systematic shifts on the measured value of $\sin^2 2\theta_{23}$. For methods that fit using a grid search the shifts are quantised in multiples of the grid spacing.

Method		2Dfit	Matrix	Far/ Near	NDfit
Grid spacing $/10^{-4}\,\mathrm{eV}^2$		0.05	0.05	N/A	0.025
Systematic: $/10^{-4}\,\mathrm{eV}^2$					
CCMA	+	−0.10	0.25	0.055	0.050
	−	0.05	−0.05	−0.130	−0.100
kno_r112122	+	−0.15	0.05	−0.121	−0.125
	−	0.10	0.15	0.076	0.050
kno_r113123	+	−0.05	0.05	−0.045	−0.050
	−	0.0	0.05	−0.024	−0.025
Shower energy scale	+	−0.15	0.05	0.574	0.0
	−	0.25	−0.45	−0.681	−0.075
NC background	+	0.15	0.90	0.573	0.800
	−	−0.15	−0.85	−0.665	−1.050
Absolute calibration	+	−0.10	0.20	0.165	−0.025
	−	0.05	−0.05	−0.160	−0.025
ID cut	+	−0.15	−0.05	−0.091	−0.250
	−	−0.05	0.15	−0.010	0.075
Track energy	+	−0.10	0.45	0.456	0.125
	−	0.05	−0.25	−0.478	−0.100
Beam modelling	+	0.0	−0.25	−0.222	−0.550
	−	0.0	0.25	0.112	0.375
Bias $/10^{-4}\,\mathrm{eV}^2$		0.10	−0.10	−0.323	−0.025
Sum in quad. $/10^{-4}\,\mathrm{eV}^2$		0.34	1.01	1.084	1.061
Stat. sensitivity $/10^{-4}\,\mathrm{eV}^2$		3.73	3.64	3.63	3.75

Table 5.3: Systematic shifts on the measured value of $|\Delta m^2_{\mathrm{atm}}|$. For methods that fit using a grid search the shifts are quantised in multiples of the grid spacing.

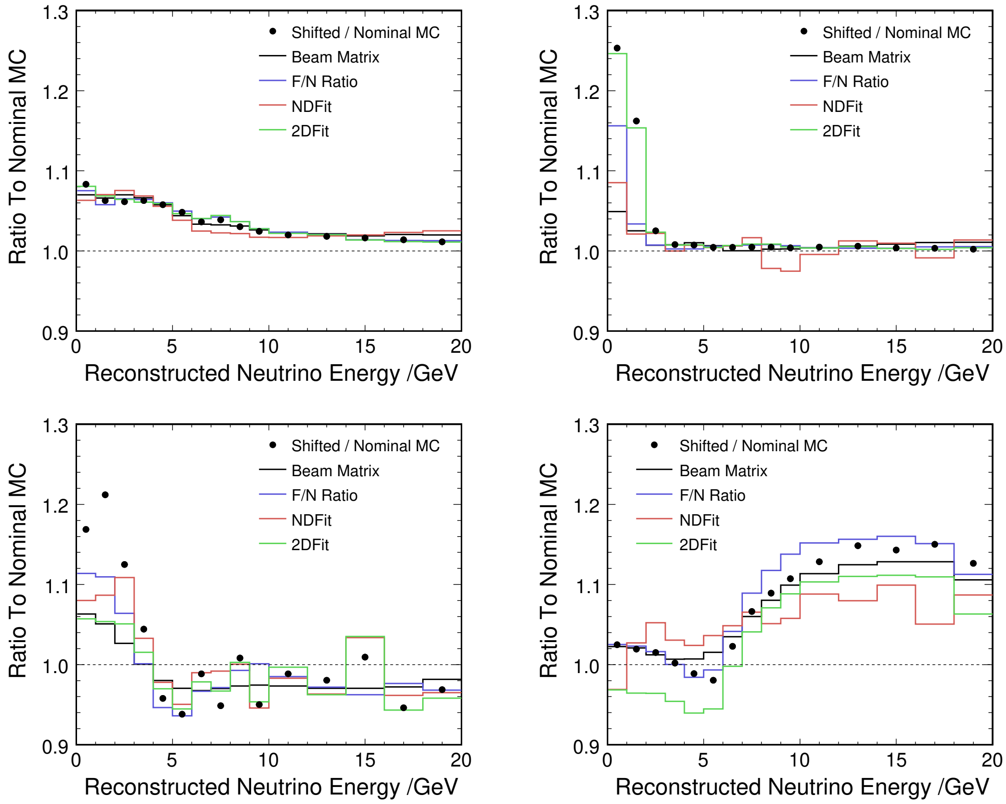

Figure 5.3: Differences in the Far Detector spectrum caused by systematic uncertainties. **(TL)** Increase CCMA parameter by 10%. **(TR)** Increase NC contamination by 50%. **(BL)** Scale shower energy down by 10%. **(BR)** Incorrect beam simulation, as described in the text.

systematic errors, and are independent of the extrapolation method, so all methods are likely to produce similar results when applied to data.

The most obvious specific result is that for both $\sin^2 2\theta_{23}$ and $|\Delta m^2_{\mathrm{atm}}|$ the 2Dfit method appears to have a markedly lower sensitivity to systematic uncertainties. In contrast, the statistical sensitivities are generally consistent to the precision used, with the 2Dfit method possibly a little worse. However, the superior performance of the 2Dfit method against systematic errors may well be illusory, and due to what can be described as the the problem of known unknowns.

The 'problem of known unknowns' is an important consideration in a study such as this. Essentially it arises because the parameterisations used by the fitting methods to tune their extrapolations are, in many cases, the same as used in the evaluation of systematic errors. For instance, in the NDfit method, the shower energy uncertainty is represented by a multiplicative scaling factor, which is exactly the same way that it is parameterised in generation of fake data. This sort of problem can occur for several reasons, but is often unavoidable as sometimes

Sensitivities	2Dfit	Matrix	Far/Near	NDfit
Systematic η				
$\sin^2 2\theta_{23}$	0.010	0.042	0.033	0.038
$\lvert\Delta m^2_{\mathrm{atm}}\rvert$ /$10^{-4}\,\mathrm{eV}^2$	0.34	1.01	1.08	1.06
Statistical σ				
$\sin^2 2\theta_{23}$	0.120	0.115	0.116	0.115
$\lvert\Delta m^2_{\mathrm{atm}}\rvert$ /$10^{-4}\,\mathrm{eV}^2$	3.73	3.64	3.63	3.75
Combined measure t	0.089	0.316	0.289	0.307

Table 5.4: Summary of the different methods' sensitivity to systematic errors. Statistical sensitivities are also shown.

there is only one appropriate way to parameterise an uncertainty. It is therefore possible for a fit-based extrapolation method to reproduce the parameter change used to generate a fake data set almost exactly, and therefore predict the Far Detector spectrum very accurately. In reality, however, the differences between the simulation and data are not so easily parameterised, and the fitting method will not necessarily adjust its parameters in the the most appropriate way. A fitting method can therefore give an artificially low estimate of its sensitivity to systematic uncertainties in this procedure.[1]

A possible way round this is to forbid the extrapolation fit from changing the same parameter as was altered in the fake data sets, but this is often too restrictive and leaves the fit methods completely unable to account for the observed difference. This is less problematic for the 2Dfit (which uses a larger number of parameters than the NDfit) as there is generally more overlap between parameters.

The problem of known unknowns therefore makes it much harder to obtain reliable estimates of the effect of systematic uncertainties on fit-based methods. Extrapolation type methods, which have a prescribed way to deal with all discrepancies, can be evaluated much more reliably.

Because of the uncertainty regarding the reliability of the systematic sensitivity estimates, the Matrix Method (which performed best for $|\Delta m^2_{\mathrm{atm}}|$) was chosen as the primary method for the first ν_μ CC analysis. The NDfit method was chosen to be used as a cross-check as it is the most conceptually different.

The other result of this study deals with the uncertainties rather than the methods. Table 5.2 and Table 5.3 show that uncertainties on the charged current cross sections are fairly unimportant, as extrapolating from the Near Detector data corrects for these quite well. Similarly, the remaining uncertainties in the neutrino flux do not have a large impact on the result. The most important systematic uncertainties remaining are the amount of contamination by neutral current events, and the shower energy scale, which can be ganged together with the absolute hadronic calibration as an 11% uncertainty. The fits to data included

[1]The problem of known unknowns is therefore related, and in some sense opposite, to the problems with fitting methods described in Section 5.4

these, along with the 4% relative normalisation uncertainty, as nuisances parameters in the analysis of data. Their values are marginalised against in the fit which uses a modified likelihood/χ^2 function:

$$\chi^2(\boldsymbol{\zeta}, N, \boldsymbol{\alpha}) = 2 \sum_{i=1}^{n_{\text{bins}}} N y_i(\boldsymbol{\zeta}, \boldsymbol{\alpha}) - n_i + n_i \ln \frac{n_i}{N y_i(\boldsymbol{\zeta}, \boldsymbol{\alpha})} + \frac{(N-1)^2}{\sigma_N^2} + \sum_{j=1}^{2} \frac{\alpha_j^2}{\sigma_j^2}, \tag{5.13}$$

where the two alphas are parameters for the deviation of the NC contamination and shower energy scale from their nominal values. The penalty scales σ_j are as described earlier: 4% for relative normalisation, 50% for NC contamination, and 11% for shower energy scale.

5.6 Results of the first ν_μ CC analysis

The results of the ν_μ charged current analysis are measurements of $|\Delta m^2_{\text{atm}}|$ and $\sin^2 2\theta_{23}$. For the lead (Beam matrix) method the best fit values and 68% C.L. errors (stat. + syst.) are:

$$\left|\Delta m^2_{\text{atm}}\right| = 2.74^{+0.44}_{-0.26} \qquad \text{and} \qquad \sin^2 2\theta_{23} > 0.87$$

The (90% C.L.) joint probability regions from each analysis are shown in Figure 5.4. As expected, there is no significant discrepancy between the different extrapolation methods.

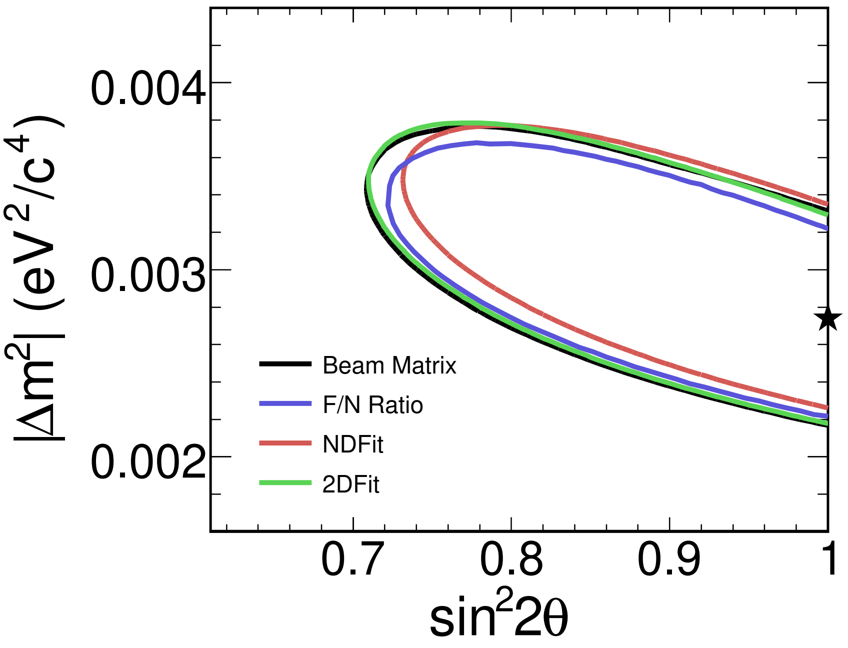

Figure 5.4: Comparison of 90% C.L. regions obtained using different extrapolation methods. The best fit point of the Beam Matrix method is shown by the star.

Chapter 6

NC/CC oscillation analysis

This analysis is primarily concerned with measuring f_{s}, the fraction of the 'disappeared' neutrinos that are sterile. The survival probability for active neutrinos was given in Equation (2.64), and is reproduced here, along with the ν_μ survival probability, for convenience:

$$P(\nu_\mu \to \nu_\mu) = 1 - \sin^2(2\theta_{23})\sin^2\left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right), \tag{6.1}$$

$$P(\nu_\mu \nrightarrow \nu_s) = 1 - f_{\mathrm{s}}\sin^2(2\theta_{23})\sin^2\left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right). \tag{6.2}$$

To measure f_{s} we need to look at the spectrum of events from neutral current interactions. Because an unknown fraction of the neutrino energy is not deposited in the detector the observed spectrum is 'washed out', which significantly degrades its usefulness for constraining $P(\nu_\mu \to \nu_s)$. As there are three parameters in Equation (2.64), the spectrum of NC events does not by itself provide a good constraint on f_{s}. However, the other two parameters are shared by the ν_μ survival probability, and can be much better constrained by the spectrum of ν_μ charged current events, so that the spectrum of neutral current events is effectively only used to get information on f_{s} itself. For this reason the analysis uses events identified as originating from both CC and NC interactions. All events are used, and are divided between two classes: 'CC-like' and 'NC-like'.

Events originating from ν_τ-CC and all ν_e interactions should ideally be excluded from the analysis. This is not done as they are extremely difficult to identify and do not constitute a significant fraction of events. Both ν_e and ν_τ-CC events predominantly end up in the NC-like sample, and must be accounted for in the simulation.

6.1 Reconstructed energy spectra

For analysis, both CC-like and NC-like samples are histogrammed as a function of their reconstructed energy. The definition of reconstructed energy differs between the samples, however.

For events in the CC-like sample (which always have a reconstructed track) the reconstructed energy is the sum of the track energy plus the shower energy. The track is assumed to be a muon, and the energy is obtained from range if it appears to stop in the detector, or from the curvature if it appears to exit. The shower energy is obtained from the summed pulse height of all strips in the shower, having first corrected for the presence of strips shared by the track. The reconstructed energy then provides an estimate of the energy of the incident neutrino.

In the the sample of NC-like events, around half the events contain a reconstructed track. If it exists, the track is ignored and the reconstructed energy is obtained from just the summed pulse height of all strips in the shower. Strips that are shared between the track and shower contribute their full pulse height to the shower. This estimates the energy of hadronic system for NC events, but the total energy of the neutrino for CC events, because the muon is buried in the shower.

6.2 Fit procedure

The oscillation parameters $\boldsymbol{\xi} = (|\Delta m^2_{\mathrm{atm}}|\,, \sin^2 2\theta_{23}, f_{\mathrm{s}})$ are extracted from the data by using the Monte Carlo simulation to generate predictions of the Far Detector reconstructed energy spectra for various values of the parameters, and comparing these predictions with data.

The reconstructed energy spectra for both NC- and CC-like samples are histogrammed between 0 and 45 GeV with a uniform bin width of 0.5 GeV, and overflow bins for events with reconstructed energies greater than 45 GeV. If the index i runs over all 182 bins in both NC- and CC-like spectra, then in the i^{th} bin there are n_i data events, and a prediction of $y_i(\boldsymbol{\xi})$ events. A binned maximum likelihood fit is used to find the best set of parameters. The fit attempts to minimise the quantity [113]

$$\chi^2(\boldsymbol{\xi}) = -2\ln\mathcal{L}(\boldsymbol{\xi}) = 2\sum_{i=1}^{n_{\mathrm{bins}}} y_i(\boldsymbol{\xi}) - n_i + n_i \ln\frac{n_i}{y_i(\boldsymbol{\xi})}, \tag{6.3}$$

which is equivalent to maximising the likelihood ratio

$$\mathcal{L}(\boldsymbol{\xi}) = \prod_{i=1}^{n_{\mathrm{bins}}} \frac{\mathrm{Po}\,(n_i; y_i(\boldsymbol{\xi}))}{\mathrm{Po}\,(n_i; n_i)}, \qquad \text{where} \quad \mathrm{Po}\,(x;\lambda) = \frac{\mathrm{e}^{-\lambda}\lambda^x}{x!}. \tag{6.4}$$

If *all* the $y_i(\boldsymbol{\xi})$ are large then the quantity $\chi^2(\boldsymbol{\xi})$ will follow a chi-squared distribution and can be used for estimating the goodness-of-fit. [114] However, wider bins would result in a loss of information and reduced statistical sensitivity. Since the expected number of events is in the region of 500 (without oscillations) and there are 182 bins used the typical values of $y_i(\boldsymbol{\xi})$ are not large and (despite the nomenclature) this statistic should not be used for evaluating goodness-of-fit. This does not affect its suitability for determining the best fit parameters and errors.

