



Goodman, Evan Arthur Gerald (2024) *Combining T2K with other experiments to better constrain oscillation parameters*. PhD thesis.

<https://theses.gla.ac.uk/84120/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

Combining T2K With Other Experiments to Better Constrain Oscillation Parameters

Evan Arthur Gerald Goodman

Submitted in fulfilment of the requirements for the
Degree of Doctor of Philosophy

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow



University
of Glasgow

September 2023

Abstract

This thesis presents an analysis of T2K data using a new external reactor constraint from Daya Bay instead of the regular one-dimensional Gaussian provided by the Particle Data Group (PDG). Both the PDG and Daya Bay data sets can be used to update the prior of given parameters in the T2K analyses. Applying Daya Bay's two-dimensional constraint on the mixing angle θ_{13} and mass splitting Δm_{32}^2 improves the constraint on the mass splitting parameter by 25% in normal hierarchy and 18% in inverted hierarchy compared to using the PDG external prior. Furthermore, denoted with a Bayes factor value which compares two hypotheses using the posterior results, it was found that there was a small increase in the preference for normal hierarchy over inverted hierarchy, $\mathcal{B}(NH/IH)$: PDG = 2.77 and Daya Bay = 2.79. There was a slightly larger increase for the upper octant in the octant degeneracy, $\mathcal{B}(UO/LO)$: PDG = 2.27 and Daya Bay = 2.38. The thesis also describes development work towards the first full joint-fit between two long baseline experiments, T2K and NOvA, showcasing the increase in statistical sensitivity for the oscillation parameters and the potential to solve some of the current degeneracies limiting the sensitivity of both experiments. Finally, there is an introductory insight into an alternate parameterisation of neutrino oscillations that could be used to better understand the constraint from the T2K data.

Contents

Abstract	i
Acknowledgements	xxii
Declaration	xxvi
1 Introduction	1
1.1 Experimental ν History	1
1.2 Neutrino Oscillations	4
1.3 3ν Oscillation Formalism and MSW Effect	7
1.3.1 Matter Effects	13
1.4 Current Status of Neutrinos	14
1.4.1 Solar Neutrinos	14
1.4.2 Atmospheric Neutrinos	15
1.4.3 Reactor Neutrinos	15
1.4.4 Accelerator Neutrinos	16
1.5 Neutrino Interactions	18
2 An Alternate PMNS Parameterisation	23
2.1 Neutrino Oscillations and the PMNS paradigm	23
2.2 T2K Oscillation Spectrum	29
2.3 NO ν A Oscillation Spectrum	30

3	Long-Baseline Neutrino Experiments	32
3.1	T2K Experiment Overview	32
3.2	T2K Beamline	33
3.3	INGRID	39
3.4	WAGASCI-BabyMIND	42
3.5	ND280	44
3.5.1	Magnet and Side Muon Range Detector	44
3.5.2	π^0 Detector	46
3.5.3	Time Projection Chambers (TPCs)	47
3.5.4	Fine Grained Detectors (FGDs)	49
3.5.5	Electromagnetic Calorimeters (ECal)	49
3.5.6	ND280 Upgrade	50
3.6	Super-Kamiokande	52
3.6.1	Water Cherenkov Detection	54
3.6.2	Analysis Samples	55
3.6.3	Gadolinium Doping	56
3.7	NOvA Experiment Overview	57
4	Statistical Inference	61
4.1	Bayes' Theorem	61
4.2	Markov Chain Monte Carlo	62
4.2.1	Monte Carlo	63
4.2.2	Markov Chain Monte Carlo	63
4.2.3	Metropolis-Hastings Algorithm	65
4.2.4	Diagnostics	67
4.3	Posterior Analysis	71
4.3.1	Credible Intervals	73
4.3.2	Prior Reweighting	75
4.3.3	Bayes Factor	75

<i>CONTENTS</i>	iv	
4.3.4	Posterior Predictive Distributions	77
4.3.5	Goodness-of-Fit Metric	78
4.4	MaCh3 Fitter algorithm	79
4.5	Likelihood and Priors	81
4.5.1	Calculation of Likelihood	81
4.5.2	Systematic Uncertainties	84
4.6	Combining T2K-NOvA	85
5	T2K-NOvA joint fit	88
5.1	Introduction to the Joint Fit	89
5.2	Joint-Analysis Technique and Shared Software	89
5.3	Container Validation	92
5.3.1	Oscillation Probability Calculation	94
5.3.2	Log-Likelihood Scans	96
5.3.3	Comparison with Other Fitters	100
5.4	Fit Results	105
5.4.1	Asimov 0	107
5.4.2	Asimov 1	116
5.4.3	Asimov 4	124
5.5	Goodness-of-Fit Metrics	132
6	Daya Bay Constraint	135
6.1	The Daya Bay Experiment	137
6.2	Daya Bay's Log-Likelihood Surface	142
6.3	Data Fit Results	143
6.3.1	One-Dimensional Posteriors	144
6.3.2	Two-Dimensional Posteriors	153
7	Conclusion	161

Appendices	164
A Full Set of Asimov Oscillation Parameter Sets	165
A.1 Asimov 0	165
A.2 Asimov 1	166
A.3 Asimov 2	166
A.4 Asimov 3	166
A.5 Asimov 4	167
Bibliography	168

List of Tables

2.1	A table of oscillation parameters and their values to create Figures 2.2, 2.4 and 2.5. Both Δm_{31}^2 and θ_{23} values were chosen for maximal mixing and Δm_{31}^2 , θ_{12} and θ_{13} were based on best-fit values listed in [69]. . . .	25
4.1	The Jefferies scale for the Bayes factors and their corresponding grades of evidence.	77
5.1	Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 0 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Normal Hierarchy and Upper Octant.	108
5.2	Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 1 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Normal Hierarchy and Upper Octant.	116
5.3	Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 4 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Inverted Hierarchy and Upper Octant.	124
6.1	A comparison of the one-dimensional $\sin^2(\theta_{13})$ priors for The PDG and Daya Bay constraint used in this analysis and the respective data releases used to create the constraint.	144

6.2	A comparison of the highest posterior density points for $\sin^2(\theta_{13})$ in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint. . . .	145
6.3	A comparison of the highest posterior density points for Δm_{32}^2 in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint. . . .	146
6.4	A comparison of the highest posterior density points for $\sin^2(\theta_{23})$ in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint. . . .	147
6.5	A comparison of the highest posterior density points for δ_{CP} in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint. . . .	147
6.6	Fraction of posterior probability lying in different combinations of mass hierarchy and θ_{23} -octant from fit to T2K data for T2K without a reactor constraint (green), T2K with PDG reactor constraint (blue), and T2K with 2D Daya Bay constraint (red).	148
6.7	A comparison of Bayes factors for no reactor constraint, PDG constraint and Daya Bay looking at mass hierarchy hypotheses and octant hypotheses.	148

List of Figures

1.1	The Electron energy spectrum recorded by James Chadwick.	2
1.2	A schematic depicting the normal hierarchy, $\nu_1 < \nu_2 < \nu_3$ (left), and inverted hierarchy, $\nu_3 < \nu_1 < \nu_2$ (right). The colours represent the flavour components of the mass eigenstate (red represents ν_e , blue for ν_μ and green for ν_τ), with the PMNS mixing angles determining the size of the flavour component. Figure taken from [37].	12
1.3	Feymann diagrams showing the Charged Current interactions for electron neutrinos on the left, electron antineutrinos in the middle, and Neutral Current interactions on the right.	12
1.4	A comparison of the joint Solar experimental constraint and the KamLAND constraint for the solar parameters, θ_{12} and Δm_{21}^2	15
1.5	Comparisons of measurements of θ_{23} and Δm_{32}^2	17
1.6	Left: neutrino Charged Current cross-section as function of energy showing the contributions from Quasi-Elastic, Resonant and Deep Inelastic Scattering. Right: antineutrino Charged Current cross-section as function of energy.	19
1.7	Tree level Feynman diagrams for the four charged current neutrino-nucleon channels. Far left: electron neutrino CC quasi-elastic, middle left: electron antineutrino Charged Current quasi-elastic, middle right: electron neutrino Charged Current resonant and far right: electron neutrino Deep Inelastic Scattering.	20

1.8 An illustration of potential FSI interactions that can occur in neutrino-nucleus scattering. 22

2.1 T2K’s current Bi-rate plot depicting the predicted number of ν_e and $\bar{\nu}_e$ events whilst varying some of the PMNS parameters. Solid (dashed) represents Normal (Inverted) ordering. The red, blue, black and green colours show the effect of altering the value of θ_{23} and the ellipse show the effects of varying δ_{CP} 24

2.2 Figure 2.2a shows the bi-probability for T2K at energies below, on and above the average energy peak. Figures 2.2b and 2.2c are the separated effects of contributions for $\sin(\delta)$ and $\cos(\delta)$. Relating to Equation 2.1, at the energy range of 500 – 700 GeV, there is a dominance of the $\sin \delta$ term. At the oscillation maximum for T2K, 600 MeV, there is minimal $\cos \delta$ contribution, suggesting limited sensitivity. 27

2.3 The effect on the oscillation probability at T2K for (anti)neutrinos on the (right) left whilst changing the cartesian coordinates representing $\sin(\delta)$ and $\cos(\delta)$. A visual demonstration shows that the new implementation works, given the perfect overlap between the PMNS paradigm (blue) and the altered parameterisation (green) of Equation 2.4. The red line shows the impact of extending the parameter space by using both a maximal CP conserving, $\chi_C = -1$, and maximal CP violating, $\chi_S = -1$, coordinate. 28

2.4 The effect on the oscillation probability at T2K for (anti)neutrinos on the (right) left at a fixed value of δ , but allowing ρ to be larger (red), equal to (green), and less than (blue) 1. 29

2.5 The effect on the oscillation probability at NOvA for (anti)neutrinos on the (right) left at a fixed value of δ , but allowing ρ to be larger (red), equal to (green), and less than (blue) 1. 31

3.1	An illustration of the T2K experiment, highlighting J-PARC, the near detectors and the far detector.	32
3.2	A schematic of the primary and secondary beamlines of the J-PARC neutrino beam.	34
3.3	A comprehensive breakdown of the target region. The top schematic depicts the graphite target located inside the first magnetic horn. The second segment shows the locations of the other two horns. Finally, the bottom image identifies the tertiary beamline located after the third magnetic horn. This is followed by the beam dump and MUMON detector. The solid line signifies the off-axis beam path, the dashed is the beam axis.	36
3.4	The left hand side shows the predicted flux in FHC mode, dominated by ν_μ but showing the wrong sign and intrinsic ν_e background. The right hand side shows the flux prediction for the RHC mode, dominated by $\bar{\nu}_\mu$	37
3.5	Top: The muon neutrino disappearance oscillation probability with the listed oscillation parameters. Middle: The appearance probability for the muon neutrino to electron neutrino, with varying values of δ_{CP} . Bottom: The neutrino flux and how it changes with various off-axis values showing that 2.5 degrees aligns the peak with the appearance maximum.	40
3.6	Layout of INGRID. With the vertical and horizontal formation, the middle detectors are located on the beam's axis. This figure also includes the unused off-axis modules.	41
3.7	A schematic of the INGRID module showing the breakdown of iron and scintillator planes and the veto planes located on the outer segment of the module.	42

3.8	Left is the full detector schematic for WAGASCI-BabyMIND. One can see the layout with WAGASCI being upstream and the surrounding Wall MRDs followed by BabyMIND. Also included is the location of a separate emulsion-based experiment that shares the housing, NINJA [82]. On the right, the 3D grid layout of the WAGASCI module, showing the alternating orientation of the scintillator planes and the grid cells containing the water target.	44
3.9	An exploded view of the ND280 detector, with the listed sub-detectors visible and the surrounding magnet yoke.	45
3.10	Schematic of scintillator layout in the PØD.	47
3.11	A schematic of the two box design of the TPC. All three time projection chambers are identical.	48
3.12	The Monte Carlo predictions for the energy loss against particle momenta drawn over the data obtained for one the TPCs. Left: positive particles traversing through the TPC. Right: Negative particles traversing through the TPC.	48
3.13	Deposited energy against the track range in FGD1. The comparison with the Monte Carlo predictions for protons, muons and pions, solid, dash and dash-dot respectively, allow for good particle identification.	50
3.14	Exploded view of the ND upgrade, replacing the PØD with a super fine grained detector and two high-angle time projection chambers. Not shown here is the addition of time-of-flight panels that will surround the new upstream sub-detectors, attached to the basket.	51
3.15	The Super-Kamiokande detector, showing its depth inside of Kamioka mine.	53

3.16	Examples of muon(a) and electron(b) like Cherenkov rings captured by SK. The colour of represents the energy of the incident photon and the hit timing information is used to recreate the interaction vertex, denoted with a white cross.	55
3.17	A diagram of the NuMI beamline including the target, decay pipe and the various monitors and absorbers.	58
3.18	A to-scale schematic showing NOvA's near and far detector sizes with a small figure showing the grid formation for the PVC cells. To the far right one can see the layout of the wavelength-shifting fibre looping through the cell.	59
4.1	The combined flux and cross-section correlation matrix used in the 2021 oscillation analysis, showing some of the parameters within the model. The red, heavily correlated, bins are the flux normalisation parameters, and those named are the cross section parameters within the model. Randomly sampling around this subset of parameters alone would be and inefficient sampling technique.	64
4.2	A flow chart depicting the step acceptance algorithm used in the MCMC.	66
4.3	A comparison of autocorrelation functions and parameter variations for a poorly tuned MCMC (a) and a well tuned MCMC (b), specifically looking at the axial mass cross-section parameter M_A^{QE}	69
4.4	An example of a set of chains converging towards a stationary distribution into a low log-likelihood, favourable region. The long decline feature at low step values is known as the "Burn-In" phase.	70
4.5	Two marginalised posteriors for flux normalisation parameters. Left: A Gaussian posterior showing good agreement between all three point estimate methods. Right: a non-Gaussian emphasising the disagreement between the point estimates. Blue is the arithmetic mean method, Gold is fitting a Gaussian and black shows the highest posterior density.	74

4.6	A schematic of the MaCh3 package, showing the various inputs required when running an oscillation analysis. These inputs come as covariance matrices that are used for the step proposal, parameter error definition and correlations.	80
4.7	Left: The posterior probability density for θ_{13} with no reactor constraint applied. Right: The same but with a reactor constraint applied. Reweighting to the reactor constraint produces a better constraint on the measurement of θ_{13} , notice the change of the x -axis.	83
5.1	A comparison of T2K (blue), NOvA (red) and SK (green) contours for $\delta_{CP} - \sin^2(\theta_{23})$ assuming normal (left) and inverted (right) hierarchy and marginalising over all other parameters.	90
5.2	Schematic of the MaCh3 fitting framework as used for this T2K+NOvA joint analysis. The dashed boxes represent the NOvA container.	91
5.3	Schematic of the CAFAna fitting framework as used for this T2K+NOvA joint analysis. The dashed boxes represent the T2K container.	91
5.4	PROB3++ and PMNSOpt comparisons at T2K's baseline in the $\nu_\mu \rightarrow \nu_e$ oscillation channel using the Asimov 1 oscillation parameter set defined in Appendix A. The grey banded area represents values of P_{T2K} and P_{NOvA} that are ≤ 0.02 . The tan-like component at 0.3 GeV is caused by the oscillation probability being close to zero, and the minimum occurring at slightly different energies for the two oscillation calculators. The red dashed line depicts the oscillation maximum at 0.6 GeV. Importantly, the percentage difference between the probability calculated by PROB3++ and PMNSOpt is minimal, at that oscillation maximum.	95
5.5	Likelihood scans for a selection of T2K systematic parameters. The blue line represents the log-likelihood output via the T2K container, with red representing the log-likelihood distribution obtained using the bare MaCh3 analysis package.	98

- 5.6 Likelihood scans for a selection of NOvA systematic parameters. The blue figure represents the log-likelihood output via the NOvA container, with red representing the log-likelihood distribution obtained using the bare CAFAna analysis package, provided by NOvA collaborators. . . . 99
- 5.7 A comparison of one-dimensional marginal posteriors for δ_{CP} , produced by MaCh3 (blue), Ptheta (red), Aria (green) and MaCh3 post-BANFF (black) at the Asimov 0 oscillation point, showing the one (dashed) and two (dotted) sigma credible regions. The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. Figures taken from [135]. 101
- 5.8 A comparison of one-dimensional marginal posteriors for δ_{CP} , produced by MaCh3 (blue), Ptheta (red), Aria (green) and MaCh3 post-BANFF (black) at the Asimov 1 oscillation point, showing the one (dashed) and two (dotted) sigma credible regions. The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. Figures taken from [135]. 102
- 5.9 Marginal posterior distributions of δ_{CP} for Asimov 0 in both hierarchies (top), inverted hierarchy (left) and normal hierarchy (right). The NOvA (pink) posterior were generated using Aria, the T2K (red) posterior was generated using the post-BANFF model, and NOvA+T2K used combining Aria and the T2K container. Figures taken from [135]. 104
- 5.10 Bi-Event plots for T2K (top) and NOvA (bottom) with a marker depicting the location of three Asimov parameter sets described in Appendix A. 106
- 5.11 Marginal posterior distributions of θ_{13} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 109

- 5.12 Marginal posterior distributions of θ_{23} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 110
- 5.13 Marginal posterior distributions of δ_{CP} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 111
- 5.14 Marginal posterior distributions of Δm_{32}^2 for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 112
- 5.15 Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 113
- 5.16 Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 114
- 5.17 Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 115

- 5.18 Marginal posterior distributions of θ_{13} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 117
- 5.19 Marginal posterior distributions of θ_{23} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 118
- 5.20 Marginal posterior distributions of δ_{CP} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 119
- 5.21 Marginal posterior distributions of Δm_{32}^2 for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 120
- 5.22 Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 121
- 5.23 Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 122

5.24 Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 123

5.25 Marginal posterior distributions of θ_{13} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 125

5.26 Marginal posterior distributions of θ_{23} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 126

5.27 Marginal posterior distributions of δ_{CP} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 127

5.28 Marginal posterior distributions of Δm_{32}^2 for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 128

5.29 Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 129

5.30 Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 130

5.31 Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 131

5.32 A posterior predictive p -value distribution generated from the posterior of a T2K+NOvA joint fit with the Asimov 4 data set, before statistical fluctuations. The severe pull in the distribution is due to the lack of statistical uncertainties in the Asimov fit. 134

5.33 A posterior predictive p -value distribution generated from the posterior of a T2K+NOvA joint fit posterior, using a statistically fluctuated version of the Asimov 4 data set. Adding statistical variation into the data sets produces a more appropriate p -value of 75.5%. 134

6.1 Left: The marginalised posterior probability density for $\sin^2(\theta_{23})$ with no reactor constraint applied. Right: The same but with the reactor constraint applied. Reweighting to the reactor constraint produces a better constraint on the measurement of $\sin^2(\theta_{23})$ 136

6.2 A global comparison of experimental constraints on $\sin^2(\theta_{13})$. This includes T2K, Daya Bay, NOvA, MINOS and RENO, using their 2018 oscillation results. The blue shaded region is the 2018 PDG best-fit constraint. 137

- 6.3 A global comparison of experimental constraints on $|\Delta m_{32}^2|$. This includes T2K, Daya Bay, NOvA, MINOS and RENO, using their 2018 oscillation results. The blue shaded region is the 2018 PDG best-fit constraint. 137
- 6.4 A schematic of the Daya Bay experiment. The figure shows the nuclear reactor cores (D1-D2, L1-L4), the Antineutrino Detectors (AD1-AD8) and the corresponding experiment halls, EH1-EH3. Halls EH1 and EH2 are the location of the near detectors and EH3 is the far detector hall. 138
- 6.5 The upper panel contains a comparison of the $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ survival probability using the full Δm_{32}^2 and Δm_{ee}^2 formulas. Bottom pad: the ratio of the two expressions. 141
- 6.6 Left: A 2D histogram surface plot containing the $\Delta\chi^2$ value (colour coded z-axis) for a given set of θ_{13} and Δm_{32}^2 values in Inverted Hierarchy. Right: The same but in Normal Hierarchy. Here, $\Delta\chi^2 = -2 \ln \left(\frac{\mathcal{L}}{\mathcal{L}_{max}} \right)$, where \mathcal{L} is the likelihood value of the current entry, and \mathcal{L}_{max} is the maximum likelihood in the distribution. Showing the difference in correlation between the two hierarchies hypotheses and the mixing angle, θ_{13} . Data taken from [148]. 142
- 6.7 Posterior distributions for $\sin^2(\theta_{13})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 149
- 6.8 Posterior distributions for Δm_{32}^2 , showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 150

- 6.9 Posterior distributions for $\sin^2(\theta_{23})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 151
- 6.10 Posterior distributions for δ_{CP} , showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 152
- 6.11 Posterior distributions for $\delta_{CP} - \sin^2(\theta_{13})$, showing the change in credible contours for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 155
- 6.12 Posterior distributions for $\delta_{CP} - \sin^2(\theta_{23})$, showing the change in credible contours for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 156
- 6.13 Posterior distributions for $\Delta m_{32}^2 - \sin^2(\theta_{23})$, showing the change in credible contours for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 157

- 6.14 Posterior distributions for $\Delta m_{32}^2 - \sin^2(\theta_{13})$, showing the change in credible contours for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. . . . 158
- 6.15 Posterior distributions for $\Delta m_{32}^2 - \delta_{CP}$, showing the change in credible contours for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. 159

Acknowledgements

Firstly, to my supervisors Phill and Paul. I will be forever thankful for this opportunity that you have given me and since starting this journey 4 years ago you have both been nothing but hugely supportive. I am indebted to the time and effort you have both given to getting me here today. The neutrino group at Glasgow has been instrumental to my PhD. To John, Lucas, Paul M, David and Beans on Roast, thank you. The lunch time chats/arguments about absolutely everything and anything really kept work fun. I enjoyed learning and tasting some incredible Brazilian food and exploring the various lunch times spots around the West End. Deeper into the experiment I had a great time working in groups MaCh3 and T2K-NOvA, so I'd like to show my appreciation for some seriously talented people. Patrick, Ed, Kevin, Zoya and the wider group, I remain grateful for the pearls of wisdom you all gave me during our trips to CERN, Japan and California.

All work and no play would makes Jack a dull boy. Outside of this PhD, when I wasn't buried in writing the thesis, I had the incredible pleasure of playing squash for SSRC. The members of SSRC welcomed me with open arms and constantly provided serotonin in the shape of hitting balls against a wall. Especially, a Mr Paul Paton. Paul, thank you listening to my awful patter and always being there when I needed a breather, you really are a bald gem.

There is a very fine establishment that deserves a mention. One that I have spent many hours and probably 90% of my stipend in. The Park Bar. Nina, Winnie, Anna and crew have been like a family away from home. I entered that place in a huff with

physics so many times and I never left without a smile, and not just because it's the best Guinness in Glasgow. Thank you all, truly.

To my family. To Mum and Kay, thank you for the facetimes, the Christmases, the Friday night drinks and more importantly, the distraction. To Dad and Karen, you gave me a vital lifeline to get this done and I owe you more than I can ever repay. Thank you for fuelling my caffeine, my food intake, letting me work all hours in the garage with so much support and letting me vent my frustration on a squash court, even if it does kill you that I win! I hope I have made you all proud.

To John, Alison, Jemma and Thomas, thank you for letting me crash the facetimes with Alex, picking me up from the ferry port, allowing me to come to Portugal and taking my mind off neutrinos with endless talk of sports. I hope it can continue for a long time.

To Adam, Ella and Gav, we made it. The time in 12A PGT was by far the best time I could have ever asked for. You became so much more than friends, you became family. Thank you for the trips to Park Bar, the nights on the PS4, the awful carbonara, red velvet cake and sausage sandwiches. You're irreplaceable, and I will never be able to show you how thankful I am for you just being yourselves, because in doing so you made life so fun.

Finally, to my rock, Alex. I struggle to find the words that show you just how important you have been in making this dream a reality. Your endless love, support and confidence in me has allowed me to get here, and this thesis is just as much yours as it is mine. You're an awe-inspiring woman who has motivated me to be the best version of myself. Thank you for being you, thank you for putting up with the silent facetimes as I sit and type, thank you for holding me up when all I wanted to do was lay down. You deserve the world of praise and love and I'm forever thankful to have you by my side.

*Dedicated to the memory of my grandmother, Nana Mo. An
incredible pillar of love and support throughout the most
testing times.*

*I am the wisest man alive,
I know one thing,
and that is that I know nothing.*

– Socrates

Declaration

Chapter 1 is an introduction to neutrino physics, based on relevant literature. Chapter 2 showcases the author’s analysis work with an alternate parameterisation and the ability to utilise this inside of an oscillation probability calculator used in the T2K analysis. Chapter 3 contains an overview of current long-baseline neutrino experiments, namely T2K and NOvA, and the detector technologies used within them, based on information available in the literature. Chapter 4 is a review of the statistical inference techniques used in this thesis’ analyses. Chapter 5 presents the analysis workflow and techniques for the first long-baseline experiment joint fit between T2K and NOvA. This chapter contains the collaborative work carried out by the author and members of the T2K-NOvA joint working group, with the author being the containerisation of the T2K Mach3 analysis package that is used by the NOvA analysers for the joint fit effort. Other contributions by the author are the validation work for said containers, running the MaCh3-postBANFF simulations and creating the joint fit posterior predictive p -value analysis. Chapter 6 is entirely the author’s own work, with the implementation of Daya Bay’s public two-dimensional constraint in T2K’s oscillation analysis.

I certify that the thesis presented here for examination for a PhD degree of the University of Glasgow is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it) and that the thesis has not been edited by a third party beyond what is permitted by the University’s PGR Code of Practice.

I declare that the thesis does not include work forming part of a thesis presented successfully for another degree. I declare that this thesis has been produced in accordance with the University of Glasgow's Code of Good Practice in Research.

Evan Arthur Gerald Goodman

Chapter 1

Introduction

*Not all questions are answered,
but fortunately some answers
are questioned.*

Terry Pratchett

1.1 Experimental ν History

Neutrinos are incredibly elusive particles. They are the most abundant massive particle in the universe, yet their low mass and lack of electric charge make them exceedingly difficult to study, as they only interact via the weak force. The neutrino was initially postulated by Wolfgang Pauli in the 1930s to explain the results of an electron emission experiment carried out by James Chadwick [1]. Initially, the beta decay (β decay) process was believed to be a two body decay, ${}_Z X \rightarrow {}_{Z+1} Y + e$, thus the electron should have been monoenergetic upon emission. Instead a spectrum of energy was found, depicted in Figure 1.1. In a letter to a radioactivity conference, Pauli proposed a spin $\frac{1}{2}$ electrically neutral fermion called the "neutron", that requires a mass of the same order of magnitude to that of the electron [2]. In 1931, James Chadwick later confirmed a somewhat massive particle in the nucleus that became the neutron that we know and understand today [3], thus Pauli's proposed particle was renamed, the neutrino.




Figure 1.1: The Electron energy spectrum recorded by James Chadwick. Taken from [4].

The initial detection of the ‘invisible’ neutrino did not arise until 1956 (initially 1953 but the background was far too large), when Reines and Cowan carried out "Project Poltergeist" [5]. The experiment was designed to measure the products of the inverse beta-decay process denoted in equation 1.1:



where p and n are protons and neutrons respectively, e^+ is a positron and $\bar{\nu}$ is an antineutrino. The experimental set-up used a tank of liquid scintillator sandwiching a layer of water doped with cadmium chloride (CdCl_2) and surrounded by 110 photomultiplier tubes (PMTs). When the neutrino interaction occurs, the positron would quickly undergo positron-electron annihilation, emitting two photons that can be detected by the PMTs. The neutron is then captured by the CdCl_2 and a further pair of photons are emitted with a delayed signal of $\sim 20 \mu\text{s}$ [6]. With significant evidence to support the existence of neutrinos, the progression of understanding was rapid.

Only six years after the detection of the (anti)neutrino, a team of researchers working on the Alternating Gradient Synchrotron (AGS) [7], at the Brookhaven National Laboratory, created the first beam of neutrinos. They collided 15 GeV protons with a beryllium target, creating pions that further decayed into muons and neutrinos with a high branching ratio. These neutrinos travelled a short distance to a 10 ton aluminium spark chamber where the neutrino interactions with aluminium would be captured by surrounding cameras. This allowed a cohort of researchers led by Leon Lederman, Melvin Schwartz and Jack Steinberger to probe if the neutrinos produced via this decay method were different to those produced by the inverse beta-decay process. If the neutrinos were the same nature then there should have been an equal number of electron and muon events. However, Lederman *et al* found many more muon events than electron events [8], confirming the discovery of the muon neutrino.

Understanding that neutrinos were related to corresponding isospin doublets (ν_e for the electron and ν_μ for the muon), it was inevitable for the hunt to begin again upon the discovery of the third and heaviest lepton, the tau, τ [9], for the ν_τ . It was not until the year 2000 that the tau neutrino was discovered by the Direct Observation of Nu-Tau (DONUT) experiment. Using a nuclear emulsion technique and producing a predominantly tau neutrino beam, DONUT was able to measure and identify tau neutrino interactions by capturing the short lived tau lepton produced by the tau neutrino interaction, decaying into a muon or electron [10].

It would be natural for one to think about the existence of more generations of neutrinos and leptons after the continuous discoveries, but with a joint-fit of four detectors in the Large Electron-Positron collider measuring the mass of the Z^0 boson, it was concluded that the best fit was three generations [11].

1.2 Neutrino Oscillations

In the 1960's an experiment was proposed by Ray Davis using methods suggested by Pontecorvo [12] to measure the neutrino flux emitted by the nuclear fusion processes, namely the proton-proton and the Carbon-Nitrogen-Oxygen chains from the sun's core. These neutrinos are now commonly called, solar neutrinos. The Homestake Experiment [13], located in South Dakota, would measure the Sun's neutrino flux using a 380 cubic litre tank filled with tetrachloroethylene (C_2Cl_4), and measured the product of the radiochemical inverse beta decay, $^{37}Cl + \nu_e \rightarrow ^{37}Ar + e^-$. The Argon atoms were extracted and the neutrino flux was calculated using the known cross-section. When comparing the results to that predicted by the Standard Solar Model (SSM) [14], Davis *et al* found that the measured flux was actually a third smaller than predicted. Initially, both the model and experiment were blamed for the discrepancy and this birthed the "Solar Neutrino Problem". Further experiments then supported the results of the Homestake experiment, SAGE [15] and GALLEX [16] used an equivalent radiochemical process as Homestake, but substituted the use of Chlorine for Gallium. Kamiokande [17] used a water Cherenkov detector but found the same result, solidifying this discrepancy. The main issue was that the comparisons were model dependent, relying on the Standard Solar Model predictions. In the early 2000's, the Sudbury Neutrino Observatory (SNO) [18] was the first model-independent measurement of the neutrino flux, and this used a heavy water detector that was sensitive to not only Charged Current interactions, but also Neutral Current. Thus, it could measure the total flux of neutrinos from the sun. When the data were collected, SNO released the analysis that the measured total flux agreed with the SSM and deemed conclusively that neutrinos change favour state as they propagate [19].

Cosmic Rays interact with nuclei in the earth's atmosphere to produce secondary particles such as pions and muons. These, in turn, decay to produce neutrinos of varying

flavour

$$\begin{aligned}\pi^\pm &\rightarrow \mu^\pm + \bar{\nu}_\mu \\ \mu^\pm &\rightarrow e^\pm + \bar{\nu}_e + \bar{\nu}_\mu.\end{aligned}\tag{1.2}$$

These are known as atmospheric neutrinos. Unlike solar neutrinos, these were not expected to be a specific flavour but due to the decays taking place (above), there should have been roughly 2 ν_μ s to every one ν_e . However, similar to the solar neutrino problem, there was a significant difference in that the measured muon-neutrino flux was smaller than the expected value [20]. A double ratio was calculated with the total number of muon-neutrino N_μ and electron-neutrino N_e events over the Monte Carlo simulated data that accounted for known physics and detector limitations

$$R = \frac{(N_\mu/N_e)_{\text{Data}}}{(N_\mu/N_e)_{\text{Simulated}}}.\tag{1.3}$$

One would expect the result of R to be ~ 1 if there were no oscillations to account for, but the measurements and calculation by Super-K [21], Soudan2 [22], IMB [23] and Kamiokande [24] obtained a value of $R \sim 0.6$. This result would suggest that the expected number of events N_μ was lower than the recorded value. Super-K also performed an analysis by comparing the flux of down-going atmospheric neutrinos, i.e those produced above the detector with a baseline of $\sim 50\text{--}100$ km, to up-going atmospheric neutrinos, i.e baseline length of $\sim 100\text{--}13000$ km, depending on the zenith angle. It was found that there was a clear flux difference, with fewer upward-going muon neutrinos than downward-going. This was a definitive measurement of muon neutrino disappearance and that neutrinos oscillate [25].

At the time of initial postulation, the hypothesis of neutrino oscillations was not supported due to the belief that neutrinos were indeed massless¹. However, Bruno Pon-

¹The Standard Model still suggests that this is the case!

tecorvo had been devising a theory that included the idea of mixing and oscillation [26]. This was not a new idea and had previously been postulated in the meson sector with $\bar{K}^0 \leftrightarrow K^0$ by Gell-mann *et al* [27]. Pontecorvo posed the possibility as to whether or not such an event could happen with any other set of neutral particles.

The standard model predicts that symmetries should be conserved and have been observed to do so under the electromagnetic force and the strong nuclear but not the weak force. Such symmetries include parity, P , and charge conjugation, C . Parity is the quantum mechanical operator that inverts a particle's spatial coordinates such that $\vec{x} \rightarrow -\vec{x}$. In essence, reflecting a particle into its mirror image. Charge conjugation transforms a particle to an anti-particle and vice versa. In the 1950's, a problem arose known as the Tau-Theta puzzle which found two particles, tau and theta, that were believed to be the same mass and lifetime but decayed into different products. These particles later became known as charged kaons. Lee and Yang proposed that parity may not be conserved under the weak force and an experiment carried out by C. S. Wu *et al* that looked at the beta decay of a Cobalt isotope, ^{60}Co , which is mediated by the weak interaction, found as a result that parity is maximally violated [28].

It was then believed that applying a product of the two symmetries, creating charge-parity symmetry, CP , would be conserved in weak interactions. However, this was shown to be violated by Cronin and Fitch *et al* when studying the decay products of Kaons [29]. CP gives a definitive symmetry between matter and antimatter, if CP is conserved then there is no difference between them. This becomes special for particle physicists as it satisfies the conditions laid out by Andrei Sakharov in order to explain the matter-antimatter asymmetry observed in the universe [30]. The three conditions required for matter and antimatter to be produced at different rates include:

- Baryon number violation;

- Charge and Charge-Parity symmetry violation;
- The universe fell out of thermal equilibrium early in its evolution.

The CP violation observed in the meson sector was not sufficiently large to support the second condition. As such, other avenues like B factories and neutrinos are being explored to find if CP symmetry is also violated through propagation and mixing as opposed to just in decay. Advancing to the late 1990's, the idea of neutrinos changing flavour was gaining acceptance and if neutrinos can undertake time-dependent neutrino oscillations, then it would imply that they must have mass.