The minimisation is performed using the MINUIT minimisation package. [115] Confidence intervals ('contours') are also found using MINUIT, using the increase in the χ^2 function (the 'up-value', defined as $\Delta\chi^2 = \chi^2 - \chi^2_{\mathrm{min}}$). The physical limits at 0 and 1 on $\sin^2 2\theta_{23}$ and f_{s} are imposed by fixing the χ^2 function at the boundary value and adding half-parabolic penalty terms of the form $+k\delta^2$, where δ is the distance the parameter has strayed outside the physical region, and k is an arbitrary scale factor. Although this procedure results in contours that are meaningless outside the physical region, the best fit point must stay in the physical region (to the within the precision of the fit), and the physical part of the contours are valid.

6.2.1 Additional fits

In the ν_μ CC result, MINOS's ability to resolve the shape of the oscillation provides a much stronger constraint on the oscillation parameters than simply observing the reduced rate of neutrino interactions. In the NC analysis, the shape measurement is likely to be much less powerful than for a CC analysis, so it interesting to ask whether rate or shape measurements contribute more to the overall sensitivity of the measurement. For both the CC and NC like spectra the fit function can be switched (independently) to use:

Rate information only, achieved by using a single bin instead of the 91 bins in the standard fit.

Shape information only, achieved by normalising the predicted spectrum to the observed one:

$$y_i(\boldsymbol{\xi}) \to y_i'(\boldsymbol{\xi}) = \frac{\sum_i n_i}{\sum_i y_i(\boldsymbol{\xi})} \times y_i(\boldsymbol{\xi}).$$

Full fit, i.e. shape and rate information (the standard fit).

None of the information, i.e. do not use this sample in the fit.

For the results presented here, the CC-like spectrum is always fitted using both rate and shape information, and the only difference is whether a full, rate-only or shape-only fit is made to the NC-like spectrum. Because a full fit to the CC-like spectrum is common to all three types of fit, $\chi^2_{\mathrm{both}} \neq \chi^2_{\mathrm{rate}} + \chi^2_{\mathrm{shape}}$.

6.2.2 Predicted spectra

To obtain the predicted spectrum for a given set of oscillation parameters $\boldsymbol{\xi}$ the approach take is to weight events according to their oscillation probability. Two sets of simulated (Monte Carlo) events, are used. One is the standard beam set, which represent the events at the Far Detector in the absence of oscillations, and contain both ν_μ and the beam ν_e (and antineutrinos). The other set is the tau set, in which all neutrinos (antineutrinos) are forced to be ν_τ ($\overline{\nu}_\tau$).[1]

Oscillation weights are then applied. Because neutral current events are unaffected by the $\nu_\mu \to \nu_\tau$ transition it is convenient to use ν_μ neutral current events to represent ν_τ neutral current events, so the tau set is only used for ν_τ CC events. Thus the 'oscillation' weights are not true oscillation probabilities, but are applied both as a function of the neutrino flavour and of the interaction type. Table 6.1 shows the oscillation weights for each type of neutrino and interaction type.

	Beam set		Tau set	
	ν_e	ν_μ	ν_τ (orig. ν_e)	ν_τ (orig. ν_μ)
CC	1	$1 - D$	0	$(1 - f_\mathrm{s})D$
NC	1	$1 - f_\mathrm{s}D$	0	0

Table 6.1: Weights applied to events to reproduce the effects of $\nu_\mu \to \nu_\tau$ and $\nu_\mu \to \nu_s$ oscillations. The factor $D = \sin^2 2\theta_{23} \sin^2 (\Delta m^2_\mathrm{atm} L/4E)$ is common to all modes of oscillation.

The simplest way to create spectra with appropriate oscillation weights is to add an entry for each event with the appropriate weight. This would require the fitting code to loop over every event each time the parameters are changed. However it can be seen in Table 6.1 that the oscillation weights for all events are the sum of two terms, one that is a constant and one that is proportional to $\sin^2 (\Delta m^2_\mathrm{atm} L/4E)$. The proportionality constants for these terms are simple functions of f_s and $\sin^2 2\theta_{23}$. If a pair of component spectra are made for each type of event (ν_μ/ν_τ NC, ν_μ CC, ν_τ CC and ν_e), one with an oscillation weight of 1, and one with an oscillation weight of $\sin^2 (\Delta m^2_\mathrm{atm} L/4E)$, then spectra can be made for any combination of f_s and $\sin^2 2\theta_{23}$ by simply adding all these component spectra in the correct proportions.

Providing the correct weights for different values of $|\Delta m^2_\mathrm{atm}|$ is somewhat more complicated as there is no (useful) subset of events where it acts as a scale factor. However between two sufficiently close values of $|\Delta m^2_\mathrm{atm}|$ the oscillation probability for any intermediate value of $|\Delta m^2_\mathrm{atm}|$ can be approximated by linearly interpolating between the two. The question is: how close is close enough? It is interesting to remember at this point that oscillations in the atmospheric sector are really

[1] Precisely, the flavour of the neutrino *flux* is altered, so that the tau data set includes the effect of the lower ν_τ CC cross-section near the $\tau^\pm$ threshold.

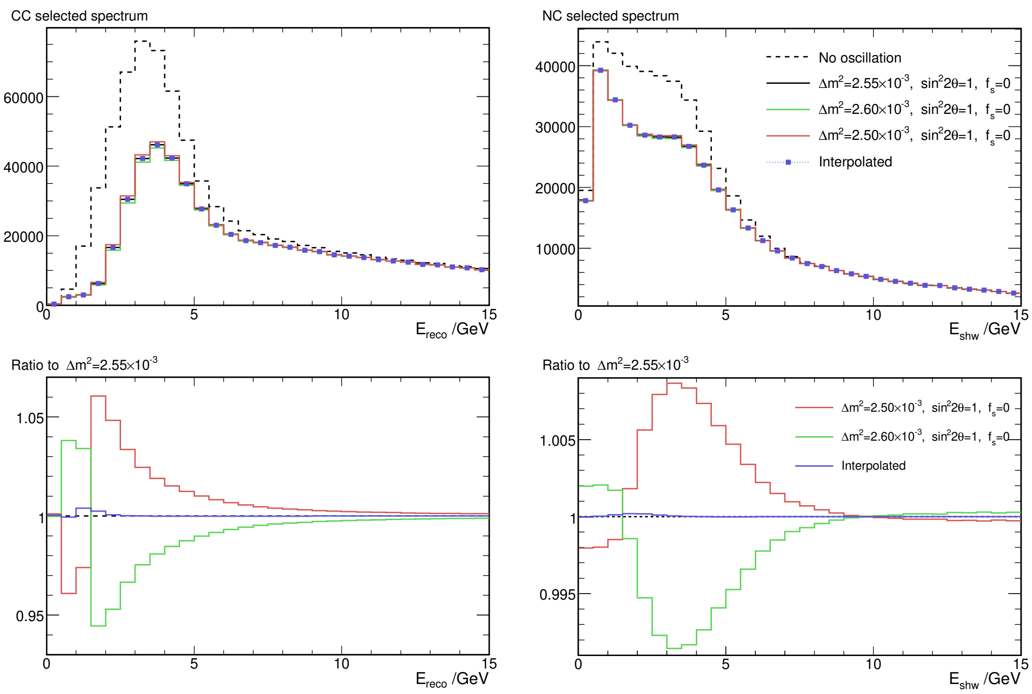

Figure 6.1: Errors introduced by interpolating $|\Delta m^2_{\rm atm}|$. The CC spectrum is affected more strongly than the NC spectrum, but the largest discrepancy is less than 0.5%, and most bins are affected much less than this.

the sum of two oscillations, one at Δm^2_{23}, and one at Δm^2_{13}. This implies that the linear interpolation should work for spacings of order Δm^2_{12}. Moreover, the error introduced by doing a linear interpolation over this range will be of a similar size to that introduced by using the 2-mass approximation. Therefore it should be sufficient to use spacings in $|\Delta m^2_{\rm atm}|$ of $10^{-4}\,{\rm eV}^2{\rm c}^{-4}$. Figure 6.1 shows an example of the error introduced by this interpolation, for $|\Delta m^2_{\rm atm}| = 2.55 \times 10^{-3}\,{\rm eV}^2{\rm c}^{-4}$. Errors due to the interpolation get smaller for larger values of $|\Delta m^2_{\rm atm}|$ whilst for smaller values the discrepancy drops to lower energies, where the neutrino flux is low.

By interpolating in $|\Delta m^2_{\rm atm}|$ it then possible to make predicted energy spectra for any set of oscillation parameters by summing nine component spectra (the oscillated spectra for ν_e and unoscillated spectrum for ν_τ are not needed). The histograms for the component spectra must be generated first but this can be done in a single pass through the simulated events, and is much faster than looping over every event each time a new set of oscillation parameters is needed. Figure 6.2 and Figure 6.3 show the different components of the Far Detector spectrum for two sets of oscillation parameters.

Finally, to take full advantage of the the data and minimise the impact of systematic uncertainties a Near to Far extrapolation is implemented. The extrapola-

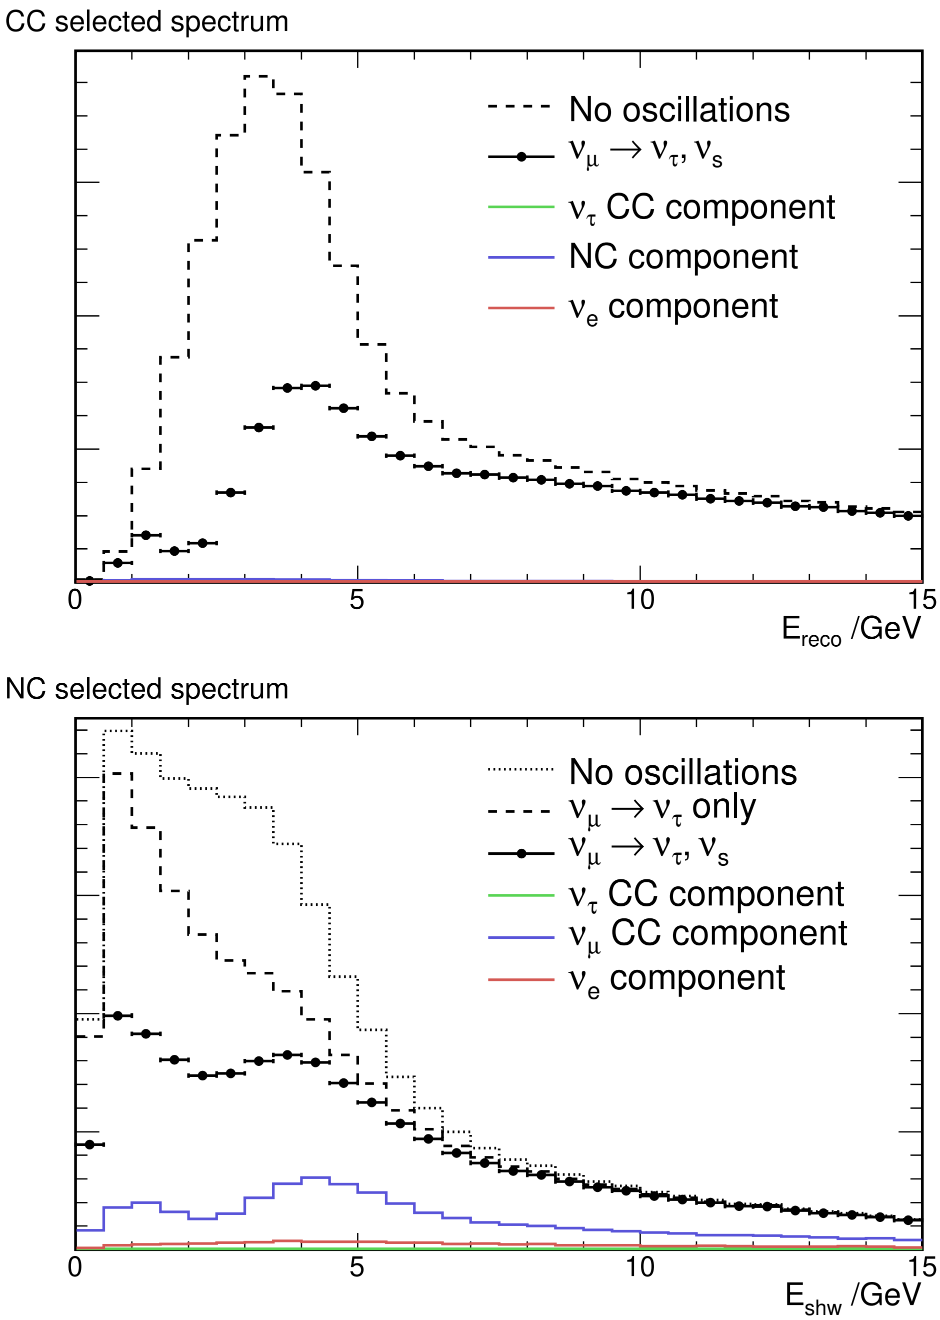

Figure 6.2: Example Far Detector spectra showing background contributions, with a set of parameters ($|\Delta m^2_{\mathrm{atm}}| = 3.5 \times 10^{-3}\,\mathrm{eV^2c^{-4}}$, $\sin^2 2\theta_{23} = 1$ and $f_s = 1$) chosen to demonstrate clearly the effect of a sterile signal.

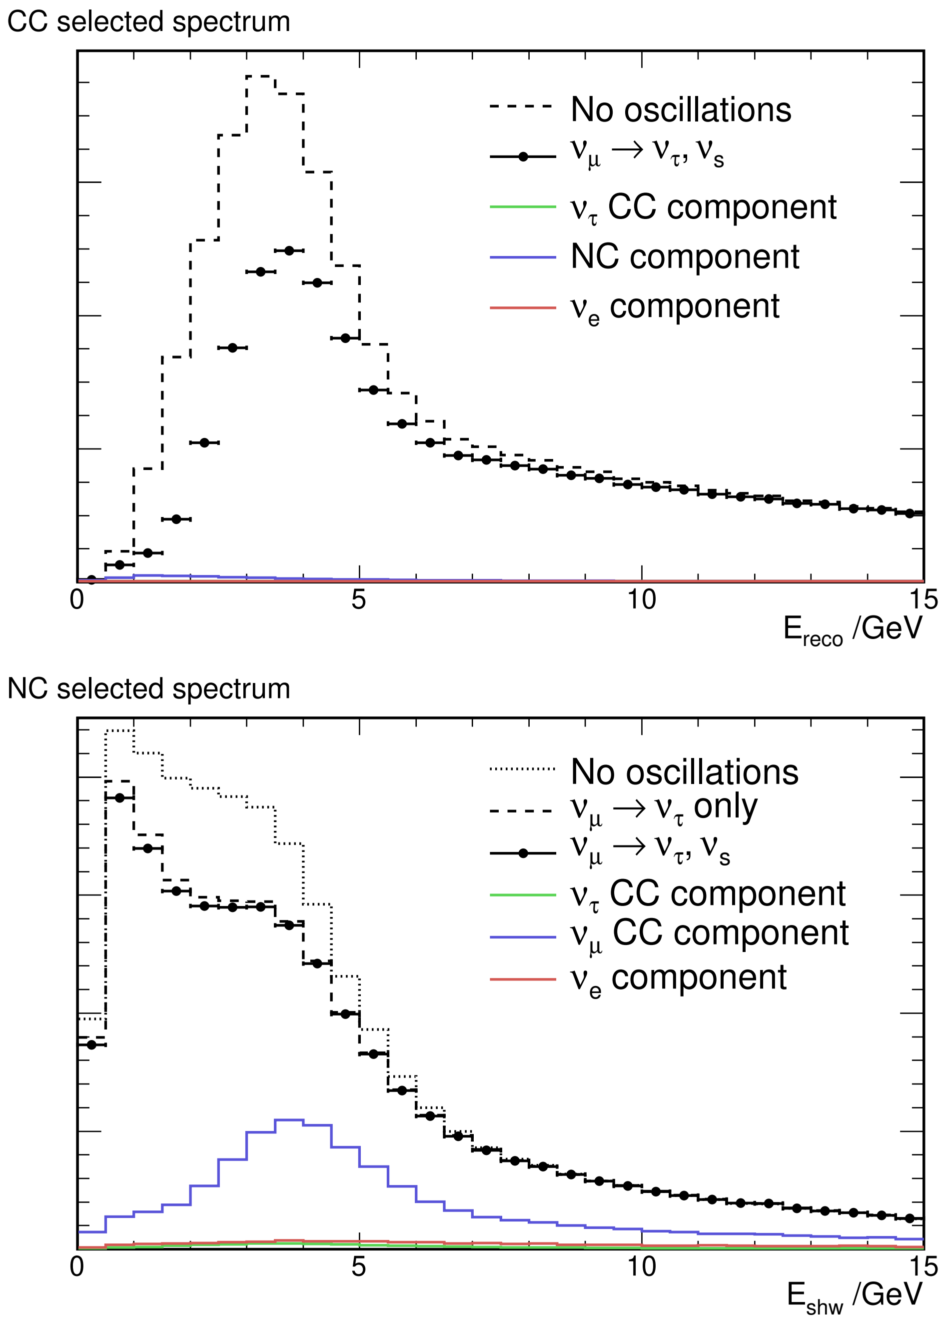

Figure 6.3: Example Far Detector spectrum showing background contributions, with a set of parameters ($|\Delta m^2_{\mathrm{atm}}| = 2.5 \times 10^{-3}\,\mathrm{eV^2c^{-4}}$, $\sin^2 2\theta_{23} = 0.9$ and $f_{\mathrm{s}} = 0.1$) that are consistent with current knowledge.

tion used is an extended version of the Far/Near method described in Section 5.3.3 As with the ν_μ CC analysis, the predicted number of events in bin i of the Far Detector energy spectra F_i is given by

$$F_i(\boldsymbol{\xi}) = \frac{N_i\, f_i(\boldsymbol{\xi})}{n_i}, \tag{6.5}$$

with f_i being the equivalent bin in the predicted Far Detector spectrum without any extrapolation, and N_i and n_i being bins in the Near Detector data and simulation respectively. To extend this to the neutral current analysis, the index i simply runs over both CC-like and NC-like spectra. This method is fairly robust; in comparison, a similar extension to the Beam Matrix method to extrapolate the NC-like spectrum runs into difficulties associated with the mapping of reconstructed event energies to neutrino energies. The Near Detector data is separated into CC- and NC-like samples in the same way as done in the Far Detector, but the set of cuts for removing the non-neutrino background is necessarily different, and described in [91].