1.3 3ν Oscillation Formalism and MSW Effect

The statement that neutrinos have mass and that there are three individual flavours we see that there is a spectrum of fixed mass for each of the three flavours, $|\nu_i\rangle$ corresponding to a fixed flavour state, $|\nu_\alpha\rangle$, with $\alpha = [e, \mu, \tau]$. It could therefore be stated that the superposition of mass eigenstates make neutrino flavour α [31], creating the equation,

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad |\bar{\nu}_\alpha\rangle = \sum_i U_{\alpha i} |\bar{\nu}_i\rangle \quad (1.4)$$

where $U_{\alpha i}$ is an individual component from the unitary Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. This matrix determines the amplitude of mixing between the flavour and mass eigenstate components, through the following form,

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix}. \quad (1.5)$$

This can be decomposed into four individual rotation matrices,

$$U = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{LBL + Reactor}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar}} \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\phi_2} & 0 \\ 0 & 0 & e^{i\phi_3} \end{pmatrix}}_{\text{Majorana Phases}} \quad (1.6)$$

where $s_{ij} = \sin(\theta_{ij})$ and $c_{ij} = \cos(\theta_{ij})$ and θ_{ij} are the mixing angles between the mass states. A further unanswered question in the sector of neutrinos is whether or not they are of the Dirac or Majorana type. A Majorana neutrino particle would be its own anti-particle and a more detailed report of this case can be found in [32]. Experiments such as CUORE and KamLAND-Zen [33] are searching for the infamous neutrinoless double beta decay, $0\nu\beta\beta$, that can only be present if neutrinos have Majorana mass. For the purpose of this thesis, the Majorana nature of the neutrino does not affect neutrino oscillations, as the Majorana phases cancel, so the assumption will be made that they are Dirac and that the complex phases ϕ_2 and ϕ_3 are $= 0$, yielding the identity matrix, $\mathbb{1}$, with $diag[1,1,1]$, which can be ignored. Therefore the PMNS matrix becomes

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}. \quad (1.7)$$

The complex phase δ_{CP} is the charge-parity violating phase within neutrino mixing, quantifying the phase difference between the flavour neutrinos. The measurement of the CP phase comes from an analysis of the subsequent difference between neutrino and antineutrino oscillations. If $\sin(\delta_{CP}) = 0$, then $P(\nu_\mu \rightarrow \nu_e)$ would be equal to

$P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ in a vacuum. However, as will be explained later in the thesis, this might not be the case and this is a further glimpse into the matter-antimatter asymmetry problem.

As the neutrino propagates through a vacuum, one can assume that $t = cL$, where L is distance travelled and c is the speed of light, which we assume to be $c = 1$ in natural units. The time evolution of the neutrino's mass states as a plane-wave solution becomes,

$$|\nu_i(L)\rangle = e^{i(E_i t - p_i \cdot L)} |\nu_i(0)\rangle. \quad (1.8)$$

Focusing on the phase term,

$$E_i t - p_i \cdot L, \quad (1.9)$$

the assumption can be made that plane waves have a fixed energy and momentum relationship, such that, $E_i^2 = m_i^2 + |p_i|^2$. But with the nature of neutrinos it can be stated that all mass states are ultrarelativistic, $m_i \ll E$. In order to equate the average energy across all mass states, a binomial expansion must be performed [34]

$$p_i = \sqrt{E_i^2 - m_i^2} \simeq E - \frac{m_i^2}{2E}. \quad (1.10)$$

There are two assumptions involved in the following steps, that $E \approx p$ and that $t \approx L$. Substituting equation 1.10 into equation 1.9, the phase is now defined as:

$$\left(E - \frac{m_i^2}{2E}\right)L - EL = \frac{-m_i^2}{2E}L. \quad (1.11)$$

Therefore the time evolution of the mass states for a neutrino that has travelled distance L becomes,

$$|\nu_i(L)\rangle = e^{\frac{-im_i^2 L}{2E}} |\nu_i(0)\rangle, \quad (1.12)$$

thus giving the neutrino a flavour at L ,

$$|\nu_\alpha(L)\rangle = \sum_i U_{\alpha i}^* e^{\frac{-im_i^2 L}{2E}} |\nu_i\rangle. \quad (1.13)$$

Applying the conjugate to equation 1.4 yields the flavour of the neutrino after propagation,

$$\langle \nu_\beta | = \sum_j U_{\beta j} \langle \nu_j |, \quad (1.14)$$

where, β and j , are flavour and mass eigenstates respectively. Then combining equations 1.13 and 1.14 produces the amplitude of oscillation,

$$\begin{aligned} \langle \nu_\beta | \nu_\alpha(L) \rangle &= \sum_i \sum_j \langle \nu_j | U_{\beta j} U_{\alpha i}^* e^{\frac{-im_i^2 L}{2E}} |\nu_i\rangle \\ &= \sum_i U_{\beta i} U_{\alpha i}^* e^{\frac{-im_i^2 L}{2E}}. \end{aligned} \quad (1.15)$$

The final step is to square the amplitude to find the probability of a neutrino oscillating from flavour α to flavour β having travelled distance L ,

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= |\langle \nu_\beta | \nu_\alpha(L) \rangle|^2 \\ &= \sum_{ij} U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j} e^{-i(m_i^2 - m_j^2) \frac{L}{2E}}. \end{aligned} \quad (1.16)$$

Expanding equation 1.16 with the use of the Kronecker delta $\delta_{\alpha\beta}$ and some exponential to trigonometric identities we yield the full probability equation,

$$\begin{aligned}
P(\nu_\alpha \rightarrow \nu_\beta) = & \delta_{\alpha\beta} - 4 \sum_{i>j} \operatorname{Re}\{U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j}\} \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \\
& + 2 \sum_{i>j} \operatorname{Im}\{U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j}\} \sin\left(\frac{\Delta m_{ij}^2 L}{2E}\right)
\end{aligned} \tag{1.17}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$ ($i, j = 1, 2, 3$). A key aspect of equation 1.17 is that the oscillation probability is dependent on the square of the mass splittings between the mass eigenstates and not the absolute mass values, which are currently unknown but being evaluated by experiments such as KATRIN [35]. There is an understanding on the difference between the first two states, $m(\nu_1)$ and $m(\nu_2)$, such that $m(\nu_1) < m(\nu_2)$ [36]. However, a topical question currently surrounding neutrino mass is the order in which ν_3 fits. Figure 1.2 illustrates the two mass ordering options, the left hand side shows normal ordering, where $m(\nu_1) < m(\nu_2) < m(\nu_3)$. The right hand side depicts inverted ordering where $m(\nu_3) < m(\nu_1) < m(\nu_2)$. Despite being unknown, the normal hierarchy generated the name given the assumption that the ν_1 is the lightest mass and given its flavour composition is two-thirds electron, the lightest lepton, this makes sense. Long-baseline and atmospheric neutrino experiments are trying to better probe and understand the mass hierarchy degeneracy: as well as being interesting in itself, this will allow for a better understanding of neutrino interactions and provide a constraint on δ_{CP} .

Interestingly, the CP violating phase can only be observed in appearance channels where, $P(\nu_\alpha \rightarrow \nu_\beta)$ if $\alpha \neq \beta$. This occurs due to the imaginary part of equation 1.17 becoming zero due to the multiplication of complex conjugates of the same term.

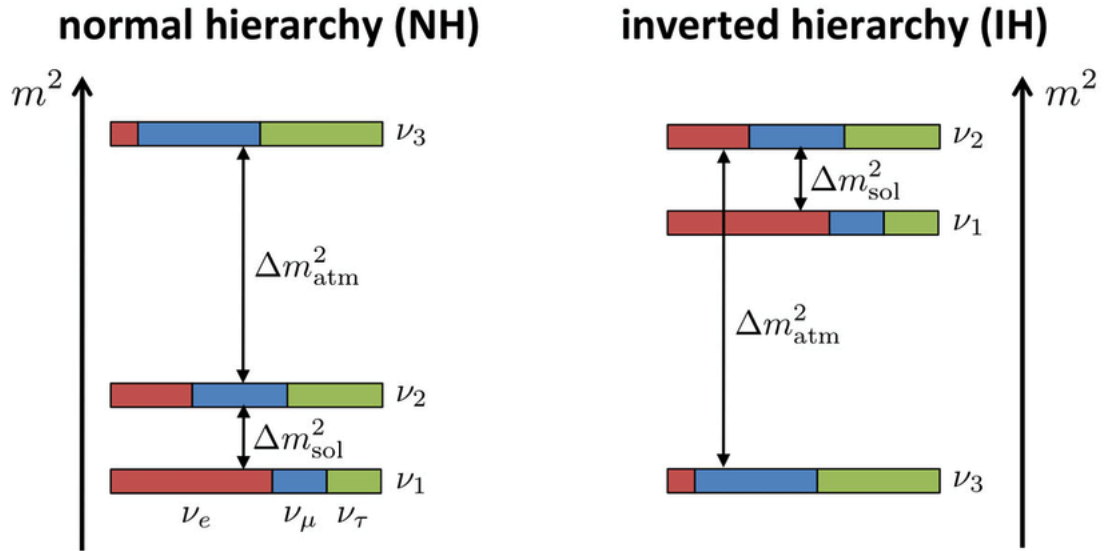


Figure 1.2: A diagram depicting the normal hierarchy, $\nu_1 < \nu_2 < \nu_3$ (left), and inverted hierarchy, $\nu_3 < \nu_1 < \nu_2$ (right). The colours represent the flavour components of the mass eigenstate (red represents ν_e , blue for ν_μ and green for ν_τ), with the PMNS mixing angles determining the size of the flavour component. Figure taken from [37].




Figure 1.3: Feynman diagrams showing the Charged Current interactions for electron neutrinos on the left, electron antineutrinos in the middle, and Neutral Current interactions on the right.

1.3.1 Matter Effects

The formalism above is true for neutrinos propagating through a vacuum, but for many instances this is not the case. Certainly in long-baseline neutrino oscillation experiments, neutrinos will propagate through a medium, the Earth’s crust, which is composite with non-exotic matter such as protons, neutrons and electrons. This can have profound effects on the neutrinos oscillation probability depending on the flavour of neutrino. There are two main interaction types, Charged Current (CC) scattering and Neutral Current (NC) scattering, depicted in Figure 1.3. CC scattering occurs when an electron neutrino or anti-neutrino interacts with an electron and exchanges a W^\pm , depending on the neutrino type. NC interactions are independent of neutrino flavour and will exchange a Z^0 boson. This is often referred to as the matter effect, or MSW effect, named after the postulators, Mikheyev–Smirnov–Wolfenstein [38, 39]. The matter effect is accounted for by a perturbation to the vacuum Hamiltonian matrix, adding a potential term, V_e , for the Charged Current scattering on the electron (anti)neutrinos,

$$V_e = \pm\sqrt{2}G_F N_e, \quad (1.18)$$

where G_F is the Fermi coupling constant and N_e is the electron density of the matter that the neutrinos are propagating through. The sign present at the front of equation 1.18 is to define the potential for electron neutrinos (+) and antineutrinos (−). One can find the full derivation of the matter term by combining the vacuum Hamiltonian, \mathcal{H}_v , and the matter effect Hamiltonian, \mathcal{H}_m , which includes the perturbation defined in Equation 1.18 in [40] and [41]. Ignoring the influence of matter will lead to “fake CP-violation” due to the difference in potential between neutrinos and antineutrinos [42]. The strength of the matter effect term will vary depending on the energy of the neutrino, or the density of the matter that the neutrino is propagating through, and importantly it changes the probability’s dependence on the mass splitting terms. The change in matter potential gives an experiment more sensitivity to the sign of Δm_{32}^2 ,

giving insight into the aforementioned mass hierarchy degeneracy.

1.4 Current Status of Neutrinos

Neutrino oscillations have gathered a lot of interest in the scientific community over the last few decades. Experiments past, present and future are dedicated to understanding and constraining the mixing parameters defined in the PMNS matrix. These experiments are built with the purpose of being sensitive to components of this matrix and can be categorised into groups. No individual experiment has good sensitivity to all parameters and so the best-fits are a collection of decades of work across multiple collaborations. Both solar and atmospheric neutrinos were mentioned in section 1.2, but two other important sources are Accelerator neutrinos and Reactor neutrinos.

1.4.1 Solar Neutrinos

Briefly mentioned in section 1.2, solar neutrinos are those that originate as a product of the ongoing nuclear reactions inside of the core of the Sun, specifically the CNO cycle and the p - p chain. Knowing the predicted flux of electron neutrinos, the main oscillation channel for solar neutrino experiments is the ν_e disappearance channel. This gives sensitivity to the θ_{12} mixing angle and the Δm_{21}^2 mass splitting PMNS parameters. The best measurement of Δm_{21}^2 to date was ironically made by a nuclear reactor experiment, KamLAND. Figure 1.4 shows the comparison for the Δm_{21}^2 and $\tan^2(\theta_{12})$ two dimensional constraint between KamLAND and joint solar constraints, where the larger uncertainty for the solar experiments is present, unlike KamLAND which has a much smaller uncertainty on the mass splitting. A global best-fit to the data sets provides a good constraint of solar PMNS parameters.




Figure 1.4: A comparison of the joint Solar experimental constraint and the KamLAND constraint for the solar parameters, θ_{12} and Δm_{21}^2 . Figure taken from [43].

1.4.2 Atmospheric Neutrinos

Atmospheric neutrinos are caused by cosmic rays interacting with the Earth's atmosphere creating a hadronic shower of charged pions and muons that further decay into neutrinos. Discussed in section 1.2, neutrino oscillations were hinted at due to the difference between the upgoing and downward going neutrino flux. Those with a large L/E , traversing through the Earth, had more time to oscillate compared to those made above a detector. Within the PMNS model, the ν_μ disappearance channel that is measured in atmospheric neutrinos, is dominated by θ_{23} and Δm_{32}^2 , commonly referred to as the atmospheric parameters.

1.4.3 Reactor Neutrinos

The initial postulation of neutrinos, described in Section 1.1, originated from a nuclear process. Recent experiments, such as Daya Bay [44] and RENO [45], used existing

nuclear reactors to measure the flux of emitted antineutrinos via the beta decay process in nuclear fission reactions. Reactor experiments commonly have a short baseline length $\mathcal{O}(1 \text{ km})$, measuring the $\bar{\nu}_e$ disappearance channel. This oscillation channel gives sensitivity to the θ_{13} oscillation parameter, better than that of long-baseline accelerator neutrino experiments. In 2012, Daya Bay published a non-zero θ_{13} result with a 5.3σ confidence [46]. This was almost simultaneously confirmed by other experiments [47, 48]. This is further discussed in chapter 4 and chapter 6.

1.4.4 Accelerator Neutrinos

Neutrinos can be produced via the same decay channels to atmospheric neutrinos. The aforementioned AGS experiment in Brookhaven national lab created a source of muon neutrinos with the understanding of the pion decay chain and this led to the foundation for future experiments. The first long-baseline accelerator neutrino experiment was called KEK-to-Kamioka (K2K) [49], based in Japan. K2K had a baseline length of 250 km from the neutrino beam source to the Super-Kamiokande far detector and an average neutrino energy of 1.4 GeV. K2K had a smaller water Cherenkov detector situated $\sim 300 \text{ m}$ downstream of the target to give better constraints on neutrino interactions before propagation. MINOS [50] was another long-baseline experiment that had a larger baseline and higher peak median neutrino energy, 735 km and 3 GeV, respectively. Although the two experiments had similar setups, the detectors were not identical and MINOS utilised a magnetised steel-scintillator design, as opposed to K2K’s water Cherenkov technique. These experiments agreed with the atmospheric findings of Δm_{32}^2 .

K2K and MINOS were primarily focused on the muon neutrino disappearance channel, emphasising searches on the atmosphere mass splitting and θ_{23} mixing angle. At the time, “next-generation” experiments focused on the electron neutrino appearance channel. As the result of a non-zero θ_{13} allowed new long-baseline experiments such as T2K

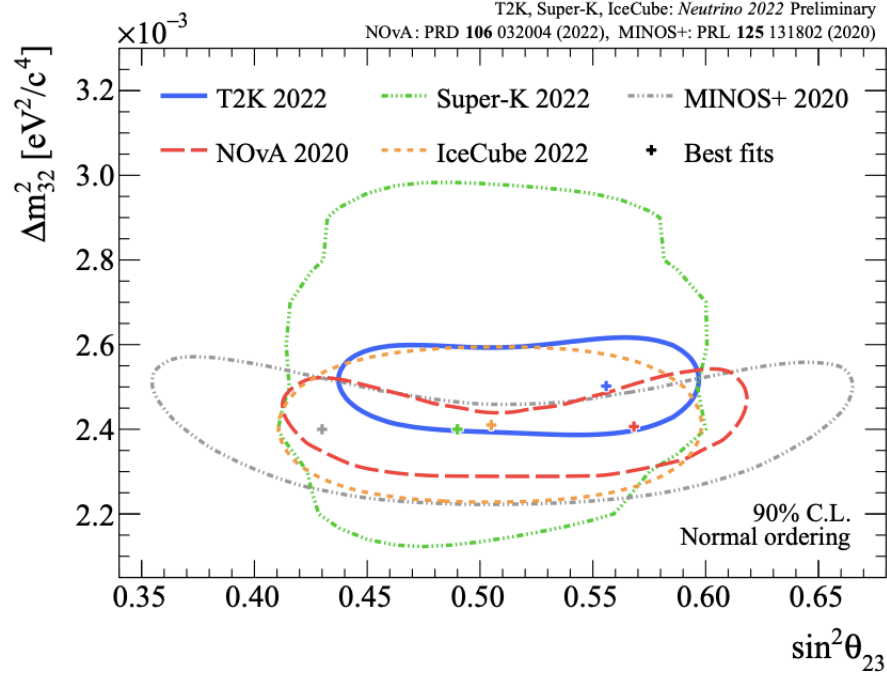


Figure 1.5: Comparison of measurements of θ_{23} and Δm_{32}^2 . Figure taken from [51].

[52] and NOvA [53] to probe into measuring δ_{CP} through the electron (anti)neutrino appearance in a muon (anti)neutrino beam. Both experiments have sensitivity to θ_{13} , $|\Delta m_{32}^2|$, θ_{23} and δ_{CP} and the experimental setups are further explained in Chapter 3. Figure 1.5 shows the 2D contours for the atmospheric mixing parameters, which is measured via the $\bar{\nu}_\mu$ survival channels. Comparing the long baseline and atmospheric experiments, T2K has a world leading constraint. The other main measurements of long-baseline experiments are the electron (anti)neutrino appearance channels. When adding the matter effect to Equation 1.17, the appearance probability can be approximated via

$$\begin{aligned}
P(\nu_\mu \rightarrow \nu_e) \approx & T_{\theta\theta} \sin^2(2\theta_{13}) \frac{\sin^2([1-V]\Delta)}{[1-V]^2} + T_{\alpha\alpha} \alpha^2 \frac{\sin^2(V\Delta)}{V^2} \\
& - T_{\alpha\theta} \alpha \sin(2\theta_{13}) \frac{\sin([1-V]\Delta)}{(1-V)} \frac{\sin(V\Delta)}{V} \sin(\Delta) \sin(\delta) \\
& + T_{\alpha\theta} \alpha \sin(2\theta_{13}) \frac{\sin([1-V]\Delta)}{(1-V)} \frac{\sin(V\Delta)}{V} \cos(\Delta) \cos(\delta)
\end{aligned} \tag{1.19}$$

where,

$$\begin{aligned}
 T_{\theta\theta} &= \sin^2(\theta_{23}), & T_{\alpha\alpha} &= \cos^2(\theta_{23}) \sin^2(2\theta_{12}), & T_{\alpha\theta} &= \cos(\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{23}) \\
 V &= \pm 2\sqrt{2}G_F n_e \frac{E}{|\Delta m_{31}^2|}, & \Delta &= \frac{\Delta m_{31}^2 L}{4E}, & \alpha &= \frac{\Delta m_{21}^2}{\Delta m_{31}^2}.
 \end{aligned}
 \tag{1.20}$$

The sign of the third term containing $\sin(\delta)$ will change if calculating the probability for neutrinos ($-$) or antineutrinos ($+$). In the approximate formula, this term is proportional to $\sin(2\theta_{13}) \sin(\delta_{CP})$, which changes sign for antineutrinos which is the hallmark of charge-parity violation. Equation 1.19 shows a degeneracy between $\sin(2\theta_{13})$ and $\sin(\delta_{CP})$ that allows for a large range of values for δ_{CP} if θ_{13} is poorly constrained. So because of their high sensitivity to θ_{13} , the reactor experiment results are often used as an external constraint by long-baseline experiments, typically in the form of the Particle Data Group (PDG) average [54]. When using this tighter constraint on θ_{13} , long-baseline experiments have better sensitivity to δ_{CP} through the electron appearance channels. This is commonly referred to as the reactor constraint or sometimes the PDG constraint.

1.5 Neutrino Interactions

Neutrino-nucleon interactions are dependent on the neutrino's energy, and as such they need to be carefully modelled due to the range of energies that a source of neutrinos can have. These can be categorised in four channels: Charged Current Quasi-Elastic, Single Pion production, Multiple Pion production and Deep-Inelastic Scattering. Illustrated in Figure 1.6 is the neutrino and antineutrino cross section with respect to neutrino energy, also displaying the dominant interaction channel for a given neutrino energy. Neutrino interactions are an extremely complicated topic in their own right. For the purpose of this thesis, only a brief outline will be given of the interaction channels and models used to parameterise them.

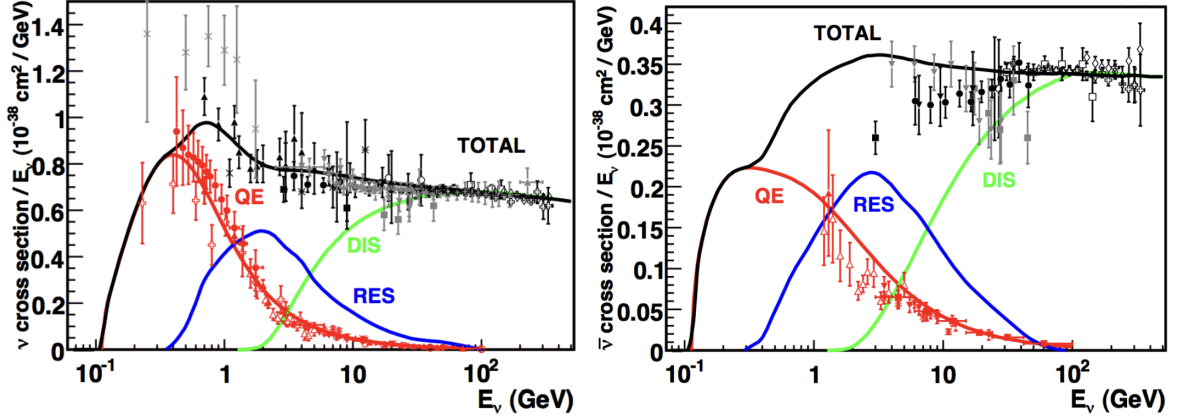


Figure 1.6: Left: neutrino Charged Current cross-section as function of energy showing the contributions from Quasi-Elastic, Resonant and Deep Inelastic Scattering. Right: antineutrino Charged Current cross-section as function of energy. Figure taken from [55].

CCQE

For current generation long-baseline experiments that operate with an average neutrino energy in the range $0.1 \text{ GeV} < E_\nu < 2 \text{ GeV}$, the dominant interaction type is the Charged Current *quasi-elastic* (CCQE) interaction. The Charged Current interaction exchanges a W^\pm boson, producing the corresponding flavour lepton, depicted in the left hand panel of Figure 1.7. The ‘one particle one hole’ interaction is defined as a neutrino interacting with a single bound nucleon quasi-elastically and the nuclear ground state ‘hole’ is modelled using the Benhar spectral function (SF) nuclear model defined in [56]. This model features a distribution of momenta and removal energies of nucleons inside the nucleus within the shell model picture. The physics of the nucleus structure is based on the results of bubble chamber data, which informs the Llewellyn Smith model [57].

The two-particle-two-hole (2p2h) interaction mode, which involves a neutrino interacting with a bound or correlated pair of nucleons, is generated via the Nieves model [59]. Unfortunately, for a water Cherenkov detector, the observable final state is virtually indistinguishable from the final state of CCQE interactions. This can be problematic




Figure 1.7: Tree level Feynman diagrams for the four charged current neutrino-nucleon channels. Far left: electron neutrino CC quasi-elastic, middle left: electron antineutrino Charged Current quasi-elastic, middle right: electron neutrino Charged Current resonant and far right: electron neutrino Deep Inelastic Scattering. Figure taken from [58].

for neutrino energy reconstruction, inducing a bias due to incorrectly accounting for momentum transfer and interaction products. For higher energy interactions, there is better agreement with the initial Nieves model, however for lower energy interactions there is a better agreement with the Martini model [60].

Single Pion Production

Single pion production (SPP) is the second most abundant interaction channel at the T2K (the experiment on which this thesis is based) beam's average neutrino energy and the result of Charged Current Resonant (CCRES) interactions. It is important to measure and constrain this interaction given it constitutes the largest background for the CCQE measurement. These processes are modelled in the Rein-Sehgal formalism [61], describing the neutrino interaction with the nucleon and the resulting excitation to a resonant state. As de-excitation occurs, a Delta baryon can be produced that decays into a neutral pion and nucleon combination. The tree level Feynman diagram for this channel is shown in the third panel of Figure 1.7.

Corrections to the model are required in order to account for coherent scattering occurring at higher neutrino energies with low momentum transfer. Both coherent and resonant interactions produce a pion in the final state and differentiating them is im-

portant but difficult. The low momentum transfer is modelled with the Berger-Sehgal model [62].

Multiple Pion Production and Deep Inelastic Scattering

There can be instances in which the de-excitation of a resonant state can produce multiple pions in the final state, if the momentum transfer is large enough. However, one can see illustrated in Figure 1.6, that as the neutrino energy reaches 3 GeV, the dominant interaction mode is Deep Inelastic Scattering (DIS). In this channel, the interaction is no longer modelled as occurring with the nucleon as a whole, but instead it is better understood at quark level. This interaction causes changes in the structure of the nucleon given the high energy impact, leading to the production of multiple pions or mesonic particles, a tree level Feynman diagram illustrates this process in Figure 1.7 (far right). The boundary between the multiple pion production and deep inelastic is an incredibly difficult one to model, but a comprehensive overview can be found in [63].

Final State Interaction

Final State Interactions (FSI) are a combination of nuclear effects within the nucleus, rather than a physical interaction, but worth noting as it is important for constraining neutrino interaction topologies. FSI occurs after an initial neutrino interaction: the hadronic final state further interacts as it exits the nucleon. This can happen via elastic scattering, absorption, charge exchange and pion production, example of these channels are depicted in Figure 1.8. Relating this to the T2K analysis, if a neutrino interaction produces a resonant baryon that de-excites and produces a pion that is reabsorbed inside of the nucleus, then the final state particle would just be the isospin doublet of the neutrino, leading to a false CCQE signal. To model this, a cascade model is used, with a description of the current status being found in [64].




Figure 1.8: An illustration of potential FSI interactions that can occur in neutrino-nucleus scattering, Figure taken from [65].

Chapter 2

An Alternate PMNS Parameterisation

*Right or wrong, it's very
pleasant to break something from
time to time.*

Fyodor Dostoevsky

2.1 Neutrino Oscillations and the PMNS paradigm

The PMNS formalism is widely accepted amongst the oscillation community as it describes the behaviour of neutrinos as they propagate. A common practice for long-baseline experiment is plotting the number of ν_e and $\bar{\nu}_e$ candidates against the PMNS prediction, a style known as Bi-rate plots, as shown in Figure 2.1. The ellipse shape occurs due to varying the value of δ_{CP} ¹, where the smallest difference in oscillation probability between neutrino and antineutrino occurs at maximal CP-conserving values, 0 and π . The maximal CP-violating values of $\frac{\pi}{2}$ and $\frac{3\pi}{2}$ suppress and enhance the electron neutrino event rate respectively. As seen in Figure 2.1, the T2K data point lies outside of the PMNS prediction, albeit that there is overlap when including the 68% statistical error and 68% systematic error regions, giving an opportunity to look at a

¹The ellipse can also be altered by other factors such as matter effect, flux and cross-section influence. However, the major factor is the value of δ_{CP} .




Figure 2.1: T2K’s current Bi-rate plot depicting the predicted number of ν_e and $\bar{\nu}_e$ events whilst varying some of the PMNS parameters. Solid (dashed) represents Normal (Inverted) ordering. The red, blue, black and green colours show the effect of altering the value of θ_{23} and the ellipse show the effects of varying δ_{CP} . Values of θ_{13}, θ_{12} and Δm_{21}^2 are kept constant at the reactor and solar best-fit values outline in [54]. It is interesting to note that T2K’s data point sits outside of the PMNS paradigm prediction of the ellipses, creating rationale to look at alternate parameterisations. Figure taken from [66].

non-standard PMNS parameterisation.

It was common practice to use an approximate formula to find the probability of a neutrino propagating through a constant matter profile, which also reduced to a simple variation. Yokomakura *et al* [67], showed that, even when using the exact formalism for the probability of oscillation, it could still be broken down into dependent and independent terms. A relevant example being that a muon to electron neutrino oscillation probability formula could be approximated into,

$$P_{\mu e}(\delta) = A_{\mu e} \cos(\delta) \pm B_{\mu e} \sin(\delta) + C_{\mu e}, \quad (2.1)$$

where $P_{\mu e}(\delta)$ is the probability for a muon neutrino to oscillate to an electron neutrino,

$A_{\mu e}$, $B_{\mu e}$ and $C_{\mu e}$ are CP-violating phase independent terms, finally $\cos(\delta)$ and $\sin(\delta)$ are the dependent terms. A full approximation for the probability formula was shown in Equation 1.19 [68], where one can see that the first two terms are equivalent to $C_{\mu e}$, the third to $B_{\mu e}$, and the fourth to $A_{\mu e}$. This will still be true for probability calculations in non-approximated formulas.

To switch between neutrino and antineutrino one would just flip the sign of the matter potential $V \rightarrow -V$ and $\delta \rightarrow -\delta$. This describes a parameterisation that differs from the PMNS model as it treats $\sin(\delta)$ and $\cos(\delta)$ as independent parameters, as opposed to just δ_{CP} . Depicted in Figure 2.2a is a T2K bi-probability plot. This is similar in nature to the bi-rate plot shown in Figure 2.1, however the former plots the oscillation probabilities for two channels, namely $P(\nu_\mu \rightarrow \nu_e)$ and $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$. The ellipse changes depending on the average neutrino energy used and the mass ordering, where the other oscillation parameters are kept at the constant values, which are shown in Table 2.1.

Parameter	Value
θ_{12}	33.5°
θ_{13}	8.51°
θ_{23}	45.0°
Δm_{21}^2	$7.53 \times 10^{-5} \text{ eV}^2/c^4$
Δm_{31}^2	$2.53 \times 10^{-3} \text{ eV}^2/c^4$

Table 2.1: A table of oscillation parameters and their values to create Figures 2.2, 2.4 and 2.5. Both Δm_{31}^2 and θ_{23} values were chosen for maximal mixing and Δm_{21}^2 , θ_{12} and θ_{13} were based on best-fit values listed in [69].

The blue, red and green ellipses of figure 2.2a represent 0.5 GeV, 0.6 GeV and 0.7 GeV, and the solid (dashed) represent normal (inverted) ordering. The major axis are nearly aligned to the $\sin(\delta)$ term shown in Figure 2.2b, the length does not change much as we go either side of the 0.6 GeV peak energy. However, Figure 2.2c shows a greater dependence in the $\cos(\delta)$ term as one changes the neutrino energy. Given the difference in behaviour of $\sin(\delta)$ and $\cos(\delta)$, they can be treated as independent parameters, now

denoted, χ_S , and, χ_C , respectively. Treating χ_S and χ_C as coordinates in a parameter space allows one to break unitarity that is otherwise conserved by the PMNS model, when $\chi_S^2 + \chi_C^2 \neq 1$. Implementing this relationship requires a polar transformation of the Cartesian coordinates to ensure that the correct values of δ_{CP} are being simulated. Fulfilling this transformation one creates polar coordinates, δ and ρ , via equations 2.2 and 2.3,

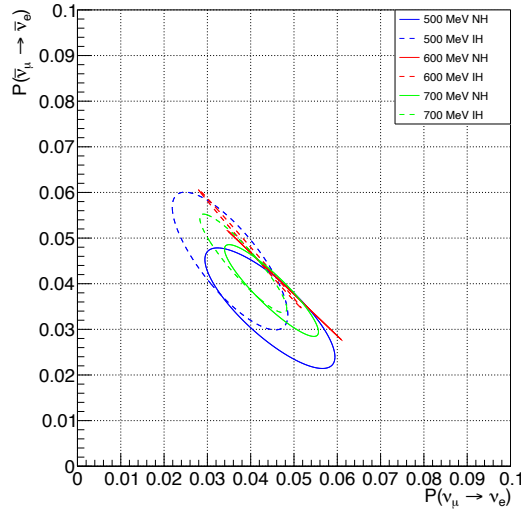
$$\delta_{CP} = \begin{cases} \arctan \frac{\chi_S}{\chi_C}, & \text{if } \chi_C \geq 0 \\ \arctan \frac{\chi_S}{\chi_C} + \pi, & \text{if } \chi_C < 0 \text{ and } \chi_S > 0 \\ \arctan \frac{\chi_S}{\chi_C} - \pi, & \text{if } \chi_C < 0 \text{ and } \chi_S \leq 0, \end{cases} \quad (2.2)$$

$$\rho = \sqrt{\chi_S^2 + \chi_C^2}. \quad (2.3)$$

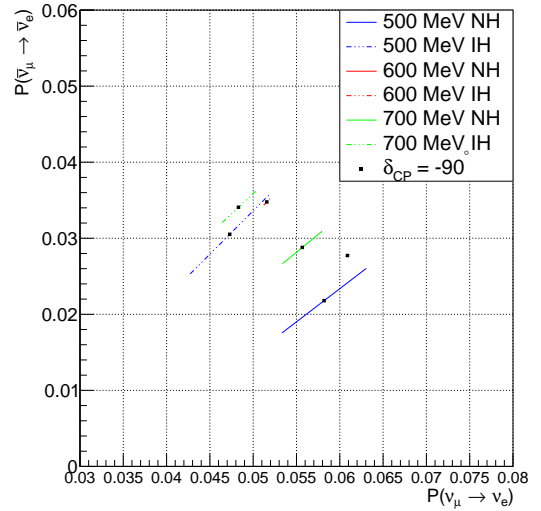
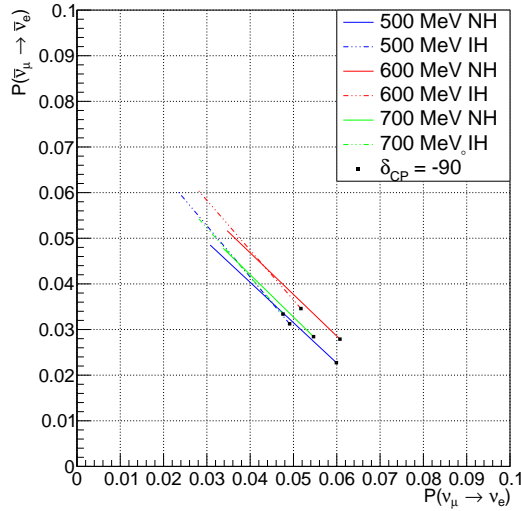
The original probability formula is then modified to a linear combination form, ensuring the use of the polar coordinates throughout,

$$P(\nu_\mu \rightarrow \nu_e) = \frac{1 + \rho}{2} P_{\mu e}(\delta) + \frac{1 - \rho}{2} P_{\mu e}(\pi + \delta). \quad (2.4)$$

The main motivation for this analysis was, instead of using the approximate formula stated in equation 1.19, to fit for this extended paradigm inside of a neutrino oscillation propagator package, `PROB3++` [70]. This includes a non-approximated MSW effect inside of the Hamiltonian.