6.3 Statistical sensitivity

One of the first things that can be done with the fit is to establish the statistical sensitivity of the data set. This is done by using the simulated events (with a known oscillation applied) as the data, as well as in their normal role of generating predicted spectra. The Far Detector spectra are scaled to an appropriate number of POT so that the correct statistical errors (implicit in Equation (6.3)) are used. The statistical sensitivities are studied at the same test point as was used in Figure 6.3:

$$\left|\Delta m^2_{\mathrm{atm}}\right| = 2.5\times 10^{-3}\,\mathrm{eV}^2\mathrm{c}^{-4}, \qquad \sin^2 2\theta_{23} = 0.9, \qquad f_{\mathrm{s}} = 0.1,$$

and the sensitivities are calculated for 1.20×10^{20} POT. Single parameter errors are calculated for a confidence level of 68% (a Gaussian 1σ error, hereafter 68% C.L.) using a $\Delta\chi^2$ of +1. Joint probability contours for two parameters are plotted for 68% C.L. and 90% C.L. using $\Delta\chi^2$ values of +2.30 and +4.61 respectively.

Asymmetric single parameter errors are quoted. In the case of $\sin^2 2\theta_{23}$ and f_{s}, the presence of physical boundaries can limit the parameter values, resulting in reported errors that are simply the distance of the best fit point to to the physical boundary. (For instance, if $\sin^2 2\theta_{23} = 0.9$, the upper error is often precisely 0.1.) For comparative purposes, only the errors that are not constrained by the physical boundary are reported here.

The single-parameter statistical sensitivities, for a simulated exposure of 1.2×10^{20} POT are given in Table 6.2. The joint sensitivity for each combination of two oscillation parameters is shown in Figure 6.4. The allowed regions show that f_{s} is

Parameter		Sensitivity		
		Rate	Shape	Full
f_s	+ only	0.531	0.610	0.411
$\lvert\Delta m^2_\mathrm{atm}\rvert$	+	0.470	0.468	0.460
$/10^{-3}\,\mathrm{eV^2c^{-4}}$	−	0.393	0.391	0.384
$\sin^2 2\theta_{23}$	− only	0.174n	0.175	0.173

Table 6.2: Statistics-only sensitivities for 1.2×10^{20} POT. Errors are obtained using $\Delta\chi^2 = +1$, and marginalising against the other parameters.

not strongly correlated with either $|\Delta m^2_\mathrm{atm}|$ or $\sin^2 2\theta_{23}$.

6.4 Systematic errors

An important consideration is the systematic errors on the measurement of f_s. This is approached in a similar way to the systematic error analysis described in Section 5.5. Sets of Monte Carlo are generated with modified input parameters. Each parameter was shifted by its estimated one standard error ('1σ') uncertainties. These were compared with fake data sets generated using the nominal Simulation[2]. The Far/Near extrapolation is done and the fit is performed to obtain the best-fit oscillation parameters. The shift in the best fit f_s is taken as an measure of the sensitivity of the measurement to the systematic uncertainty. All systematics were studied at the same test point as was used in Figure 6.3, which was also used for the figures in this section. The values of $\sin^2 2\theta_{23}$ and f_s are offset from the physical limit so that shifts in the best fit points were not obstructed by the physical limits. For some systematics this is not sufficient and the best fit points were still obstructed by the physical limits, so for these studies the limits are not applied until $\sin^2 2\theta_{23}$ and f_s move more than 0.2 outside the physically allowed regions.

The systematic uncertainties tested are summarised in Table 6.3. Included are the uncertainties expected to give the largest errors on f_s, [116] described below, plus the two Far Detector cleaning systematics described in Section 4.2.5.

6.4.1 Relative normalisation

This uncertainty refers to the number of neutrino events observed at the Far Detector, relative to the number seen at the Near Detector. This is different from the absolute normalisation (or proton counting) which is well corrected by the Near to Far extrapolation. The value of 4% is taken from the charged current analysis [57], and arises from several sources, which are listed in Table 6.4.

[2]This is different from the CC analysis where the systematic shifts were applied to the fake data spectra.

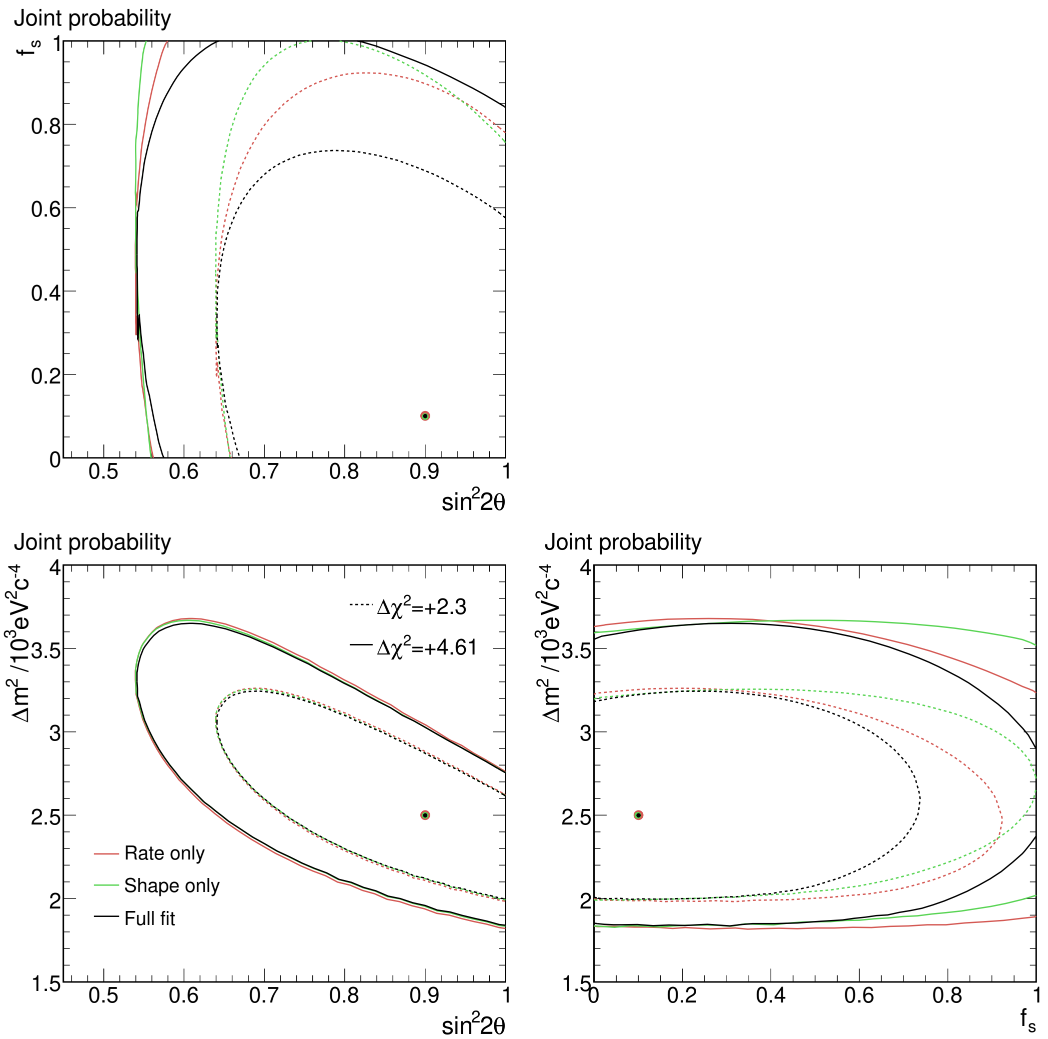

Figure 6.4: Statistics-only joint sensitivity contours for 1.2×10^{20} POT. The dashed and solid lines are 68% C.L. and 90% C.L. contours respectively.

Systematic	'1σ' value
Relative normalisation	4%
Relative shower energy scale	2.6%
Absolute shower energy scale	10%
CC background in NC spectrum	5%
NC background in CC spectrum	50%
Far Detector cleaning 1	see Section 4.2.5
Far Detector cleaning 2	see Section 4.2.5

Table 6.3: Systematic uncertainties considered, and their magnitudes.

Source	Contribution
Livetime	1.0%
Detector composition	0.3%
Fiducial volume	2.1%
Reconstruction efficiency	3.0%
Total (sum in quadrature)	4%

Table 6.4: Contributions to the 4% relative normalisation systematic. Summarised from [117].

6.4.2 Relative shower energy

This is another relative uncertainty that refers to a difference between the Far and Near detectors. In this case the difference is in the calorimetric response of the detectors. This uncertainty reflects residual differences between Near and Far detectors, and reducing it is perhaps the most important reason for calibrating the detectors. The most recent estimate of this uncertainty puts it at 3.8% [118], but part of this includes an uncertainty at each detector covering the use of an old version of the simulation of cosmic rays. The difference between this and the more recent simulations is now thought to be an overall scale factor [119] and would probably not contribute to a relative error. The relative shower energy uncertainty should properly be recalculated from scratch, but this has not yet been done and for now these contributions are removed, leaving other contributions unchanged. This results an overall uncertainty of 2.6%. Because this was uncertain at the time of calculating the resulting shifts, uncertainties of 2% and 3% were used, and interpolation used to obtain the shift from an intermediate 2.6% uncertainty.

As this is a relative uncertainty it is applied by modifying the reconstructed shower energy (in simulation) in one detector only. For this evaluation it applied in the Far Detector as this is computationally easier. The effect of a 3% relative shower energy systematic uncertainty on the predicted CC-like (NC-like) spectrum is shown in Figure 6.5 (6.6).

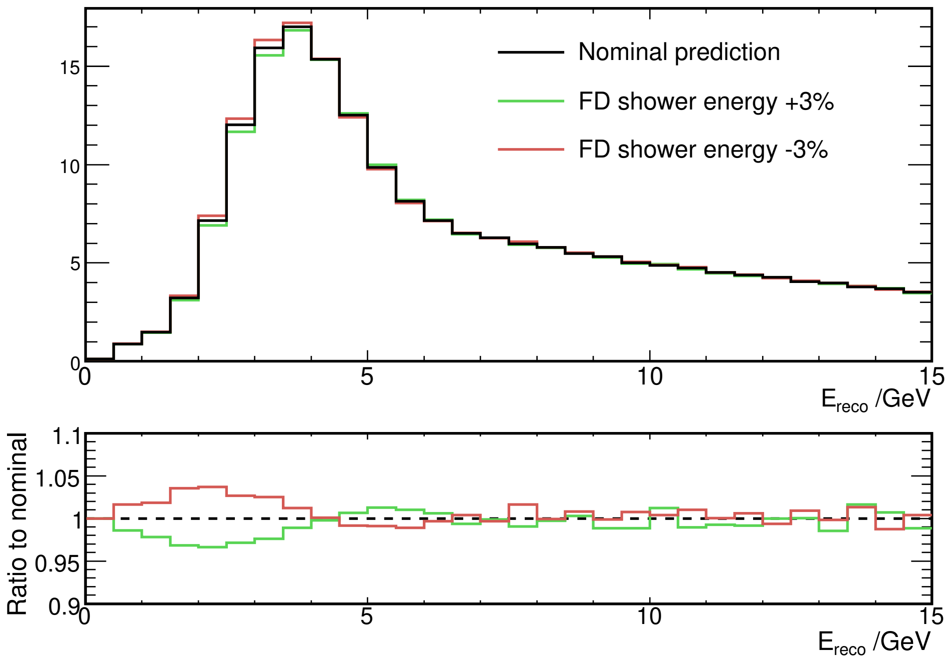

Figure 6.5: Effect of a 3% variation of the shower energy scale in the Far Detector on the predicted CC-like energy spectrum.

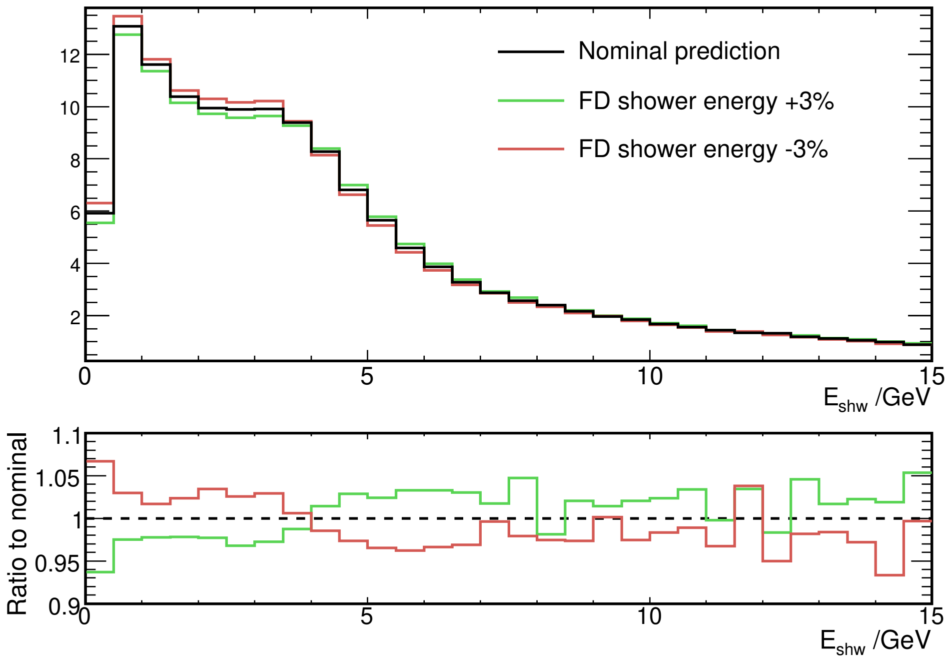

Figure 6.6: Effect of a 3% variation of the shower energy scale in the Far Detector on the predicted NC-like energy spectrum.

6.4.3 Absolute shower energy

This uncertainty reflects the detector response after calibration (5.7% [118]), and the uncertainty on the fraction of the hadronic system's energy absorbed by the struck nucleus (8% [120]). A different model of the hadronic system would also result in a different NC/CC classification for some events; this is not taken into account in this systematic, but does form part of other systematic uncertainties, such as the two background systematics below.

Although the effect on the Far Detector simulation is substantially larger than with the relative shower energy uncertainty, the impact of this uncertainty is mitigated by the Near to Far extrapolation. In the ideal case of a perfect extrapolation method this uncertainty would have no effect on the predicted Far Detector spectra.