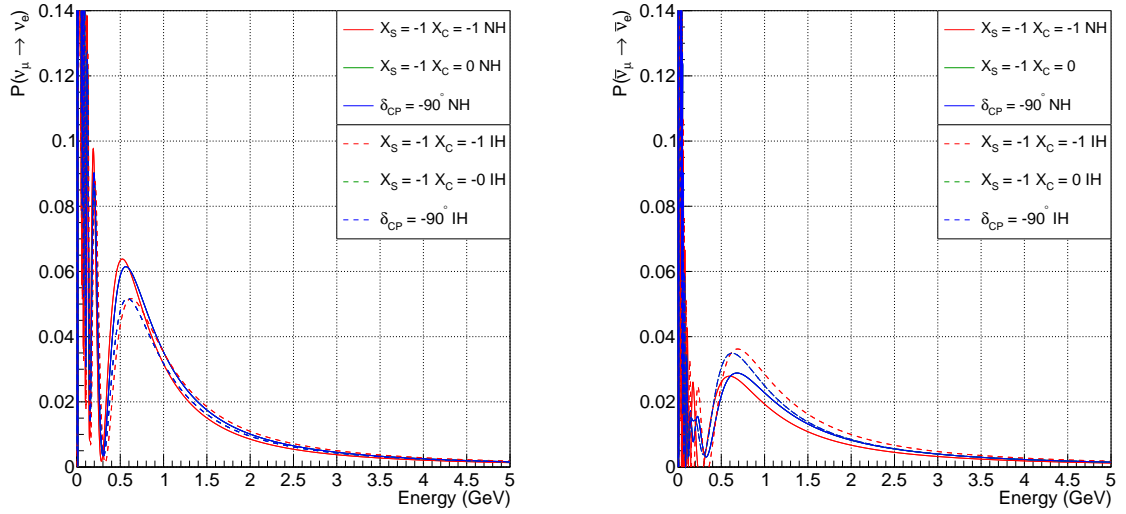
(a) T2K bi-probability with varying average neutrino energy.



(b) The independent contribution of $\sin(\delta)$ at varying energy values to the bi-probability plot, keeping $\cos(\delta)$ constant at, $\chi_c = 0$.

(c) The independent contribution of $\cos(\delta)$ at varying energy values to the bi-probability plot, keeping $\sin(\delta)$ constant at, $\chi_s = -1$.

Figure 2.2: Figure 2.2a shows the bi-probability for T2K at energies below, on and above the average energy peak. Figures 2.2b and 2.2c are the separated effects of contributions for $\sin(\delta)$ and $\cos(\delta)$. Relating to Equation 2.1, at the energy range of 500 – 700 GeV, there is a dominance of the $\sin \delta$ term. At the oscillation maximum for T2K, 600 MeV, there is minimal $\cos \delta$ contribution, suggesting limited sensitivity.



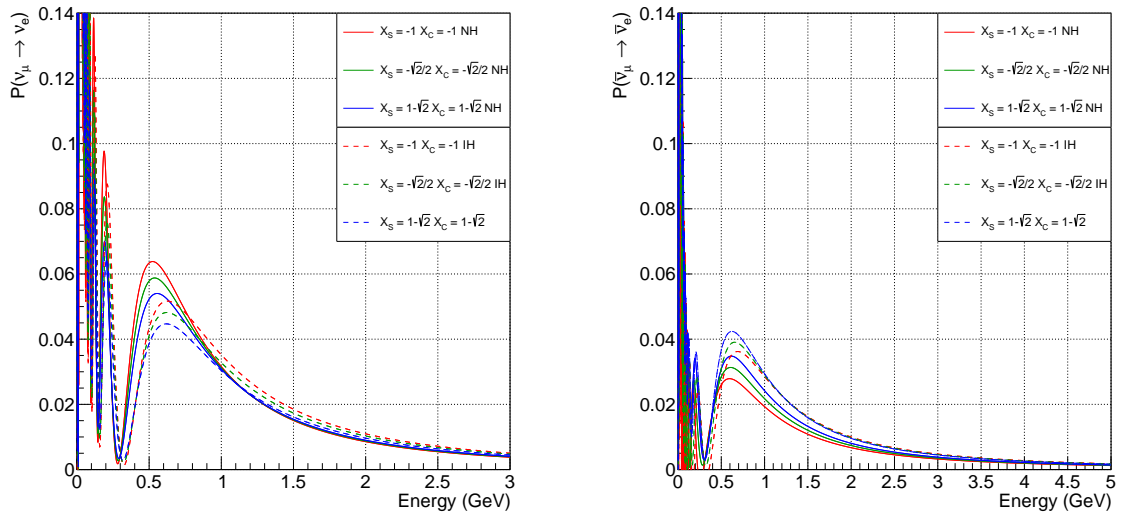
(a) Overlap of PMNS paradigm and the parameterisation of Equation 2.4, for neutrinos in T2K.

(b) Overlap of PMNS paradigm and the parameterisation of Equation 2.4, for antineutrinos in T2K.

Figure 2.3: The effect on the oscillation probability at T2K for (anti)neutrinos on the (right) left whilst changing the cartesian coordinates representing $\sin(\delta)$ and $\cos(\delta)$, as stated in Equation 2.1. When values of χ_S and χ_C are substituted into Equation 2.3 and equal 1, then unitarity is conserved. A visual demonstration shows that the new implementation works, given the perfect overlap between the PMNS paradigm (blue) and the altered parameterisation (green) of Equation 2.4. The red line shows the impact of extending the parameter space by using both a maximal CP conserving, $\chi_C = -1$, and maximal CP violating, $\chi_S = -1$, coordinate.

2.2 T2K Oscillation Spectrum

Figure 2.3 shows the oscillation for $P(\nu_\mu \rightarrow \nu_e)$ (left) and $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ (right). Similar to that in earlier figures, the solid (dashed) represents normal (inverted) ordering. In this formalism, if one selects coordinates, χ_s and χ_c , such that $\rho = 1$, then the second term of Equation 2.4 is zero and one is left with an analysis equal to the PMNS formalism. Used as a validation tool, the coordinates $\chi_s = -1$ and $\chi_c = 0$, represent a $\delta_{CP} = -\frac{\pi}{2}$. The overlap between the green and blue spectra in Figure 2.3 suggests a correct implementation into the PROB3++ software, as one cannot see the green figure due to the perfect agreement. As well as validating that the alternate parameterisation was correctly implemented, it was also interesting to see the effect of adding a $\cos(\delta)$ to the coordinate system. The red and red-dashed lines represent a coordinate system of $\chi_s = -1$ and $\chi_c = -1$, and it is interesting to see the shift in peak with an increase (decrease) in normal (inverted) for the peak probability around the average neutrino energy for T2K. The opposite occurs for antineutrinos as seen in Figure 2.3b.



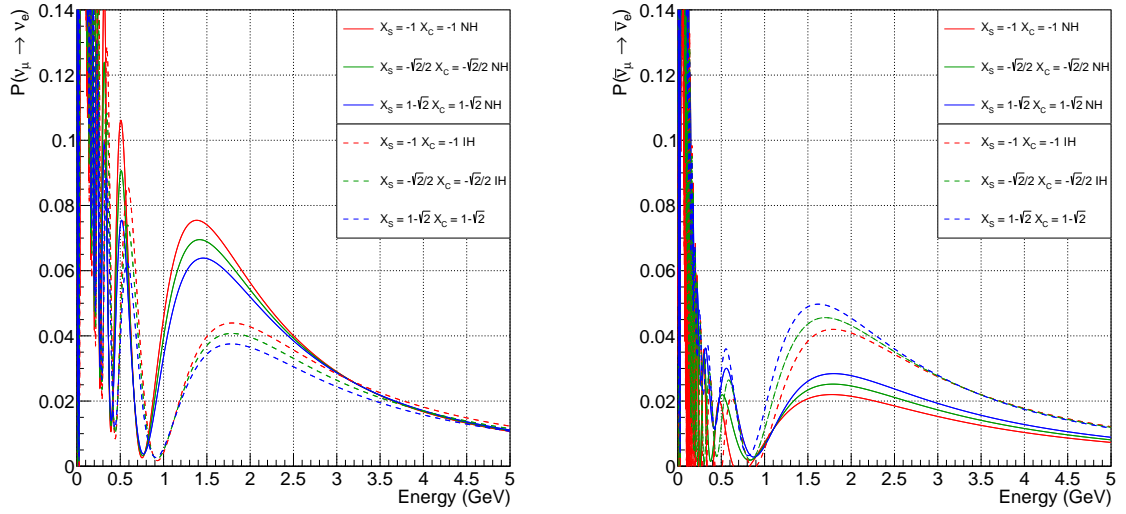
(a) The oscillation spectrum for $P(\nu_\mu \rightarrow \nu_e)$ with varying energy for T2K. (b) The oscillation spectrum for $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ with varying energy for T2K.

Figure 2.4: The effect on the oscillation probability at T2K for (anti)neutrinos on the (right) left at a fixed value of δ , but allowing ρ to be larger (red), equal to (green), and less than (blue) 1.

Figure 2.4 depicts the oscillation spectrum for T2K's baseline length with three sets of coordinates. The green, with coordinates $\chi_s = -\sqrt{2}/2$ and $\chi_C = -\sqrt{2}/2$, conserves unitarity. This is confirmed by substituting these coordinates into Equation 2.3 to find $\rho = 1$. The other two coordinates violate unitarity. These were chosen to violate unitarity in equal measure to better see the effects. The red and red-dashed coordinates, $\chi_S = -1$ and $\chi_C = -1$, violate unitarity above 1, whereas the blue and blue-dashed line has coordinates $\chi_S = 1 - \sqrt{2}$ and $\chi_C = 1 - \sqrt{2}$ and violate unitarity below 1. All three sets of coordinates represent the same value of δ_{CP} , substituting them into Equation 2.2 will show $\delta_{CP} = -3\pi/4$. In Figure 2.4, one can see visible shifts at the oscillation probability peak, around 600 MeV and a difference in probability between the red (red-dashed) and blue (blue-dashed). The red line has an increased peak probability compared to the blue. For antineutrino, the effect is the opposite and there is an increase in the peak probability for blue, and a decrease for red. This result is not unexpected, as the coordinates used for the blue (blue-dashed) line stem towards a ρ value of zero, suggesting no CP violation and the probability of electron neutrino appearance would be similar to the electron antineutrino appearance. This does show that T2K is not very sensitive to the spectral distortion and a full fit would be required to identify if T2K is sensitive to unitarity violations.

2.3 NOvA Oscillation Spectrum

The extended parameterisation can also be applied to other experiment parameters. Changing the length of the baseline to 810 km and the electron density for the matter to 2.84 g/cm^3 and keeping the same parameters as those listed in Table 2.1, one can see the basic effect of extending the oscillation parameter space for NOvA. The wider beam produced by the NuMI beam gives a better sensitivity to $\cos(\delta)$. This can be seen in NOvA's bi-rate plot illustrated later in Figure 5.10b, as the width of the ellipse



(a) The oscillation spectrum for $P(\nu_\mu \rightarrow \nu_e)$ with varying energy for NOvA. (b) The oscillation spectrum for $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ with varying energy for NOvA.

Figure 2.5: The effect on the oscillation probability at NOvA for (anti)neutrinos on the (right) left at a fixed value of δ , but allowing ρ to be larger (red), equal to (green), and less than (blue) 1.

is determined by the strength of the $\cos \delta$ term. Looking at the oscillation spectra in Figure 2.5, the effect of violating unitarity becomes more apparent at the oscillation peak and ~ 1.5 GeV. One can also see in the antineutrino spectrum that the peak shifts to a higher (lower) energy for the blue (blue-dashed) line. Continuing the search for distortions in the neutrino spectra has interesting prospects in finding experimental sensitivity to unitarity violations, especially in experiments which use a wide-band beam.

Chapter 2 has given an introduction into a parameterisation that breaks unitarity violation within the PMNS paradigm by treating δ_{CP} as two independent parameters, $\sin(\delta)$ and $\cos(\delta)$. It goes on to show a successful implementation of this alternate parameterisation into a probability calculator, **Prob3++**. Initial studies showcase the effect of this parameterisation when applied to two long baseline neutrino experiments, T2K and NOvA, highlighting the impact of the individual components of $\sin \delta$ and $\cos \delta$ and potential unitarity violation.

Chapter 3

Long-Baseline Neutrino Experiments

Neutrino physics is largely an art of learning a great deal by observing nothing.

Haim Harari

3.1 T2K Experiment Overview

T2K (Tokai-to-Kamioka) [52] is a long baseline experiment situated in Japan that studies neutrino oscillations through the disappearance channels ($\nu_\mu \rightarrow \nu_\mu, \bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$), and the appearance ($\nu_\mu \rightarrow \nu_e, \bar{\nu}_\mu \rightarrow \bar{\nu}_e$). Protons are accelerated to 30 GeV at the Japan Pro-




Figure 3.1: An illustration of the T2K experiment, highlighting J-PARC, the near detectors and the far detector. Figure taken from [71]

ton Accelerator Research Complex (J-PARC) and collide with a graphite target. This in turn will produce a decay chain resulting in a muon (anti)neutrino beam. The experiment studies the oscillation channels by comparing the data collected at the complex of near detectors, situated 280 m downstream of the beam target, to what is observed in the far detector, Super-Kamiokande [21], after a 295 km propagation through the Earth's crust. A simple schematic of this is depicted in Figure 3.1. The goal of the experiment is to measure and constrain the parameters of the PMNS matrix, namely θ_{13} , θ_{23} and δ_{CP} . T2K also has some sensitivity to constraining the mass hierarchy given the propagation of the neutrinos through the Earth's crust and the matter effect that occurs.

This chapter will give an outline into the experimental setup for T2K. This includes: the neutrino beam, the near detectors as well as their future upgrades and the far detector. There will be a short overview of the NOvA experiment and its layout to give better context into the joint-fit that is discussed later in this thesis.

3.2 T2K Beamline

The process of creating the neutrino beam for T2K starts within the linear accelerator (LINAC) segment of the J-PARC's accelerator complex. Negatively charged hydrogen ions, H^- , are accelerated to 400 MeV/ c and passed through charge stripping foils, turning the ions into protons. These protons then pass into the Rapid Cycling Synchrotron (RCS) and are further accelerated to 3 GeV/ c before being injected into the Main Ring (MR) to accelerate to a final momentum of 30 GeV/ c as a train of eight bunches [72]. When the protons are at this stage then two extraction methods can be applied, slow extraction and fast extraction. Fast extraction uses kicker magnets to extract all of the bunches present in the MR and direct them toward the hadron beamline. The slow extraction method uses an electrostatic separator to direct individual bunches in one

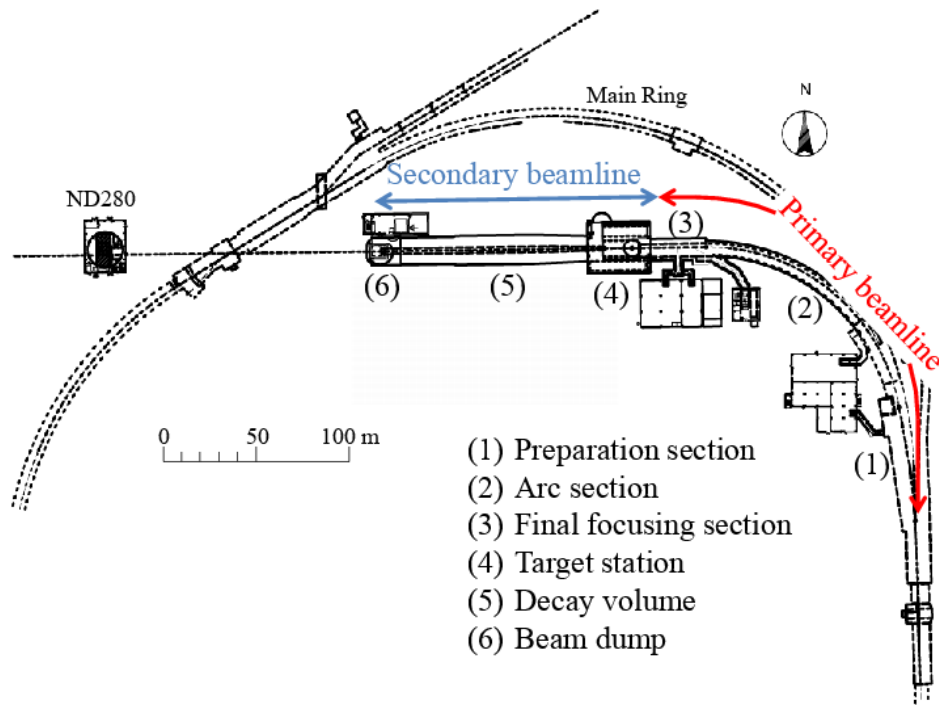


Figure 3.2: A schematic of the primary and secondary beamlines of the J-PARC neutrino beam. Figure taken from [52].

cycle toward the neutrino beamline.

Following the extraction for the MR, the beamline comprises two segments. Following the schematic in Figure 3.2, the primary segment's initial stage is the preparation of the extracted protons. Focused by an assortment of dipole and quadrupole magnets the beam is monitored closely for its profile, position and intensity to ensure acceptance by the second segment within the primary beamline, the Arc section. The Arc section, labelled 2 in Figure 3.2 is pivotal in ensuring that the direction of the beam is toward the target. Surrounded by 14 doublet superconducting combined function magnets (SCFM's) [73] the beam is directed toward the final focusing area. This segment consists of 10 normal magnets that focus the beam downward 3.36° with respect to the horizontal on to the target.

The secondary beamline segment starts with a baffle and an optical transition radiation

monitor (OTR) upstream from the target housing. The target comprises of a graphite rod that is 90 cm in length, with a 2.6 cm diameter. The high energy protons collide with the carbon of the target, producing mainly pions and kaons. These charged particles will be focussed by the magnetic horn with an applied pulsed current of 250 kA located downstream of the target, this is depicted in Figure 3.3. With a 250 kA current, the horns can have a maximum magnetic field of 2.1 Telsa, increasing the neutrino flux at the far detector by a factor of roughly 16. The polarity of the magnetic horn will define the particle charge being focused depending on the requirement for a muon neutrino or muon antineutrino beam. For a neutrino beam the horns will have a positive current, and this is referred to as the Forward Horn Current, FHC. In antineutrino mode, the horn will be pulsed with a negative current, known as the Reverse Horn Current, RHC. Having passed through these three magnetic horns the hadrons will enter a 96 meter decay volume, shown in the bottom schematic of Figure 3.3.

Focusing the positively charged hadrons will allow for the production of the muon neutrino beam via the π^+ and K^+ decays:

$$\begin{aligned}\pi^+ &\rightarrow \mu^+ + \nu_\mu \\ K^+ &\rightarrow \mu^+ + \nu_\mu.\end{aligned}\tag{3.1}$$

If running in RHC, the polarity of the current supplied to the magnetic horn reverses and instead the negatively charged particles are focused toward the decay volume. Given this, the decays then become:

$$\begin{aligned}\pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \\ K^- &\rightarrow \mu^- + \bar{\nu}_\mu.\end{aligned}\tag{3.2}$$

An issue that arises for this set-up is the contamination of wrong sign neutrinos and

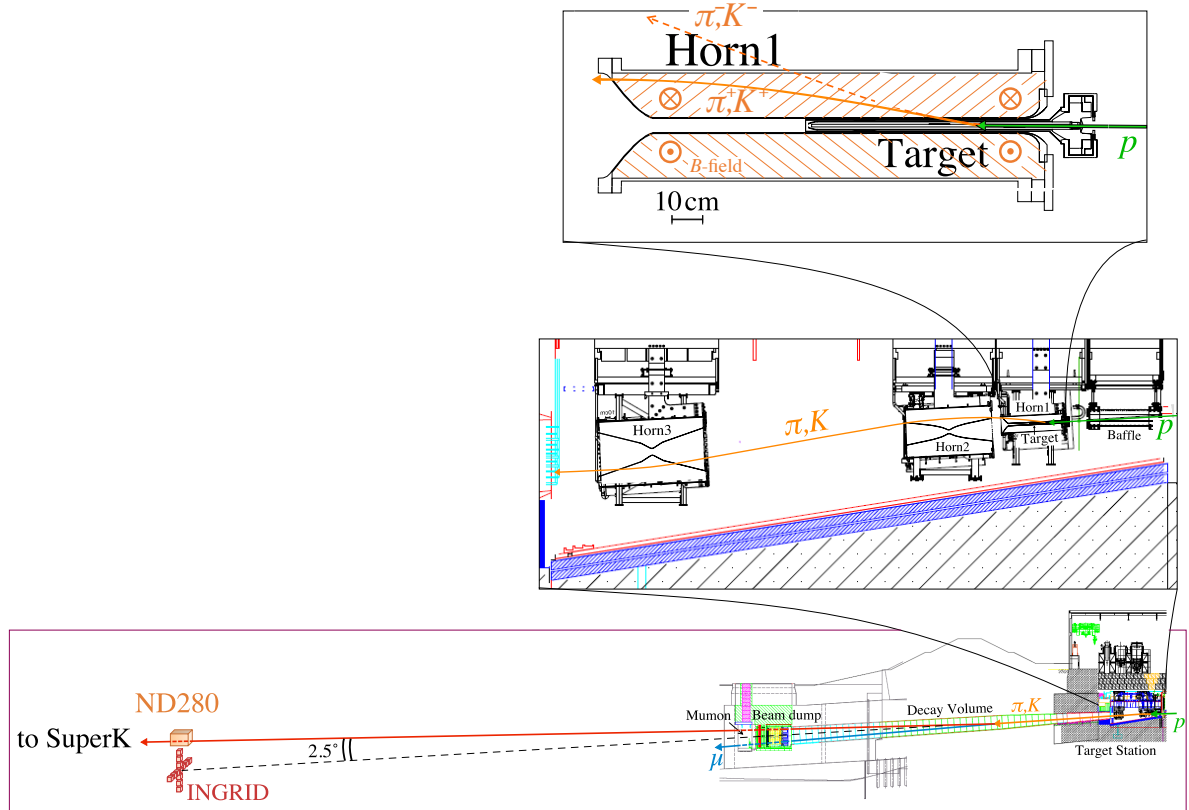


Figure 3.3: A comprehensive breakdown of the target region. The top schematic depicts the graphite target located inside the first magnetic horn. The second segment shows the locations of the other two horns. Finally, the bottom image identifies the tertiary beamline located after the third magnetic horn. This is followed by the beam dump and MUMON detector. The solid line signifies the off-axis beam path, the dashed is the beam axis. Figure taken from [74].

intrinsic wrong flavour neutrinos. Unfortunately, the horns are not 100% efficient at only focussing the correct charged particles. This creates a flux of antineutrinos in what should be a neutrino beam, known as “wrong sign” background. Referring to Figure 3.4, one can see the wrong sign flux for both FHC and RHC. Another form of contamination is known as the “intrinsic background”. In the instance of T2K, this refers to having an electron neutrino flux in a muon neutrino beam, through the decays:

$$\begin{aligned}\mu^+ &\rightarrow e^+ + \bar{\nu}_\mu + \nu_e \\ K^+ &\rightarrow \pi^0 + e^+ + \nu_e.\end{aligned}\tag{3.3}$$

Understanding and constraining the intrinsic background that exists in the flux is important for estimating the physics of oscillations and expectations in the far detector. If these neutrinos propagate to Super-Kamiokande and interact, it can create false oscillation signals leading to fake CP violation.

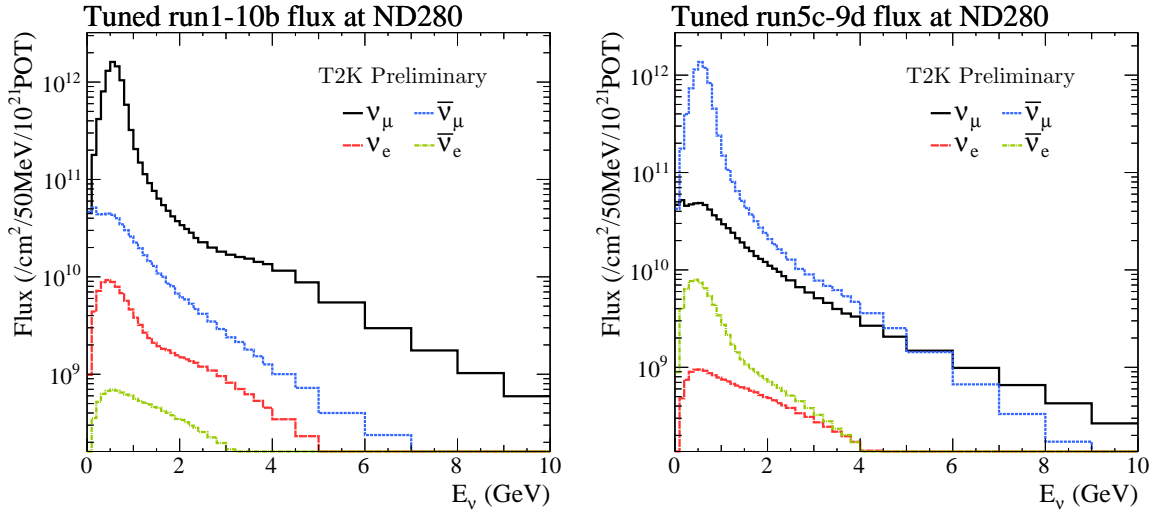


Figure 3.4: The left hand side shows the predicted flux in FHC mode, dominated by ν_μ but showing the wrong sign and intrinsic ν_e background. The right hand side shows the flux prediction for the RHC mode, dominated by $\bar{\nu}_\mu$. Figure taken from [75].

Figure 3.4 shows the flux for the various neutrino types, noting the logarithmic scale,

the wrong sign background and intrinsic background is less than 3% and 1% respectively, in both FHC and RHC run modes. One can see from Figure 3.4 that the wrong sign background is larger in the RHC mode. This is due to the proton producing a majority of positive particles after collision with the graphite target.

Downstream of the decay volume is the beam dump and the muon monitor detector, MUMON. The beam dump is used to stop remnants of the primary beam produced within the decay volume. Muons below an energy of ~ 5 GeV/c are stopped by the beam dump [52]. The muon monitor is designed to measure the direction and intensity of the beam on very short timescales.

T2K has specifically designed the experiment around the J-PARC beam to utilise what is known as the ‘off-axis technique’. Initially proposed by the Brookhaven National Laboratory [76], it works by situating a consortium of detectors at a small angle off axis with respect to the beam’s direction. Looking at the two-body pion decay, the neutrino energy can be calculated via equation 3.4, given the defined direction of the pion on the decay volume axis:

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi \cos(\theta_{OA}))}, \quad (3.4)$$

where E_ν and E_π are the neutrino and pion energy respectively, m_π^2 and m_μ^2 are the masses squared of the pion and muon, p_π is the momentum of the pion and $\cos(\theta_{OA})$ is the angle of the outgoing neutrino with respect to the parent pion. Referring to the bottom image in Figure 3.5, measuring the neutrino flux at differing off-axis ranges will change the peak of the average neutrino energy. Measuring the neutrino flux on-axis has a wide range of neutrino energies, peaking at roughly 1 GeV. Moving to 2.0° off-axis one can see that the peak shifts to a more narrow peak situated at 0.8 GeV. This reduces the spectrum of oscillation probabilities given the reduced range of neutrino energies. T2K is designed to use a beam 2.5° off-axis. This results in a narrower band to the

aforementioned values and importantly peaking at 0.6 GeV, represented in red in the bottom panel of Figure 3.5.

Looking at the top and middle segment of Figure 3.5, a neutrino energy at 0.6 GeV is the minimum for the muon neutrino disappearance and also a peak for the electron neutrino appearance. Having a peak neutrino flux at this energy optimises the sensitivity to these two oscillation channels. This also highlights why it is important to understand and constrain the wrong sign and intrinsic background in the neutrino flux due to the indistinguishable interaction products at SK. The distance of the Super Kamiokande detector from T2K sits on the oscillation peak for maximal electron neutrino appearance. This creates a ‘narrow band’ beam, creating an environment that has maximum sensitivity towards the measurement of δ_{CP} .

3.3 INGRID

The Interactive Neutrino GRID (INGRID) [78] is T2K’s only on-axis detector, located 280 metres downstream from the target. Its primary purpose is to measure the beam intensity and the beam direction, playing quite an important role in reducing the systematic uncertainty in the flux given the off-axis nature of the experiment. Comprising of 14 identical modules, stacked 7 vertically and 7 horizontally as shown in Figure 3.6, the detector has a $10 \times 10 \text{ m}^2$ total $x - y$ coverage. The detector was originally designed with two additional modules located above the horizontal, slightly off-axis, with the purpose of monitoring the beam’s axial symmetry.




Figure 3.5: Top: The muon neutrino disappearance oscillation probability with the listed oscillation parameters. Middle: The appearance probability for the muon neutrino to electron neutrino, with varying values of δ_{CP} . Bottom: The neutrino flux and how it changes with various off-axis values showing that 2.5 degrees aligns the peak with the appearance maximum. Figure taken from [77].




Figure 3.6: Layout of INGRID. With the vertical and horizontal formation, the middle detectors are located on the beam's axis. This figure also includes the unused off-axis modules. Figure taken [52].

The modules themselves consist of 9 iron and 11 tracking scintillator planes sandwiched together. The iron planes act as a target for the neutrinos, with a mass of 7.1 tonnes per module and the scintillator planes consist of 24 bars placed vertically and horizontally. As the neutrino interaction occurs, the final state particle will pass through the bar and create scintillator light that is transported via wavelength-shifting fibres (WLS). At one end of the WLS one will find the Multi-Pixel Photon Counter (MPPC) [79] that is used as a read out mechanism. The size of the photon read-out is proportional to the energy lost by the particle as it passes through the scintillator bars and this allows for reconstruction of the particle's energy, which in turn will allow one to find the neutrino energy. To decrease the background and improve the interaction reconstruction, the modules are surrounded by veto planes that contain scintillator bars segmented in the direction of the beam [78]. An exploded view of the INGRID module can be seen in Figure 3.7.

In the centre of the cross configuration of the INGRID detector where the vertical and




Figure 3.7: A schematic of the INGRID module showing the breakdown of iron and scintillator planes and the veto planes located on the outer segment of the module. Figure taken from [80].

horizontal arms overlap sat a sub-detector of INGRID called the proton module. The proton module was very similar to the modules that exist on the arms of INGRID, however it used more finely-segmented scintillator bars and had no iron layers. This module was used to monitor the charge current quasi-elastic interactions directly on-axis, and as a comparison for the Monte Carlo predictions of the beam and interactions. The proton module had much better tracking capabilities than the INGRID modules with the finer scintillator bars and lack of iron planes, allowing for the reconstruction of shorter tracks.

3.4 WAGASCI-BabyMIND

WATER-Grid-And-SCIntillator-Detector (WAGASCI) is a near detector located 280 m downstream of the beam and sits at 1.5° off axis with respect to the beam. Assessing Figure 3.5, it is visible to see the trend in the peak average neutrino energy as one tends to on-axis, even though 1.5° is not present, the trend shows that WAGASCI will have access to slightly higher neutrino energies than ND280. The objective of this detector is to help measure and constrain the neutrino interaction cross-sections and aid in reducing the uncertainties. The two components of the WAGASCI detector include

the main target area, containing the relocated proton module and two target modules. These target modules consist of a 3-dimensional scintillator grid format along with alternating $x - y$ scintillator planes. The cells located within the scintillator grid are filled with the target material, water and scintillator, as shown in Figure 3.8b. One key component to the detector's construction is the ability to drain the water module, known as water-in/water-out. The water-out mode allows one to make a measurement of the neutrino-carbon interaction, which is useful for constraining the expected background on the water-in measurements. This allows for a better constraint on the oxygen-neutrino cross-section, reducing the uncertainty at the far detector. Surrounding the target area are muon range detectors known as, Wall-MRDs. These are made up of sandwiched scintillator and iron planes, used to measure muons with a large angle track but also situated at a distance from the target area allowing for a time of flight measurement of the particle's trajectory.

Immediately downstream of WAGASCI is the Baby Magnetised Iron Neutrino Detector (BabyMIND) [81], as seen in Figure 3.8a. BabyMIND is made from alternating 33 magnetised iron plates that are 30 mm thick with 18 plastic scintillator planes. The magnetic field inside of BabyMIND is 1.5 T, giving excellent charge identification and momentum measurements for particles exiting the WAGASCI target volume. Muons produced by the beam's muon neutrinos interacting with the target material via CCQE will lead to a decrease in the background uncertainty at the near detector. It is not the primary muons produced, but depending on the interaction type, but a pion can decay into a muon, muon antineutrino and electron neutrino. This effects the wrong sign background depending on the horns polarity, so BabyMIND will also aid in reducing the uncertainty in the flux model. BabyMIND has been optimised to measure the range of muon energies that are expected to be produced by neutrino interactions with the detector being 1.5° off-axis.




Figure 3.8: Left is the full detector schematic for WAGASCI-BabyMIND. One can see the layout with WAGASCI being upstream and the surrounding Wall MRDs followed by BabyMIND. Also included is the location of a separate emulsion-based experiment that shares the housing, NINJA [82], which focuses on neutrino cross-section measurement and interaction physics. On the right, the 3D grid layout of the WAGASCI module, showing the alternating orientation of the scintillator planes and the grid cells containing the water target. Figures from [83, 84].

3.5 ND280

The Near Detector at 280 m (ND280) immediately follows downstream from INGRID. As described in Section 3.2, ND280 is situated 2.5° off-axis from the target with the intention of measuring various aspects of the analysis such as the flux, cross-section measurements, as well as wrong-sign and intrinsic contamination before oscillations have occurred. ND280 is composed of a group of sub-detectors contained within a metal “basket” as well as two sub-detectors outside of the basket, and their roles will be discussed below.

3.5.1 Magnet and Side Muon Range Detector

The ND280 magnet encases all but one of the various sub-detectors and is a refurbished magnet previously used by the UA1 [86] and NOMAD [87] experiments. It consists of two symmetric yokes¹ that contain water-cooled aluminium coils operated with a 3 kA

¹This is allow the detector to be opened up more conveniently.




Figure 3.9: An exploded view of the ND280 detector, with the listed sub-detectors visible and the surrounding magnet yoke. Figure taken from [85].

current that produces a uniform magnetic field of 0.2 T. The magnetic field allows for particle charge determination and high resolution of the particle's momentum within the sub-detectors by tracking the path of the particles. The magnetic yoke contains all of the various sub-detectors shown in Figure 3.9; the Electromagnetic Calorimeter, the π^0 detector, Fine Grained Detectors and Time Projection Chambers, except for the Side Muon Range Detector (SMRD).

The SMRD [88] consists of polystyrene scintillator bars which sit inside the air gaps of the magnetic yoke. A total of 192 horizontal modules and 248 vertical modules exist inside of the ND complex, each embedded with a wavelength-shifting fibre. The main purposes of the SMRD are to tag escaping muons that have a high angle with respect to the beam, but also to identify tracks of particles that originate outside the detector, such as those from cosmic muons.

3.5.2 π^0 Detector

The π^0 detector (PØD) [89] is the most upstream sub-detector in the ND280 complex. The primary function of the detector is to measure the rate of neutral current neutrino interactions that produce a single π^0 in the final state, hence the name. The PØD is a scintillator detector with water target layers split into two segments, the central water target and the upstream water target, shown in Figure 3.10. In between these interaction targets exist the scintillator layers, comprising of two perpendicular layers of triangular scintillator bars that are read out via WLS fibres (the same as those in INGRID) and brass sheets. The water target is contained in bags and is easily removable by draining the bags. Measurements are taken with and without water to allow for a better constraint on the cross section on water and the intrinsic ν_e background. The π^0 final state in SK is the largest background given the decay into photons, which are almost indistinguishable from electrons in a water Cerenkov detector, so it is important to understand them. Surrounding the central target are the ECals, also split by central and upstream segments. These are used as a veto for cosmic showers and to monitor interactions that have occurred in other sub-detectors. Importantly, they also measure high angle particle tracks that may only pass through few X and Y planes before leaving the detector, pivotal for correct event reconstruction.