Figure 6.7 (6.8) shows the effect of this systematic on the predicted CC-like (NC-like) spectrum. It is interesting to note there is a more obvious effect on the CC-like spectrum than on the NC-like spectrum, despite the shower energy being a more important quantity for NC-like events. This is possibly because neutrino oscillations have a greater effect on the shape of the CC-like spectrum, making the Far and Near spectra more dissimilar, and therefore reducing the effectiveness of the extrapolation method.

6.4.4 CC background

The charged current uncertainty affects the NC-like spectrum only, in both detectors. The charged current background is scaled (by $\pm 5\%$) to represent the uncertainty in the number of charged current events that appear NC-like. ν_e and ν_τ events are also scaled, although their contribution is negligible.

This uncertainty is intended to represent a number of factors: reconstruction effects, such as the curvature direction uncertainty examined Section 4.3.1; or interaction and hadronisation model uncertainties. It is unclear whether a corresponding change should be made to the CC-like spectrum. The change in the number of charged current events in the NC-like spectrum is small in comparison to the total number of charged current events in the CC-like spectrum, however, so this is not a great concern.

6.4.5 NC background

The NC background uncertainty plays a similar role to the CC background uncertainty, but for neutral current events in the CC-like spectrum. Although the scaling factor is much greater ($\pm 50\%$) than for the CC background uncertainty the overall effect is small because of the much lower neutral current contamination in the CC-like spectrum. The magnitude of this uncertainty is taken from the ν_μ charged current analysis, which uses a different selection. It is assumed that this is still a reasonable estimate.

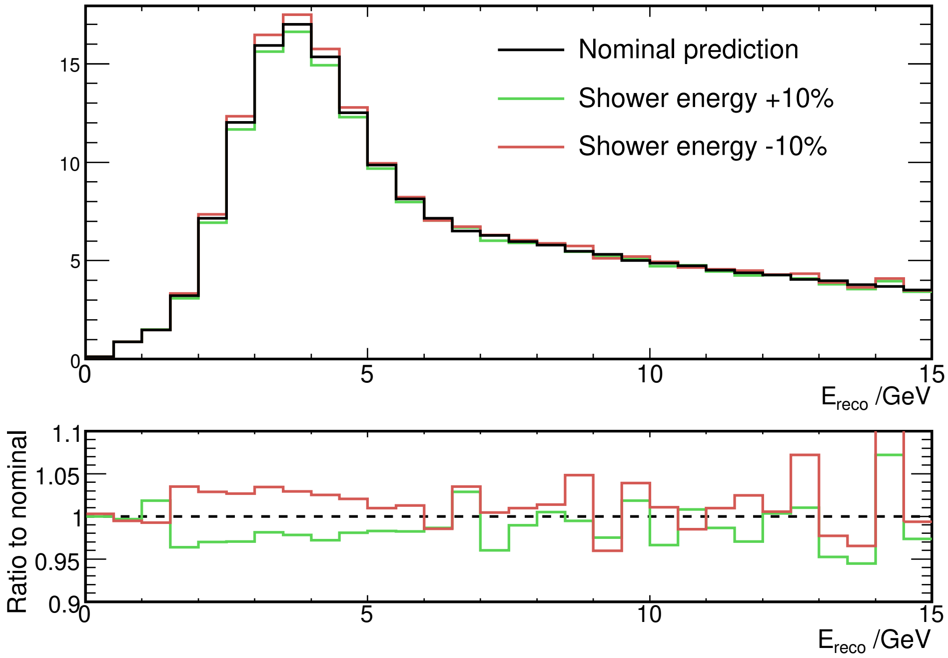

Figure 6.7: Effect of a 10% variation of the shower energy scale in both detectors on the predicted Far Detector CC-like energy spectrum.

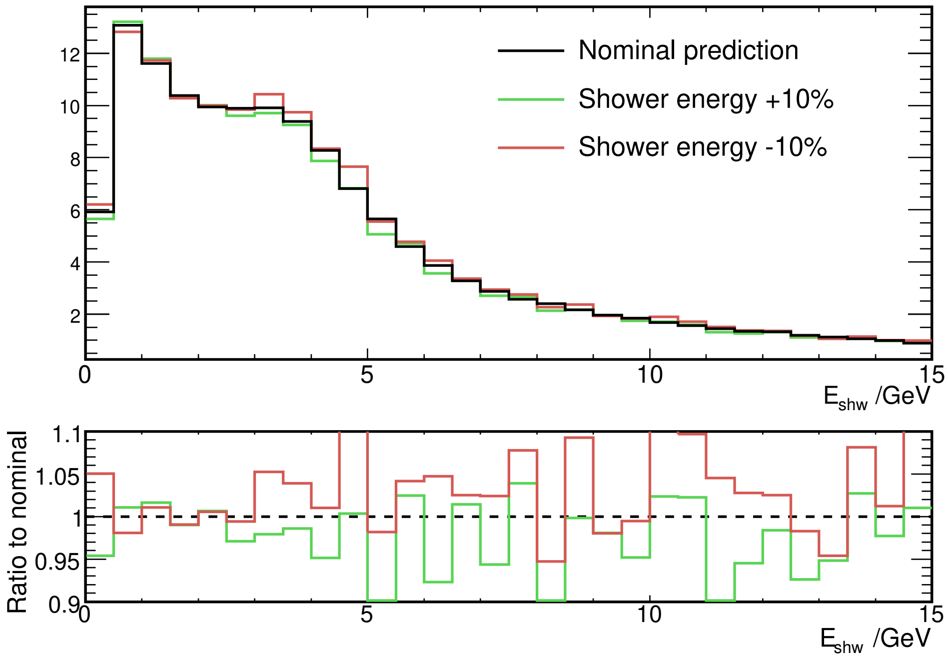

Figure 6.8: Effect of a 10% variation of the shower energy scale in both detectors on the predicted Far Detector NC-like energy spectrum. Excursions are less than ±0.15

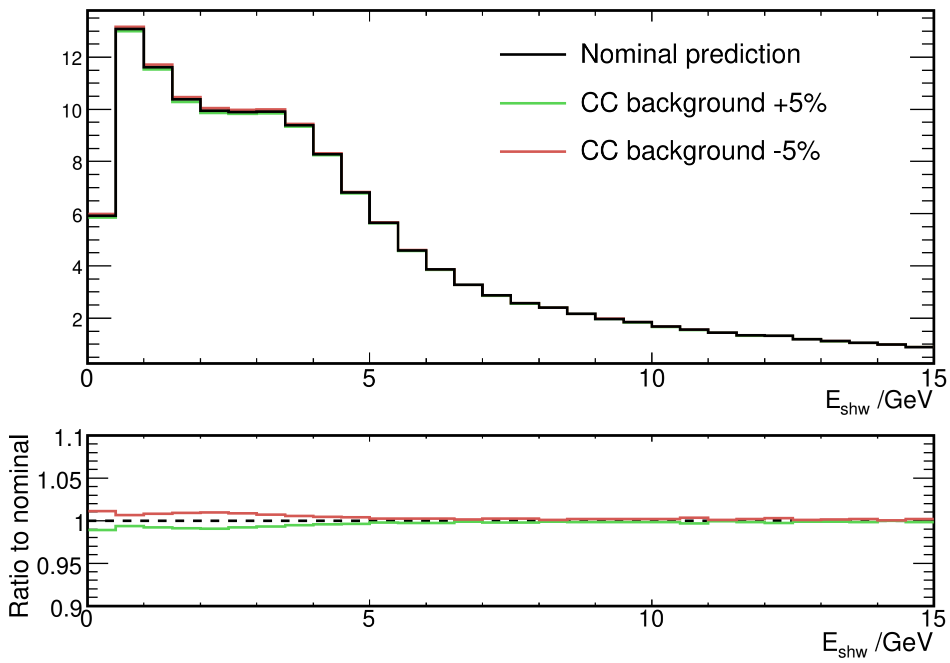

Figure 6.9: Effect of a 5% variation of the charged current background on the predicted Far Detector NC-like energy spectrum.

6.4.6 Far Detector cleaning

The two Far Detector cleaning systematics are described in Section 4.2.5. The systematic uncertainty affecting events with shower energy below 0.75 GeV is referred to as (Far Detector cleaning) systematic 1, whilst the uncertainty arising from the cosmic ray removal cuts is systematic 2.

Both uncertainties were obtained without making any distinction between NC-like or CC-like events, although it is reasonable to expect that the NC-like events are more severely affected. As a result, the systematic errors were tested with shifts applied to both the NC-like and CC-like spectra, or just to the NC-like spectrum. The resulting errors on f_s were worse in the latter case, so that is what is presented here. Figure 6.11 shows the effect of the (low-energy) Far Detector cleaning systematic 1, and Figure 6.12 shows the effect of the (cosmic ray removal) Far Detector cleaning systematic 2.

6.4.7 Results of systematic study

Each systematic uncertainty parameter was varied in both directions. The positive and negative shifts of the best fit parameter (η_i^+ and η_i^-) were combined by taking the average of the magnitude of the shift.

The resulting average shift (η_i) is signed to show the direction of the shift when

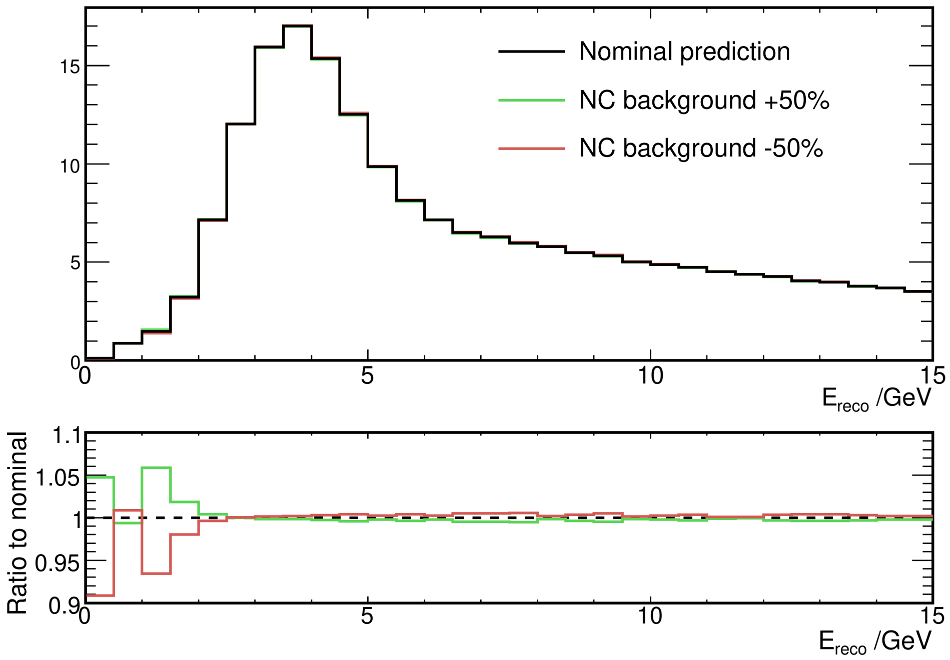

Figure 6.10: Effect of a 50% variation of the neutral current background on the predicted Far Detector CC-like energy spectrum.

the systematic parameter is increased, i.e.

$$\eta_i = \frac{1}{2}\mathrm{sign}(\eta_i^+)\left(\left|\eta_i^+\right| + \left|\eta_i^-\right|\right). \tag{6.6}$$

This allows the effect of a given systematic on the rate and shape fits to be compared, to see if there are competing effects, but the sign of the shift is not otherwise used.

Table 6.5 summarises the results of the systematic study. Several interesting results are apparent:

- The statistical errors on the value of f_s are around 2–3 times larger than sum in quadrature of these systematic shifts. This analysis will therefore be dominated by statistical error, although systematic errors may become important for future analyses with higher statistics.

- The systematic shifts on f_s are well behaved, in that the the two opposite-sign shifts in each parameter result in opposite-sign shifts in the value of f_s. The magnitude of each of a pair of shifts is also similar. To a slightly lesser extent this is true for the rate- and shape-only fits (not shown). This contrasts with the ν_μ CC analysis described in Section 5.5.3 where some systematic shifts were highly asymmetric.

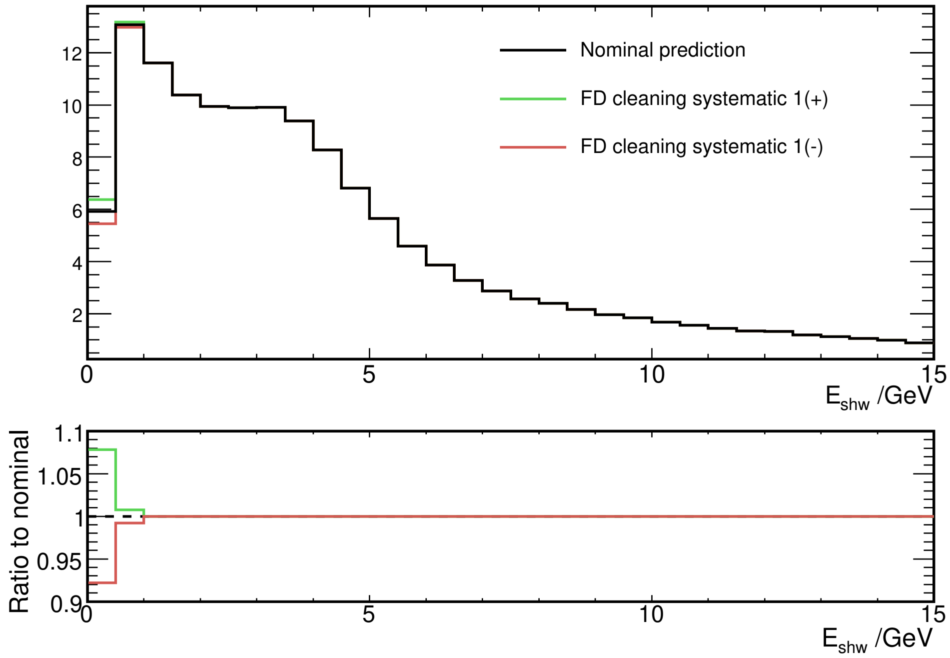

Figure 6.11: Effect of '1σ' variations of the the first Far Detector cleaning systematic on the NC-like energy spectrum.

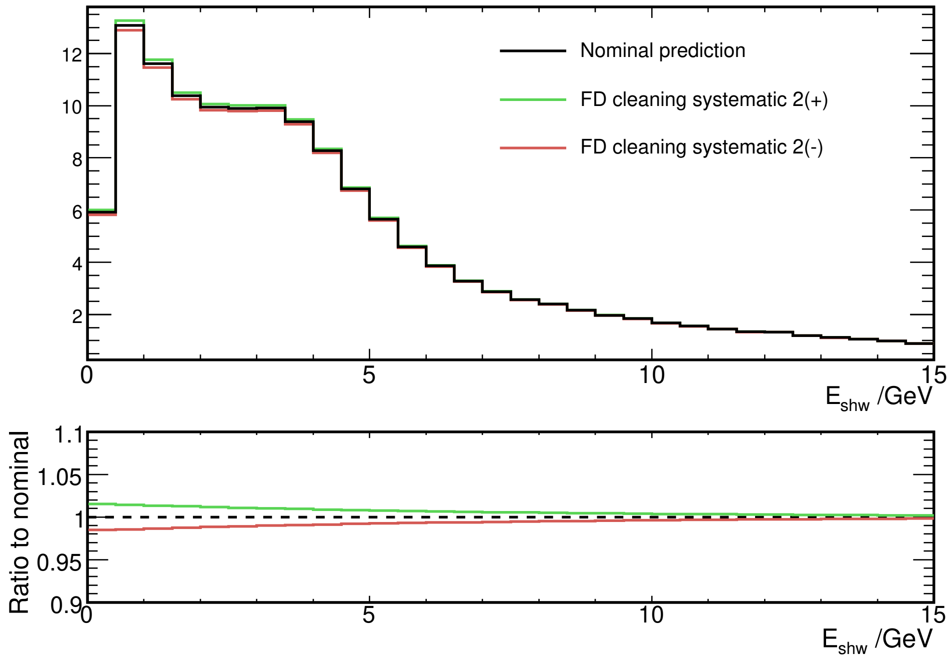

Figure 6.12: Effect of '1σ' variations of the the second Far Detector cleaning systematic on the NC-like energy spectrum.

Systematic		Shift in f_s	Signed average	Rate only	Shape only
Relative normalisation	+	0.125	0.137	0.222	0.007
	−	0.045			
Relative shower energy	+	−0.058	−0.058	0.017	−0.176
	−	0.058			
Absolute shower energy	+	−0.028	−0.030	−0.086	0.056
	−	0.033			
CC background	+	−0.028	−0.026	−0.032	−0.023
	−	0.024			
NC background	+	−0.008	−0.008	−0.006	−0.011
	−	0.007			
Far Detector cleaning 1	+	0.033	0.034	0.022	0.046
	−	−0.034			
Far Detector cleaning 2	+	0.045	0.046	0.054	−0.266
	−	−0.047			
Bias			< 0.001	< 0.001	< 0.001
Total (sum in quadrature)			0.168	0.248	0.305

Table 6.5: Shifts in the best fit value of f_s due to various systematic uncertainties

- The relative normalisation uncertainty is clearly the most important. This is not unsurprising, given the rate-only fit was more sensitive to f_s than the the shape-only fit.

- The relative normalisation uncertainty also affects the shape-only fit. This seems counter-intuitive, but is due to the fact that CC-like spectrum part of the fit still makes use of rate information. Although this has very little direct impact on f_s, the best fit values of $|\Delta m^2_{\mathrm{atm}}|$ and $\sin^2 2\theta_{23}$ change and this indirectly causes a shift in the best fit value of f_s. In a similar way the two shower energy scale uncertainties affect the rate-only fit.

- In many cases the shape-only fit is more susceptible to systematic uncertainty than the rate-only fit. The total systematic uncertainty is also larger. This underlines the importance of the rate measurement to the sterile neutrino analysis.

- In some cases the systematic shifts obtained with the rate- and shape-only fits are in the same direction. The systematic shifts with the combined fit are then intermediate between these values.

- In other cases the systematic shifts obtained with the rate- and shape-only fits pull in different directions. In these cases the competition between these two aspects tends to keep the systematic shifts in the combined fit under control.

6.4.8 Fit with variable normalisation

Given the large systematic error resulting from an incorrect estimate of the relative normalisation, it is interesting to consider the possibility of allowing this parameter to vary in the fit. This approach was used to handle the three main systematics in the ν_μ CC analysis [57].

This approach is somewhat akin to performing shape-only fits to both the CC-like and NC-like spectra. There is an important difference, however: here the normalisation factor must be common to both CC-like and NC-like spectra. This is an important distinction,as we can expect the CC-like spectrum to constrain the normalisation factor.

An additional constraint can be put on the normalisation, representing our prior belief about the accuracy of the normalisation used in the prediction. This is implemented by extending Equation (6.3) to include a Gaussian penalty term based upon the normalisation value:

$$\chi^2(\boldsymbol{\xi}, N) = 2\sum_{i=1}^{n_{\mathrm{bins}}} N y_i(\boldsymbol{\xi}) - n_i + n_i \ln \frac{n_i}{N y_i(\boldsymbol{\xi})} + \frac{(N-1)^2}{\sigma_N^2}. \quad (6.7)$$

The quantity σ_N is our estimate of how large the (1 standard deviation) uncertainty on the normalisation is, i.e. 0.04 or 4%.

The results of fitting with a variable normalisation are shown in Table 6.6. It can be seen that the error introduced by the relative normalisation systematic is somewhat reduced. However in the combined fit the total systematic sensitivity does not reduce by as much as might be hoped. The table shows the reason for this: when the normalisation is allowed to vary the analysis becomes more sensitive to the relative shower energy systematic.

The increased susceptibility of the floating normalisation fit to the relative shower energy uncertainty can be understood by looking again at Figure 6.6. The oscillation fit will attempt to 'correct' the distorted ratio shown in the figure. With the value of $|\Delta m^2_{\mathrm{atm}}|$ constrained by the CC-like spectrum there is no way an oscillation signal can affect the high energy tail so this aspect of the distortion is in some sense ignored by the fixed normalisation fit. But varying the normalisation allows a better fit in the high energy tail, at the cost of increasing the distortion at low energies. The increased low energy distortion is then reflected by a larger shift in f_{s}.

Although it increases the susceptibility to other sources of systematic error, including a variable normalisation in the fit allows it to take into account the largest source of systematic error, and therefore provides a better estimate of the overall error on the measurement of f_{s}. For the remainder of this chapter, and the

Systematic		Shift in f_s	Signed average	Rate only	Shape only
Relative normalisation	+	0.094	0.101	0.174	0.006
	−	0.054			
Relative shower energy	+	−0.071	−0.070	0.007	−0.177
	−	0.070			
Absolute shower energy	+	−0.022	−0.022	−0.048	0.057
	−	0.022			
CC background	+	−0.028	−0.028	−0.032	−0.023
	−	0.028			
NC background	+	−0.006	−0.006	0.002	−0.011
	−	0.005			
Far Detector cleaning 1	+	0.033	0.034	0.024	0.049
	−	−0.035			
Far Detector cleaning 2	+	0.037	0.038	0.043	0.033
	−	−0.039			
Bias			< 0.001	< 0.001	< 0.001
Total (sum in quadrature)			0.138	0.190	0.196

Table 6.6: Shifts in the best fit value of f_s due to various systematic errors, with the relative normalisation allowed to float under constraint.

measurement of f_s, we continue to use the variable normalisation fit.

6.5 Sensitivity with variable normalisation

Having decided to include the normalisation as a fit parameter, it is necessary to recalculate the sensitivities as the will necessarily be worse than with a fixed normalisation. Table 6.7 shows these recomputed values. In comparison with Table 6.2 the allowed parameter ranges do not increase by much, as the measurement is still dominated by statistical errors. A similar result can be seen in the joint sensitivity contours in Figure 6.13, which are not significantly larger than those shown in Figure 6.4.

6.5.1 Fitting without f_s

Before turning to result, there is another potential use for the NC-like spectrum. Because of the large number of ν_μ CC events in the NC-like sample it has some sensitivity to $\sin^2 2\theta_{23}$ and $|\Delta m^2_{\rm atm}|$. If we fit for f_s we lose much of this additional sensitivity because of the additional freedom in the fit, but if we assume the standard 3-neutrino picture is correct and $f_s = 0$ then the NC-like spectrum can be used (in principle) to improve the ν_μ disappearance analysis.

Parameter		Sensitivity		
		Rate	Shape	Full
f_s	+ only	0.552	0.609	0.423
$\lvert\Delta m^2_{atm}\rvert$	+	0.476	0.471	0.466
$/10^{-3}\,eV^2c^{-4}$	−	0.406	0.402	0.398
$\sin^2 2\theta_{23}$	− only	0.174	0.175	0.174

Table 6.7: Sensitivities for 1.2×10^{20} POT ,where the normalisation has been allowed to vary. Errors are obtained using $\Delta\chi^2 = +1$, and marginalising against the other parameters.

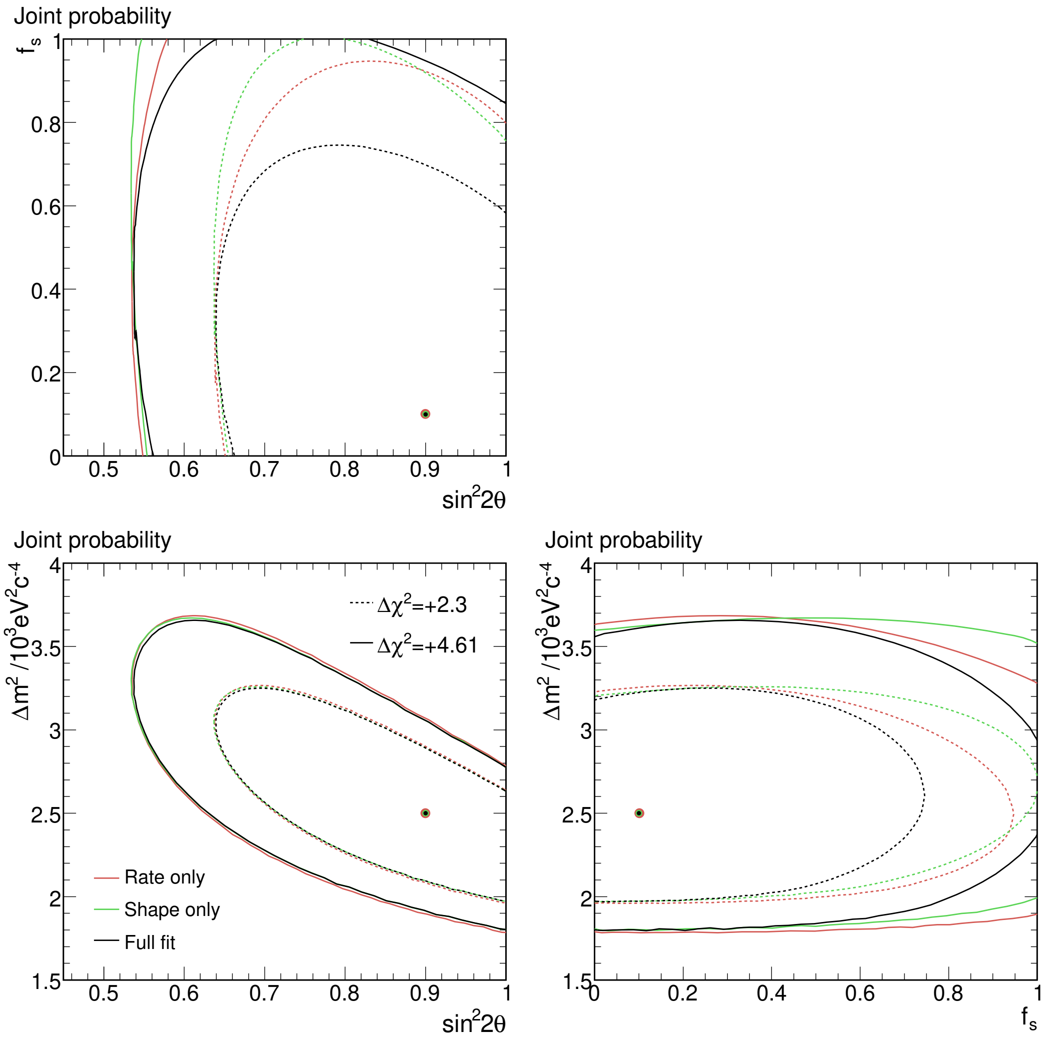

Figure 6.13: Joint sensitivity contours for 1.2×10^{20} POT, where the normalisation has been allowed to vary. The dashed and solid lines are 68% C.L. and 90% C.L. contours respectively.

A simple test of this idea is shown. The input values are taken from the latest MINOS ν_μ disappearance analysis [121]: $|\Delta m^2_{\rm atm}| = 2.38\times10^{-3}\,\rm eV^2c^{-4}$, $\sin^2 2\theta_{23} = 1$, with a similar Far Detector exposure of 3.0×10^{20} POT. The resulting joint sensitivity contours are shown in Figure 6.14. The improvement in the sensitivity is around 3%. So from a statistical standpoint the use of the NC-like events is not a great benefit to the ν_μ disappearance analysis. It should also be noted that the selection of events is not optimised for measuring $\sin^2 2\theta_{23}$ and $|\Delta m^2_{\rm atm}|$, so making use of the NC-like spectrum will quite possibly have even less impact on a dedicated ν_μ disappearance analysis.

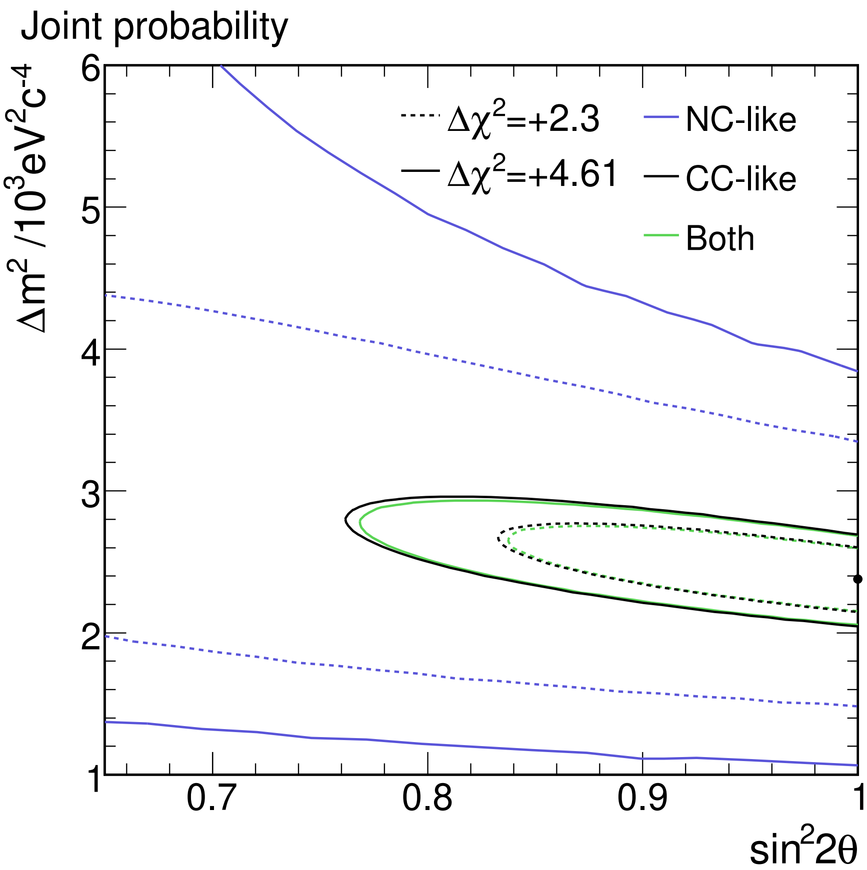

Figure 6.14: Joint sensitivity contours for 3-neutrino model ($f_{\rm s} = 0$). Contours making use of the CC-like, NC-like or both energy spectra are shown. The best fit point for the 'normal' CC-like only analysis is shown with a black circle.

The merit of making use of the NC-like events will therefore depend on how they impact the systematic errors. It is possible that this could go in either direction. If the systematic uncertainties affecting the NC-like spectrum are well understood then the use of such events can help constrain common uncertainties, such as the relative normalisation and shower calibration. On the other hand, if the additional systematic uncertainties are not well understood then the the NC-like spectrum will introduce unwanted biases and negatively impact the analysis. Doing such an analysis properly is a significant task and is beyond the scope of this thesis.

6.6 Analysis of the first year of MINOS data

The Far Detector data set used for this analysis is was taken between June 2005 and February 2006 and is essentially the same as used in the ν_μ CC analysis described in [57, 121]. Basic data quality cuts such as having the beam and detector operating normally are implemented in the software framework and are largely unchanged from that analysis. The one exception is the use of slightly tighter cuts on the NuMI horn current. Because of this the data set represents a marginally (around 2%) lower exposure of 1.245×10^{20} POT.

The Near Detector data set is taken from a similar period, although the livetime of the Near Detector is somewhat lower, as it does not limit the sensitivity of the experiment. As a result the Near Detector data set corresponds to a slightly lower beam exposure, 1.23×10^{20} POT. This is somewhat lower than the available 1.33×10^{20} POT equivalent of simulated Near Detector events, and more than enough to ensure the statistical error on the Near Detector data/simulation ratio is not a significant contribution to the overall error.

Figure 6.15 shows the stability of the Near Detector ν_μ CC-like spectrum over the data taking period. Given this lack of variation, which reflects the stability of the beam flux, the small difference in exposure periods at each detector will not have a significant systematic effect.

6.6.1 Near Detector data

As noted previously, the Far/Near extrapolation effectively multiplies the Far Detector simulation by the data/simulation spectrum ratio from the Near Detector. These ratios give a measure of how much the simulation differs from data and are shown in Figure 6.16. The CC-like ratio departs furthest from unity in the region corresponding to the rising part of the peak, suggesting that the errors on the neutrino flux are possibly a large component of the remaining difference. The fact that the NC-like distribution does not show such divergences is probably due to the smearing effect of observing an unknown fraction of the neutrino energy.

6.6.2 Far detector spectra and fits

The Far Detector CC-like and NC-like spectra are shown in Figure 6.17 and Figure 6.18 respectively. The number of events in data and simulation is recorded in Table 6.8.