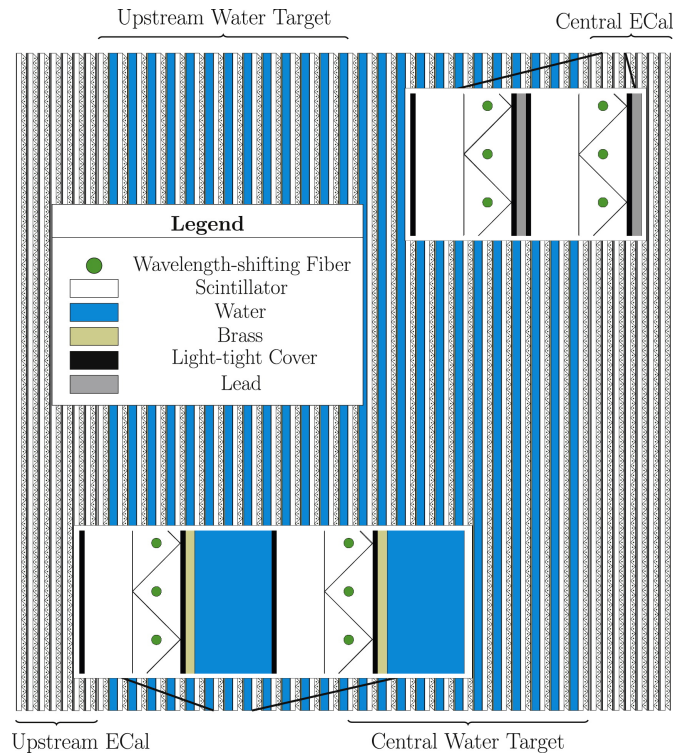


Figure 3.10: Schematic of scintillator layout in the PØD. Figure taken from [52].

3.5.3 Time Projection Chambers (TPCs)

Following on from the PØD sub-detector is the ND’s ‘Tracker’ sub-system. The tracker consists of three time projection chambers and two fine grained scintillator detectors sandwiched together, visualised in Figure 3.9. The three TPCs [90] are each built with a double box system, as seen in Figure 3.11. The outer box is grounded and filled with carbon dioxide to act as insulator for the inner box which forms a field cage and filled with an Argon based gas mixture. A central cathode splits the inner box with two anodes residing on either side of the TPC. The cathode has a high voltage supply creating an electric field parallel to the magnetic field produced by the surrounding magnet. As a charged particle traverses through the TPC, it ionises the TPC gas. Given the electric field, the electrons will drift toward the anode which contains MicroMegas [91] readout panels. Each TPC contains twelve MicroMegas modules, each with 1728 pads that read out the charge pulses. Using the read-out from the MicroMegas, one can find the energy loss ($\frac{dE}{dx}$) as a function of the particle true momentum. Comparing




Figure 3.11: A schematic of the two box design of the TPC. All three time projection chambers are identical [90].

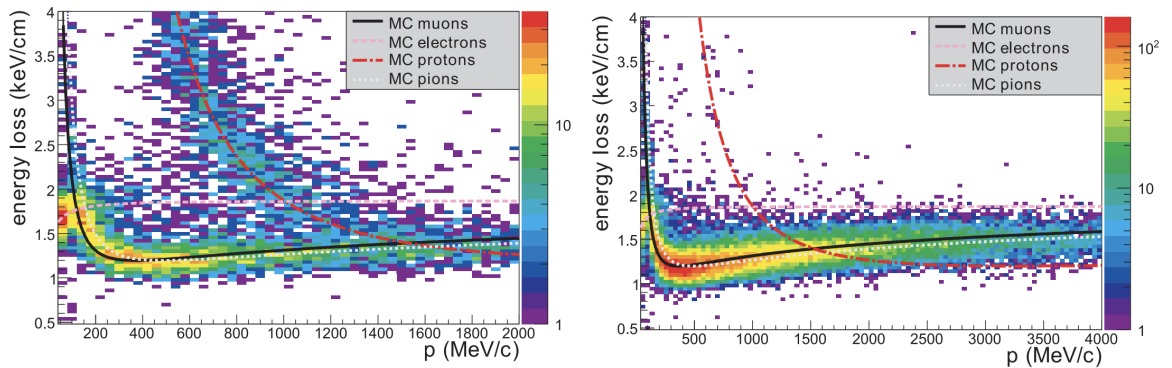


Figure 3.12: The Monte Carlo predictions for the energy loss against particle momenta drawn over the data obtained for one the TPCs. Left: positive particles traversing through the TPC. Right: Negative particles traversing through the TPC. Figure taken from [90].

this to a prediction from Monte Carlo gives a high efficiency for particle identification, depicted in Figure 3.12. Alongside this, the magnetic field from the enclosing solenoid causes the particles to have a circular projection on the read out phase which is used to calculate the particle's momentum. The direction of the curve allows for particle charge determination, giving better event reconstruction and discrimination between a neutrino and antineutrino interaction.

3.5.4 Fine Grained Detectors (FGDs)

Sandwiched between the three TPCs are the two Fine Grained Detectors [92], FGD1 and FGD2. These sub-detectors act as the target mass for neutrino interactions with a mass of 1.1 tonnes each. FGD1 is the more upstream detector of the two and comprises of fifteen XY planes, each plane consists of 192 horizontal and 192 vertical scintillator bars made from extruded polystyrene. Each bar is coated with titanium dioxide, TiO_2 , which reflects the scintillator light back to the WLS to reduce the cross-talk between fibres. FGD2 only contains seven of these XY modules and the space in between is filled with a 2.5 cm thick layer of water, allowing a comparison of the interaction rates in the two FGDs permitting cross-section measurements on carbon and water. Although it is important to constrain the interaction cross-section, this depends on successfully defining the interaction type. The FGDs have fine granularity to track particles produced at the interaction vertex that do not penetrate the TPCs, such as low-momentum protons. Measuring the energy loss in the FGD and comparing it to the MC prediction allows for good particle discrimination between protons and minimum ionising particles, shown in Figure 3.13.

3.5.5 Electromagnetic Calorimeters (ECal)

ND280 contains three electronic calorimeter components, with a total of thirteen ECal modules. Figure 3.9 shows the PØD ECal, which as the name suggest surrounds the PØD², the Barrel ECal (BrECal) that surrounds the tracker and the most downstream sub-detector, the downstream ECal (DsECal). The barrel and downstream ECals both enclose the tracker region and is often referred to as the tracker ECal (TrECal) [93]. The BrECal contains 31 scintillator-lead interleaved layers that alternate in orientation compared to the DsECals 34 layers and all scintillator bars contain WSL fibres, read out via MPPCs. It is also used to track particles ejected with a high-angle from

²Not to be confused with the central and upstream ECals that exist in the PØD.




Figure 3.13: Deposited energy against the track range in FGD1. The comparison with the Monte Carlo predictions for protons, muons and pions, solid, dash and dash-dot respectively, allow for good particle identification. Figure taken from [92].

the interaction vertex with respect to the beam, aiding the full reconstruction with the information from the FGDs and TPCs. Finally, it is also used to measure outgoing photons which is crucial for reconstructing π^0 events. The lead sheets act as an absorber to encourage electromagnetic showers and further improve the event reconstruction.

The PØDECal is slightly different to the TrECal, both in its construction and function due to the PØD sub-detector already containing the high-angle tracks. The PØDECal contains just six scintillator layers interleaved with thicker lead layers that are ~ 4 mm as opposed to TrECal's, 1.75 mm. Therefore, the aim of the PØDECal is to ensure high-angle acceptance of the PØD by catching photons that exit a water layer without passing through a scintillator layer.

3.5.6 ND280 Upgrade

Since 2022, T2K is entering a new era in its operational timeline with multiple upgrades underway, the first of which are planned upgrades on the beam power at J-PARC and




Figure 3.14: Exploded view of the ND plans, replacing the PØD with a super fine grained detector and two high-angle time projection chambers. Not shown here is the addition of time-of-flight panels that will surround the new upstream sub-detectors, attached to the basket. Schematic taken from [95].

the components to handle this increase [94]. In parallel with this upgrade, ND280 will also be revamped. Figure 3.14 is an illustration of the upgrade plans for the near detector complex [95]. The most upstream sub-detector, PØD, is being removed and replaced with three different sub-detectors. The middle layer sub-detector is the Super Fine Grained Detector (SFGD) (coloured grey in Figure 3.14) and surrounding that above and below are the High-Angle Time Project Chambers (HA-TPCs). The SFGD is comprised of two million plastic scintillator cubes, each with a volume of 1 cm^3 . Providing a target mass for the upstream sub-detector to two tonnes. All cubes will have three wavelength-shifting fibres passing through them giving an improved 3 dimensional reconstruction ability. This increase in spatial resolution means better particle identification, certainly for protons and electrons with tests currently ongoing to understand the sensitivity to neutrons. If it proves fruitful then this should lead to a better understanding of antineutrino interactions.

Either side of the SFGD sit the High Angle Time Projection Chambers. The HA-TPCs

are very similar to the existing TPCs further downstream, but the new versions lie horizontal in the basket. These will have a very similar internal layout, with a key difference being the use of resistive MicroMegas [96]. This reduces the number of readout planes that exist on the anode but adds a small resistive film which smooths out the charge distribution and improves spatial resolution allowing for excellent particle tracking. The final upgrade is the addition of six time-of-flight (TOF) panels. These will surround the new sub-detectors and attach to the metal basket. The panels consist of 20 scintillator bars each that will provide better track direction determination, further improving the interaction reconstruction. Overall, the upgrades bring exciting new particle tracking and identification potential and impressive particle discrimination. Another impact of the upgrade is the combination of the three new sub-detectors means that the angular coverage will be nearly the same as the far detector. SK has an angular coverage of 4π , giving the ability to reconstruct tracks in any direction. Prior to the upgrade, ND280 had a large preference to forward-going tracks. The new sub-detectors give accessibility to high-angle and backward-going tracks, which were previously difficult to reconstruct. Having a similar angular coverage will also lead to a reduction in systematic uncertainty for future analyses.

3.6 Super-Kamiokande

The Super-Kamiokande [21] (Super-K) detector is operated as an independent collaboration on its own but it is also used as the far detector for the T2K experiment. Located 295 km away and 2.5° off-axis from the neutrino beam in J-PARC for the purpose of residing on the first oscillation maximum for a 600 MeV beam. Super-K is a 50 kt cylindrical water Cherenkov detector located 1 km inside of the Mozumi Mine, Kamioka, giving sufficient shielding from cosmic muons to reduce the background from said cosmics. Super-K is split into two optically isolated volumes, the outer detector (OD) and the inner detector (ID). Figure 3.15 depicts the 41 m tall tank that has a 39 m diameter.




Figure 3.15: The Super-Kamiokande detector, showing its depth inside of Kamioka mine, taken from [97].

When a neutrino interacts via charged current with the water target, the produced lepton traverses the water at faster than the speed of light in the medium, which creates Cherenkov radiation. This radiation is then registered by the photomultiplier tubes (PMTs) that reside on the walls of the OD and the ID. The OD comprises of 1,885 sparsely spaced outward facing PMTs that are eight inches in diameter. This allows the detector to identify particles originating from a background source such as the aforementioned cosmic muons or neutrino interactions with the matter surrounding the detector. The ID contains 11,129 inward facing twenty inch PMTs that covers roughly 40% of the inner detector surface area. This large coverage allows for better resolution within the target region. Alongside the OD vetoing background noise, the ID also has a fiducial volume boundary of which any event vertex that sits outside of said boundary is rejected, increasing the purity of the samples.

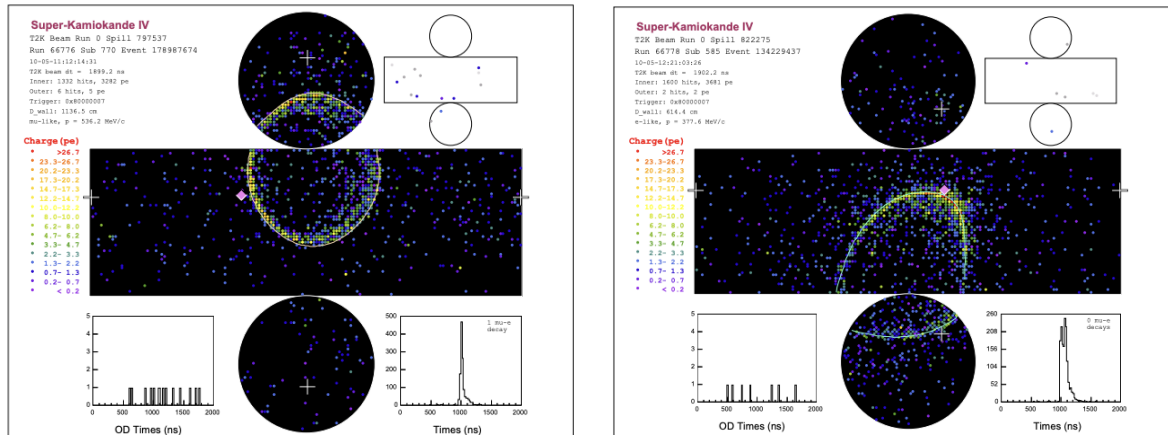
The photomultiplier tubes are a photosensor designed to measure an incident photon by utilising the photoelectric effect of the photon striking a photocathode. This releases an electron which is accelerated into a dynode that then emits a group of roughly three electrons. The group of electrons are then accelerated creating a multiplying cascade of electrons striking the chain of dynodes until finally a few million electrons reach the anode which generates an electric signal that can be read out. The size of the charge pulse is proportional to the number of incident photons. Events can be reconstructed using the energy and timing information alongside the surrounding PMTs.

3.6.1 Water Cherenkov Detection

As mentioned above, Super-K is a Cherenkov light detector. Cherenkov light occurs as a charge particle traverses through a dielectric medium at a greater velocity than the phase velocity of light. Analogous to a sonic boom, the traversing particle emits a shock wave of photons whilst its momentum is above the Cherenkov threshold. Equation 3.5 shows the angle at which the electromagnetic radiation is emitted due to the relativistic charged lepton, commonly referred to as the ‘‘Cherenkov angle’’;

$$\cos(\theta_c) = \frac{1}{\beta n}; \quad (3.5)$$

where θ_c is the Cherenkov angle, n is the refractive index for the medium (in this case water, $n = 1.34$), and β is the ratio of the particle’s velocity over the speed of light, $\beta = v/c$. The assumption is made that the charged leptons are travelling relativistically, giving $\beta = 1$ and thus the Cherenkov angle is $\sim 42^\circ$. With the information from the PMTs and the understanding of Cherenkov radiation, SK has an accurate selection method for particle identification despite not having a magnetic field like the near detector complex, which allows for particle charge identification, and thus neutrino and antineutrino discrimination. Using the PMTs readout and hit timing, one can



(a) Muon-like Cherenkov ring.

(b) Electron-like Cherenkov ring.

Figure 3.16: Examples of muon(a) and electron(b) like Cherenkov rings captured by SK. The colour of represents the energy of the incident photon and the hit timing information is used to recreate the interaction vertex, denoted with a white cross. Figure taken from [52].

work out the particle's flavour as well as the momentum and interaction vertex. The particle identification comes from the "fuzziness"³ of the reconstructed Cherenkov light ring, as seen in Figure 3.16, with the noticeable difference between Figures 3.16a and 3.16b being the sharpness of the ring. Muons have a relatively large mass and thus do not scatter very much as they traverse the water. They do lose energy which is why the ring dissipates towards the centre, as the particle's velocity reduces so does the rate of Cherenkov radiation emission. On the other hand, electrons are much lighter and scatter much more, causing electromagnetic showers that create this "fuzzy" appearance. With the knowledge of the particle leptonic flavour and our understanding of CCQE interactions, we can determine the flavour of neutrino that interacted with the water target in the far detector.

3.6.2 Analysis Samples

Briefly mentioned in the previous section was the distinguishable nature of the Cherenkov radiation produced by the electron and muon, giving information about the neutrino

³Yes, this is the official term.

involved in the interaction. Unfortunately the other products, such as the proton cannot be observed in SK due to their low velocity. There is the ability to detect an outgoing neutron via capture on hydrogen from an electron antineutrino interaction, however this has a very low efficiency. Therefore, the neutrino energy is reconstructed via lepton kinematics alone. The assumption is made that the neutrino interacts with a nucleon at rest, via CCQE and thus the neutrino energy can be reconstructed via:

$$E_{\nu}^{rec} = \frac{2(m_n - E_b)E_l + m_l^2 + m_p^2 - (m_n - E_b)^2}{2(m_p - E_l + p_l \cos(\theta_{beam}))}, \quad (3.6)$$

where m_n, m_p and m_l are the rest masses for the neutron, proton and lepton respectively. E_b is the binding energy for the nucleon, E_l and p_l are the respective energy and momentum of the lepton, $\cos(\theta_{beam})$ is the angle of lepton with respect to the neutrino direction and E_{ν}^{reco} is the energy of the reconstructed neutrino. For the newer analyses, SK can deduce that a charged pion was produced in the initial interaction via charged current resonance. The reconstruction equation becomes:

$$E_{\nu}^{rec} = \frac{2m_p E_l + m_{\Delta^{++}}^2 - m_p^2 - m_l^2}{2(m_p - E_l + p_l \cos(\theta_{beam}))}, \quad (3.7)$$

where E_{ν}^{rec} , m_n, m_p, m_l, p_l, E_b and $\cos(\theta_{beam})$ are the same as those described for equation 3.6. Finally, $m_{\Delta^{++}}^2$ is the mass squared of the Delta Baryon.

3.6.3 Gadolinium Doping

The lack of magnetic field⁴ means that SK does not have sensitivity to the charge of the outgoing lepton, so constraining the wrong sign background at the near detector is very important. In previous analyses, this was included in the simulation using the neutron capture on a hydrogen nucleus which would de-excite and release a single photon of 2.2 MeV, which is just on the threshold for SK reconstruction and thus has a low efficiency. The incoming electron antineutrino would interact with a proton via CCQE,

⁴SK contains magnetic coils to ensure there is no interference from the Earth's magnetic field, so enough to affect a particle's momentum.

producing a positron and a neutron. SK would see an electron-like event but with a large inefficiency in observing the neutron information. Thus the electron antineutrino would be indistinguishable from an electron neutrino. To aid this, the SK tank was opened and drained in 2018 to replace broken PMTs, perform general cleaning and to seal leaks that had occurred. After this, SK could be doped with Gadolinium sulphate, $\text{Gd}_2(\text{SO}_4)_3$, which allows neutrons to be captured in the Gadolinium nuclei with 75% efficiency [98]. This produces a gamma cascade signal of 8 MeV on a time scale of 20 μs after the initial lepton signal [99]. This will help with neutrino-antineutrino discrimination in future analyses. SK will also benefit from the Gd doping with searches for low energy (\mathcal{O} MeV) electron antineutrinos that may have originated from pre-supernova sources [100].

3.7 NOvA Experiment Overview

Chapter 5 shows a joint-fit between T2K and NOvA. Therefore it seems sensible to give an overview into the NOvA experiment and its detector set up to highlight the differences between the two experiments. This section contains a general overview, but a much more detailed description can be found in the Technical Design Report [53].

The neutrinos from the Main Injector (NuMI) [101] Off-Axis ν_e Appearance (NOvA) Experiment [53] is another long-baseline neutrino oscillation experiment, situated in the United States of America. Like T2K, the experimental set-up consist of a beam target, near detector and a far detector and is optimised for ν_e appearance and ν_μ disappearance.

Starting at the NuMI facility in Fermi National Accelerator Laboratory (Fermilab), Illinois, protons are accelerated using a multi-ring system to 120 GeV energy before




Figure 3.17: A diagram of the NuMI beamline including the target, decay pipe and the various monitors and absorbers, taken from [53].

colliding with a 0.95 m long graphite target. Secondary pions and kaons are focused using a pair of magnetic horns in the same way as T2K. The charged hadrons are focused towards a 675 m decay volume where the products of the decays are the same as described in section 3.2. Similarly to J-PARC, the polarity of the horns defines the charge of hadrons focused toward the decay volume and the NuMI neutrino beam can run in a Reverse Horn Current (RHC) or Forward Horn Current (FHC) mode. Further down the beamline one will find the hadron absorber, designed to absorb the remaining protons and mesons from the beam whilst the sequential muon monitors and ‘earth shields’ are designed to absorb the remaining muons. The schematic of the beamline process is illustrated in Figure 3.17.

Like T2K, the goal is to measure and constrain the values of the PMNS matrix and an off-axis configuration allows NOvA to have a peak neutrino energy near the first oscillation maximum, although the neutrino energy spectrum is slightly wider than T2K’s. In order to measure neutrinos before oscillation, there is a 300 tonne segmented liquid scintillator detector located 1.015 km downstream of the graphite target and 0.8° off-axis, situated 100 m underground in the Fermilab facility with the dimensions of $14.3\text{ m} \times 2.9\text{ m} \times 4.2\text{ m}$. The ND comprises of a grid-like structure made of extruded PVC spanning 16 cells wide. This is grouped with another 16 cells to form a plane-like module. The modules are placed perpendicular to the beam and alternate sequentially




Figure 3.18: A to-scale schematic showing NOvA's near and far detector sizes with a small figure showing the grid formation for the PVC cells. To the far right one can see the layout of the wavelength-shifting fibre looping through the cell [103].

in the vertical (Y) and horizontal direction (X), giving the ability for three dimensional tracking. The near detector contains 214 planes of these alternating modules giving a total of 21,192 cells [102]. For the near detector the extrusion reaches a length of 3.9 m long. The cells themselves are filled with liquid scintillator which also houses a looped wavelength-shifting fibre leading to an avalanche photo diode (APD) which is used as the read out mechanism. This is rendered in Figure 3.18.

NOvA's far detector is located 810 km away from the proton target in Ash River, Minnesota. Unlike T2K, NOvA's near and far detector are functionally identical with the main contrast being the size of the target volume. The extruded PVC length for the far detector cells is 15.6 m as opposed to the near detector's 3.9 m, creating a target mass of 14 kt. The far detector contains 896 planes and houses a total of 344,054 cells, each with a wavelength-shifting fibre looped through it leading to an APD. Interestingly, the far detector is situated on the surface, which is unusual for a particle detector due to the very high background from cosmic particles. In order to combat the high level of background noise from cosmic muons, neutrinos and other higher energy particles,

the detector is housed under an overburden comprising 1.2 m of concrete and 15 cm of barite [104].

Chapter 4

Statistical Inference

Statistics: the only science that enables different experts using the same figures to draw different conclusions.

Evan Esar

This chapter will discuss the Bayesian approach used in the oscillation analysis produced by the **Markov Chain Monte Carlo 3**, MaCh3, package. MaCh3 uses a Markov Chain Monte Carlo (MCMC) [105] method to produce a simultaneous fit between the near detector and far detector. Markov chains sample the parameter space of a model, calculating likelihood values for each iteration, and build a posterior distribution using Bayes' theorem. Bayesian methodologies such as marginalisation and credible intervals will also be related to the fitter results and their conclusions. An analysis overview, including a discussion on systematics, will also be given in the later parts of the chapter.

4.1 Bayes' Theorem

Bayesian Inference is a statistical method which uses Bayes' theorem to update a hypothesis as more data or information becomes available. Bayes' theorem itself can be

described as finding the conditional probability of an event happening with a set of hypothesised data, testing a hypothesis, H , with prior information, B , and data set, D . Combining the various strands of information, one can find the probability of the posterior distribution for a given parameter, using Bayes' Theorem:

$$p(H|D, B) = \frac{p(H|B)p(D|H, B)}{p(D|B)}, \quad (4.1)$$

where $p(H|B)$ is the prior probability and $p(D|H, B)$ is the probability of the data set being observed given the hypothesis and set of priors, B , also known as the likelihood function, and this will be discussed further in section 4.5. The denominator, $p(D|B)$, is a normalisation factor used to keep the posterior probability from violating unitarity [106]. Finally, $p(H|D, B)$, is the posterior probability, the probability of the hypothesis being true, given a set of data and prior information. Relating equation 4.1 to the oscillation analysis; the hypothesis, H , are the oscillation parameters and B is the collection of priors placed on the model parameters used to define the prediction. When simplifying Bayes theorem, it is easier to see the relationship between the posterior distribution, likelihood function and prior distributions given that:

$$p(H|D, B) \propto p(D|H, B) \times P(H|B), \quad (4.2)$$

and how the choice of prior distribution can have an effect on the posterior distribution, this is depicted in [107].

4.2 Markov Chain Monte Carlo

It is important to give context behind Monte Carlo methodologies, and why a simple Monte Carlo would not be effective for the T2K oscillation analysis due to naive sampling. Section 4.2.1 gives an introduction to the methods, and Section 4.2.2 describes

what is carried out in the analysis.

4.2.1 Monte Carlo

To build a suitable distribution and gather information about best fit values, lots of parameter values would need to be sampled. This can be accomplished via a Monte Carlo (MC) method [108]. The MC algorithm uses complete random sampling of the model points within a distribution. If this distribution lies within a set of limits then this is accepted and another random point is assessed. When the number of dimensions¹ is low or very well constrained, i.e the boundaries on the parameters are tight, then this would be an effective technique. T2K's full oscillation analysis contains 779 systematics (100 Flux, 75 cross-section, 552 ND280 detector systematics, 46 SK detector systematics and 6 oscillation parameters). Figure 4.1 shows one of the correlation matrices used in the T2K analysis, housing the flux and cross-section parameters. Other matrices also contain the Near detector specific systematics and the Far detector specific systematics. As one can guess, randomly sampling this space would be extremely inefficient and computationally expensive and there will end up being simulations where there is more sampling happening in low likelihood regions compared to more 'favourable' parameter spaces. Therefore, an alternate approach is required.

4.2.2 Markov Chain Monte Carlo

MCMC is a stochastic process in which the points that are sampled are dependent on the current point in parameter space and accepted or rejected, but independent of any other point in the chain. This allows the simulation to sample regions in parameter space that are of better interest to the analyser, depending on the condition set. In the case of this analysis, when the parameter space of the model is sampled, it is desirable to find the regions in which the likelihood is highest for obtaining the data. There are conditions that must be satisfied for the MCMC to successfully converge to a stationary

¹The number of parameters that exists in the model being sampled is the number of dimensions.




Figure 4.1: The combined flux and cross-section correlation matrix used in the 2021 oscillation analysis, showing some of the parameters within the model. The red, heavily correlated, bins are the flux normalisation parameters, and those named are the cross section parameters within the model. Randomly sampling around this subset of parameters alone would be an inefficient sampling technique. Figure taken from [109].

region:

- **Irreducibility:** Given any state, there is a non-zero probability that any other state will be sampled in a finite number of steps.
- **Recurrence:** When the chain is within the stationary distribution, all subsequent steps must be within that stationary distribution.
- **Aperiodicity:** The chain will not visit a given point with any fixed regularity.

A more thorough discussion on the conditions described above can be found in [110]. As the chain runs, it will build this multi-dimensional posterior probability, allowing an analyst to understand what parameter values are most probable within the model given some observed data set.

4.2.3 Metropolis-Hastings Algorithm

The Metropolis-Hastings algorithm [111] is an algorithm used within the MCMC to decide if a sample is accepted or rejected, that utilises the proposed steps probability, and compares it to the current state. Relating this to the analysis, the point in parameter space will undergo a likelihood calculation, $\mathcal{L}_{\vec{x}_n}$, where, \vec{x}_n , is the vector of parameter values in the current step. The methodology can be simplified into 5 steps:

1. Begin the MCMC with a proposed step, consisting of a set of values for each parameter in the model, \vec{x}_n , the initialisation.
2. Step into a region of parameter space and create a vector of parameter values, \vec{y}_n , the proposal. It is important to note that \vec{y}_n is in the vicinity of \vec{x}_n , otherwise this would inefficiently explore the parameter space.
3. Calculate the likelihood function for the parameter values of \vec{y}_n .

4. Compare to the likelihoods of \vec{x}_n and \vec{y}_n and accept or reject the step. The condition for acceptance or rejection is based on the likelihood comparison between the proposed, \vec{y}_n , and the current step, \vec{x}_n , such that:

$$\alpha(\vec{x}_n, \vec{y}_n) = \min \left(1, \frac{P(\vec{y}_n|D)}{P(\vec{x}_n|D)} \right). \quad (4.3)$$

A random number is then generated from a uniform distribution between 0 and 1. If this value is smaller than α , then the step is accepted and $\vec{x}_{n+1} = \vec{y}_n$. If the step is rejected then $\vec{x}_{n+1} = \vec{x}_n$ for the next iteration.

5. Repeat steps 2-4 for N iterations.

Looking at the steps outlined above, step 4 is very important for building a posterior. When the step is proposed and the likelihood is calculated, this will be compared to the current steps likelihood and if the new proposed values in parameter space are more probable then the step is accepted and the parameter values are saved. However, if the likelihood is smaller, i.e $\mathcal{L}_{\vec{x}_n} > \mathcal{L}_{\vec{y}_n}$, the step will have to undergo a conditional basis to determine acceptance. Figure 4.2 shows a flowchart of the acceptance/rejection method outlined above.

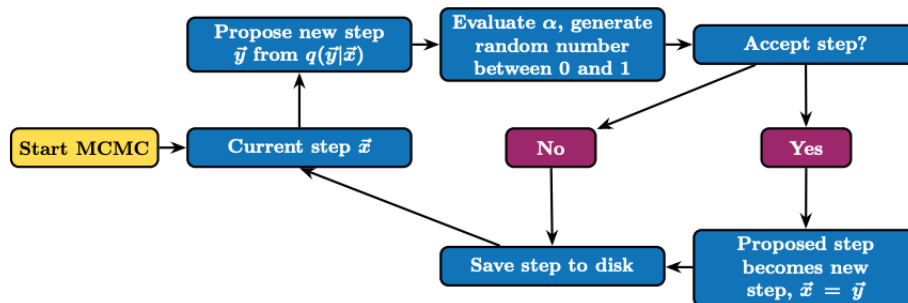


Figure 4.2: A flow chart depicting the step acceptance algorithm used in the MCMC. Figure taken from [112].

The chain will then move and propose a new step in the parameter space. If the proposed step is rejected, then step $\vec{x}_{n+1} = \vec{x}_n$, and \vec{x}_n is saved to disk. This will build

the posterior probability density histogram, where a region that has more saved points will have a higher posterior probability than that of regions with less saved points. The rejected step must not be saved to disk as this will skew the posterior probability distribution. It is important that if the likelihood for the proposed step is smaller than the current state that the step may still be accepted, otherwise this would just work as a minimizer. Allowing less probable steps to be accepted ensures that every region in the parameter space has a non-zero probability, fulfilling the aforementioned irreducibility criteria.

4.2.4 Diagnostics

It is important to check the chain progression in order to improve the efficiency at which the chain will converge to its stationary range. The tuning features of the MCMC include: the step proposal function, the step-sizes of the parameters and the total number of steps determined for the chain, N . The total step count does not need to be pre-determined and the chain can always be continued, but due to computational expense one would prefer the convergence and posteriors to be reached in the lowest number of steps. The proposal function itself comprises of a multi-variate Gaussian for each of the parameters with the prior error on the parameter defining the sigma on the Gaussian, multiplied by a scale factor that is individually tuned, as well as a global covariance scale that defines a common scale for all parameters at the same time. These step sizes will be changed as the model changes with correlations and anti-correlations accounted for prior to running the MCMC. However, this becomes increasingly difficult with high dimensional models. Fortunately, regardless of an initial position in parameter space the chain will always arrive at the stationary distribution, but the sooner the better [113].

The step sizes are tuned to ensure that once the chain has converged upon its station-

ary region, that the parameter space is effectively sampled. Step size tuning becomes important because having a small step size will increase the acceptance rate of the steps but increase the time taken to find and map out the posterior likelihood [114]. Having a small step-size also decreases the chains ability to step outside of high density regions, stopping the MCMC from sampling all of the parameter space. The opposite effect can happen if the step size is too large. The accepted steps ratio will be drastically reduced as the MCMC is more likely to enter an unfavourable region. There is an optimum acceptance ratio for the MCMC, suggested to be 0.23 [115] for the probability of accepting a step. Unfortunately, there is no golden equation for calculating the correct step-size and proposal function for a given model and is more often a recursive process with trial and error for each parameter. The right hand side of Figure 4.3a depicts a poor step variation for a tested step scale. The ‘jumpy’ range in parameter variation signifies a poorly sampled parameter space, affecting the number of accepted steps and parameter space that is visited by the MCMC and suggesting that the step size is too small. Alternatively, one can see the coverage of the parameter space that is achieved with an appropriately tuned chain in Figure 4.3b. As this process has been carried out for previous analyses, there are benchmarks than can be compared against. For example, if one runs a fit using a known ND constraint, then the expected parameter variation range would be very similar to the prior error range. It is worth noting that the step acceptance rate alone is not sufficient to determine the stability of the chain, as a poorly mixing chain can have a good acceptance ratio and efficiency, so other diagnostic criteria are also used.

Burn-In

As the MCMC chain embarks on its path to convergence, it may have an initial starting point in parameter space that is highly unfavourable for the given data. It may take some iterations before the chain finds its stationary range and correctly samples from

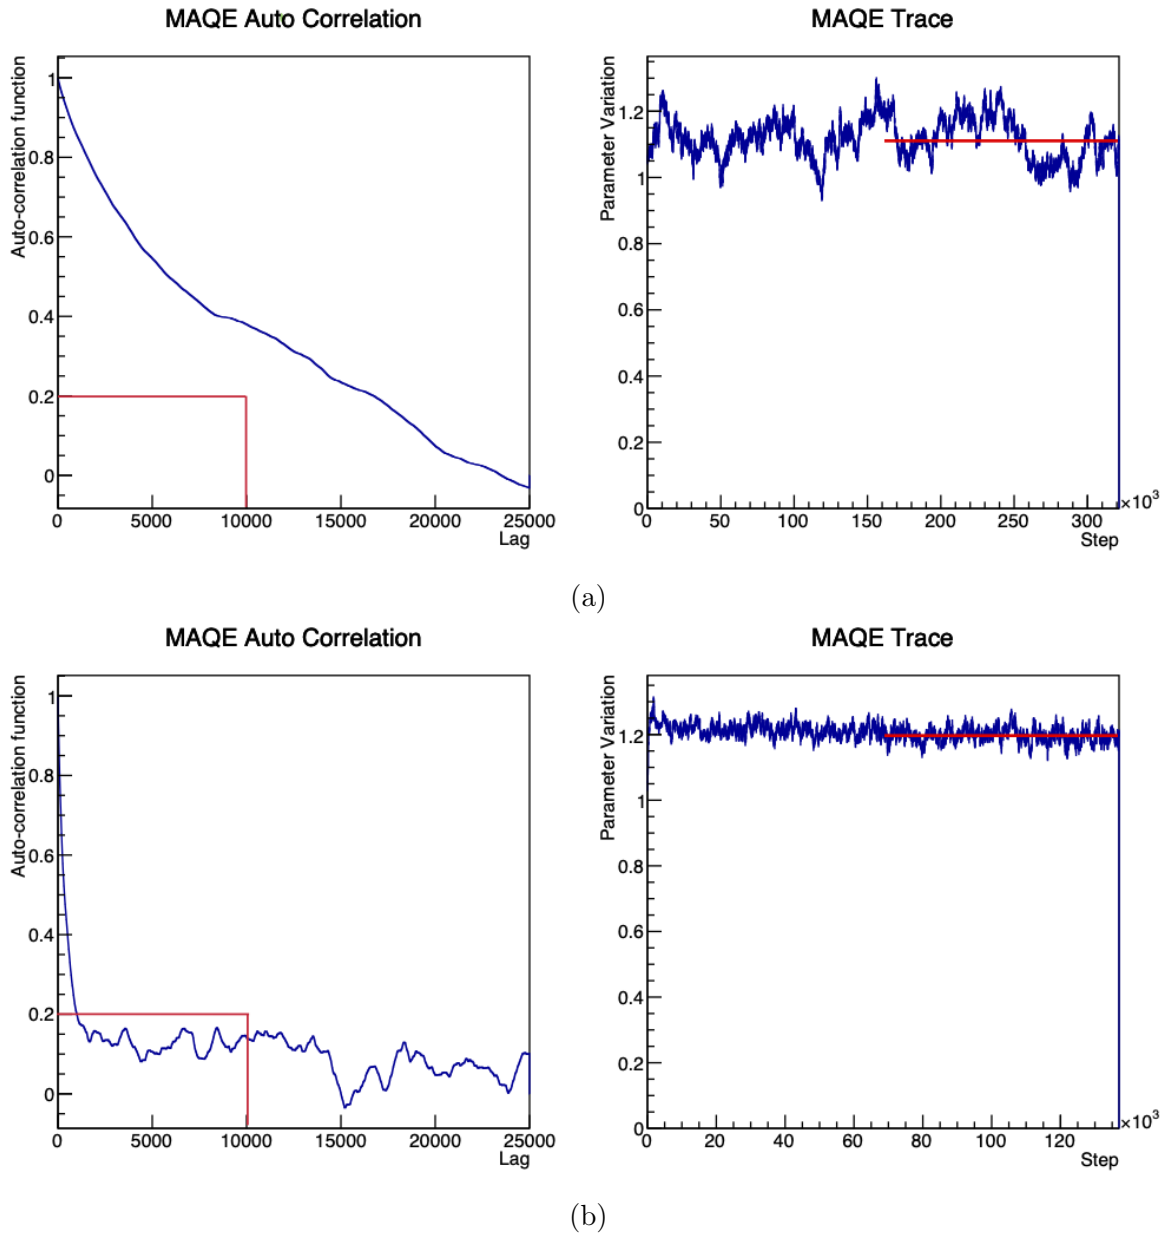


Figure 4.3: A comparison of autocorrelation functions and parameter variations for a poorly tuned MCMC (a) and a well tuned MCMC (b), specifically looking at the axial mass cross-section parameter M_A^{QE} . The red box on the left hand side shows the acceptance criteria for the autocorrelation functions. The red line on the trace plots represents the stationary range for the Markov Chain Monte Carlo, and represents the post burn-in phase. Lag is defined as an iteration of the value, i , through the chain of length, N , as described in Equation 4.4.

a favourable likelihood region. This initial phase of the chain is known as the Burn-In phase. Ideally, the step-size and proposal function will be well tuned so that this Burn-in is as small as possible. It is important to cut these points from the finalised posterior distribution as this will poorly represent the model's posterior probability function. One can see in Figure 4.4 how the chain travels to the stationary distribution through the chain steps. The Burn-In phase for Figure 4.4 would be set at steps less than one hundred thousand. This will vary on each iteration of a fit and is a good way of diagnosing the stability of the chain.




Figure 4.4: An example of a set of chains converging towards a stationary distribution into a low log-likelihood, favourable region. The long decline feature at low step values is known as the "Burn-In" phase. Figure taken from [116].

Auto-Correlation

Aside from looking at the Burn-In phase, another diagnostic tool is the auto-correlation function. The auto-correlation function will show a parameter's correlation with itself, k steps further along in the chain. Using the auto-correlation formula:

$$A_k = \frac{\sum_{i=1}^{N-k} (X_i - \bar{X})(X_{i+k} - \bar{X})}{\sum_{i=1}^N (X_i - \bar{X})^2}, \quad (4.4)$$

where A_k is the auto-correlation value, N is the total number of steps, k is known as the lag where k is an iteration of the value through, i , through the chain of length, N .

\bar{X} is the mean of the parameter distribution from the chain and X_{i+k} is the parameter value for X that is k steps ahead of iteration, i . The general rule of thumb for an acceptable auto-correlation is less than 0.2, at lag k of 10,000 as we want the steps to be as random as possible. The left hand panel of Figure 4.3a shows the effect of a poorly tuned step size on the auto-correlation function. Comparing this to the left pad of Figure 4.3b, one can see that the chain converges on the aforementioned criteria for the auto-correlation function.

Batch Mean

A final method for chain diagnostics is the Batch Mean method. This will take the stored parameter values from a given chain, dividing them into equal groups and the mean of these groups are compared. If a chain is covering the parameter space correctly then there should be minimal variation in the chain's subset mean. However, one should not use this method alone. Depicted on the upper right hand plot of Figure 4.3a, the variation in parameter value is minimal which would suggest a similar batched mean result to that produced by the same chain used in Figure 4.3b, but as mentioned it does not scan the parameter space well. A variation of the batch mean method is the acceptance batch method, rather than taking the mean of the steps in intervals, one can look at the acceptance rate in said intervals which should be also be consistent with each other.

4.3 Posterior Analysis

When the MCMC concludes the required number of steps, the full posterior distribution function is humanly impossible to interpret. With a full ND and FD joint-fit, the number of dimensions of the posterior would exceed 800. Most of these are parameters such

as ND detector systematics, flux, cross-section and SK detector systematic parameters. These are often referred to as nuisance parameters. They are important for the model, and analysis and are utilised in the MCMC via a correlation matrix shown in Figure 4.1.

Given the difficulty of interpreting such a large dimension posterior, a process called marginalisation is applied. Each systematic can be marginalised to assess systematic constraints and behaviour. We can remove these by integrating the posterior over the nuisance parameters allowing one to visualise the more parameter of choice on a one dimension or two dimensional scale. These are known as ‘parameters of interest’. This is to account for the effects that the nuisance parameters have on the final posterior distribution. Marginalising over the nuisance parameters via:

$$P(\theta_{\text{poi}}|D) = \int P(\eta, \theta_{\text{poi}}|D) d\eta, \quad (4.5)$$

will find the marginal posterior. Here, η is the vector of nuisance parameters and θ_{poi} are the parameters of interest. Often a one or two dimensional histogram showing the parameter entries accepted in the MCMC chain is the easiest method of interpreting MCMC outputs to create conclusions on the fit outcomes. One can compare the marginalised posterior distributions of all the parameters to their prior in order to understand how well the model describes the data. It is almost always the case that two of the oscillation parameters, θ_{12} and Δm_{21}^2 , are marginalised over due to T2K’s low sensitivity and the external constraint that is available from solar neutrino experiments as a prior.

Point Estimates

Given the marginal posteriors, one would aptly look for a point estimate, error or a value of best fit for a parameter in an imposed model given some data. This can be done in a number of ways, with none being the ‘correct’, but all three being used:

- **Highest Posterior Density (HPD)** - The highest posterior point, the modal value of the distribution.
- **Mean and RMS** - Taking the arithmetic mean of the parameter values in the chain and calculating the Root Mean Square to equate as the error,
- **Gaussian** - Fitting a symmetric Gaussian to the posterior distribution, using the Gaussian's mean and sigma as the point estimate and error respectively.

The HPD is created by binning the marginalised posterior into a 1D or 2D histogram, depending on the number of parameters of interest, and using the central value of the highest bin as the point estimate. This is the most common method used in MaCh3, as the other two options depend on the posterior being a Gaussian, which is not the case for some cross-section parameters, illustrated in Figure 4.5. If the marginalised posterior is Gaussian, the agreement is good, however, when the marginalised posterior distribution is non-Gaussian, the agreement can be poor. Following from this, there is also a feature known as marginalisation bias. When marginalising over non-Gaussian posteriors one can find a 'shift' in the point estimate for the posteriors, however this is a known and usually small effect.

4.3.1 Credible Intervals

Credible intervals are a way of showing a degree of belief that a given range of the model will contain the true value of an unknown parameter. Utilising the marginal posteriors it is possible to build a Bayesian credible interval from the binned output for individual or pairs of parameters. An X% credible interval is a statement of belief that there is X% probability that the true value of this unknown parameter lies within the interval range. This differs from the frequentist confidence interval, an X% confidence level suggests that if one was to repeat the experiment N times then the true value of the parameter would sit within the interval range for X% of N . The credible interval can be calculated using equation 4.6,

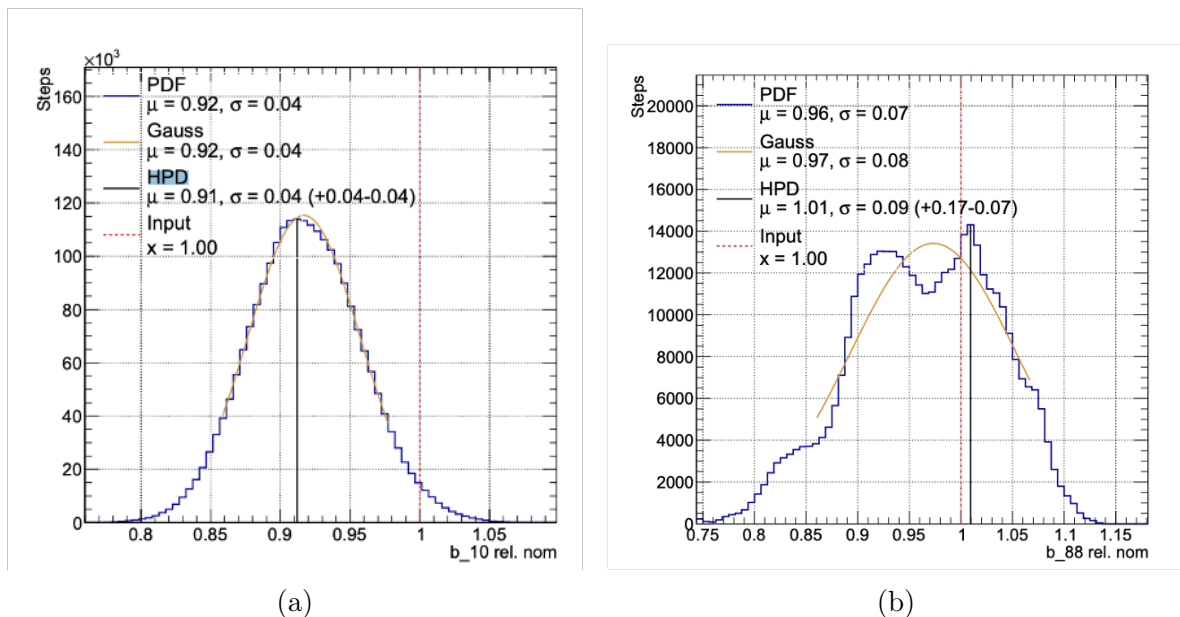


Figure 4.5: Two marginalised posteriors for flux normalisation parameters. Left: A Gaussian posterior showing good agreement between all three point estimate methods. Right: a non-Gaussian emphasising the disagreement between the point estimates. Blue is the arithmetic mean method, Gold is fitting a Gaussian and black shows the highest posterior density.

$$\alpha = \int_{\theta_L}^{\theta_U} \pi(\theta|D) d\theta, \quad (4.6)$$

where, $\pi(\theta|D)$, is the posterior density, θ_U and θ_L are the parameters upper and lower limits respectively, and α multiplied by 100% is the percentage of the posterior distribution. An $\alpha = 0.68$ equates to a 68% credible interval. Using the HPD method, one starts with the highest populated bin and progressively adds the next highest bin until a value, α , of the posteriors total integral is found.

Sequentially counting from the most populated bin under the desired credible interval is trivial for a Gaussian posterior, however for a non-Gaussian posterior, this can produce unintuitive results. If one was to look at the posteriors for Δm_{32}^2 , illustrated in Figure 5.14a, the degeneracy that exists related to the mass ordering shows two peaks in the posterior. This creates a disjointed contour which should not be neglected when looking at model preferences.

4.3.2 Prior Reweighting

Credible intervals are a tool for building conclusions of the unknown parameter values of the model in the posterior, but it is also important to look at the effects of changing the priors. Looking at Equation 4.1, the reliance on a prior to build the posterior distribution used to analyse a model and data set is pivotal. As stated before, the posterior is proportional to the likelihood multiplied to the prior and a very beneficial tool in the Bayesian approach is the ability to alter the prior probability of a parameter and see the effect on the posterior, without needing to re-run the MC by creating a weight. If one was to replace an original prior, $f(x)$, for parameter value x , with a new prior, $f(x')$, for updated parameter value x' , then finding the weight is simply taking the ratio of the two using Equation 4.7,

$$\omega(x) = \frac{f(x')}{f(x)}. \quad (4.7)$$

Each step in the MCMC will have an individual weight calculated and this technique is often used to see the effects on the posterior when updating the model. This could be the application of external constraints on specific parameters, such as the reactor constraint or checking the effect when enhancing or suppressing a cross-section parameter in the Monte Carlo, without the need of generating a new MC with the updated prior, as the results will be the same. This saves a large amount of computational time and resources.

4.3.3 Bayes Factor

It is often desirable to test different models or hypotheses on the same data set to see which hypothesis is more probable. For example, one can find if the model being fit to a given data set has a preference for the normal ordering or inverted ordering in the mass

ordering degeneracy. This is achieved by calculating a Bayes factor. If one refers back to Equation 4.1, the global posterior is proportional to the global likelihood multiplied by the model's prior information. If we marginalise this for a specific model we can find the marginal posterior as we would for an individual parameter. The Bayes' factor is in essence the ratio of the posterior probability and prior probability and as such can be represented in the following way:

$$\frac{P(H_1|D, B)}{P(H_2|D, B)} = \frac{P(D|H_1, B)}{P(D|H_2, B)} \times \frac{P(H_1|B)}{P(H_2|B)}, \quad (4.8)$$

The components to the left hand side of Equation 4.8 represent the marginal posterior probability for hypotheses H_1 and H_2 . The middle term denotes the ratio of marginal likelihoods for the two hypotheses and the far right is the prior information. Specifically in this thesis we look at the octant degeneracy (the sign of $\theta_{23} - \pi/2$) and the mass hierarchy (the sign of Δm_{32}^2) as independent hypotheses to see which of those the model finds more probable. We set the prior probability for these hypotheses to be equal, meaning that the probability of stepping into a positive or negative Δm_{32}^2 parameter space is 50%, and the same for the lower and upper octant of θ_{23} . Thus, the Bayes factor will be equivalent to the posterior odds ratio:

$$\frac{P(D|H_1, B)}{P(D|H_2, B)} = B.F_{1,2}, \quad (4.9)$$

where $B.F_{1,2}$ represents the Bayes' factor quantity. However, it is important to quantify the Bayes factor and its strength, rather than just stating that there is a preference because this can induce a false conclusion, albeit at the reader's discretion. In order to aid this, the significance of the Bayes factor is categorised via the Jeffreys scale [117]. Table 4.1 denotes a modified version of the original scale taken from [118]. Relating this to the mass hierarchy problem, where normal hierarchy is H_1 and inverted hierarchy is H_2 , one can simply take the ratio of entries in the MC chain for $\Delta m_{32}^2 > 0$ over $\Delta m_{32}^2 < 0$, to get a Bayes factor for NH/IH hierarchy, given the equal prior probability.

Bayes Factor	Grades of Evidence
1 to 3	Anecdotal in favour of H_1
3 to 10	Moderate in favour of H_1
10 to 30	Strong in favour of H_1
30 to 100	Very Strong in favour of H_1
> 100	Extreme in favour of H_1

Table 4.1: The Jefferies scale for the Bayes factors and their corresponding grades of evidence.

4.3.4 Posterior Predictive Distributions

An aspect that is important to testing the model against the data is the posterior predictive approach. As mentioned in Section 4.3, it is impossible for one to interpret and conceptualise the full posterior probability density function output of the MCMC and so we marginalise over all parameters in order to aid this. Unfortunately this comes as a trade off for understanding the full pdf. This can be overcome by sampling from the posteriors and obtaining post-fit predictions of the data in order to assess whether the model and data are in compatible regions of parameter space. One can create a “best-fit spectrum” from the MCMC by taking random throws from the MC chain N times, and build a Gaussian distribution of the predicted spectra from each throw. This will also derive an “error band” on the prediction. The method used for this consists of:

- Draw 2500 random steps from the Markov chain.²
- Each draw corresponds to a vector of parameter values. Reweight the MC to each draw to generate a new distribution of predicted events. Iterate through the bins of the SK samples and get reweighted event rate for the given neutrino energy bin.
- Draw a Gaussian around the distribution of event rates. Use the mean and sigma

²Important to note that the samples should be drawn from the post Burn-In region of the chain.

of the Gaussian to establish the central value and uncertainty of the posterior predictive.

Monitoring the comparison of the data and the distribution of fake data sets is a visual check. If the data point is extreme compared to the predicted distribution, i.e on or beyond the tails of the Gaussian, then the model will need to be reassessed and this will need to be understood.

4.3.5 Goodness-of-Fit Metric

In conjunction to what was outlined above, it is just as important to test the goodness of fit that the model provides to the data obtained. The approach used in this thesis utilises the Bayesian posterior predictive p-value derived from [119]. The majority of methods for calculating a p-value require that the fit be run with many different fake data studies, which for a high parameter model would be problematic and computational expensive. However, an existing Markov chain allows one to create these fake data sets and find a point likelihood using a similar method to the one above. To find the p-value one needs to calculate two likelihoods: \mathcal{L}_1 and \mathcal{L}_2 , via:

- Draw a random step from the Markov chain, ensuring the step is after the Burn-In phase.
- Calculate the likelihood of obtaining the data given the parameter values of the draw, this is \mathcal{L}_1 .
- Reweight the Monte Carlo prediction to the parameter values, this creates a prediction for the given draw.
- Iterating through each sample bin of the reweighted Monte Carlo, taking the number of events as the mean of a Poisson distribution and generating a random value from said distribution and setting that as the bin content.

- Calculate the likelihood of obtaining this Poisson fluctuated data set against the original MC prediction with the parameter values of the draw, \mathcal{L}_2 .
- Repeat this cycle for a desired number of iterations, N .

One can then draw a 2D distribution of \mathcal{L}_1 against \mathcal{L}_2 and find the p-value using:

$$p = \frac{N(\mathcal{L}_1 < \mathcal{L}_2)}{N_{Draws}}. \quad (4.10)$$

Importantly, one should use the same parameter vector from each draw to calculate both likelihoods. This is a variation of the traditional p -value test, where the data set is compared to many statistical variation of the ‘best fit’ model. Instead, the data set is compared to an ensemble of models by taking multiple draws corresponding to the full posterior probability. This way there is an account for the systematic uncertainty in the associated fit. The results of this hypothesis testing indicate the model’s ability to fit the data better than with fake data, which is just a statistical fluctuation of the observed data. Ideally, one would want a p-value of 0.5, as if the model is a good fit then there should be an even chance for the data or fake to better the prediction.

4.4 MaCh3 Fitter algorithm

MaCh3 is an oscillation fitter program that utilises the methods mentioned above to perform a joint-fit between the Near Detector and the Far detector simultaneously to produce a posterior probability density function. Depicted in Figure 4.6, MaCh3 uses a MCMC to semi-randomly traverse through the parameter space and build the posterior distribution by implementing the algorithm described in section 4.2.3. T2K also has two other main oscillation fitters: **P-Theta** [120] and **Valor** [121]. Both fitters are quasi-frequentist that utilise an input from a separate near-detector fit, then they propagate it to the far detector and use this as an input for the far-detector-only fit. MaCh3 has the added benefit of also performing a near-detector-only fit which can be used in

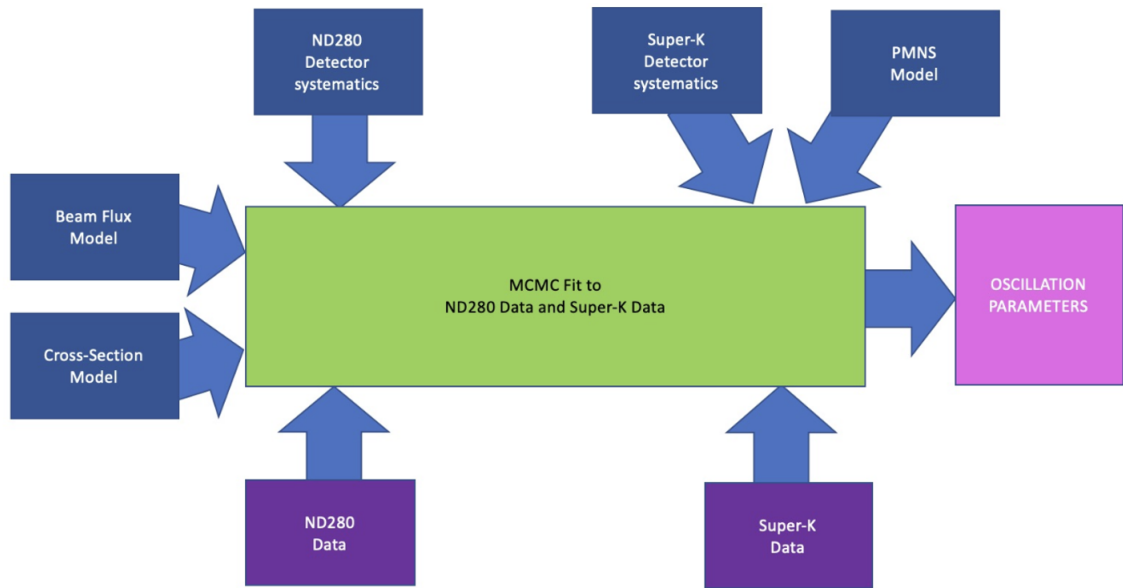


Figure 4.6: A schematic of the MaCh3 package, showing the various inputs required when running an oscillation analysis. These inputs come as covariance matrices that are used for the step proposal, parameter error definition and correlations. Taken from [124].

comparison with T2K’s other near detector fitter, **BANFF** (Beam and Near detector Flux Task Force). BANFF uses a gradient minimization technique with MINUIT [122] to find a global minimum within a parameter space. This allows BANFF to constrain the cross-section and flux parameters, passing this information to P-Theta and VALOR via a covariance matrix to be used as the aforementioned input. One benefit to using the gradient method is the ability to find a point of best fit. As discussed in section 4.3, this cannot be done efficiently in MCMC and is only an estimate. P-Theta and MaCh3 are used as comparison aids to ensure consistent results between the fitters, these comparisons can be found in [123].

4.5 Likelihood and Priors

4.5.1 Calculation of Likelihood

As discussed in section 4.4, as the MCMC steps through parameter space, it will calculate the likelihood for the given parameters proposed to obtain the given data. Equation 4.11 shows the Poisson likelihood function used to find the ratio of predicted and observed events.

$$\begin{aligned}
-\log P(\vec{\theta}|D) = & \sum_{\text{samples}}^{ND} \sum_i^{NDbins} \left[N_i^{ND,p}(\vec{\theta}_f, \vec{\theta}_x, \vec{\theta}_{ND}) \right] - N_i^{ND,d} \log \left(N_i^{ND,d} / N_i^{ND,p}(\vec{\theta}_f, \vec{\theta}_x, \vec{\theta}_{ND}) \right) \\
& + \sum_{\text{samples}}^{SK} \sum_i^{SKbins} \left[N_i^{SK,p}(\vec{\theta}_f, \vec{\theta}_x, \vec{\theta}_{SK}) \right] - N_i^{SK,d} \log \left(N_i^{SK,d} / N_i^{SK,p}(\vec{\theta}_f, \vec{\theta}_x, \vec{\theta}_{SK}) \right) \\
& + \frac{1}{2} \sum_{i,j}^{osc} (\theta_{o,i} - \bar{\theta}_{o,i})(V_o^{-1})_{ij}(\theta_{o,j} - \bar{\theta}_{o,j}) \\
& + \frac{1}{2} \sum_{i,j}^{flux} (\theta_{f,i} - \bar{\theta}_{f,i})(V_f^{-1})_{ij}(\theta_{f,j} - \bar{\theta}_{f,j}) \\
& + \frac{1}{2} \sum_{i,j}^{xsec} (\theta_{x,i} - \bar{\theta}_{x,i})(V_x^{-1})_{ij}(\theta_{x,j} - \bar{\theta}_{x,j}) \\
& + \frac{1}{2} \sum_{i,j}^{ND} (\theta_{ND,i} - \bar{\theta}_{ND,i})(V_{ND}^{-1})_{ij}(\theta_{ND,j} - \bar{\theta}_{ND,j}) \\
& + \frac{1}{2} \sum_{i,j}^{SK} (\theta_{SK,i} - \bar{\theta}_{SK,i})(V_{SK}^{-1})_{ij}(\theta_{SK,j} - \bar{\theta}_{SK,j}).
\end{aligned} \tag{4.11}$$

Breaking down the numerous variables in Equation 4.11, one has the subscript variables, o, f, x, ND, SK at the base of the $\vec{\theta}$, and these represent oscillation, flux, cross-section, near detector and far detector parameters respectively. The V_y , where y is just a generic substitution for the subscripts, terms are covariance matrices, individual to each of the previously mentioned subscripts containing the constraints on these systematic parameters. These are also important for applying penalty terms based on the prior; the

further from the prior central value, $\bar{\theta}$, for a given parameter, the larger the penalty term. For parameters that have a physical boundary, then an extremely large likelihood penalty is applied so that the probability of accepting that step is close to zero, reducing the number of steps accepted in these regions.

The terms $N^{ND,d}$ and $N^{SK,d}$ are the number of selected events at the near and far detector, where $N^{ND,p}$ and $N^{SK,p}$ are the number of predicted events for ND280 and Super-K respectively. The events are placed into histogram bins depending on the reconstructed energy of the neutrino and the bins are read into the analysis. The ND data histograms are binned in terms of two-dimensional lepton kinematics of $p_\mu - \cos(\theta_\mu)$ which is the lepton momentum and angle of the muon track with respect to the beam. For the SK data histogram, three of the muon-like samples are binned in one-dimension reconstructed energy histograms. The electron-like samples are binned in two-dimensional neutrino and $\cos(\theta)$ binning, where θ is the angle of the track with respect to the angle of the beam. The number of predicted events is calculated using:

$$N_i^{ND,p} = \Phi_\alpha^{ND}(E_i, \vec{\theta}_f) \sigma_\alpha(E_i, \vec{\theta}_x) \chi^{ND}(E_i, \vec{\theta}_{ND}) M^{ND}, \quad (4.12)$$

$$N_i^{SK,p} = P(\nu_\alpha \rightarrow \nu_\beta, E_i, \vec{\theta}_o) \Phi_\alpha^{SK}(E_i, \vec{\theta}_f) \sigma_\alpha(E_i, \vec{\theta}_x) \chi^{SK}(E_i, \vec{\theta}_{SK}) M^{SK}, \quad (4.13)$$

where, Φ is the neutrino flux, χ is the detector efficiency, M is the number of targets, P is the oscillation probability for a neutrino to oscillate from flavour α to flavour β , σ is the cross section for neutrino interactions and E_i is the incident neutrino energy in bin i . The events are obtained by running a neutrino interaction simulation program, NEUT [125] and tuned by T2K's Neutrino Interaction Working Group (NIWG). The neutrino interaction processes and nuclear effects for T2K were mentioned briefly in Section 1.5, but much more information can be found in [126] and [127].

As stated in Section 4.3.2, it is possible to apply a new prior for a given parameter

after the chains have been run, without requiring the analysis to be run again. This is done via a process known as reweighting. This allows one to alter the prior used in the chain by taking the ratio of the calculated likelihood for the original prior value and the likelihood for the new prior value at each step, creating a new weight for each event. This is done for each iteration of the chain and the weights are then stored and can be applied accordingly. Reweighting is utilised in two main ways within the MaCh3 analysis. The first way is to reweight events according to the effect of systematics for each proposed step. Relating this to Equation 4.11, this reweighting indirectly impacts the likelihood calculation by changing the predicted number of samples in $N^{ND,p}$ and $N^{SK,p}$ as each throw of parameters will need a new prediction.

The other way is to compare the effects of applying different priors by directly reweighting the posterior probability itself. This also effects the likelihood of Equation 4.11 by altering the penalty term applied in $(V_o^{-1})_{ij}$. This is most common on θ_{13} . Reactor experiments have a larger sensitivity to the oscillation parameter θ_{13} via the electron antineutrino disappearance channel. One can see from Figure 4.7 the effect of applying the reactor constraint as the prior on θ_{13} . A shift in the HPD and a much tighter constraint on the parameter shows the benefit of this reweighting.




Figure 4.7: Left: The posterior probability density for θ_{13} with no reactor constraint applied. Right: The same but with a reactor constraint applied. Reweighting to the reactor constraint produces a better constraint on the measurement of θ_{13} , notice the change of the x -axis. Taken from [128].

4.5.2 Systematic Uncertainties

Having outlined how the fitter works towards building results, the dominant systematic uncertainties that exist in the fit will be discussed further. Section 4.5.1 briefly touches upon the numerous flux, cross section and detector systematic parameters that are present in the model and how they influence the likelihood calculation.

Flux Model

The first term inside of Equation 4.12, $\Phi_{\alpha}^{ND}(E_i, \vec{\theta}_f)$, relates to the flux model. The flux model used inside of the MaCh3 analysis is tuned using external NA61/SHINE [129] data which utilises a replica T2K 90 cm graphite target to better map out hadronic production, whilst also including a better physical model for secondary interactions [130]. This tuning has a significant impact on the constraint of the flux parameters, increasing the π^{\pm} statistics and reducing the uncertainty on higher energy neutrinos from kaon production from $\sim 10\%$ down to $\sim 5\%$, compared to previous iterations of the analysis [75].

Overall there are 100 flux systematic parameters, 50 for the near detector and 50 for the far detector and these encapsulate the uncertainties on the hadron production model, secondary interactions, beam profile and alignment. They are separated into 8 groups with the four neutrino modes; ν_{μ} , ν_e , $\bar{\nu}_{\mu}$ and $\bar{\nu}_e$ and repeated for FHC and RHC run modes. The parameters are then further split into bins defined by neutrino energy with varying width, the closer to the oscillation maximum, the finer the binning, and the opposite for energies further away.

Cross-Section Systematics

T2K's cross-section model is the largest source of uncertainty within the analysis, with the simulation of neutrino interactions being a big focus of work. Taking a consortium of external models to apply to neutrino-nuclear interactions induces a large uncertainty

with models not completely agreeing across data space. With a total of 75 cross-section systematics, these can be grouped into individual interaction modes that are generated in the NEUT package [125]. A brief introduction to T2K's cross section model was given in Section 1.5. Uncertainties about the correctness of these models is encapsulated by systematics included in the model.

Some of these parameters are defined as normalisation or shape parameters. Normalisation parameters act as a weight that affect all parameters of a given interaction type, i.e CCQE parameters or 2p2h parameters. Shape parameters are sensitive to event kinematics and as such the weight will vary on an event by event basis. These are described by response functions which specify how the data prediction scales up and down as the systematic changes and stored as cubic splines. These splines are used for reweighting when calculating the number of events as cross-section parameters change in the MCMC.

Detector Systematics

A large number of the systematics included in the fit are the ND280 and SK detector systematics. These uncertainties cover specific sub-detector efficiencies, such as the TPCs tracking efficiency and charge identification efficiency. They also include observable-like systematics such as the TPC momentum resolution uncertainty and TPC particle identification. For SK, the detector systematics largely surround the ring reconstruction efficiency and the discriminating power of selections used for the data samples.

4.6 Combining T2K-NOvA

The T2K-NOvA joint-fit effort will use a combination of the individual experimental likelihoods, obtained using the MCMC techniques outlined in section 4.2 to sample their respective models represented by the full-likelihood of the two experiments inde-

pendently, to better constrain the PMNS parameters thanks to the extended parameter space available to be sampled outside of the individual experimental limits. Due to the differences in degeneracies of certain parameters, namely δ_{CP} , θ_{23} and the mass ordering problem, a combined analysis may produce a much clearer picture than either analysis alone.

Adding the NOvA model parameters into the global model, and sampling through parameter space obtaining a likelihood for the NOvA data allows the joint-fit MCMC to concurrently sample two independent models and two independent data sets, finding a favourable region in parameter space. To test that the implementation of the joint fit model is correct, fits are carried out on Asimov data sets. These Asimov data sets are Monte Carlo simulations that are generated around the model, with all systematics set to their prior central value and exactly equal to the PMNS prediction for a chosen set of oscillation parameters [131]. The truth data is the Monte Carlo prediction. This allows one to test the robustness and sensitivity of the fitter [123]. If the output of the fit is not the same as the input, i.e the HPD of an oscillation parameter is not the same as that used to generate the truth data set, then there is a problem within the fitting framework. Not only does it test the robustness, but it also allows for an experiment to determine the sensitivity to parameters, as Asimov data sets have no statistical fluctuations, thus generating constraints as if the model perfectly describes the data. The next chapter shows the effect on the sensitivity to these oscillation parameters when running a joint fit with Asimov data sets, generated around different oscillation parameter points.

This chapter has presented an introduction into the statistical techniques used in the analysis presented in this thesis and one of the analyses carried out by the T2K collaboration. Using the Metropolis-Hastings algorithm in the Markov Chain Monte Carlo allows one to build a posterior distribution function in order to assess how well the

model describes the data, and make conclusions about systematic uncertainties within the analysis, or conclusions about parameter constraints. Giving a brief overview of techniques such as reweighting and how this is implemented within the analysis to apply alternate priors to the MCMC without the need to re-run the fit. Finally, discussing the goodness-of-fit analysis and how it can allow one to quantify the models ability to describe the data by comparing it to its ability to describe a statistically fluctuated data set.

Chapter 5

T2K-NOvA joint fit

*Teamwork makes the dream
work.*

John C. Maxwell

Both T2K and NOvA are trying to answer the same questions within neutrino physics, but with different approaches. With dissimilar detector technologies, experiment baselines and model sensitivities, producing a joint fit between the two experiments will hopefully lead to better constraints on the oscillation parameters and reduce the degeneracies that are otherwise present in individual analyses [132].

As the joint fit is currently in development, this chapter will only be able to highlight the framework and validation that has gone into the joint fit. Figures are shown that visualise the increase in Monte Carlo sensitivity, but do not include the far detector data. Due to the inter-collaboration nature of the joint fit, it is a rigorously monitored process, with many external and internal adaptations that must be checked with simulations before applying the joint analysis to far detector data. At the time of writing, data fits are starting to be run and hopefully the results will be published soon after digesting the outputs.

5.1 Introduction to the Joint Fit

The working group of the joint fit effort is complicated. Collaborations do not usually disclose details of the full analysis publicly. Therefore, doing so between two competing experiments required a careful balance, so much so that a Memorandum of Understanding (MOU) was agreed between the collaborations in order to formalise how much could be shared with the wider collaborations and even more so externally. In the early stages of the analysis, it was discussed how much should be shared between the collaborations and it was decided to limit this as much as possible.

Figure 5.1 shows the two-dimensional (2D) 90 % confidence level contour for the $\delta_{CP} - \sin^2(\theta_{23})$ parameter space of T2K (blue), NOvA (red) and SK (green). A distinct aspect of this is the three different best-fit locations, but noticeably there is a disagreement between T2K and NOvA as both experiments exclude the other’s best-fit point at the 90% confidence level. There is sufficient overlap in the contour that suggests that no tension is present but this points to the importance of doing a full joint fit. There is a good alignment between T2K and NOvA in $\Delta m_{32}^2 - \sin^2(\theta_{23})$ space, as shown in Figure 1.5. Current fits, such as those composed by NuFIT [133], try and recreate contours published by experiments with simplified systematic assumptions. Many ‘joint fit’ attempts will multiply the likelihood distributions across the marginalised 2D parameter space. However, this does not account for the full correlations in likelihood and only the more common two-dimensional projection and thus may not give an accurate representation of a joint fit.

5.2 Joint-Analysis Technique and Shared Software

An important aspect to external joint fit methodologies is that they cannot account for the full systematics and therefore not a full likelihood. It was discussed in the early stages of the joint fit analysis as to the extent at which a likelihood was shared. Three




Figure 5.1: A comparison of T2K (blue), NOvA (red) and SK (green) contours for $\delta_{CP} - \sin^2(\theta_{23})$ assuming normal (left) and inverted (right) hierarchy and marginalising over all other parameters. Figure taken from [51].

options were presented; a marginalised likelihood, giving the likelihood of obtaining the oscillation parameters only, a profiled likelihood, also only giving likelihood as a function of the oscillation parameters, and finally the full likelihood, accounting for all oscillation and nuisance parameters in the model. As the two experiments use different approaches (T2K uses marginalisation and NOvA use profiling), sharing a full likelihood allows either collaboration to present results as they normally would for internal use. Also, providing a full likelihood did not require any new implementation and could be easily shared.