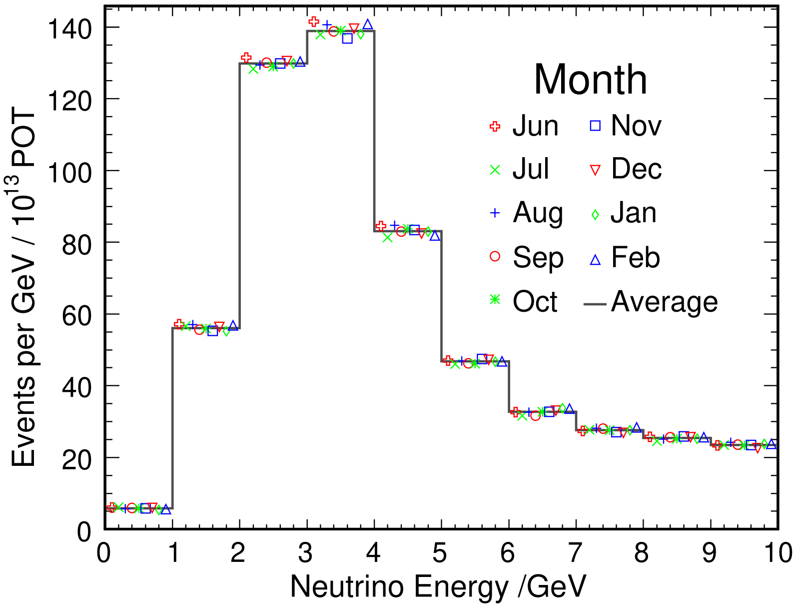

Figure 6.15: Stability of the Near Detector CC-like spectrum, modified from [121]. The lack of month-to-month variation shows that both the beam and detector were stable over the run period.

	Predicted (no osc)		Predicted (best fit)		Data	
	CC	NC	CC	NC	CC	NC
0–15 GeV	284.1	160.1	189.3	137.4	184	133
15–45 GeV	69.1	15.3	67.9	15.4	73	15
> 45 GeV	10.7	1.0	10.5	1.0	9	0
Total	363.8	176.3	267.7	153.8	266	148

Table 6.8: Numbers of CC-like and NC-like events in various energy ranges.

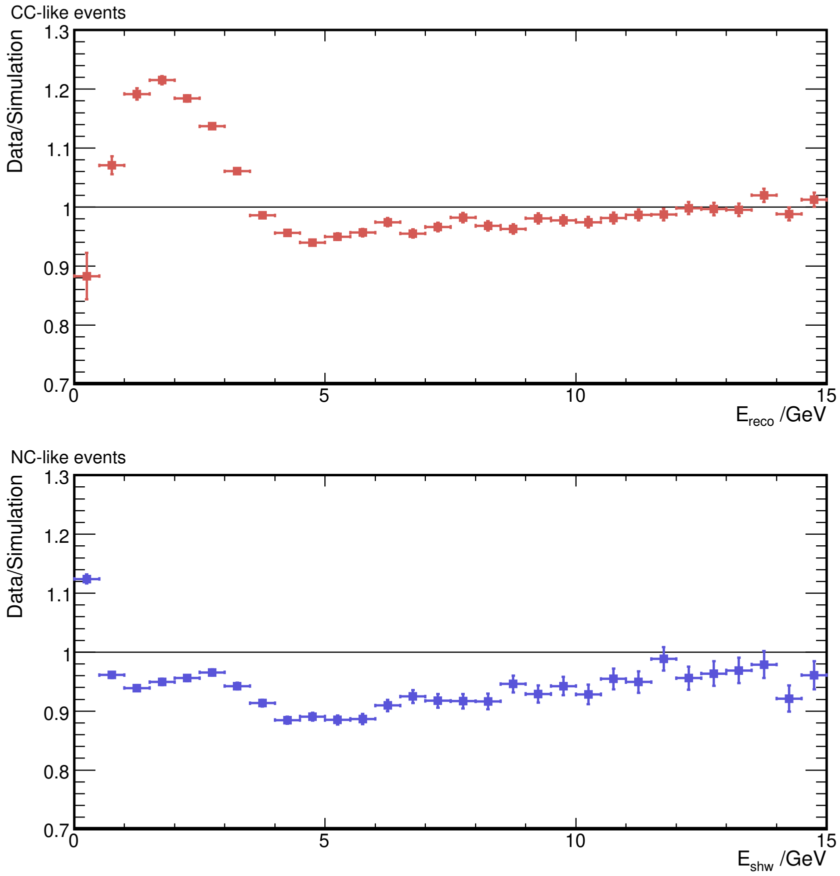

Figure 6.16: Data/simulation ratios of Near Detector reconstructed energy spectra. These ratios are used by the Far/Near method to weight the predicted Far Detector energy spectra.

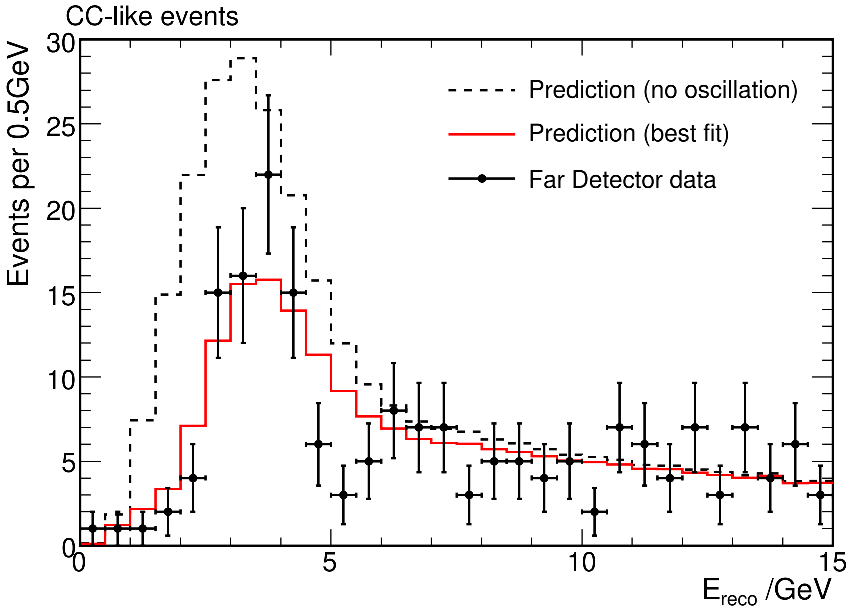

Figure 6.17: Reconstructed energy spectrum of CC-like events in the Far Detector. The red line shows the best fit to both CC-like and NC-like spectra, with the parameters given in the text.

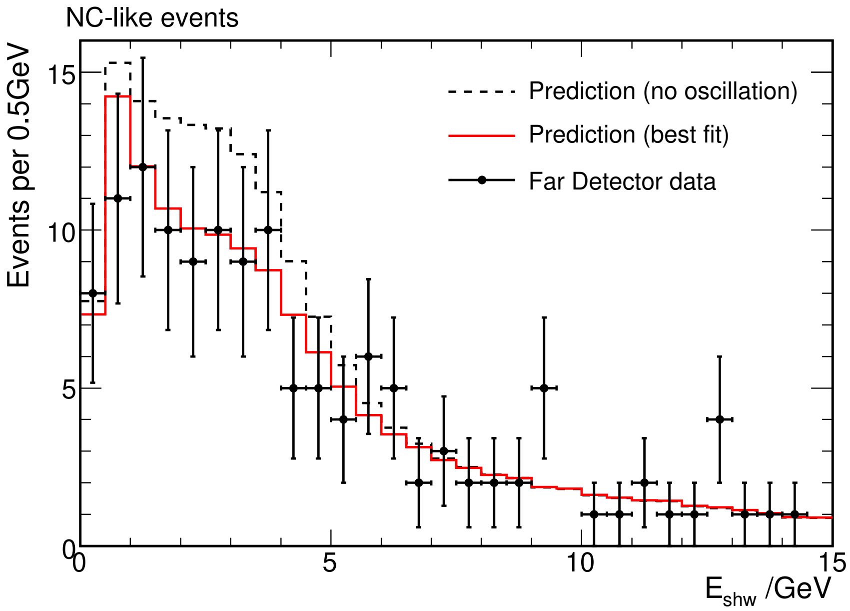

Figure 6.18: Reconstructed energy spectrum of NC-like events in the Far Detector. The red line shows the best fit to both CC-like and NC-like spectra, with the parameters given in the text.

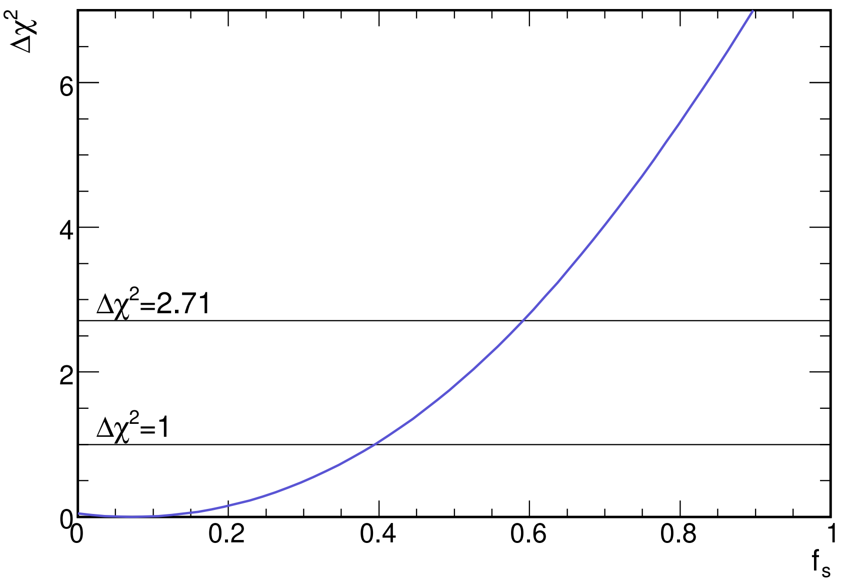

Figure 6.19: Values of $\Delta\chi^2$, as a function of f_s and marginalised against all other parameters.

Parameter	Best fit and errors		
	Rate	Shape	Full
f_s	$0.05^{+0.42}$	$0.03^{+0.47}$	$0.07^{+0.32}$
$\lvert\Delta m^2_{\mathrm{atm}}\rvert/10^{-3}\,\mathrm{eV}^2\mathrm{c}^{-4}$	$2.67^{+0.22}_{-0.21}$	$2.67^{+0.22}_{-0.21}$	$2.67^{+0.22}_{-0.21}$
$\sin^2 2\theta_{23}$	$1.00_{-0.07}$	$1.00_{-0.07}$	$1.00_{-0.07}$
Normalisation	+0.1%	+0.2%	+0.1%

Table 6.9: Best fit parameters and (single parameter) 68% errors from NC analysis of the first year of MINOS Far Detector data.

The full fit to both CC-like and NC-like spectra gives $f_s = 0.072^{+0.323}$ (68% C.L. from $\Delta\chi^2 = 1$). At 90% C.L. ($\Delta\chi^2 = 2.71$), the limit is $f_s < 0.591$. In Figure 6.19 the minimum value of $\Delta\chi^2$ is shown for different values of f_s. As can be seen from the figure, the data are quite consistent with the standard three neutrino model ($f_s = 0$) and do not require a sterile neutrino. The two-parameter joint confidence regions for are shown in Figure 6.20. Best fit parameters with associated errors are given in Table 6.9

6.6.3 Eventual sensitivity

As the measurement of f_s is still statistically limited, improvements will come as the Far Detector exposure increases. So far, the total exposure is around 3×10^{20} POT, and the final exposure will probably be something in the region of

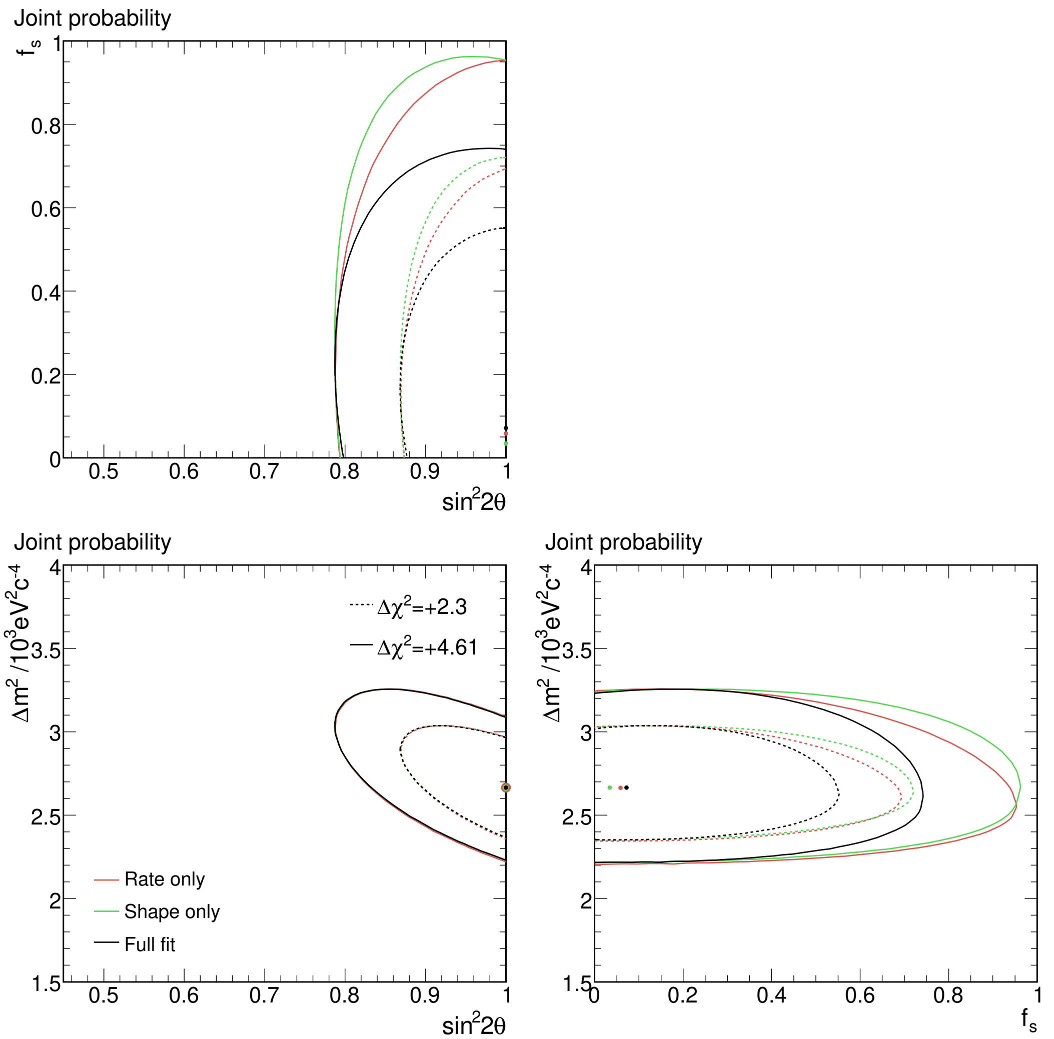

Figure 6.20: Joint confidence regions from NC analysis of the first year of MINOS Far Detector data. The dashed and solid lines are 68% C.L. and 90% C.L. contours respectively.

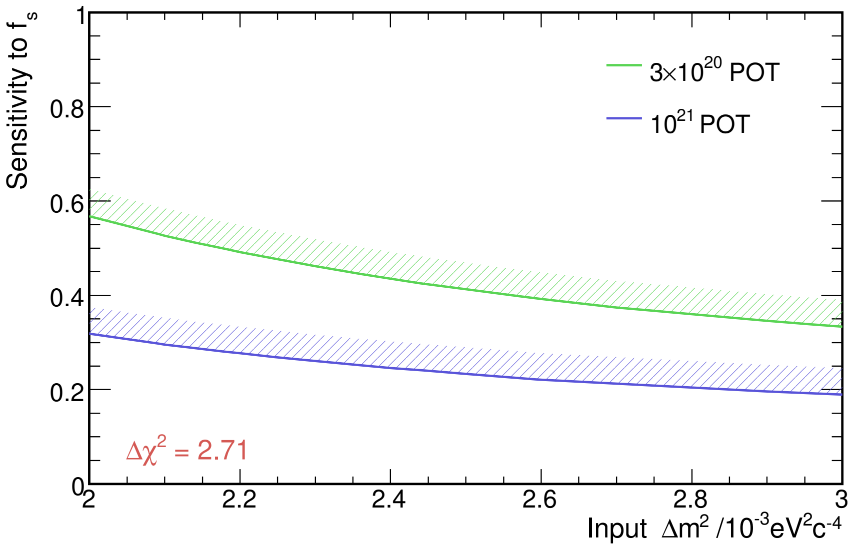

Figure 6.21: Sensitivity to f_s for various values of $|\Delta m^2_\mathrm{atm}|$, with increased Far Detector exposures.