The next pressing issue was how this likelihood function and the information was physically shared, given a strict MOU. Packages could not simply be swapped and ‘run’, and so the solution was to treat each respective collaboration’s analysis as a ‘black box’. This would take a vector of parameter values and return a log-likelihood, ultimately calculated from a comparison to experimental data, but during the development stage this is Monte Carlo predictions used as fake data. Due to the software implementation, they will be referred to as containers, where “T2K’s container” and “NOvA’s container” are from the respective experiments. T2K’s chosen package to be containerised was MaCh3 (described in section 4.4), and NOvA used an internal package know as CAFAna. Fig-

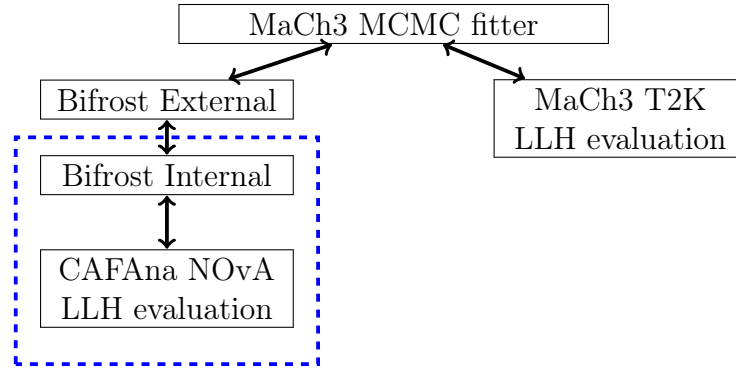


Figure 5.2: Schematic of the MaCh3 fitting framework as used for this T2K+NOvA joint analysis. The dashed boxes represent the NOvA container.

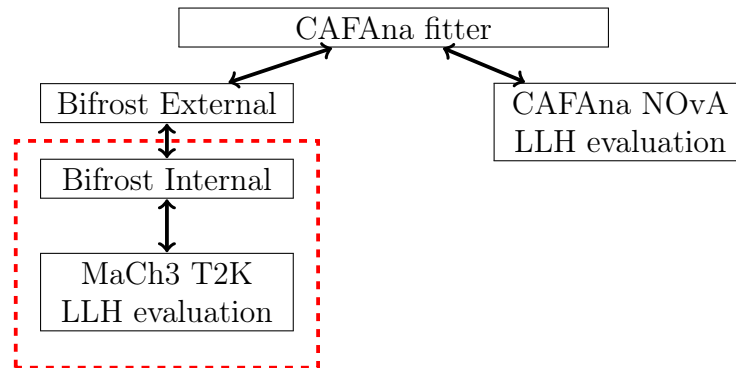


Figure 5.3: Schematic of the CAFAna fitting framework as used for this T2K+NOvA joint analysis. The dashed boxes represent the T2K container.

ures 5.2 and 5.3 depict the fitting framework implemented by both collaborations. The blue box and red box represent the CAFAna and MaCh3 containers respectively. This way, experiments only had access to their own packages, and could not see the inner workings of the other. From this, the next issue was how to communicate with the containers themselves. A specialist C++ package called, ‘Bifrost’, was built. This utilised `std::stream` to pipe in information and pipe out outputs that could be read by the respective packages, depicted in Figures 5.2 and 5.3 as Bifrost Internal and External respectively¹.

Figure 5.2 represents the framework as used by T2K. MaCh3 will run a Markov Chain Monte Carlo as described in section 4.2, inclusive of NOvA parameters. As a set of

¹It is worth noting that Bifrost is independent of either experiment and just acts as a “bridge” of information.

parameters is proposed, the NOvA parameters are piped into the NOvA container via Bifrost External, read by Bifrost Internal, and then NOvA's log-likelihood calculation will commence. Once finished, the value of log-likelihood for that given set of parameters is returned and piped out via Bifrost and picked up by the MaCh3 fitter, combining the likelihood and assessing whether or not to accept or reject the steps via the criteria outlined in sub-section 4.2.3. This combined likelihood, written in equation 5.1, depicts the joint likelihood of either experiment obtaining their respective data sets with a given set of parameters.

$$\mathcal{L}_{Total} = \mathcal{L}_{T2K} \cdot \mathcal{L}_{NOvA} \cdot \mathcal{L}_{Other} \quad (5.1)$$

where \mathcal{L}_{T2K} and \mathcal{L}_{NOvA} are the likelihoods of T2K and NOvA respectively, and \mathcal{L}_{Other} is the likelihood from external aspects such as the reactor constraint which must be utilised after running as a weight instead of during either experiments individual likelihood calculation, since it must only be used once.

5.3 Container Validation

As discussed above, neither experiment can see the inner workings of the other container, so it is important to confirm that the likelihood output from the container is correct. To do this, extensive validations have occurred and will be discussed below. It is worth noting that there is not just one container. In fact, five iterations of the T2K container have been created and validated as the physics studies and operational capability changes through requests of other collaborators or computational limitations. The five iterations consist of:

Version 1 - The initial container storing all Monte Carlo simulations, Data² and

²A very important point to highlight is that both experiments run a near detector and far detector set-up (discussed in chapter 3). Real near detector data is used during the validation process and

covariance matrices. This was a very large container that was \mathcal{O} 10 GB, accessed via a program called Docker [134].

Version 2 - Arguably the most important update, version 2 brought changes to Bifrost and the containers that allowed mountable objects. In essence, this means that the container no longer hosted the Monte Carlo, data or covariance matrices, and instead they could be loaded/‘piped’ in. This created two major advantages. Firstly, it allowed analysers to alter the inputs used for the fits. This meant one could use real near detector data, near detector Monte Carlo generated data sets or far detector data sets generated at different oscillation points (described further in section 5.4), thus changing the sensitivity and physics. Secondly, it drastically reduced the disk storage load for analysers. Those on shared resources did not need multiple containers due to permission blockages and could store both containers and container inputs in a common location.

Version 3 - This entailed a small tweak to Bifrost and the output from the containers, changing the output from the parameters physical values to distance from nominal in terms of the assigned prior uncertainty. This just created a constant benchmark when running with real data or Monte Carlo generated data at the near detector.

Version 4 - Mach3 has the ability to utilise Graphic Processing Units (GPUs) inside of its core code, while NOvA’s package does not, so everything was run on Central Processing Units (CPUs), which is slower. To aid this, T2K implemented a new system for the propagation to the far detector. Originally, for every likelihood calculation a new oscillation probability function was derived. This update brought about using a binning system to eliminate the need to calculate an oscillation weight at every step, drastically reducing the time taken to perform a fit

Monte Carlo sensitivity studies, but during development the far detector data is replaced with Monte Carlo. This will of course be swapped out for real data once the collaborations agree that the analysis is correct.

from NOvA's side using the T2K container.

Version 5 - Finally, the newest update allows one to calculate posterior predictive p -values. Described initially in section 4.3.5, the use of this test statistic and results are described in section 5.5. This capability required changes to both the core code and Bifrost.

With each new version of a container came more validations, not only for the code residing within the container, but also for the altered Bifrost mechanisms.

5.3.1 Oscillation Probability Calculation

As a neutrino propagates from the point of creation to a point of interaction, it traverses matter and a given distance. Using the PMNS matrix the probability of oscillation from a flavour state to another (or to remain the same flavour state) is calculated. Equation 4.13 includes this probability in the prediction. Unfortunately this is not a shared package between T2K and NOvA and the oscillation probability packages are PROB3++ and PMNSOpt respectively. The first aspect of validation the author was responsible for was ensuring that these oscillation calculator packages produce compatible results. To do this, various oscillation points outlined in Appendix A were used to produce an oscillation probability for the four oscillation channels, $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu)$ and $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ as a function of energy. These functions only rely on the PMNS parameters as well as the baseline length, matter density and matter profile. As in a normal fit, a constant matter density was used with an electron density value of 2.84 g/cm^3 and a baseline length of 295 km. Figure 5.4 show the percentage difference of the probability calculated by the two oscillation calculators in the $\nu_\mu \rightarrow \nu_e$ oscillation channel. The spacing of test points used for the scan is 100 MeV with a threshold for T2K being 200 MeV, equivalent to the minimum energy threshold used in T2K's analysis. One can see that the disagreements are less than 1%, with these discrepancies occurring at regions where the

sharp feature at 0.3 GeV is caused by the calculated oscillation minimum being slightly different for the two probability calculators. The grey banded areas indicate regions where the probability for the given oscillation channel is less than or equal to 0.02. The red dashed line indicates a maximum for the oscillation channel. It is important that the key take away is that there is a very minimal amount of disagreement at the oscillation maximum, as this is where both experiments focus their analyses. These differences were deemed to be negligible and a sufficient level of agreement between the two collaborations that there will be no meaningful effect on the sensitivity of the joint fit.




Figure 5.4: A comparison of PROB3++ and PMNSOpt at T2K’s baseline in the $\nu_\mu \rightarrow \nu_e$ oscillation channel using the Asimov 1 oscillation parameter set defined in Appendix A. The grey banded area represents values of P_{T2K} and P_{NOVA} that are ≤ 0.02 . The tan-like component at 0.3 GeV is caused by the oscillation probability being close to zero, and the minimum occurring at slightly different energies for the two oscillation calculators. The red dashed line depicts the oscillation maximum at 0.6 GeV. Importantly, the percentage difference between the probability calculated by PROB3++ and PMNSOpt is minimal at that oscillation maximum.

5.3.2 Log-Likelihood Scans

The main purpose of the container is to allow the sharing of the full likelihood function, without sharing the package. Log-likelihood scans are the most important validation technique used for the containers, with the main workflow being an input of parameter values and returning a corresponding likelihood for said parameters. It is important to note here why this became an integral part of validation. During the creation of version 2 outlined above, it was important to make the container more versatile with inputs by changing the way the data was handled in the container. Alongside this, there were issues with the sheer time taken to run an analysis due to the large number of parameters that exist in the T2K model. As mentioned in section 4.4, MaCh3 will simultaneously fit both the near and far detector, making the parameter space over 750 dimensions, a difficult computational task for the profiling approach. To reduce this, MaCh3 took a similar approach to that used by another T2K package, P-Theta, by using a near detector fit, created by the BANFF group (introduced in Section 4.4), to constrain the parameters that can be propagated to the far detector. This nullifies the need to vary the near detector systematics, near detector flux and near detector only cross-section parameters (the cross-section systematics that only apply to interactions on a carbon target for example, as there is no carbon in the far detector target volume) during the joint fit. This drastically reduces the parameter space to 147 parameters. Validation checks such as predicted samples and log-likelihood scans were carried out to ensure that the physics was not significantly altered as an expense to the reduction in computational load. Therefore, T2K has two modes of fitting, MaCh3-postBANFF, the reduced parameter model, and normal MaCh3 that continues to fit the near detector and far detector simultaneously. The container utilises this MaCh3-postBANFF and the log-likelihood scans shown are within this fit space.

The second validation implemented by the author was ensuring that the likelihood being returned by the container is as expected, a log-likelihood plot is generated by passing

a vector of parameter values to the container, keeping all but one parameter at the nominal value and scanning the given parameter across a range of three sigma, where sigma is defined as the prior uncertainty on the parameter. This process is repeated for each parameter in the model. Figure 5.5 shows the log-likelihood scans for six T2K systematics, obtained via the container (red) and the bare³ package (blue). The top pad shows the scans themselves, whereas the bottom pad depicts the difference between the output from the container and the MaCh3 package. The disagreement is of the order $\mathcal{O}(10^{-12})$ and thus deemed negligible. The same is done for the NOvA parameters and the NOvA container, depicted in Figure 5.6. Differences arose due to trivial issues, including offset (seen in the bottom left of Figure 5.6) and binning issues where some scans were stored in histograms as opposed to graphs leading to a shift in the minima. The largest issue found with respect to the physics implications can be seen in the bottom right of Figure 5.6, the final systematic in the NOvA model, ‘Cosmic Scale’, was found to be off by a constant factor of 7. This was later confirmed by NOvA to have been implemented in the container, and not in the bare CAFAna package that was used to create the internal NOvA LLH scans handed over for comparison. It highlights the importance to carry out these continuous verifications and make sure that every discrepancy is understood.

³Bare is used to describe the MaCh3 package outside of the container.

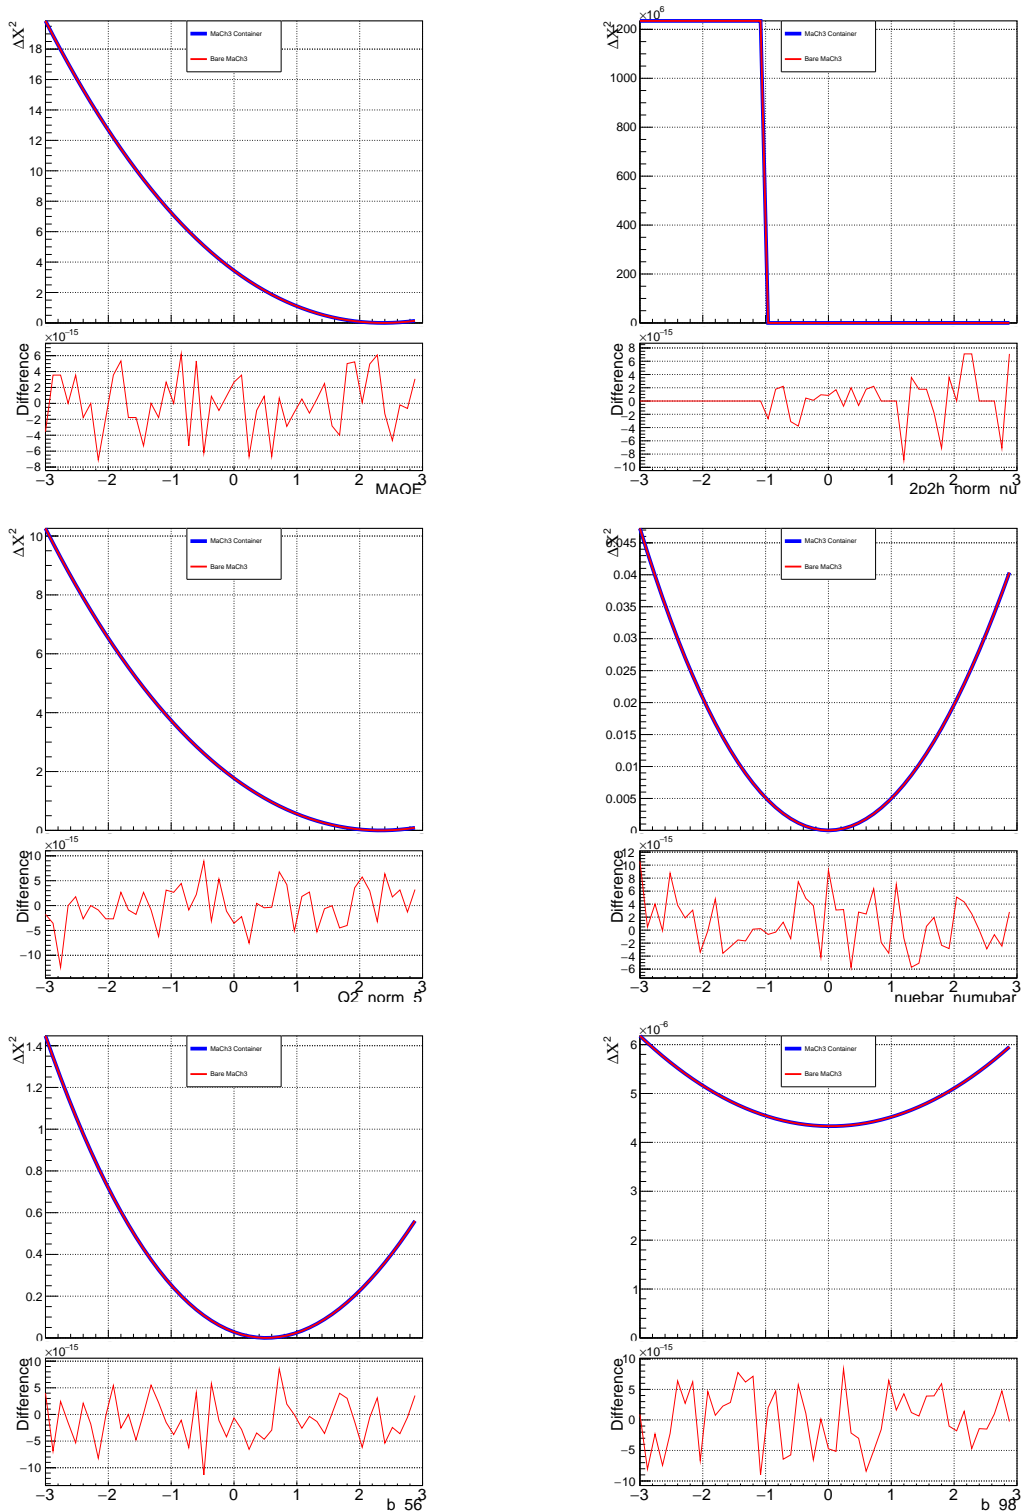


Figure 5.5: Likelihood scans for a selection of T2K systematic parameters. The blue line represents the log-likelihood output via the T2K container, with red representing the log-likelihood distribution obtained using the bare MaCh3 analysis package.

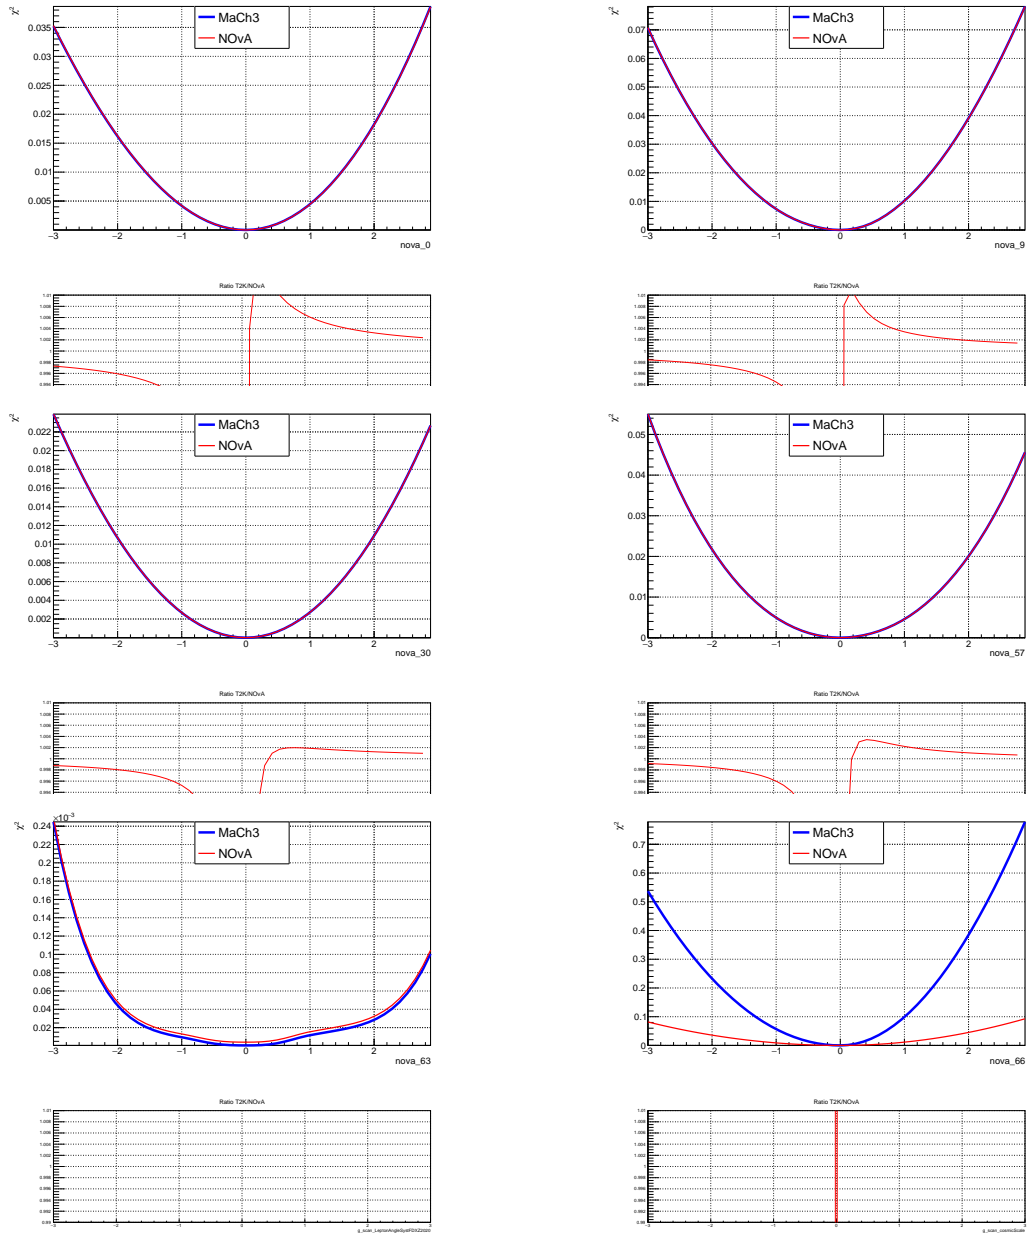



Figure 5.6: Likelihood scans for a selection of NOvA systematic parameters. The blue figure represents the log-likelihood output via the NOvA container, with red representing the log-likelihood distribution obtained using the bare CAFAna analysis package, provided by NOvA collaborators.

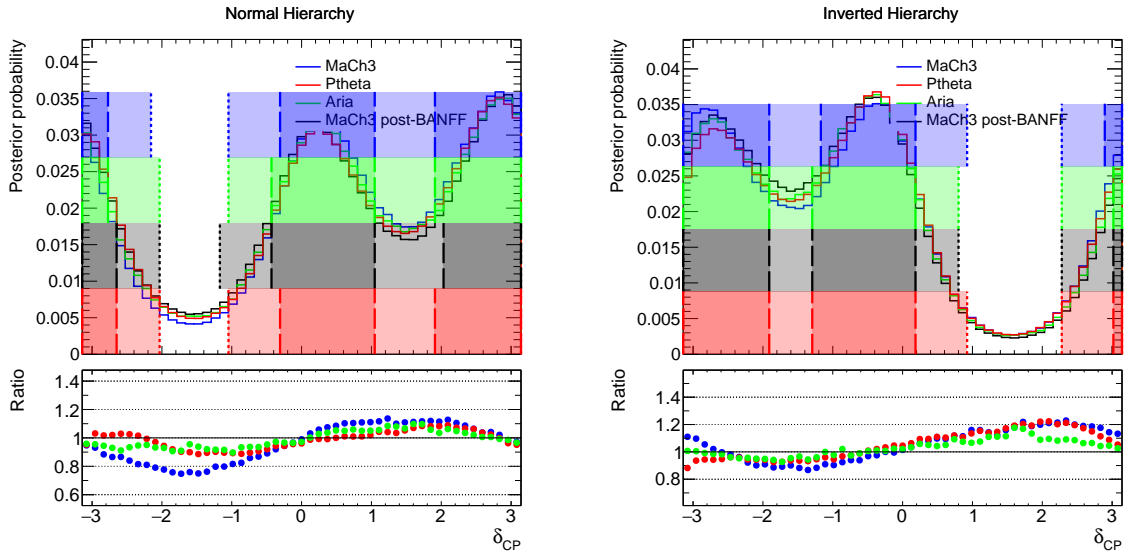
5.3.3 Comparison with Other Fitters

The final method of validation was comparing the output of an analysis to other fitters running the same analysis. It was at this stage that NOvA built a new package that used the likelihood calculator of CAFAna, but with a Markov Chain Monte Carlo fitter. This produces a Bayesian posterior probability function that can then be marginalised to find the marginal posterior distribution for a parameter of interest. This package is called Aria and utilises a Metropolis–Hastings Algorithm (discussed in section 4.2.3), just like MaCh3, making the outputs much more comparable. The marginalised posterior distributions were also compared to another T2K fitter, P-Theta (mentioned in section 4.4, a full MaCh3 fit that fits the near and far detector simultaneously and also a MaCh3-postBANFF analysis. Each of these fits were prepared by different analysers with the MaCh3-postBANFF being carried out by the author.

Figures 5.7 and 5.8 show the marginalised posterior distribution for δ_{CP} at Asimov points 0 and 1, which are the respective NOvA and T2K best fit points for the 2020 analysis, defined in Appendix A. The sub-figures depict the posteriors (a) marginalised over both hierarchies with equal prior, for (b) Normal Hierarchy and (c) Inverted Hierarchy respectively. The bottom pad illustrates the ratio between MaCh3-postBANFF and any of the other fitters. The largest difference that is visible from Figures 5.7 and 5.8 is the MaCh3 (blue) posterior. The fact that MaCh3 vs MaCh3 post-BANFF is the largest difference is expected within the working group. The good agreement between MaCh3-postBANFF, Aria and P-Theta is promising with all three analyses taking a near detector constraint from BANFF. One could question why they do not agree perfectly but this is understood to be due to different statistics (MCMC steps) and also step-tuning plays an important factor for convergence and coverage.



(a) Both hierarchies.



(b) Inverted hierarchy.

(c) Normal hierarchy.

Figure 5.7: A comparison of one-dimensional marginal posteriors for δ_{CP} , produced by MaCh3 (blue), Ptheta (red), Aria (green) and MaCh3 post-BANFF (black) at the Asimov 0 oscillation point, showing the one (dashed) and two (dotted) sigma credible regions. The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. Figures taken from [135].

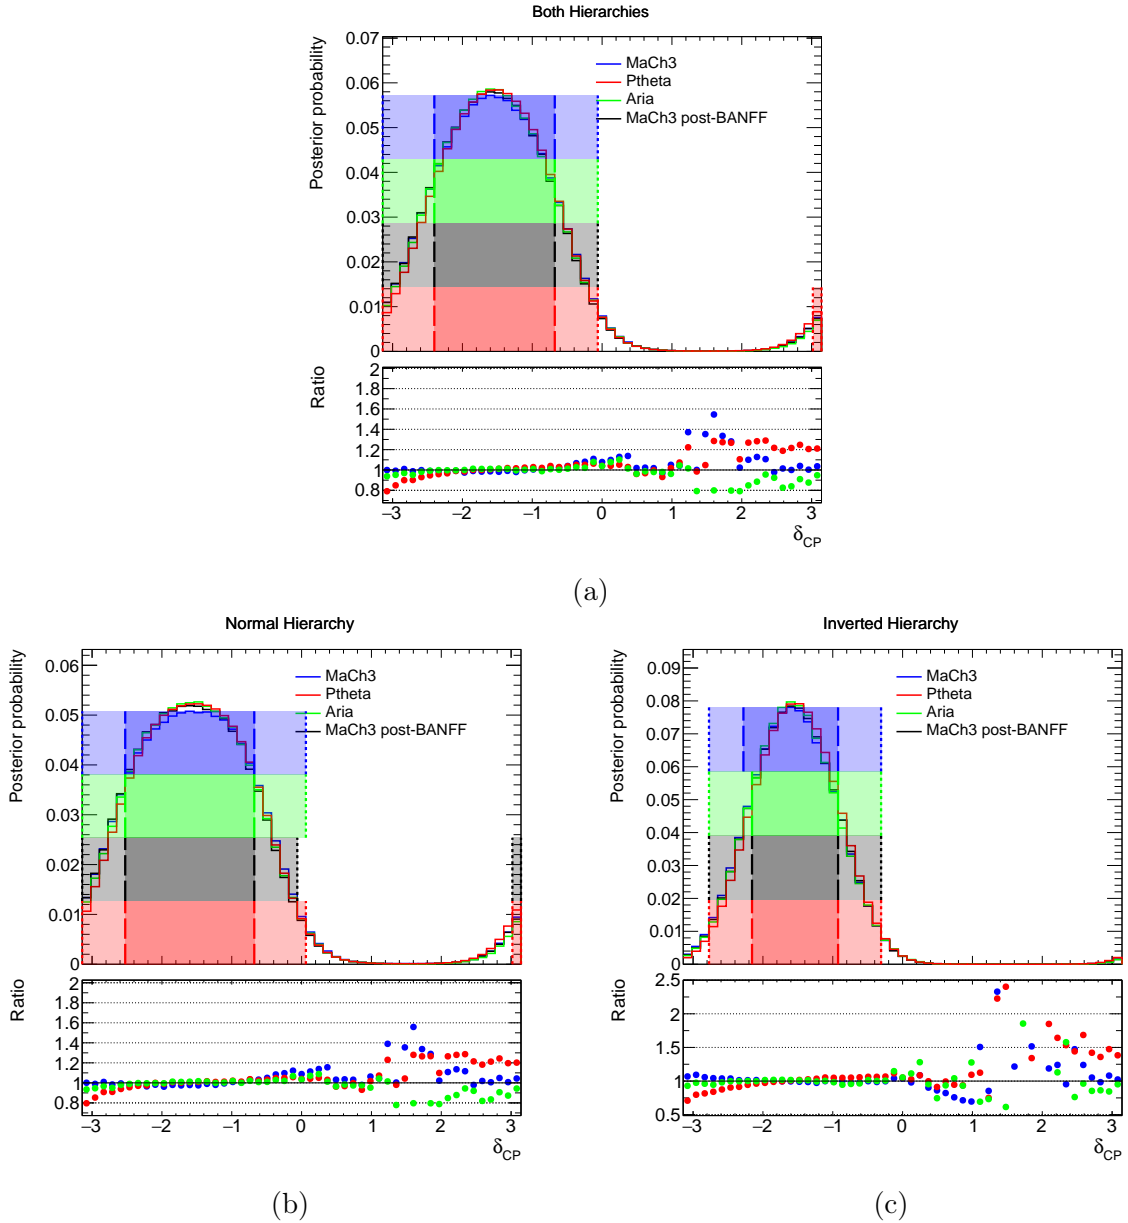


Figure 5.8: A comparison of one-dimensional marginal posteriors for δ_{CP} , produced by MaCh3 (blue), Ptheta (red), Aria (green) and MaCh3 post-BANFF (black) at the Asimov 1 oscillation point, showing the one (dashed) and two (dotted) sigma credible regions. The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range. Figures taken from [135].

Shown in Figure 5.7 and Figure 5.8 are the one-dimensional δ_{CP} sensitivity comparisons for four individual fits, in order to compare and validate. It is important to highlight that these do not use either container described in Section 5.1. Instead, Figure 5.9 shows a comparison of Aria (NOvA only), T2K (T2K post-BANFF) and the joint T2K+NOvA (using the T2K container) one-dimensional marginalised posterior for δ_{CP} at Asimov 0, showcasing the increase in sensitivity of the joint fit compared to either experiment's individual fit. The T2K posteriors and the T2K container were generated by the author of this thesis and due to the MOU not all figures for the various Asimov points have been made official. The purpose of this segment is to highlight the author's contribution to the joint fit effort. Section 5.4 shows the marginalised posterior comparisons between T2K, NOvA and a joint T2K+NOvA fit which was generated using the NOvA container, the counterpart to the T2K container. These plots have been created by the author using the available Markov Chain Monte Carlo chains open to the working group in order to show the change in sensitivity to oscillation parameters at different Asimov data points. Following from this, Section 5.5 shows the validation work for version 5 of the T2K container, carried out by the author.

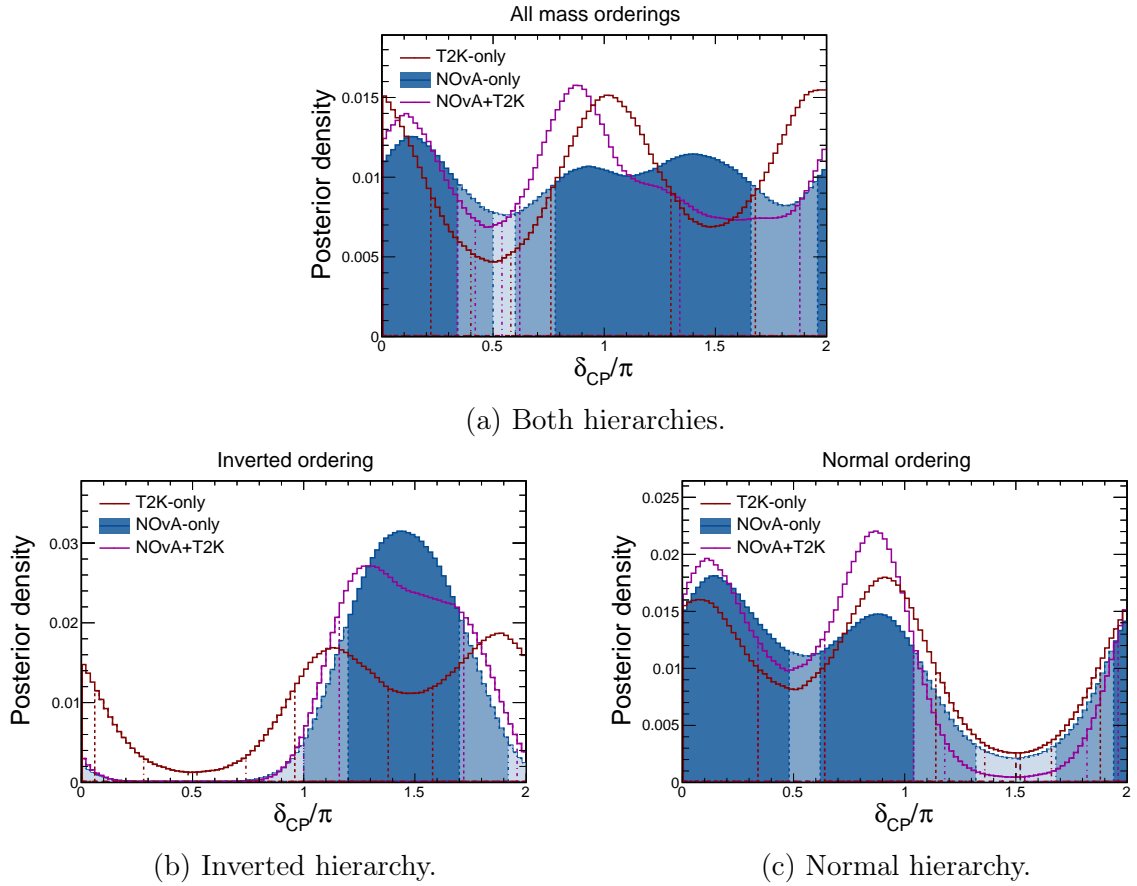



Figure 5.9: Marginal posterior distributions of δ_{CP} for Asimov 0 in both hierarchies (top), inverted hierarchy (left) and normal hierarchy (right). The NOvA (pink) posterior were generated using Aria, the T2K (red) posterior was generated using the post-BANFF model, and NOvA+T2K is joint fit using Aria and the T2K container. Figures taken from [135].

5.4 Fit Results

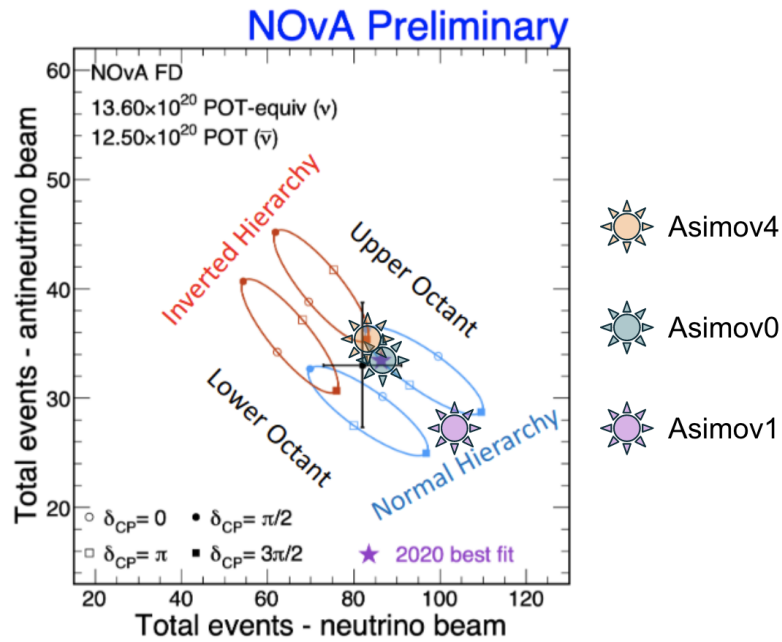
Utilising the framework and containers outlined in the previous sections and implementing them together gives the ability to run a joint fit between T2K and NOvA. Although a list of five oscillation parameters were given in Appendix A and initially used for validation, the joint fit narrowed this down to three with focus being on Asimov 0 - NOvA 2020 best fit [136], Asimov 1 - T2K 2020 best fit [137] and Asimov 4 - NuFit 5.1 Global fit [133]. These points were selected as they cover a good range of parameter space, allowing the joint fit to probe the degeneracies that exist with an individual experiment's analysis. The three Asimov points are depicted as markers on the bi-event rates in Figure 5.10, and were used to generate Asimov data sets, described in Section 4.6.

Asimov 0 is degenerate for T2K given the insensitivity to the mass hierarchy (at this oscillation point, values for normal hierarchy and inverted hierarchy exist that give identical values for the e-like samples on the x -axis). Asimov 1 is the only non degenerate Asimov oscillation set used for both experiments. The combination of a near maximal δ_{CP} with the assumption of normal hierarchy being true is a very non-degenerate region for both experiments, illustrated in Figure 5.10. Asimov 4 is similarly degenerate to Asimov 0, but although it resembles Asimov 0 at NOvA it is quite different in T2K, being intermediate between Asimov points 0 and 1.

The following section shows the one and two dimensional posteriors for each Asimov point in turn, illustrating the individual and (green) joint sensitivity for (blue) T2K and (red) NOvA. There are four sets of one-dimensional posteriors for $\sin^2(\theta_{13})$, $\sin^2(\theta_{23})$, δ_{CP} , Δm_{32}^2 and three sets of two-dimensional posteriors for $\sin^2(\theta_{13}) - \delta_{CP}$, $\sin^2(\theta_{23}) - \delta_{CP}$ and $\sin^2(\theta_{23}) - \Delta m_{32}^2$ which are commonly used to show the main degeneracies in the the individual experiments. Both experiments use real near detector data and Monte Carlo generated far detector data, reweighted to each Asimov oscillation



(a)



(b)

Figure 5.10: Bi-Event plots for T2K (top) and NOvA (bottom) with a marker depicting the location of three Asimov parameter sets described in Appendix A.

point. The PDG $\sin^2(\theta_{13})$ Gaussian constraint is applied throughout. It is worth noting here that the first iteration of the joint fit only has a few correlations between the models. Notably the oscillation parameters are fully correlated with some correlations between influential systematics, more detail can be found in [138].

5.4.1 Asimov 0

Starting with Asimov 0, the NOvA best-fit point, the joint fit creates tighter credible intervals across all parameters, with one notable exception being the sensitivity in $\delta_{CP} - \sin^2(\theta_{13})$ parameter space marginalised over Inverted hierarchy, depicted in Figure 5.15b. This could have occurred due to the pseudo-degenerate region of δ_{CP} at this Asimov point. One can see from the blue star marker in Figure 5.10b, the input value is near values of $\delta_{CP} = \pi/2$ and $-\pi/2$ which suppresses and enhances the e-like samples respectively. This creates this triple peak feature in the posterior, visualised in Figure 5.13a.


It is expected that sensitivity to the PMNS parameters will increase. The main motivation behind the joint fit is resolving degeneracies, where the mass hierarchy (NH or IH) and octant degeneracy (whether $\sin^2(\theta_{23})$ less than or greater than 0.5) are separate hypotheses that can be tested. As described in section 4.3.3, one can take the ratio of the probabilities given two hypotheses, H_1 and H_2 . Relating this to the degeneracies, H_1 is the Normal Hierarchy hypothesis, or $\sin^2(\theta_{23}) < 0.5$ and H_2 are the alternatives. Thankfully, having the MCMC posteriors probability, one does not need to run more analyses with the assumption that a given hypothesis is true. Instead, one can take the integral of the marginalised posterior, assuming a given hypothesis. For example, referring to Figure 5.14a, one can find the integral, i.e integration range $0 < \Delta m_{32}^2 < +\infty$, of the marginalised posterior and take the ratio of the posterior probability in each hypothesis, divided by the ratios of integrals on the prior. This gives one an indication

to the preference of hypotheses for the mass hierarchy degeneracy.


Table 5.1 indicates the fraction of the Asimov 0 posterior probability for a given hypothesis, where Sum is the sum of a given row or column, indicating the total integral for one hypothesis. To find the Bayes factor for the mass hierarchy as an example, one would simply take a ratio of the Sum of NH over the Sum of IH, finding the preference of mass hierarchy hypotheses. The preference for Normal Hierarchy and upper octant is expected given the oscillation parameters used to simulate the far detector data.

	$\sin^2 \theta_{23} < 0.5$			$\sin^2 \theta_{23} > 0.5$			Sum		
NH	0.202	0.172	0.227	0.531	0.324	0.462	0.733	0.496	0.689
IH	0.084	0.170	0.079	0.183	0.334	0.232	0.267	0.504	0.311
Sum	0.286	0.342	0.306	0.714	0.658	0.694	1.000		


Table 5.1: Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 0 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Normal Hierarchy and Upper Octant.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.11: Marginal posterior distributions of θ_{13} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.12: Marginal posterior distributions of θ_{23} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.13: Marginal posterior distributions of δ_{CP} for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.14: Marginal posterior distributions of Δm_{32}^2 for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.15: Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.16: Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.



(b) Inverted hierarchy



(c) Normal hierarchy.

Figure 5.17: Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 0 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.


5.4.2 Asimov 1

Looking at Asimov 1, T2K best-fit point, the lack of degeneracy allows one to really see the influence of performing a joint fit. The sensitivity increase is more prominent throughout the posteriors. Again, the only questionable credible region comes in the $\delta_{CP} - \sin^2(\theta_{13})$ parameter space, illustrated in Figure 5.22b. This ‘wobbly’ contour is the result of low MCMC steps in the Inverted Hierarchy region, visualised by the non-smooth one dimensional marginalised posterior evident in Figure 5.18b.


For Asimov 1, there are also no standout features shown in table 5.2. The fraction of posterior in Inverted Hierarchy and lower octant being zero is a clear indication of the lack of degeneracy at this Asimov point. There is a very similar sensitivity to the octant degeneracy in both experiments, with no real increase from the joint fit.

	$\sin^2 \theta_{23} < 0.5$			$\sin^2 \theta_{23} > 0.5$			Sum		
NH	0.270	0.233	0.284	0.728	0.530	0.713	0.998	0.763	0.997
IH	0.000	0.057	0.001	0.002	0.180	0.002	0.002	0.238	0.003
Sum	0.270	0.290	0.285	0.730	0.710	0.715	1.000		


Table 5.2: Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 1 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Normal Hierarchy and Upper Octant.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.18: Marginal posterior distributions of θ_{13} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.19: Marginal posterior distributions of θ_{23} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.

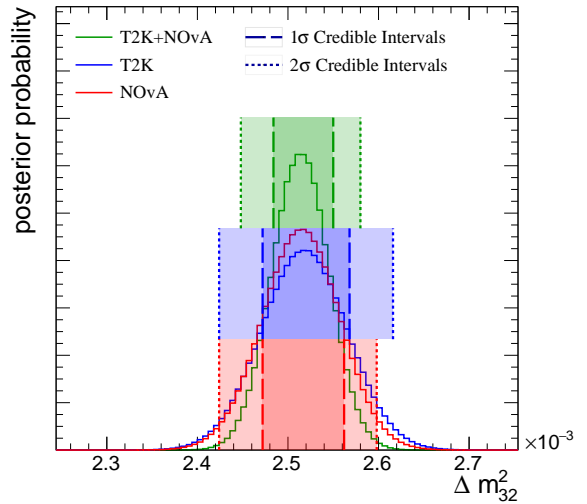
Figure 5.20: Marginal posterior distributions of δ_{CP} for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.21: Marginal posterior distributions of Δm_{32}^2 for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.22: Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.23: Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.



(b) Inverted hierarchy



(c) Normal hierarchy.

Figure 5.24: Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 1 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.


5.4.3 Asimov 4

Finally, Asimov 4 is a very degenerate oscillation point for both experiments. Focusing on δ_{CP} , T2K only has sensitivity to negative values of δ_{CP} regardless of the assumed mass hierarchy. However, this is the opposite for NOvA as they can only obtain negative values of δ_{CP} if Inverted hierarchy was true, so this creates this large difference in posterior probability visible in Figure 5.27a. As the posteriors are normalised in each figure, the contribution of δ_{CP} in the Normal Hierarchy (Figure 5.27c) looks larger than it actually is. The triple peak feature seen in Figure 5.27c is created by CP conserving values of $\delta_{CP} = -\pi/\pi$ better describing the data if normal hierarchy is assumed, but due the narrow band beam of T2K these are nearly indistinguishable. For the mass hierarchy sensitivity itself, T2K has none and NOvA is also limited, illustrated in Figure 5.28a showing the joint fit sensitivity to the correct assumed mass hierarchy. Overall, there is an increase in sensitivity that can be seen in the marginal posteriors when combining the two experiments for all Asimov points, highlighting the impact this joint fit can have on lifting degeneracies.


Finally, Asimov 4 has one interesting feature. In table 5.3, both experiments show an individual preference for the incorrect Normal Hierarchy hypothesis, albeit insignificant, but the joint fit sensitivity gives preference for the correct Inverted Hierarchy. The incorrect preference is understood to be a result of the small change in event rate, where a bigger fraction of the prior gives similar results in normal hierarchy than in inverted.

	$\sin^2 \theta_{23} < 0.5$			$\sin^2 \theta_{23} > 0.5$			Sum		
NH	0.097	0.199	0.212	0.227	0.346	0.377	0.324	0.545	0.589
IH	0.184	0.136	0.123	0.492	0.319	0.288	0.676	0.455	0.411
Sum	0.281	0.335	0.335	0.719	0.665	0.665	1.000		


Table 5.3: Fraction of posterior probability lying in different combinations of hierarchy and θ_{23} -octant from fit to Asimov 4 simulated data for T2K+NOvA (green), T2K-only (blue), and NOvA-only (red). The true Asimov point lays in Inverted Hierarchy and Upper Octant.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.25: Marginal posterior distributions of θ_{13} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.26: Marginal posterior distributions of θ_{23} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.27: Marginal posterior distributions of δ_{CP} for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.28: Marginal posterior distributions of Δm_{32}^2 for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.29: Marginal posterior distributions of $\delta_{CP} - \theta_{13}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.




(b) Inverted hierarchy




(c) Normal hierarchy.


Figure 5.30: Marginal posterior distributions of $\delta_{CP} - \theta_{23}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both hierarchies.



(b) Inverted hierarchy



(c) Normal hierarchy.

Figure 5.31: Marginal posterior distributions of $\Delta m_{32}^2 - \theta_{23}$ for Asimov 4 in both hierarchies with equal priors (top), and assuming inverted hierarchy (left) or normal hierarchy (right). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.

5.5 Goodness-of-Fit Metrics

A key aspect to the joint fit is ensuring that the joint-model fits the data well. The author’s final task was through a posterior predictive p -value described in section 4.3.5 in which one fraction of fake data sets with log-likelihood (\mathcal{L}_2) worse than the ‘real’⁴ data (\mathcal{L}_1) with a random throw of parameter values from the Markov Chain Monte Carlo. Fortunately, the MCMC means one can create a fake data set by reweighting the Monte Carlo with a vector of parameter values (the throw), and apply a Poisson fluctuation to that data set, then repeating many times to find the fraction. This test statistic allows one to find if the ensemble of models represented by the prior gives a good fit or not. With respect to the joint fit, this is achieved by randomly throwing from the MCMC posterior chain, including both the T2K and NOvA parameters and generating log-likelihoods for T2K and NOvA and comparing them to randomly generated statically fluctuated Monte Carlo. Summing the respective likelihoods for T2K+NOvA gives the 2 dimensional distribution defined in section 4.3.5.

However, in version 4 of the NOvA container it was not possible to run the NOvA software in such a way to take a vector of parameter values and generate a fake data set for lots of throws; this was remedied in version 5. Implementing the changes into Bifrost brought new challenges given the inability to view inside the respective containers. The major issue was actually getting the container to apply statistical fluctuations and set this as a ‘true’ data set, but then allowing one to reset this data file to calculate a normal likelihood, \mathcal{L}_1 , for the next throw.

In order to test the implementation of the Bifrost and container updates in version 5 for the posterior predictive p -value, a T2K+NOvA joint fit was run using the Asimov 4 data set. From this posterior, a posterior predictive p -value was generated, illustrated

⁴It should be emphasised that the ‘real’ data in this analysis is truth data generated via Monte Carlo prediction as the joint fit is still in development.

in Figure 5.32. The diagonal line represents $\mathcal{L}_1 = \mathcal{L}_2$, defining equal likelihood for the model to describe the real data and the statistically fluctuated fake data. The y -axis is the χ^2 , where $\chi^2 = -2 \ln \mathcal{L}$, between the Asimov 4 truth data and the MC prediction, and the x -axis is the χ^2 between fake data and MC prediction. The p -value itself is calculated by the ratio of $\mathcal{L}_1/\mathcal{L}_2$, or the number of throws below that diagonal line over the total number of throws, signifying the probability to better describe the data over fake data. Figure 5.32 shows a p -value of 1 which suggests that the model describes the data incredibly well. Thankfully, the realisation that the reason this skew exists and the low χ^2 distribution in the y -axis is present, is because of the use of Asimov data, generated via the Monte Carlo. An Asimov data set, by definition does not include any of the statistical fluctuations expected in the real data, thus the log-likelihood contribution, \mathcal{L}_1 , is minimal. Therefore, it is also difficult to tell if the Bifrost changes are indeed working. In order to test the posterior predictive method with a data set that (realistically) includes statistical fluctuations, a joint fit was carried out using a Poisson fluctuation around the Asimov 4 data. This includes a realistic χ^2 contribution from \mathcal{L}_1 and therefore a better representation of the test applied when using real data. Building the posterior distribution in this analysis adds more statistical and systematic uncertainty. Using the MCMC chains from this analysis and the steps mentioned in section 4.3.5, the p -value depicted in Figure 5.33 was obtained with a value of 0.76, showing around 3/4 of the distribution below the $\mathcal{L}_1/\mathcal{L}_2$ line. One may notice the difference in magnitude of the χ^2 data distribution between Figures 5.32 and 5.33, and this is due to the statistical and systematic uncertainty generated when using a statistical fluctuated Asimov data set. Of course since the test is not yet applied to the real data, this is not reflective of the true goodness-of-fit of the joint fit model, as that can only be obtained when it is run with the real data samples. However, the fact that the returned value is not extreme in the range [0,1] indicates that the implementation is working and now validated, ready for the real data fits.




Figure 5.32: A posterior predictive p -value distribution generated from the posterior of a T2K+NOvA joint fit with the Asimov 4 data set, before statistical fluctuations. The severe pull in the distribution is due to the lack of statistical uncertainties in the Asimov fit.




Figure 5.33: A posterior predictive p -value distribution generated from the posterior of a T2K+NOvA joint fit posterior, using a statistically fluctuated version of the Asimov 4 data set. Adding statistical variation into the data sets produces a more appropriate p -value of 0.76.

Chapter 6

Daya Bay Constraint

*If I have seen further, it is by
standing on the shoulders of
giants.*

Isaac Newton

Throughout this thesis, there have been multiple mentions of using an external reactor constraint, commonly referred to as the PDG constraint. A one-dimensional constraint of $\sin^2(\theta_{13})$ used by experiments to increase the sensitivity to not only θ_{13} , but also other correlated parameters. The PDG reactor constraint comes as a combination of results from a group of reactor experiments: Daya Bay, Double Chooz and Reno, creating a one-dimensional Gaussian distribution. The effects of applying this constraint can be seen in the marginalised posterior probability for θ_{13} , exhibited in Figure 4.7. The left panel shows the posterior distribution without the application of the constraint and the right side shows the distribution with this weighting applied. One can see the reduction in the 3σ interval range and much tighter constraint in the 1σ interval. Furthermore, applying this constraint gives MCMC steps in a less favourable region a very small weight, which also has an impact on the posterior of other parameters, as noted previously in Section 1.4.4 it would have an impact on the posterior probability for δ_{CP} . But given T2K's data, this also has a substantial effect on $\sin^2(\theta_{23})$. The left panel of




Figure 6.1: Left: The marginalised posterior probability density for $\sin^2(\theta_{23})$ with no reactor constraint applied. Right: The same but with the reactor constraint applied. Reweighting to the reactor constraint produces a better constraint on the measurement of $\sin^2(\theta_{23})$. Figures taken from [123].

Figure 6.1 shows the marginalised posterior for $\sin^2(\theta_{23})$, a bi-modal distribution creating a disjointed 1σ credible interval. When applying the reweighting, the unfavourable steps of $\sin^2(\theta_{13})$ that receive a heavy weight from the constraint mostly lay within the lower octant segment and thus change the shape of the posterior, visible on the right hand side of Figure 6.1. This reduces the bi-modal feature and eliminates the disjoint feature, creating a preference for the upper octant hypothesis. This happens due to the PDG’s lower best-fit value of θ_{13} , as shown in Figure 4.7. Referring to Equation 1.19, the leading order term for the $\nu_\mu \rightarrow \nu_e$ channel is determined by the magnitude of $\sin^2(2\theta_{13})$ and $\sin^2(\theta_{23})$. Therefore, if the PDG pulls the value of θ_{13} down, then the T2K fit must increase the value of θ_{23} to compensate, thus increasing the preference for the upper octant.

Diving into the PDG constraint itself, one can see from Figure 6.2, that the 2018 constraint is dominated by Daya Bay with a very similar best-fit value and error region. Using this and the other listed reactor experiments, PDG creates the averaged one-dimensional Gaussian. Two experiments, Daya Bay and RENO have also released constraints on $|\Delta m_{32}^2|$ as shown in Figure 6.3. More recently, the Daya Bay collab-




Figure 6.2: A global comparison of experimental constraints on $\sin^2(\theta_{13})$. This includes T2K, Daya Bay, NOvA, MINOS and RENO, using their 2018 oscillation results. The blue shaded region is the 2018 PDG best-fit constraint. Figure taken from [139].




Figure 6.3: A global comparison of experimental constraints on $|\Delta m_{32}^2|$. This includes T2K, Daya Bay, NOvA, MINOS and RENO, using their 2018 oscillation results. The blue shaded region is the 2018 PDG best-fit constraint. Figure taken from [139].

oration have also released a two-dimensional constraint, on both the amplitude (θ_{13}) and the location (Δm_{32}^2) of the oscillation dip. This analysis investigates whether the additional constraint on Δm_{32}^2 has noticeable effects on the final results when used by the T2K analysis.

6.1 The Daya Bay Experiment

The Daya Bay experiment [140, 141] is a nuclear reactor neutrino experiment, situated in Daya Bay, China, that is designed for sensitivity to the $\sin^2(\theta_{13})$ mixing angle through the electron antineutrino disappearance channel. The experiment consists of three main components: the neutrino sources, the experiment halls and the antineu-




Figure 6.4: A schematic of the Daya Bay experiment. The figure shows the nuclear reactor cores (D1-D2, L1-L4), the Antineutrino Detectors (AD1-AD8) and the corresponding experiment halls, EH1-EH3. Halls EH1 and EH2 are the location of the near detectors and EH3 is the far detector hall. Figure taken from [141].

trino detectors themselves. The antineutrinos are sourced from three nuclear power plants, the Daya Bay nuclear power plant, the Ling Ao nuclear power plant and the Ling Ao-II nuclear power plant. Each plant comprises of two identical cores spaced 88 m apart. Ling Ao is located 1100 m from Daya Bay and Ling Ao-II is separated by another 500 m, depicted in Figure 6.4. The experiment halls (labelled E1-3) house the antineutrino detectors and act as near detectors (EH1, EH2) and a far detector (EH3), located ~ 500 m and ~ 1600 m respectively. The experiment uses eight identical antineutrino detectors (AD1-8), each consisting of 20 tonnes of liquid scintillator doped with Gadolinium and surrounded by 192 20 cm photo-multiplier tubes. The ADs are sensitive to low-energy antineutrinos produced by the products of nuclear fission reaction from the aforementioned power plants, observing these via the Inverse β -Decay (IBD) process. The doped scintillator gives excellent discrimination between electron antineutrinos and electron neutrinos, which are otherwise indistinguishable, due to the neutron capture prompt signal.

One can see the oscillation probability spectrum for the electron antineutrino survival

channel at the bottom of Figure 6.4. The far detector is located near the dip of the survival probability which is driven by the $\sin^2(\theta_{13})$ mixing angle, giving the experiment sensitivity to θ_{13} . The frequency however is governed by the mass-splitting term. This is shown in the electron antineutrino survival probability formula,

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \cos^4(\theta_{13}) \sin^2(2\theta_{12}) \sin^2 \Delta_{21} - \sin^2(2\theta_{13}) (\cos^2(\theta_{12}) \sin^2 \Delta_{31} + \sin^2(\theta_{12}) \sin^2 \Delta_{32}), \quad (6.1)$$

where,

$$\Delta_{ij} = \frac{\Delta m_{ij}^2 L}{E}. \quad (6.2)$$

Given Daya Bay's short baseline of ~ 1.6 km and the low energy of the reactor antineutrinos, the second term of Equation 6.1 is very small and thus negligible, meaning the leading term is governed by the θ_{13} mixing angle. Because of this, the two mass splittings, Δm_{31}^2 and Δm_{32}^2 , are unresolvable from a single $\sin^2(\Delta_{3j})$ dependence, where $j = 1, 2$. Therefore, Daya Bay and other short-baseline experiments released best-fits of an effective mass splitting, Δm_{ee}^2 , essentially averaging between the two existing mass splittings [142]. Relating to equation 6.1, one can rewrite the electron disappearance probability as:

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \cos^4(\theta_{13}) \sin^2(2\theta_{12}) \sin^2 \Delta_{21} - \sin^2(2\theta_{13}) \sin^2(\Delta_{ee}). \quad (6.3)$$

This simplification works if the following equation is true,

$$\sin^2(\Delta_{ee}) \approx \cos^2(\theta_{12}) \sin^2 \Delta_{31} + \sin^2(\theta_{12}) \sin^2 \Delta_{32}, \quad (6.4)$$

which is best approximated if:

$$\Delta m_{ee}^2 \equiv \cos^2(\theta_{12})\Delta m_{31}^2 + \sin^2(\theta_{12})\Delta m_{32}^2 \quad (6.5)$$

for $L/E < 0.84 \text{ km/MeV}$. Daya Bay satisfies this criteria with the low energy neutrinos and short baselines. This is a non-trivial simplification and unfortunately a non-standard one also. Different definitions have arisen by both Daya Bay, RENO and theorists in varying literature [143–145]. However, Parke et al. show in [142] that for experiments such as Daya Bay, Equation 6.5 is most accurate.

More generally, it is shown in [146] that the oscillation probability can be expressed with a more complicated form:

$$\begin{aligned} P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = & 1 - 2s_{13}^2 c_{13}^2 \\ & + 2s_{13}^2 c_{13}^2 \sqrt{1 - 4s_{12}^2 c_{12}^2 \sin^2(\Delta_{21})} \cos(2\Delta_{32} \pm \phi) \\ & - 4c_{13}^4 s_{12}^2 c_{12}^2 \sin^2(\Delta_{21}), \end{aligned} \quad (6.6)$$

where,

$$\phi = \arctan \left(\frac{\sin(2\Delta_{21})}{\cos(2\Delta_{21}) + \tan^2(\theta_{12})} \right). \quad (6.7)$$

However, this prescription is not necessary for Daya Bay and Equation 6.5 is accurate enough. Daya Bay and current reactor experiments have a small L/E ratio, leading to the phase term (Equation 6.7) tending to zero and the root contained in the third term of Equation 6.6 tending to 1, shown in [146], due to the small magnitude of the solar mixing parameters. The oscillation probability stated in Equation 6.6 does become important for future experiments with a larger L/E ratio. The next-generation reactor experiment, JUNO, will use the same energy reactor neutrinos but a baseline ~ 50 times longer, so the more exact expression is necessary because ϕ is large and thus the $\cos(2\Delta_{32} \pm \phi)$ term gives access to the neutrino mass hierarchy in a way that is

independent of the matter effect [144]. Figure 6.5 shows a comparison of the formulae for Daya Bay’s baseline length and average neutrino energy. The largest difference is of the order $\sim 1\%$. Similarly to the T2K-NOvA comparisons described in Section 5.3.1, where a 1% difference was shown to be negligible at current precision and does not change the interpretation of the data.

Although not as sensitive as the effect in JUNO, the combination of T2K and the measurement of Δm_{ee}^2 has some slight sensitivity to the mass hierarchy, due to the known sign of Δm_{21}^2 , in that $m_2 > m_1$. In T2K’s analysis, the constraint is expressed on Δm_{32}^2 , and in Equation 6.5, for a given value of Δm_{ee}^2 the value of $|\Delta m_{32}^2|$ depends on the mass ordering and hence the reactor constraint is different for the two cases.




Figure 6.5: The upper panel contains a comparison of the $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ survival probability using the full Δm_{32}^2 and Δm_{ee}^2 formulas. Bottom panel: the ratio of the two expressions. Figure taken from [147].

6.2 Daya Bay's Log-Likelihood Surface




Figure 6.6: Left: A 2D histogram surface plot containing the $\Delta\chi^2$ value (colour coded z-axis) for a given set of θ_{13} and Δm_{32}^2 values in Inverted Hierarchy. Right: The same but in Normal Hierarchy. Here, $\Delta\chi^2 = -2 \ln \left(\frac{\mathcal{L}}{\mathcal{L}_{max}} \right)$, where \mathcal{L} is the likelihood value of the current entry, and \mathcal{L}_{max} is the maximum likelihood in the distribution. Showing the difference in correlation between the two hierarchies hypotheses and the mixing angle, θ_{13} . Data taken from [148].

The results from Daya Bay's analysis can be shown in two dimensional histograms representing $\Delta\chi^2$ surface plots, for both hierarchies, depicted in Figure 6.6. This constraint in $\sin^2(\theta_{13}) - \Delta m_{32}^2$ space has three interesting components. It contains an asymmetry for the constraint in either hierarchy, visible with the $\Delta\chi^2$ surface plot and has a small correlation between $\sin^2(\theta_{13})$ and Δm_{32}^2 for both hierarchies. Although minor, there is a negative correlation coefficients of 0.133 in Inverted Hierarchy and 0.113 in Normal Hierarchy. Published T2K analyses to date assume no correlations between oscillation parameter priors when running a fit. To implement this would require a modification of the fitting code. However in MaCh3, the new Daya Bay constraint can be incorporated by reweighting the existing analysis posteriors via the technique outlined in Section 4.3.2, which makes it possible to compare the effects of the 2D constraint without re-running the analyses.

Daya Bay's data release [148] is a text file containing varying data points of the θ_{13} mixing angle and effective mass splitting with a corresponding $\Delta\chi^2$, rather than being

plotted over a function. As a first attempt, the data points were simply filled into a two-dimensional histogram. The binning is determined by the data points published, and the initial reweighting was achieved by simply finding the bin content that contained the θ_{13} and $|\Delta m_{32}^2|$ values of the MCMC step. However, this meant that areas with the highest density of steps, i.e the best fit point, would receive the same reweighting. In order to increase the accuracy of the reweighting, multiple methods such as a ‘slicing’ technique, fitting two half Gaussians and also, looking into performing a MINUIT fit to obtain a function that best describes the surface plot were considered. These methods were tested in order to try and extend Daya Bay’s surface plot range. Although T2K and Daya Bay have similar sensitivities to Δm_{32}^2 , Daya Bay is much more sensitive to θ_{13} . Due to the lower sensitivity, a large integral of T2K θ_{13} posteriors sits outside of Daya Bay’s published data range. Therefore, extrapolation would be required in order to give these more extreme MCMC steps a weight. However since these points will be reweighted by a lower prior probability, they do not contribute much to the final distribution. It was important to test that the loss in MCMC statistics did not effect results and make sure they were still sufficient to produce stable contours. As there was no need to extrapolate outside of Daya Bay’s data fit range, the main focus shifting to ensuring a good interpolation of the published data. This was achieved using Delaunay triangulation to find a $\Delta\chi^2$ value between the given data points, creating smoother contours and less disjoint regions due to disproportionate reweighting.

6.3 Data Fit Results

Posterior distributions are used to build credible intervals. One can see these effects on the posterior distributions for no constraint applied (green), with the one-dimensional PDG Gaussian (blue) and the two-dimensional Daya Bay constraint applied (red), marginalised over both, for inverted ($\Delta m_{32}^2 < 0$) and normal ordering ($\Delta m_{32}^2 > 0$). The Daya Bay constraint used in this analysis contains more data than the Daya Bay

result used in the PDG constraint. Looking at the one-dimensional $\sin^2(\theta_{13})$, the updated Daya Bay constraint is slightly better as shown in Table 6.1. However, using the two-dimensional constraint trades the sensitivity in $\sin^2(\theta_{13})$ for sensitivity to Δm_{32}^2 .

	Value of $\sin^2(\theta_{13})$	Data reference
PDG	0.0220 ± 0.0007	[139, 149–152]
Daya Bay	0.0217 ± 0.0006	[148]

Table 6.1: A comparison of the one-dimensional $\sin^2(\theta_{13})$ priors for The PDG and Daya Bay constraint used in this analysis and the respective data releases used to create the constraint.

6.3.1 One-Dimensional Posteriors

Starting with the one-dimensional distribution of θ_{13} , Figure 6.7 shows, as expected, a much tighter constraint when applying either of the reactor constraints to T2K’s posterior. There is a visible shift in the interval range of the Daya Bay constraint compared to the PDG (see Table 6.2), which is interesting given the dominance of Daya Bay in the PDG constraint itself, exhibited in Figure 6.2. This PDG constraint uses an older data release from Daya Bay [151], highlighting Daya Bay’s increase in sensitivity with more data. It is visible in Figure 6.7 that using the Daya Bay constraint with T2K prefers a smaller value of $\sin^2(\theta_{13})$ compared to using the PDG external constraint, evidenced in Table 6.2. This is expected given the different priors listed in Table 6.1.

Applied Constraint	$\sin^2(\theta_{13})(\times 10^{-2})$		
	Both	NH	IH
No External Constraint	$2.44^{+0.51}_{-0.40}$	$2.36^{+0.44}_{-0.40}$	$2.63^{+0.48}_{-0.40}$
1D PDG Constraint	$2.21^{+0.07}_{-0.08}$	$2.20^{+0.07}_{-0.08}$	$2.22^{+0.07}_{-0.07}$
2D Daya Bay Constraint	$2.19^{+0.06}_{-0.07}$	$2.19^{+0.06}_{-0.07}$	$2.19^{+0.07}_{-0.06}$

Table 6.2: A comparison of the highest posterior density points for $\sin^2(\theta_{13})$ in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint.

The one-dimensional posteriors for Δm_{32}^2 , shown in Figure 6.8 show the most drastic change when applying the Daya Bay constraint. There is a tighter constraint in both normal (Figure 6.8c) and inverted (Figure 6.8b) hierarchy with a larger pull in the best fit point in the normal hierarchy due to Daya Bay's lower best-fit point, as stated in Table 6.3. One can also see the posteriors integrals in Table 6.6, indicating an increase in the fractional posterior probability for normal hierarchy when using the Daya Bay constraint. This reflects in the Bayes factor value shown in Table 6.7. Although there is an increase in the hypothesis preference compared to using the PDG constraint, it is still not substantial.

Applied Constraint	$\Delta m_{32}^2 (\times 10^{-3} eV^{-2})$		
	Both	NH	IH
No External Constraint	N/A	$2.50^{+0.06}_{-0.06}$	$-2.55^{+0.06}_{-0.06}$
1D PDG Constraint	N/A	$2.50^{+0.06}_{-0.06}$	$-2.54^{+0.05}_{-0.06}$
2D Daya Bay Constraint	N/A	$2.48^{+0.05}_{-0.04}$	$-2.56^{+0.05}_{-0.04}$

Table 6.3: A comparison of the highest posterior density points for Δm_{32}^2 in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint.

Figure 6.9 shows no real change in the constraint of θ_{23} between either of the reactor constraints. There is a bin's width different in the interval ranges, noticeable in Figure 6.9. There is a significant shift in the highest posterior density point when applying either reactor constraint, moving T2K's best-fit point from the lower octant ($\theta_{23} < 0.5$) to the upper octant ($\theta_{23} > 0.5$), as shown in Table 6.4. As stated in the beginning of the chapter, the decrease in θ_{13} causes an increase in θ_{23} to better model T2K's large electron neutrino event rate. This is also evident in the fraction of posterior table (Table 6.6), used to calculate the Bayes factors shown in Table 6.7, with an increased preference for the upper octant, $\mathcal{B}(UO/LO) = 2.38$. This is marginally larger than the PDG Bayes factor of 2.27 and the Bayes factor for no external constraint applied is 1.37.

Applied Constraint	$\sin^2(\theta_{23})$		
	Both	NH	IH
No External Constraint	$0.49^{+0.06}_{-0.02}$	$0.49^{+0.06}_{-0.02}$	$0.49^{+0.06}_{-0.02}$
1D PDG Constraint	$0.54^{+0.02}_{-0.05}$	$0.53^{+0.03}_{-0.04}$	$0.54^{+0.03}_{-0.04}$
2D Daya Bay Constraint	$0.54^{+0.02}_{-0.05}$	$0.53^{+0.03}_{-0.04}$	$0.54^{+0.02}_{-0.05}$

Table 6.4: A comparison of the highest posterior density points for $\sin^2(\theta_{23})$ in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint.

The final one-dimensional distribution is δ_{CP} . As shown in Figure 6.10a and Table 6.5, there is a minor increase to the constraint when comparing the PDG (blue) with the Daya Bay constraint (green). Both constraints still allow a CP -conserving value of π inside of the 2σ interval.

Applied Constraint	δ_{CP}		
	Both	NH	IH
No External Constraint	$-1.79^{+1.10}_{-0.97}$	$-2.42^{+5.56}_{-0.72}$	$-1.16^{+0.66}_{-0.91}$
1D PDG Constraint	$-1.85^{+0.85}_{-0.79}$	$-2.17^{+0.97}_{-0.66}$	$-1.29^{+0.54}_{-0.72}$
2D Daya Bay Constraint	$-1.85^{+0.91}_{-0.72}$	$-2.17^{+1.04}_{-0.60}$	$-1.29^{+0.53}_{-0.66}$


Table 6.5: A comparison of the highest posterior density points for δ_{CP} in both hierarchies with equal priors, and assuming normal or inverted hierarchy when using no external constraint, one-dimensional PDG external constraint and the two-dimensional Daya Bay external constraint.

	$\sin^2 \theta_{23} < 0.5$			$\sin^2 \theta_{23} > 0.5$			Sum		
NH	0.257	0.237	0.229	0.369	0.498	0.507	0.626	0.735	0.736
IH	0.165	0.069	0.067	0.209	0.196	0.197	0.374	0.265	0.264
Sum	0.422	0.306	0.296	0.578	0.694	0.704	1.000		


Table 6.6: Fraction of posterior probability lying in different combinations of mass hierarchy and θ_{23} -octant from fit to T2K data for T2K without a reactor constraint (green), T2K with PDG reactor constraint (blue), and T2K with 2D Daya Bay constraint (red).

$\mathcal{B}(H1/H2)$	No RC	PDG RC	Daya Bay RC
$\mathcal{B}(NH/IH)$	1.67	2.77	2.79
$\mathcal{B}(UO/LO)$	1.37	2.27	2.38


Table 6.7: A comparison of Bayes factors for no reactor constraint, PDG constraint and Daya Bay looking at mass hierarchy hypotheses and octant hypotheses.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.7: Posterior distributions for $\sin^2(\theta_{13})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy

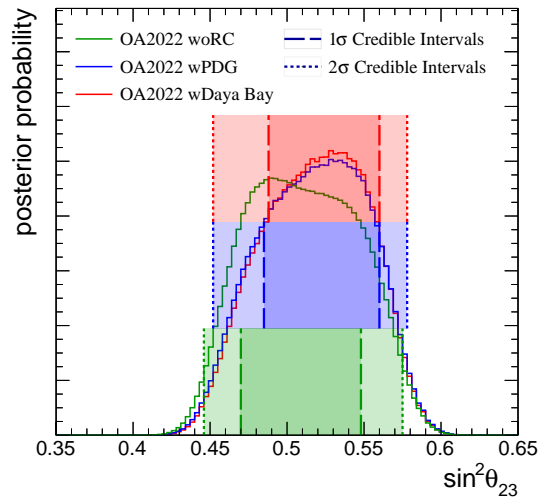
Figure 6.8: Posterior distributions for Δm_{32}^2 , showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.9: Posterior distributions for $\sin^2(\theta_{23})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies



(b) Inverted Hierarchy



(c) Normal Hierarchy

Figure 6.10: Posterior distributions for δ_{CP} , showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.