10^{21} POT.

If the analysis were to remain unchanged, with a 4% normalisation uncertainty as the only source of systematic error, the eventual limit on f_s will improve as shown in Figure 6.21. The figure shows how the sensitivity is dependent on the true value of $|\Delta m^2_\mathrm{atm}|$, assuming that $\sin^2 2\theta_{23} = 1$ and $f_s = 0$. As the number of ν_μ oscillation to ν_s is proportional to $f_\mathrm{s} \sin^2 2\theta_{23}$, the sensitivity to f_s goes roughly as $1/\sin^2 2\theta_{23}$, although this ignores the effects of a reduced ν_μ disappearance probability.

Chapter 7

Summary

MINOS is currently the only accelerator neutrino experiment able to probe the atmospheric sector of neutrino oscillations. The design of the experiment is optimised for studying ν_μ disappearance, and makes it sensitive to $|\Delta m^2_{\mathrm{atm}}|$. MINOS should provide the most precise measurement of this quantity within the next few years. The experiment also provides the most sensitive cross check of the $\sin^2 2\theta_{23}$ measurement from Super-Kamiokande.

One of the major concerns with this kind of experiment are the systematic errors introduced by uncertainties on the neutrino production and interaction rates. MINOS is explicitly designed to minimise these uncertainties by having a Near Detector with similar construction to the Far Detector. The similarity of the detectors makes it possible to use a diverse range of methods to reduce these systematic errors. The optimal approach depends on which uncertainties are largest.

The experiment can be used to look for exotic models of neutrino disappearance. This thesis concentrates on testing for oscillations into a fourth, sterile, flavour eigenstate. This can be done by studying events resulting from neutral current neutrino interactions. Because the only visible part of these interactions is the recoiling hadronic system it is more challenging than the ν_μ disappearance measurements.

The first step in this analysis is to distinguish neutrino events from the non-neutrino background. In the Far Detector these come mainly from cosmic rays and instrumental backgrounds. The absence of any distinctive structure in the neutral current events means that many of the backgrounds arise as artifacts of the reconstruction, and a range of approaches must be used. Estimates are made of the systematic uncertainties resulting from this selection.

The events are divided into CC- and NC-like samples, using a straightforward method. The most important discriminating characteristic is the presence of a muon track. Refinements are made based on the direction of curvature of the track, and identification of tracks that exit the detector.

Using the CC-like and NC-like samples, a model of sterile oscillations is tested.

The model extends the standard two-neutrino formula for atmospheric-scale oscillations by the inclusion of a parameter f_{s}, the sterile fraction, that is zero for pure $\nu_\mu \to \nu_\tau$ oscillations and one for pure $\nu_\mu \to \nu_s$. The parameter f_{s} can be cast in terms of an extension of the PMNS mixing matrix, or regarded simply as an interpolation parameter between these two pure models.

Various systematic uncertainties affecting the measurement of f_{s} are considered. The most important is the ratio of event rates at the Near and Far detectors. This is not surprising, given that around half the sensitivity of the measurement comes from counting the number of NC-like events.

Analysis of the first year of MINOS data yields a best fit value of $f_{\mathrm{s}} = 0.07^{+0.32}$. The standard three-neutrino model ($f_{\mathrm{s}} = 0$) is entirely consistent with these data. This measurement is based on a data set corresponding to 1.24×10^{20} POT. The final data set will correspond to something closer to 10^{21} POT, so should provide a tighter constraint.

Appendix A

Tuning the NC selection cuts

A.1 Defining a figure of merit

As alluded to in Chapter 4, it is usually[1] necessary to find a balance between a tight cut that removes most of the background events and a looser cut that retains as many of the signal events as possible. Ideally this could be estimated by simulating many pseudo-experiments with different cut positions and choosing the one that is most sensitive to the parameter of interest, but this is time consuming. A faster approach is to find a figure of merit that gives an indication of the overall sensitivity. As the intended analysis is at heart a counting experiment we look to minimise σ_s/s, the fractional error on the number of signal events seen. This is equivalent to maximising its inverse-squared s^2/σ_s^2, which is more convenient to evaluate. But in data only n, the total number of selected events is known. Although it is not done explicitly, the number of signal events is effectively estimated by subtracting b, the number of background events predicted by the simulation, so

$$s = n - b. \tag{A.1}$$

Turning to the errors, $\sigma_n = \sqrt{n}$ is the statistical error on the total number of events selected and σ_b is the error on the background estimate. It is assumed that systematic errors on this estimate are small, so that $\sigma_b = \sqrt{b}$. It is also assumed the error on the background estimate is independent of the total number of events selected so that

$$\sigma_s^2 = \sigma_n^2 + \sigma_b^2 = n + b. \tag{A.2}$$

From the definitions in Chapter 4, $b = n(1 - \mathbb{P})$ and $s = \mathcal{E}S = \mathbb{P}n$, so

$$\frac{s^2}{\sigma_s^2} = \frac{\mathcal{E}S\ \mathbb{P}n}{(2 - \mathbb{P})\,n}. \tag{A.3}$$

[1]If there are enough signal events then systematic errors may be a bigger concern and it will be more important to reduce these.

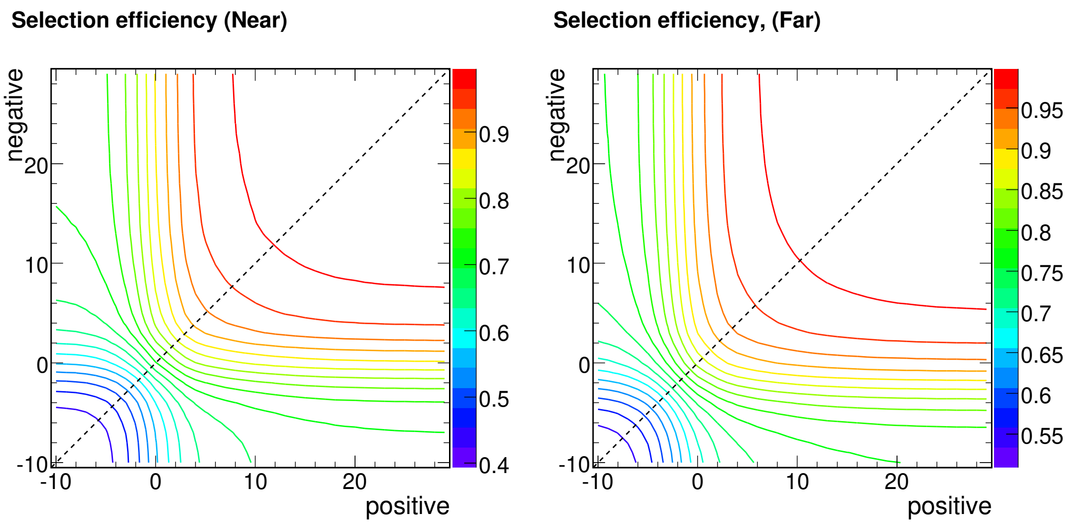

Figure A.1: Contours of efficiency of the NC selection cuts in the Near Detector, and the Far Detector. The diagonal lines show what can be achieved using the same cut for both positively and negatively curving tracks.

S is independent of the selection, so the figure of merit to be maximised is [122]:

$$\frac{\mathcal{E}\mathbb{P}}{(2-\mathbb{P})}. \tag{A.4}$$

It is interesting to compare this to two common figures of merit: $s^2/n \sim \mathcal{E}\mathbb{P}$ and $s^2/b \sim \mathcal{E}\mathbb{P}/(1-\mathbb{P})$. The figure of merit used is intermediate between these two in terms of the importance of the selection purity, and reduces to these in the limit of small backgrounds and zero signal events respectively.

A.2 Tuning procedure

The selection described in Section 4.3.1 can be tuned by varying the position of the two cuts on δ_{ts}. The track containment cut is always applied. It is possible to evaluate the efficiency and purity for the positively and negatively curving subsamples individually, but these are not directly related to the overall $\mathcal{E}$ and $\mathbb{P}$, so both cuts must be evaluated together, with the shower only events included. The efficiency contours are shown in Figure A.1, purity contours in Figure A.2 and figure of merit in Figure A.3. The optimal positions (where the figure of merit is maximised) in the Far Detector are $\delta_{\text{ts}} = 2$ for negatively curving tracks and $\delta_{\text{ts}} = 13$ for positively curving tracks. Differences in the neutrino spectrum and detector acceptance at the Near Detector lead to slightly different optimal cuts positions: $\delta_{\text{ts}} = 4$ for negatively curving tracks and $\delta_{\text{ts}} = 11$ for positively curving tracks.

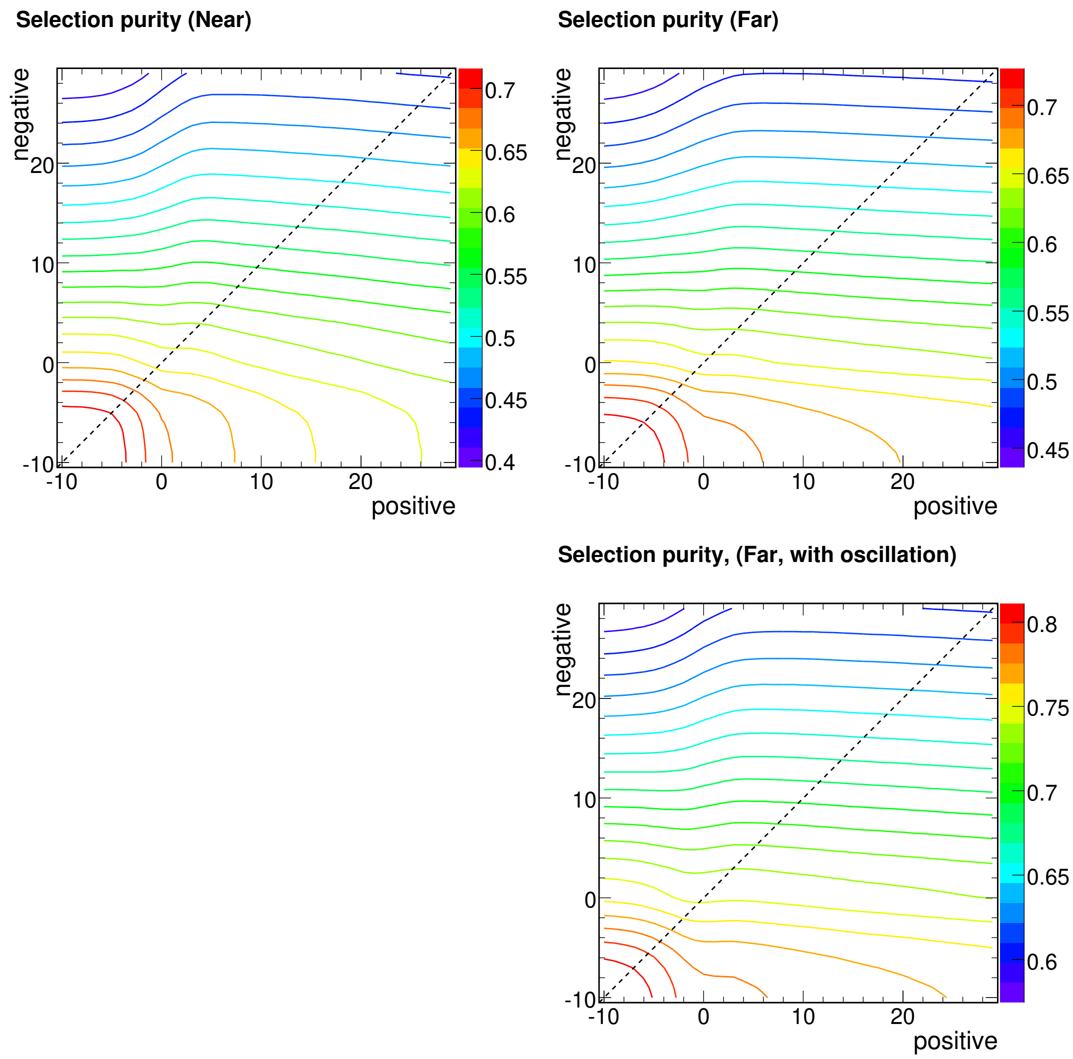

Figure A.2: Contours of selection purity after applying the NC selection cuts in the Near Detector, and the Far Detector. The lower plot shows the purity if oscillations are applied to the charged current events. The diagonal lines show what can be achieved using the same cut for both positively and negatively curving tracks.

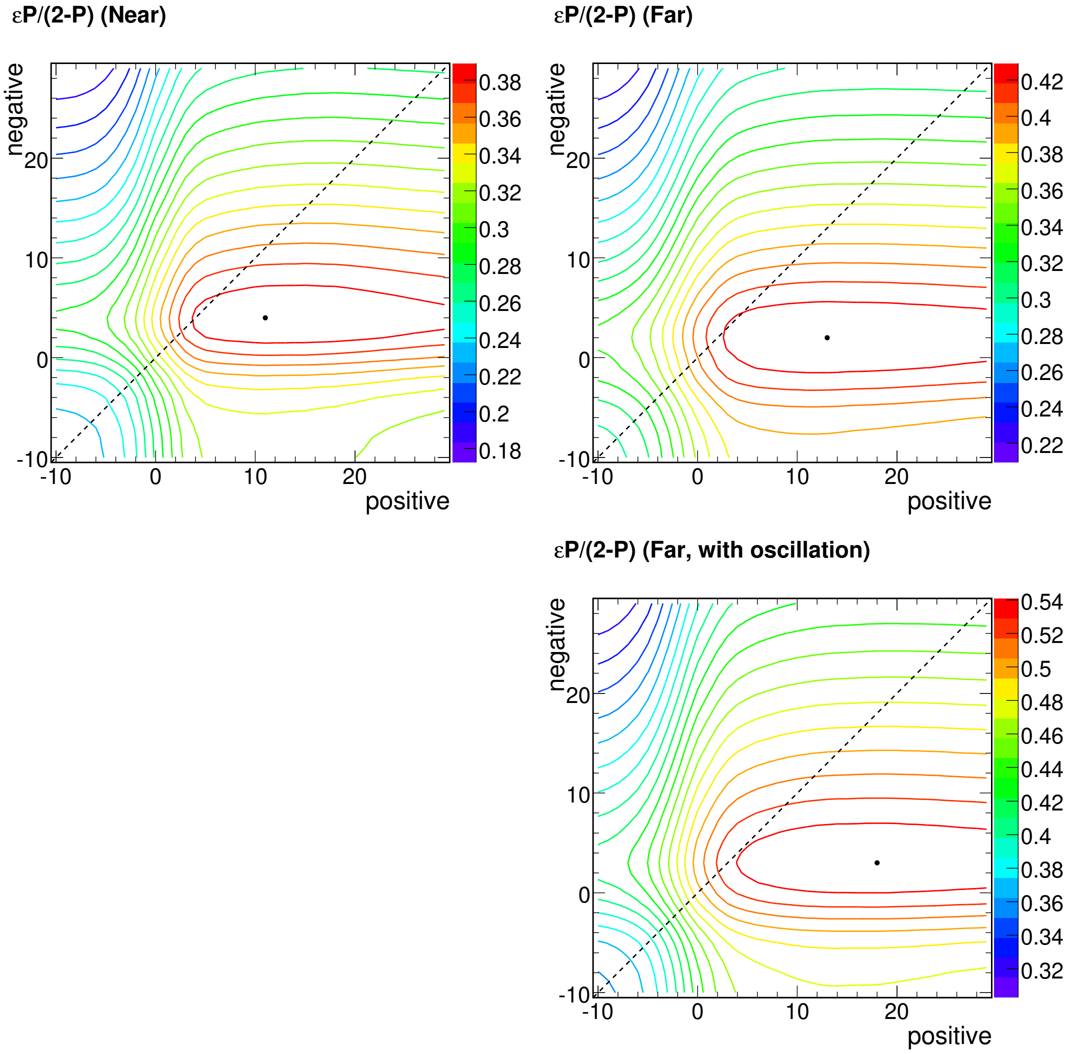

Figure A.3: Contours of selection figure of merit of the NC selection cuts in the Near Detector, and the Far Detector. The lower plot shows the figure of merit if oscillations are applied to the charged current events. The diagonal lines show what can be achieved using the same cut for both positively and negatively curving tracks.

In the context of this sterile neutrino analysis, the phenomena of ν_μ disappearance is assumed, and can be expected to affect the charged current background. Given this, one could also consider the evaluating the purity and figure of merit when the charged current background has been oscillated[2]. As there is less background to be removed, the resulting cut positions are slightly looser. The cut positions derived this way are at $\delta_{\mathrm{ts}} = 3$ for negatively curving tracks, and $\delta_{\mathrm{ts}} = 18$ for positively curving tracks.

Finally, the approach of cutting separately on events with positively and negatively curving tracks can be be equally well applied (and tuned) for any selection that uses a single 'separation parameter' This has been examined for the Neural Network based selection described in [91] and a Multivariate Discriminant Analysis based selection, similar to that described in [123]. Similar improvements are observed with these comparatively sophisticated approaches.

[2]The efficiency is not affected as it only depends on the number of neutral current events

Bibliography

[1] W. Pauli, Letter, addressed to participants of Tübingen conference on radioactivity, 1930, Available from the CERN Document server: `http://cdsweb.cern.ch/record/83282`.

[2] A. Franklin, *Are there really neutrinos? An evidential history* (Perseus, 2000).

[3] E. Fermi, Z. Phys. **88**, 161 (1934).

[4] F. L. Wilson, Am. J. Phys. **36**, 1150 (1968).

[5] S. Sakata and T. Inoue, Prog. Theor. Phys. **1**, 143 (1946), Translated from J.Phys.Math.Soc.Jpn. 16, 232 (1942).

[6] R. E. Marshak and H. A. Bethe, Phys. Rev. **72**, 506 (1947).

[7] H. Yukawa, Proc. Phys. Math. Soc. Jap. **17**, 48 (1935).

[8] C. M. G. Lattes, M. Muirhead, G. P. S. Occhialini, and C. F. Powell, Nature **159**, 694 (1947).

[9] F. Reines and C. L. Cowan, Phys. Rev. **92**, 830 (1953).

[10] R. Davis, Phys. Rev. **97**, 766 (1955).

[11] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Science **124**, 103 (1956).

F. Reines, C. L. Cowan, F. B. Harrison, A. D. McGuire, and H. W. Kruse, Phys. Rev. **117**, 159 (1960).

[12] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Phys. Rev. **109**, 1015 (1958).

[13] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962).

[14] T. D. Lee and C. N. Yang, Phys. Rev. **104**, 254 (1956).

[15] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Phys. Rev. **105**, 1413 (1957).

[16] R. L. Garwin, L. M. Lederman, and M. Weinrich, Phys. Rev. **105**, 1415 (1957).

[17] T. D. Lee, R. Oehme, and C. N. Yang, Phys. Rev. **106**, 340 (1957).

[18] E. C. G. Sudarshan and R. E. Marshak, Phys. Rev. **109**, 1860 (1958).

[19] R. P. Feynman and M. Gell-Mann, Phys. Rev. **109**, 193 (1958).

[20] S. L. Glashow, Nucl. Phys. **22**, 579 (1961).
A. Salam and J. C. Ward, Phys. Lett. **13**, 168 (1964).
S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967).
G. 't Hooft, Nucl. Phys. **B35**, 167 (1971).

[21] Gargamelle Neutrino Collaboration: F. J. Hasert *et al.*, Phys. Lett. **B46**, 121 (1973).

[22] Gargamelle Neutrino Collaboration: F. J. Hasert *et al.*, Phys. Lett. **B46**, 138 (1973).

[23] UA1 Collaboration: G. Arnison *et al.*, Phys. Lett. **B126**, 398 (1983).
UA2 Collaboration: P. Bagnaia *et al.*, Phys. Lett. **B129**, 130 (1983).

[24] M. Gell-Mann and A. Pais, Phys. Rev. **97**, 1387 (1955).

[25] A. Pais and O. Piccioni, Phys. Rev. **100**, 1487 (1955).

[26] B. Pontecorvo, Sov. Phys. JETP **6**, 429 (1957).

[27] B. Pontecorvo, Sov. Phys. JETP **7**, 172 (1958).

[28] B. Pontecorvo, Sov. Phys. JETP **26**, 984 (1968).

[29] N. Cabibbo, Phys. Rev. Lett. **10**, 531 (1963).

[30] Z. Maki, M. Nakagawa, and S. Sakata, Prog. Theor. Phys. **28**, 870 (1962).

[31] S. L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. **D2**, 1285 (1970).

[32] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).

[33] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Lett. **13**, 138 (1964).

[34] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).

[35] Kamiokande-II Collaboration: K. S. Hirata *et al.*, Phys. Rev. Lett. **63**, 16 (1989).

[36] SAGE Collaboration: A. I. Abazov *et al.*, Phys. Rev. Lett. **67**, 3332 (1991).

[37] GALLEX Collaboration: P. Anselmann *et al.*, Phys. Lett. **B285**, 376 (1992).

[38] Kamiokande-II Collaboration: K. S. Hirata *et al.*, Phys. Lett. **B205**, 416 (1988).

[39] IMB Collaboration: T. J. Haines *et al.*, Phys. Rev. Lett. **57**, 1986 (1986).
IMB Collaboration: D. Casper *et al.*, Phys. Rev. Lett. **66**, 2561 (1991).

[40] Fréjus Collaboration: C. Berger *et al.*, Phys. Lett. **B227**, 489 (1989).

[41] NUSEX Collaboration: M. Aglietta *et al.*, Europhys. Lett. **8**, 611 (1989).

[42] SNO Collaboration: Q. R. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 (2002), nucl-ex/0204008.

[43] SNO Collaboration: Q. R. Ahmad *et al.*, Phys. Rev. Lett. **87**, 071301 (2001), nucl-ex/0106015.

[44] Super-Kamiokande Collaboration: S. Fukuda *et al.*, Phys. Rev. Lett. **86**, 5651 (2001), hep-ex/0103032.

[45] L. Wolfenstein, Phys. Rev. **D17**, 2369 (1978).

[46] S. P. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).

[47] CHOOZ Collaboration: M. Apollonio *et al.*, Eur. Phys. J. **C27**, 331 (2003), hep-ex/0301017.

[48] F. Boehm *et al.*, Phys. Rev. **D64**, 112001 (2001), hep-ex/0107009.

[49] KamLAND Collaboration: T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005), hep-ex/0406035.

[50] Soudan 2 Collaboration: W. W. M. Allison *et al.*, Phys. Lett. **B391**, 491 (1997), hep-ex/9611007.

[51] Super-Kamiokande Collaboration: J. Hosaka *et al.*, Phys. Rev. **D74**, 032002 (2006), hep-ex/0604011.

[52] Super-Kamiokande Collaboration: Y. Ashie *et al.*, Phys. Rev. Lett. **93**, 101801 (2004), hep-ex/0404034.

[53] S. D. Biller and N. A. Jelley, private communication, 2007.

[54] W.-M. Yao *et al.*, Journal of Physics G **33**, 1+ (2006).

[55] MACRO Collaboration: M. Ambrosio *et al.*, Eur. Phys. J. **C36**, 323 (2004).

[56] K2K Collaboration: M. H. Ahn *et al.*, Phys. Rev. **D74**, 072003 (2006), hep-ex/0606032.

[57] MINOS Collaboration: D. G. Michael *et al.*, Phys. Rev. Lett. **97**, 191801 (2006), hep-ex/0607088.

[58] T. Schwetz, Phys. Scripta **T127**, 1 (2006), hep-ph/0606060.

[59] LSND Collaboration: A. Aguilar *et al.*, Phys. Rev. **D64**, 112007 (2001), hep-ex/0104049.

[60] KARMEN Collaboration: B. Armbruster *et al.*, Phys. Rev. **D65**, 112001 (2002), hep-ex/0203021.

[61] NOMAD Collaboration: P. Astier *et al.*, Phys. Lett. **B570**, 19 (2003), hep-ex/0306037.

[62] Y. Declais *et al.*, Nucl. Phys. **B434**, 503 (1995).

[63] E. D. Church, K. Eitel, G. B. Mills, and M. Steidl, Phys. Rev. **D66**, 013001 (2002), hep-ex/0203023.

[64] MiniBooNE Collaboration: A. A. Aguilar-Arevalo *et al.*, (2007), arXiv:0704.1500 [hep-ex].

[65] P. Minkowski, Phys. Lett. **B67**, 421 (1977).

T. Yanagida, In Proceedings of the Workshop on the Baryon Number of the Universe and Unified Theories, Tsukuba, Japan, 13–14 Feb 1979.

P. Ramond, (1979), hep-ph/9809459, Invited talk given at Sanibel Symposium, Palm Coast, Fla., Feb 25 – Mar 2, 1979.

[66] K. Zuber, Acta Phys. Polon. **B37**, 1905 (2006), nucl-ex/0610007.

[67] J. Rich, Phys. Rev. **D48**, 4318 (1993).

[68] B. Kayser, Phys. Rev. **D24**, 110 (1981).

[69] MINOS Collaboration: P. Adamson *et al.*, Phys. Rev. **D76**, 072005 (2007), arXiv:0706.0437 [hep-ex].

[70] L.-L. Chau and W.-Y. Keung, Phys. Rev. Lett. **53**, 1802 (1984).

[71] E. K. Akhmedov, (2006), hep-ph/0610064, Presented at 22nd International Conference on Neutrino Physics and Astrophysics (Neutrino06), Santa Fe, NM, 13–19 June 2006.

[72] V. D. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, Phys. Rev. Lett. **82**, 2640 (1999), astro-ph/9810121.

V. D. Barger *et al.*, Phys. Lett. **B462**, 109 (1999), hep-ph/9907421.

[73] M. Blennow, T. Ohlsson, and W. Winter, JHEP **06**, 049 (2005), hep-ph/0502147.

[74] E. Lisi, A. Marrone, and D. Montanino, Phys. Rev. Lett. **85**, 1166 (2000), hep-ph/0002053.

[75] The ALEPH, DELPHI, L3, OPAL, SLD Collaborations, the LEP Electroweak Working Group, the SLD Electroweak and Heavy Flavour Groups, Phys. Rept. **427**, 257 (2006), hep-ex/0509008.

[76] Super-Kamiokande Collaboration: A. Habig, (2001), hep-ex/0106025, Contributed to 27th International Cosmic Ray Conferences (ICRC 2001), Hamburg, Germany, 7–15 Aug 2001.

Super-Kamiokande Collaboration: T. Nakaya, (2002), hep-ex/0209036.

[77] M. Maltoni, T. Schwetz, M. A. Tortola, and J. W. F. Valle, Nucl. Phys. **B643**, 321 (2002), hep-ph/0207157.

M. Maltoni, T. Schwetz, M. A. Tortola, and J. W. F. Valle, Nucl. Phys. Proc. Suppl. **114**, 203 (2003), hep-ph/0209368.

[78] M. Maltoni and T. Schwetz, (2007), arXiv:0705.0107 [hep-ph].

[79] Double Chooz Collaboration: F. Ardellier *et al.*, (2006), hep-ex/0606025.

[80] Daya Bay Collaboration: X. Guo *et al.*, (2007), hep-ex/0701029.

[81] M. Lindner, Springer Tracts Mod. Phys. **190**, 209 (2003), hep-ph/0209083.

[82] NOνA Collaboration: D. S. Ayres *et al.*, (2004), hep-ex/0503053.

[83] T2K Collaboration: Y. Itow *et al.*, (2001), hep-ex/0106019.

[84] MINOS Collaboration: D. G. Michael *et al.*, To be published in Nucl. Instrum. Meth. A.

[85] Hamamatsu Photonics K.K., 314-5 Shimokanzo, Toyooka-mura, Iwata-gun, Shizuoka-ken, 438-0193, Japan. Multianode photomultiplier tube R5900-00-M64 Data sheet.

[86] S. Eilerts, T. Koslowski, and K. Lang, Effects of magnetic fields on Hamamatsu R5900-00-M16 PMTs, NuMI-L-568, 1999.

M. A. Barker and A. Weber, A more rigorous description of the influence of magnetic fields on Hamamatsu M64 photomultipliers in relation to the MINOS experiment, NuMI-L-650, 2000.

[87] N. Tagg *et al.*, Nucl. Instrum. Meth. **A539**, 668 (2005), physics/0408055.
K. Lang *et al.*, Nucl. Instrum. Meth. **A545**, 852 (2005).

[88] T. Cundiff *et al.*, IEEE Trans. Nucl. Sci. **53**, 1347 (2006).

[89] Ideas ASA, PO Box 1, 1330 Fornebu, Norway.

[90] J. Huth *et al.*, Prepared for 5th Workshop on Electronics for the LHC Experiments (LEB 99), Snowmass, Colorado, 20-24 Sep 1999.

[91] T. M. Raufer, *A study of neutrino oscillations in MINOS*, DPhil thesis, University of Oxford, 2007.

[92] S. Avvakumov *et al.*, Nucl. Instrum. Meth. **A545**, 145 (2005).

[93] Creative Electronic Systems (CES) Corp., The PVIC and VIC family data sheet (2004). PVIC stands for PCI Vertical Inter-Connection.

[94] A. Belias *et al.*, IEEE Trans. Nucl. Sci. **51**, 451 (2004).

[95] P. Adamson *et al.*, Nucl. Instrum. Meth. **A492**, 325 (2002), hep-ex/0204021.

[96] J. J. Hartnell, *Measurement of the calorimetric energy scale in MINOS*, DPhil thesis, University of Oxford, 2005.

[97] M. A. Kordosky, *Hadronic Interactions in the MINOS Detectors*, PhD thesis, University of Texas at Austin, 2004.

[98] P. L. Vahle, *Electromagnetic Interactions in the MINOS Detectors*, PhD thesis, University of Texas at Austin, 2004.

[99] A. Cabrera, *Systematic Comparison of the MINOS Near and Far Detector Readout Systems*, DPhil thesis, University of Oxford, 2005.

[100] C. Zeitnitz and T. A. Gabriel, Nucl. Instrum. Meth. **A349**, 106 (1994).

[101] H. Gallagher, Nucl. Phys. Proc. Suppl. **112**, 188 (2002).

[102] MINERνA Collaboration: D. Drakoulakos *et al.*, (2004), hep-ex/0405002.

[103] GEANT—detector description and simulation tool, CERN Program Library, Long Writeup, W5013,
`http://wwwasd.web.cern.ch/wwwasd/geant/index.html`.

[104] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, CERN-2005-010.

[105] J. R. Klein and A. Roodman, Ann. Rev. Nucl. Part. Sci. **55**, 141 (2005).

[106] MINOS Collaboration: P. Adamson *et al.*, arXiv:0705.3815 [hep-ex].

[107] MINOS Collaboration: M. Kordosky, AIP Conf. Proc. **896**, 185 (2007), hep-ex/0701009.

[108] R. Wigmans, *Calorimetry* (Oxford University Press, 2000).

[109] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, Eur. Phys. J. **C20**, 13 (2001), hep-ph/0101163.

[110] S. Kopp, Ž. Pavlović, and P. Vahle, Presentation at NuFact 06, UC Irvine, USA, 24–30 Aug 2006.

[111] MINOS Collaboration: P. Adamson *et al.*, arXiv:0711.0769 [hep-ex], Submitted for publication in Phys. Rev. D.

[112] J. J. Evans, private communication, 2007.

[113] S. Baker and R. D. Cousins, Nucl. Instrum. Meth. **A221**, 437 (1984).

[114] J. G. Heinrich, The log likelihood ratio of the Poisson distribution for small μ, CDF/MEMO/CDF/CDFR/5718, 2001, `http://www-cdf.fnal.gov/publications/cdf5718_loglikeratv2.ps.gz`.

[115] F. James and M. Roos, Comput. Phys. Commun. **10**, 343 (1975).

[116] A. J. G. Weber, f_{sterile} systematics for separation comparison, MINOS-doc-3479, 2007.

[117] J. K. Nelson, Notes on the normalisation systematic, MINOS-doc-2106, 2006.

[118] MINOS Calibration Working Group, Calibration position paper, MINOS-doc-3137, 2007.

[119] J. J. Hartnell, private communication, 2007.

[120] M. A. Kordosky, H. R. Gallagher, and S. Dytman, Estimated uncertainty on the shower energy scale in muon neutrino CC events, MINOS-doc-3362, 2007.

[121] MINOS Collaboration, (2007), arXiv:0708.1495 [hep-ex], Presented at 23rd International Symposium on Lepton-Photon Interactions at High Energy (LP07), Daegu, Korea, 13–18 Aug 2007.

[122] T. M. Raufer, private communication, 2006.

[123] A. B. Sousa, *Studies of $\nu_\mu \rightarrow \nu_e$ oscillation appearance in the MINOS experiment*, PhD thesis, Tufts University, 2006.