6.3.2 Two-Dimensional Posteriors

T2K also publishes sets of two-dimensional posterior distributions and their corresponding credible contours. In this section, the conventional 2D posterior plots used by T2K and NOvA will be shown, followed by other parameters in conjunction with Δm_{32}^2 to better illustrate how using the 2D Daya Bay constraint can affect other oscillation parameters.


Figure 6.11 shows the constraint of $\delta_{CP} - \sin^2(\theta_{13})$, illustrating why T2K uses a reactor constrain to improve its measurement of δ_{CP} . The width of the contour when using the Daya Bay constraint is very similar to the PDG reactor constraint due to the similar constraint on θ_{13} compared to T2K-only. Due to this better constrained range of θ_{13} , one sees a decrease in the uncertainty on δ_{CP} for both external constraints. The next set of contours are in the $\delta_{CP} - \sin^2(\theta_{23})$ parameter space which is important in comparison with NOvA due to the way the degeneracies affect their experiment, illustrated in Figure 6.12. As seen with the one-dimensional projections for the individual parameters, there is no significant difference between the PDG and Daya Bay constraint. The only notable feature exists in Figure 6.12c, with a better constraint of δ_{CP} in the normal hierarchy.

The two-dimensional contours for $\Delta m_{32}^2 - \sin^2(\theta_{23})$ show the largest differences in constraint between the PDG and Daya Bay constraint. Figure 6.13a illustrates the narrowing constraint on the normal hierarchy values of δm_{32}^2 , with visible overlap between the PDG's 1σ and Daya Bay's 2σ contour. The same is visible in the inverted hierarchy pad (bottom pad of Figure 6.13a).


Following this, the contours for $\Delta m_{32}^2 - \sin^2(\theta_{13})$ are shown in Figure 6.14. This showcases the increase in sensitivity when applying the Daya Bay constraint. Both external constraints create an increase in sensitivity compared to T2K without an external

constraint and the PDG constraint. Furthermore, one can see the small correlations mentioned in Section 6.2.


Finally, Figure 6.15 shows the two-dimensional constraint for $\Delta m_{32}^2 - \delta_{CP}$. One will notice the tighter constraint on δ_{CP} in inverted hierarchy (bottom pad of Figure 6.15a, when both hierarchies are fit simultaneously). This is due to the lifting of the mass hierarchy degeneracy for maximal CP -violation seen by T2K, as noted in Section 5.4 for Asimov point 1. Additionally, it was noted in Section B of [137] that T2K could observe a correlation between $\Delta m_{32}^2 - \delta_{CP}$ in the muon (anti)neutrino survival channel, which looks to be visible in Figures 6.15b and 6.15c. This could be further studied using the 2D Daya Bay constraint.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.11: Posterior distributions for $\delta_{CP} - \sin^2(\theta_{13})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.12: Posterior distributions for $\delta_{CP} - \sin^2(\theta_{23})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.13: Posterior distributions for $\Delta m_{32}^2 - \sin^2(\theta_{23})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies




(b) Inverted Hierarchy




(c) Normal Hierarchy


Figure 6.14: Posterior distributions for $\Delta m_{32}^2 - \sin^2(\theta_{13})$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.



(a) Both Hierarchies



(b) Inverted Hierarchy



(c) Normal Hierarchy

Figure 6.15: Posterior distributions for $\Delta m_{32}^2 - \delta_{CP}$, showing the change in credible intervals for no reactor constraint applied (green), with the PDG one-dimensional constraint (blue) and the Daya Bay two-dimensional constraint (red). The darker shade shows the one sigma interval range, and the lighter shaded region shows the two sigma interval range.

The contents of this chapter have given an overview into the reactor neutrino experiment, Daya Bay, and their two-dimensional constraint on the mixing angle and mass splitting, $\theta_{13} - \Delta m_{ee}^2$. The effective mass splitting term Δm_{ee}^2 is interesting due to being invariant under mass hierarchy selection. This thesis carried out an analysis of using Daya Bay's data release as a prior, and comparing it to the standard one-dimensional constraint on the mixing angle, θ_{13} prior, when combining it with T2K's current Bayesian analysis. Overall, there was an increase in constraint on some of the oscillation parameters, namely θ_{23}, θ_{23} and $|\Delta m_{32}^2|$, when using Daya Bays 2D prior, with no significant increase to the other three oscillation parameters. Another aspect of the analysis was looking into the Bayes Factor, which allows one to find a preference in hypotheses. The two hypotheses in question were the octant degeneracy and mass hierarchy degeneracy, with a small increase in both when compared to the current standard T2K analysis.

Chapter 7

Conclusion

*Started making it, had a
breakdown, Bon Appetit*

James Acaster

This thesis focused on how T2K’s measurement of oscillation parameters combine with external information. The first was an initial introduction into a beyond-the-standard PMNS parameterisation and how the cyclic parameter δ_{CP} has two different observable effects that can be decomposed as two terms to test consistency with the PMNS model. Although this ultimately took a back seat to the T2K-NOvA joint fit effort, it demonstrates that this formalism can successfully be carried out even with a numerical probability calculator like PROB3++ that cannot explore the expanded parameter space directly. Furthermore, there are on-going studies within T2K to implement this into other analysis fitters with results due to be published in the near future.

The second variation is through the efforts of a full joint fit between two long-baseline neutrino experiments, T2K and NOvA. Chapter 5 highlighted the workflow of the containers and the increase in sensitivity that was present in the Markov chain Monte Carlo posteriors, for a series of selected Asimov points that show a range of parameter values. A significant improvement in the sensitivity to the mass hierarchy is visible, as well as

improvements to all other oscillation parameters. Although the marginalised posterior distributions shown throughout Chapter 5 are just the MC sensitivity, the results are interesting due to the contrast in best-fit points by the two experiments. With data fits currently ongoing, T2K-NOvA will enter a review period for both collaborations, with a public release of the analysis expected shortly after.

Finally, Chapter 6 shows the effects of adding a new external input on T2K's oscillation parameter constraints. Ordinarily, T2K would use the PDG's one-dimensional Gaussian constraint on the θ_{13} mixing angle, capitalising on the increase in constraint, notably moving T2K's best-fit point for the mixing angle, θ_{23} , from the lower octant to the upper octant. Daya Bay released the first two-dimensional constraint on $\theta_{13} - \Delta m_{ee}^2$ which shows a small correlation between the two parameters. When applying this constraint with T2K's data and model, an enhanced measurement of $|\Delta m_{32}^2|$ is achieved, improving the 1σ constraint by 25% in normal hierarchy and 18% in inverted hierarchy. The Bayes factors displayed in Table 6.7 indicate a small increase in preference for both the normal mass hierarchy hypothesis and the upper octant hypothesis, although the interpretation of this is up to the reader's discretion. This reweighting has now been implemented inside of the oscillation analysis and will hopefully yield some interesting results in future iterations as T2K increases the amount of data collected. To progress this analysis, one could perform a joint likelihood surface reweighting by combining the results of Daya Bay and Reno, as shown in [153]. This is likely to become of greater interest to the community once the current generation of reactor of experiments release their final results, especially if an accepted combination is made available in the 2D parameter space.

The analysis work presented here has many potential avenues for moving forward. A measurement of Δm_{ee}^2 from reactor experiments can potentially give more information on the mass hierarchy given the increase in precision on Δm_{32}^2 when combining with

long-baseline, experiments as it relates differently in normal hierarchy or inverted hierarchy, as discussed in Section 6.3. Future experiments may have a wide-band beam that will give greater sensitivity to the CP conserving term, $\cos(\delta)$, and with greater statistical precision it will be interesting to investigate the consistency with the standard model and test alternate parameterisations for potential unitarity violations, like that shown in Chapter 2.

Looking to the future, T2K started undergoing upgrades on the Near Detector. Coupling this with the upgrades to the J-PARC beam will lead to reduction in systematics and improvements on the cross-section, flux and oscillation model. Future Experiments will pave the way for exciting conclusions within particle physics and beyond and allow a deeper understanding of neutrino physics and its implications, which could explain the unanswered questions of today.

Appendices

Appendix A

Full Set of Asimov Oscillation Parameter Sets

The full list and description of the five Asimov oscillation parameter sets used throughout the aforementioned studies [54] are presented in this Appendix. Three parameters were kept constant in all five points, $\Delta m_{21}^2, \theta_{12}, \theta_{13}$, due to the application of solar and reactor priors.

$$\begin{aligned}\Delta m_{21}^2 &= (7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2/\text{c}^4 \\ \sin^2(\theta_{12}) &= 0.307 \pm 0.013 \\ \sin^2(\theta_{13}) &= (2.18 \pm 0.007) \times 10^{-2}\end{aligned}\tag{A.1}$$

A.1 Asimov 0

This Asimov set corresponds to the **NOvA 2020 Best Fit** value.

$$\begin{aligned}
\Delta m_{32}^2 &= 2.41 \times 10^{-3} \text{ eV}^2/c^4 \\
\sin^2(\theta_{23}) &= 0.57 \\
\delta_{CP} &= 0.83\pi
\end{aligned}
\tag{A.2}$$

A.2 Asimov 1

This is equivalent to **T2K Asimov A** and very similar to T2K 2020 Best fit point.

$$\begin{aligned}
\Delta m_{32}^2 &= 2.509 \times 10^{-3} \text{ eV}^2/c^4 \\
\sin^2(\theta_{23}) &= 0.528 \\
\delta_{CP} &= 1.49\pi
\end{aligned}
\tag{A.3}$$

A.3 Asimov 2

The values for this are the same as **T2K Asimov B** and is CP conserving.

$$\begin{aligned}
\Delta m_{32}^2 &= 2.509 \times 10^{-3} \text{ eV}^2/c^4 \\
\sin^2(\theta_{23}) &= 0.45 \\
\delta_{CP} &= 0
\end{aligned}
\tag{A.4}$$

A.4 Asimov 3

This Asimov point was chosen at the **maximal-mixing value**.

$$\begin{aligned}
\Delta m_{32}^2 &= 2.45 \times 10^{-3} \text{ eV}^2/\text{c}^4 \\
\sin^2(\theta_{23}) &= 0.5 \\
\delta_{CP} &= 0.5\pi
\end{aligned}
\tag{A.5}$$

A.5 Asimov 4

Global fitters performing joint analysis prefer **IH + UO + maximum CP violation**. This Asimov point is chosen to cover this option.

$$\begin{aligned}
\Delta m_{32}^2 &= -2.45 \times 10^{-3} \text{ eV}^2/\text{c}^4 \\
\sin^2(\theta_{23}) &= 0.55 \\
\delta_{CP} &= 1.5\pi
\end{aligned}
\tag{A.6}$$

Bibliography

- [1] L. Nanni. “Fermi theory of beta decay: A first attempt at electroweak unification”. In: *arXiv:1803.07147* (2018).
- [2] W. Pauli. “Pauli letter collection: letter to Lise Meitner”. Typed copy. URL: <http://cds.cern.ch/record/83282>.
- [3] J. Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129 (1932), p. 312. DOI: 10.1038/129312a0.
- [4] C. D. Ellis and W. A. Wooster. “The average energy of disintegration of radium E”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 117.776 (1927), pp. 109–123.
- [5] F. Reines. “The neutrino: From poltergeist to particle”. In: *Reviews of Modern Physics* 68.2 (1996), p. 317.
- [6] C. L. Cowan Jr et al. “Detection of the free neutrino: a confirmation”. In: *Science* 124.3212 (1956), pp. 103–104.
- [7] R. Beth and C. Lasky. “The Brookhaven Alternating Gradient Synchrotron: Construction of a massive nuclear research machine requires ideas, men, and methods from many fields.” In: *Science* 128.3336 (1958), pp. 1393–1401.
- [8] G. Danby et al. “Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos”. In: *Physical Review Letters* 9.1 (1962), p. 36.

- [9] M. L. Perl et al. “Evidence for anomalous lepton production in $e^+ - e^-$ annihilation”. In: *Physical Review Letters* 35.22 (1975), p. 1489.
- [10] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (2001), pp. 218–224.
- [11] J. Bagger et al. “Precision electroweak measurements on the Z resonance”. In: *Physics Reports* 427.5-6 (2006), pp. 257–454.
- [12] B. Pontecorvo. “Inverse- β processes”. In: *Chalk River Laboratory Report PD-205* (1946).
- [13] R. Davis Jr, D. S. Harmer, and K. C. Hoffman. “Search for neutrinos from the sun”. In: *Physical Review Letters* 20.21 (1968), p. 1205.
- [14] J. Bahcall et al. “Solar Neutrino Flux.” In: *The Astrophysical Journal* 137 (1963), pp. 344–346.
- [15] J. Abdurashitov et al. “Measurement of the solar neutrino capture rate with gallium metal”. In: *Physical Review C* 60.5 (1999), p. 055801.
- [16] W. Hampel et al. “GALLEX solar neutrino observations: Results for GALLEX IV”. In: *Physics Letters B* 447.1-2 (1999), pp. 127–133.
- [17] K. S. Hirata et al. “Observation of B-8 solar neutrinos in the Kamiokande-II detector”. In: *Physical Review Letters* 63.1 (1989), p. 16.
- [18] Q. R. Ahmad et al. “Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 89.1 (2002), p. 011301.
- [19] J. Evans. “The MINOS experiment: results and prospects”. In: *Advances in High Energy Physics* 2013 (2013), p. 182537.
- [20] J. M. LoSecco. “The History of “Anomalous” Atmospheric Neutrino Events: A First Person Account”. In: *Physics in Perspective* 18.2 (2016), pp. 209–241.

- [21] S. Fukuda et al. “The Super-Kamiokande detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 501.2-3 (2003), pp. 418–462.
- [22] E. Peterson. “Atmospheric neutrino studies in Soudan 2”. In: *Nuclear Physics B - Proceedings Supplements* 77.1 (1999), pp. 111–116. ISSN: 0920-5632. DOI: 10.1016/S0920-5632(99)00405-3.
- [23] R. Becker-Szendy et al. “Neutrino measurements with the IMB director”. In: *Nuclear Physics B-Proceedings Supplements* 38.1-3 (1995), pp. 331–336.
- [24] K. Hirata et al. “Experimental study of the atmospheric neutrino flux”. In: *Physics Letters B* 205.2-3 (1988), pp. 416–420.
- [25] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Physical Review Letters* 81.8 (1998), p. 1562.
- [26] B. Pontecorvo. “Mesonium and antimesonium”. In: *Zhur. Eksptl'. i Teoret. Fiz.* 33 (1957).
- [27] M. Gell-Mann and A. Pais. “Behavior of neutral particles under charge conjugation”. In: *Murray Gell-Mann: Selected Papers*. World Scientific, 2010, pp. 65–67.
- [28] C.-S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Physical review* 105.4 (1957), p. 1413.
- [29] J. H. Christenson et al. “Evidence for the 2π Decay of the K^0 Meson”. In: *Physical Review Letters* 13.4 (1964), p. 138.
- [30] A. D. Sakharov. “Violation of CP-invariance, C-asymmetry, and baryon asymmetry of the Universe”. In: *In The Intermissions... Collected Works on Research into the Essentials of Theoretical Physics in Russian Federal Nuclear Center, Arzamas-16*. World Scientific, 1998, pp. 84–87.
- [31] K. Zuber. *Neutrino Physics*. CRC press, 2011.

- [32] B. Kayser. “Are neutrinos their own antiparticles?” In: *Journal of Physics: Conference Series*. Vol. 173. 1. IOP Publishing. 2009, p. 012013.
- [33] S. Abe et al. “Search for the Majorana nature of neutrinos in the inverted mass ordering region with KamLAND-Zen”. In: *Physical Review Letters* 130.5 (2023), p. 051801.
- [34] R. P. Litchfield. *Neutrino induced events in the MINOS detectors*. Tech. rep. Fermi National Accelerator Lab.(FNAL), Batavia, IL (United States), 2008.
- [35] “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. In: *Nature Physics* 18.2 (2022), pp. 160–166.
- [36] P. F. De Salas et al. “Neutrino mass ordering from oscillations and beyond: 2018 status and future prospects”. In: *Frontiers in Astronomy and Space Sciences* 5 (2018), p. 36.
- [37] S. Vagnozzi. “Cosmological searches for the neutrino mass scale and mass ordering”. In: *arXiv:1907.08010* (2019).
- [38] L. Wolfenstein. “Neutrino oscillations in matter”. In: *Physical Review D* 17.9 (1978), p. 2369.
- [39] S. Mikheyev and A. Y. Smirnov. “Resonant neutrino oscillations in matter”. In: *Progress in Particle and Nuclear Physics* 23 (1989), pp. 41–136.
- [40] S. Boyd. *Lectures on Neutrino Oscillations*. 2020. URL: https://warwick.ac.uk/fac/sci/physics/staff/academic/boyd/stuff/neutrinolectures/lec_oscillations.pdf.
- [41] A. De Santo. “An experimentalist’s view of neutrino oscillations”. In: *International Journal of Modern Physics A* 16.25 (2001), pp. 4085–4151.
- [42] A. Y. Smirnov. “The MSW effect and matter effects in neutrino oscillations”. In: *Physica Scripta* 2005.T121 (2005), p. 57.

- [43] S. Abe et al. “Precision measurement of neutrino oscillation parameters with KamLAND”. In: *Physical Review Letters* 100.22 (2008), p. 221803.
- [44] J. Cao. “Daya Bay neutrino experiment”. In: *arXiv:hep-ex/0509041* (2005).
- [45] J. Ahn, R. Collaboration, et al. “RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang”. In: *arXiv:1003.1391* (2010).
- [46] F. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Physical Review Letters* 108.17 (2012), p. 171803.
- [47] K. Abe et al. “Observation of electron neutrino appearance in a muon neutrino beam”. In: *Physical Review Letters* 112.6 (2014), p. 061802.
- [48] J. K. Ahn et al. “Observation of reactor electron antineutrinos disappearance in the RENO experiment”. In: *Physical Review Letters* 108.19 (2012), p. 191802.
- [49] S. Ahn et al. “Detection of accelerator-produced neutrinos at a distance of 250 km”. In: *Physics Letters B* 511.2-4 (2001), pp. 178–184.
- [50] D. Michael et al. “The MINOS experiment”. In: *Progress in Particle and Nuclear Physics* 48.1 (2002), pp. 99–109.
- [51] K. Yasutome et al. “Six Sample Joint Oscillation Analysis with T2K Run1-10 Data including a new multi-ring μ like sample”. In: TN-430 (2023). URL: <https://www.t2k.org/docs/technotes/430>.
- [52] K. Abe et al. “The T2K experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (2011), pp. 106–135.
- [53] D. Ayres et al. “The NOvA Technical Design Report”. In: (Oct. 2007). DOI: 10.2172/935497.
- [54] P. Zyla et al. “Particle data group”. In: *Prog. Theor. Exp. Phys.* 2020 (2020), p. 083C01.

- [55] P. Kyberd et al. “nuSTORM: neutrinos from STOREd Muons”. In: *arXiv:1206.0294* (2012).
- [56] O. Benhar and A. Fabrocini. “Two-nucleon spectral function in infinite nuclear matter”. In: *Physical Review C* 62.3 (2000), p. 034304.
- [57] C. H. Smith. *NEUTRINO REACTIONS AT ACCELERATOR ENERGIES*. Tech. rep. Stanford Univ., Calif., 1972.
- [58] B. Batell, J. L. Feng, and S. Trojanowski. “Detecting dark matter with far-forward emulsion and liquid argon detectors at the LHC”. In: *Physical Review D* 103.7 (2021), p. 075023.
- [59] R. Gran et al. “Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV”. In: *Physical Review D* 88.11 (2013), p. 113007.
- [60] M. Martini et al. “Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei”. In: *Physical review C* 80.6 (2009), p. 065501.
- [61] D. Rein and L. M. Sehgal. “Neutrino-excitation of baryon resonances and single pion production”. In: *Annals of Physics* 133.1 (1981), pp. 79–153.
- [62] C. Berger and L. M. Sehgal. “Partially conserved axial vector current and coherent pion production by low energy neutrinos”. In: *Physical Review D* 79.5 (2009), p. 053003.
- [63] J. M. Conrad, M. H. Shaevitz, and T. Bolton. “Precision measurements with high-energy neutrino beams”. In: *Reviews of Modern Physics* 70.4 (1998), p. 1341.
- [64] W. Ma et al. “Current status of final-state interaction models and their impact on neutrino-nucleus interactions”. In: *Journal of Physics: Conference Series*. Vol. 888. 1. IOP Publishing, 2017, p. 012171.
- [65] S. Gollapinni. “Neutrino cross section future”. In: *arXiv:1602.05299* (2016).

- [66] K. Abe et al. “Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target”. In: *arXiv:2303.03222* (2023).
- [67] H. Yokomakura, K. Kimura, and A. Takamura. “Overall feature of CP dependence for neutrino oscillation probability in arbitrary matter profile”. In: *Physics Letters B* 544.3-4 (2002), pp. 286–294.
- [68] L. H. Whitehead. “Neutrino oscillations with MINOS and MINOS+”. In: *Nuclear Physics B* 908 (2016), pp. 130–150.
- [69] M. Tanabashi et al. “Particle data group”. In: *Phys. Rev. D* 98.3 (2018), p. 030001.
- [70] URL: <https://webhome.phy.duke.edu/~raw22/public/Prob3++/>.
- [71] *About T2K*. URL: <https://t2k-experiment.org/t2k/>.
- [72] H. Hotchi et al. “J-PARC 3-GeV RCS: 1-MW beam operation and beyond”. In: *Journal of Instrumentation* 15.07 (2020), P07022.
- [73] T. Nakamoto et al. “Status of Superconducting Magnet System for the J-PARC Neutrino Beam Line”. In: *IEEE Transactions on Applied Superconductivity* 21.3 (2011), pp. 1700–1704. DOI: 10.1109/TASC.2010.2094174.
- [74] L. Berns. “T2K results on long-baseline oscillations”. Neutrino Oscillation Workshop. 2022. URL: <https://agenda.infn.it/event/30418/contributions/170626/attachments/95733/131672/20220904-NOW-T2K-v7.pdf>.
- [75] L. Berns et al. “Flux Prediction and Uncertainty with NA61/SHINE 2010 Replica Target Measurements”. In: TN-401 (2022). URL: <https://www.t2k.org/docs/technotes/401>.
- [76] D. Beavis, A. Carroll, and I. Chiang. *Long baseline neutrino oscillation experiment at the AGS. Physics design report*. Tech. rep. Brookhaven National Lab., 1995.

- [77] P. Fernandez. “Recent results from the long-baseline (LBL) neutrino oscillation experiments”. In: *SciPost Physics Proceedings* 1 (2019), p. 029.
- [78] K. Abe et al. “Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 694 (2012), pp. 211–223.
- [79] M. Yokoyama et al. “Application of Hamamatsu MPPCs to T2K neutrino detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 610.1 (2009), pp. 128–130.
- [80] K. Abe et al. “Measurements of $\bar{\nu}_\mu$ and $\bar{\nu}_\mu + \nu_\mu$ charged-current cross-sections without detected pions or protons on water and hydrocarbon at a mean anti-neutrino energy of 0.86 GeV”. In: *Progress of Theoretical and Experimental Physics* 2021.4 (2021), p. 043C01.
- [81] M. Antonova et al. “Baby MIND: a magnetized segmented neutrino detector for the WAGASCI experiment”. In: *Journal of Instrumentation* 12.07 (2017), p. C07028.
- [82] A. Hiramoto et al. “First measurement of $\bar{\nu}_\mu$ and ν_μ charged-current inclusive interactions on water using a nuclear emulsion detector”. In: *Physical Review D* 102.7 (2020), p. 072006.
- [83] T. Odagawa et al. “Design and performance of a scintillation tracker for track matching in nuclear-emulsion-based neutrino interaction measurement”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1034 (2022), p. 166775.
- [84] P. Giorgio et al. “T2K-WAGASCI: MIDAS-based DAQ software and online monitor for the readout of a large number of MPPCs”. In: *Journal of Physics: Conference Series*. Vol. 1468. 1. IOP Publishing. 2020, p. 012152.

- [85] E. Frank et al. “A dedicated device for measuring the magnetic field of the ND280 magnet in the T2K experiment”. In: *Journal of instrumentation* 7.01 (2012), P01018.
- [86] G. Arnison et al. “Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider”. In: *Physics Letters B* 126.5 (1983), pp. 398–410.
- [87] J. Altegoer et al. “The NOMAD experiment at the CERN SPS”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 404.1 (1998), pp. 96–128.
- [88] S. Aoki et al. “The T2K side muon range detector (SMRD)”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 698 (2013), pp. 135–146.
- [89] S. Assylbekov et al. “The T2K ND280 off-axis pi-zero detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 686 (2012), pp. 48–63.
- [90] N. Abgrall et al. “Time projection chambers for the T2K near detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 637.1 (2011), pp. 25–46.
- [91] I. Giomataris et al. “Micromegas in a bulk”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 560.2 (2006), pp. 405–408.
- [92] P. Amaudruz et al. “The T2K fine-grained detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 696 (2012), pp. 1–31.
- [93] D. Allan et al. “The electromagnetic calorimeter for the T2K near detector ND280”. In: *Journal of Instrumentation* 8.10 (2013), P10019.

- [94] K. Abe et al. “J-PARC neutrino beamline upgrade Technical Design Report”. In: *arXiv:1908.05141* (2019).
- [95] K. Abe et al. “T2K ND280 Upgrade–Technical Design Report”. In: *arXiv:1901.03750* (2019).
- [96] M. S. Dixit et al. “Position sensing from charge dispersion in micro-pattern gas detectors with a resistive anode”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 518.3 (2004), pp. 721–727.
- [97] Y. Itow et al. “The JHF-Kamioka neutrino project”. In: *hep-ex/0106019* (2001).
- [98] L. Machado. “Super-Kamiokande Pre-Supernova Alarm and Combined Monitoring with KamLAND”. XVIII International Conference on Topics in Astroparticle and Underground Physics. 2023. URL: https://indico.cern.ch/event/1199289/contributions/5447207/attachments/2705523/4696708/TAUP2023_preSN_LucasMachado.pdf.
- [99] H. Sekiya. “The Super-Kamiokande gadolinium project”. In: *Journal of Physics: Conference Series*. Vol. 1342. 1. IOP Publishing. 2020, p. 012044.
- [100] L. Machado et al. “Pre-supernova alert system for Super-Kamiokande”. In: *The Astrophysical Journal* 935.1 (2022), p. 40.
- [101] P. Adamson et al. “The NuMI neutrino beam”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (2016), pp. 279–306.
- [102] L. Kolupaeva. “Current results of the NOvA experiment”. In: *EPJ Web of Conferences*. Vol. 125. EDP Sciences. 2016, p. 01002.
- [103] Y. Kudenko. “Neutrino detectors for oscillation experiments”. In: *Journal of Instrumentation* 12.06 (2017), p. C06003.

- [104] P. Shanahan et al. “Physics with NOvA: A half-time review”. In: *The European Physical Journal Special Topics* (2021), pp. 1–15.
- [105] C. Robert and G. Casella. “A short history of Markov chain Monte Carlo: Subjective Recollections From Incomplete Data”. In: (2011).
- [106] R. Calland. “A 3 flavour joint near and far detector neutrino oscillation analysis at T2K”. PhD thesis. University of Liverpool, 2014.
- [107] G. Cowan. “Data analysis: frequently Bayesian”. In: *Phys. Today* 60.4 (2007), pp. 82–83.
- [108] P. Gregory. *Bayesian logical data analysis for the physical sciences: a comparative approach with Mathematica® support*. Cambridge University Press, 2005.
- [109] T. Doyle et al. “Constraining the flux and cross-section models for the 2021 oscillation analysis using ND280 data”. In: TN-423 (2022). URL: <https://www.t2k.org/docs/technotes/423>.
- [110] D. Gamerman and H. F. Lopes. *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. CRC press, 2006.
- [111] W. K. Hastings. *Monte Carlo sampling methods using Markov chains and their applications*. Oxford University Press, 1970.
- [112] K. E. Duffy. “Measurement of the neutrino oscillation parameters $\sin^2 \theta_{23}$, Δm_{23}^2 , $\sin^2 \theta_{13}$, and δ_{CP} in neutrino and antineutrino oscillation at T2K”. PhD thesis. University of Oxford, 2016.
- [113] W. R. Gilks, S. Richardson, and D. Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.
- [114] D. A. Spade. “Markov chain Monte Carlo methods: Theory and practice”. In: *Handbook of Statistics*. Vol. 43. Elsevier, 2020, pp. 1–66.

- [115] J. Dunkley et al. “Fast and reliable Markov chain Monte Carlo technique for cosmological parameter estimation”. In: *Monthly Notices of the Royal Astronomical Society* 356.3 (2005), pp. 925–936.
- [116] E. T. Atkin. “Neutrino Oscillations Analysis at the T2K experiment including studies of new uncertainties on interactions involving final state hadrons”. PhD thesis. Imperial College London, 2023.
- [117] H. Jeffreys. *The Theory of Probability*. Oxford Classic Texts in the Physical Sciences. OUP Oxford, 1998. ISBN: 9780191589676. URL: <https://books.google.co.uk/books?id=vh9Act9rtzQC>.
- [118] M. D. Lee and E.-J. Wagenmakers. *Bayesian cognitive modeling: A practical course*. Cambridge university press, 2014.
- [119] A. Gelman, X.-L. Meng, and H. Stern. “Posterior predictive assessment of model fitness via realized discrepancies”. In: *Statistica sinica* (1996), pp. 733–760.
- [120] C. Bronner et al. “Markov Chain Monte Carlo framework for electron neutrino appearance analysis using electron angle and momentum information”. In: TN-137 (2013). URL: <https://www.t2k.org/docs/technotes/137>.
- [121] C. Andreopoulos et al. “Valor Run 1–10 Analysis”. In: TN-394 (2022). URL: <https://www.t2k.org/docs/technotes/394>.
- [122] F. James et al. *MINUIT User’s Guide*. 2004.
- [123] E. Atkin et al. “Measuring Oscillation Parameters in a joint ND280+SK analysis using Markov Chain Monte Carlo Methods”. In: TN-429 (2023). URL: <https://www.t2k.org/docs/technotes/429>.
- [124] J. G. Walsh. “Constraining the T2K Neutrino Oscillation Parameter Results Using Data from the Off-Axis Near Detector, ND280: Implementation of a Nucleon Removal Energy Systematic Uncertainty Treatment in the BANFF Fit”. PhD thesis. Lancaster University (United Kingdom), 2022.

- [125] Y. Hayato. “A NEUTRINO INTERACTION SIMULATION PROGRAM LIBRARY NEUT.” In: *Acta Physica Polonica B* 40.9 (2009).
- [126] U. Mosel. “Neutrino interactions with nucleons and nuclei: importance for long-baseline experiments”. In: *Annual Review of Nuclear and Particle Science* 66 (2016), pp. 171–195.
- [127] A. Hillairet. “Study of neutrino interactions at the T2K near detector”. In: *Journal of Physics: Conference Series*. Vol. 650. 1. IOP Publishing. 2015, p. 012014.
- [128] E. Atkin et al. “MaCh3 Run 1-10 Analysis”. In: TN-393 (2023). URL: <https://www.t2k.org/docs/technotes/393>.
- [129] N. Abgrall, O. Andreeva, A. Aduszkiewicz, et al. “NA61/SHINE facility at the CERN SPS: beams and detector system”. In: *Journal of Instrumentation* 9.06 (2014), P06005.
- [130] N. Abgrall et al. “Measurements of π^\pm , K^\pm and proton double differential yields from the surface of the T2K replica target for incoming 31 GeV/c protons with the NA61/SHINE spectrometer at the CERN SPS: NA61/SHINE Collaboration”. In: *The European Physical Journal C* 79 (2019), pp. 1–45.
- [131] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71 (2011), pp. 1–19.
- [132] P. B. Denton et al. “Snowmass Neutrino Frontier: NF01 Topical Group Report on Three-Flavor Neutrino Oscillations”. In: *arXiv:2212.00809* (2022).
- [133] I. Esteban et al. “The fate of hints: updated global analysis of three-flavor neutrino oscillations”. In: *Journal of High Energy Physics* 2020.9 (2020), pp. 1–22.
- [134] URL: <https://docs.docker.com/reference/>.

- [135] E. Atkin et al. “Measuring PMNS parameters in a joint T2K-NOvA analysis: Fitting workstream”. In: TN-435 (2023). URL: <https://www.t2k.org/docs/technotes/435>.
- [136] M. Acero et al. “Improved measurement of neutrino oscillation parameters by the NOvA experiment”. In: *Physical Review D* 106.3 (2022), p. 032004.
- [137] K. Abe et al. “Improved constraints on neutrino mixing from the T2K experiment with 3.13×10^{21} protons on target”. In: *Physical Review D* 103.11 (2021), p. 112008.
- [138] E. Atkin et al. “Impact of Interaction Model Parameters on the NOvA-T2K Joint Oscillation Fit”. In: TN-437 (2023). URL: <https://www.t2k.org/docs/technotes/437>.
- [139] G. Bak et al. “Measurement of reactor antineutrino oscillation amplitude and frequency at RENO”. In: *Physical Review Letters* 121.20 (2018), p. 201801.
- [140] X. Guo. “A Precision Measurement of the Neutrino Mixing Angle θ_{13} using Reactor Antineutrinos at Daya Bay”. In: *hep-ex/0701029* (2007).
- [141] J. Cao and K.-B. Luk. “An overview of the Daya Bay reactor neutrino experiment”. In: *Nuclear Physics B* 908 (2016), pp. 62–73.
- [142] H. Nunokawa, S. Parke, and R. Z. Funchal. “Another possible way to determine the neutrino mass hierarchy”. In: *Physical Review D* 72.1 (2005), p. 013009.
- [143] S. J. Parke and R. Z. Funchal. “Comment on Daya Bay’s definition and use of Δm_{ee}^2 ”. In: *arXiv:1903.00148* (2019).
- [144] F. An et al. “Neutrino physics with JUNO”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.3 (2016), p. 030401.
- [145] H. Minakata et al. “Determining neutrino mass hierarchy by precision measurements in electron and muon neutrino disappearance experiments”. In: *Physical Review D* 74.5 (2006), p. 053008.

- [146] S. Parke. “What is Δm_{ee}^2 ?” In: *arXiv:1601.07464* (2016).
- [147] B. Roskovec. “Detailed Investigation of Electron Antineutrino Oscillations in the Daya Bay Experiment”. PhD thesis. Charles University, 2016.
- [148] F. An et al. “Precision measurement of reactor antineutrino oscillation at kilometer-scale baselines by Daya Bay”. In: *Physical Review Letters* 130.16 (2023), p. 161802.
- [149] “Double Chooz θ_{13} measurement via total neutron capture detection”. In: *Nature Physics* 16.5 (2020), pp. 558–564.
- [150] C. Shin et al. “Observation of reactor antineutrino disappearance using delayed neutron capture on hydrogen at RENO”. In: *Journal of High Energy Physics* 4 (2020), pp. 1–27.
- [151] D. Adey et al. “Measurement of the electron antineutrino oscillation with 1958 days of operation at Daya Bay”. In: *Physical Review Letters* 121.24 (2018), p. 241805.
- [152] F. An et al. “New measurement of θ_{13} via neutron capture on hydrogen at Daya Bay”. In: *Physical Review D* 93.7 (2016), p. 072011.
- [153] M. A. Acero, A. A. Aguilar-Arevalo, and D. J. Polo-Toledo. “Neutrino Oscillation Analysis of 217 Live Days of Daya Bay and 500 Live Days of RENO”. In: *Advances in High Energy Physics* 2020 (2020), pp. 1–13.