

Durham E-Theses

Probing leptonic flavour with future long-baseline neutrino oscillation experiments

BALLETT, PETER ALEXANDER

How to cite:

BALLETT, PETER ALEXANDER (2013) *Probing leptonic flavour with future long-baseline neutrino oscillation experiments*, Durham theses, Durham University. Available at Durham E-Theses Online:
<http://etheses.dur.ac.uk/7745/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Academic Support Office, Durham University, University Office, Old Elvet, Durham DH1 3HP
e-mail: e-theses.admin@dur.ac.uk Tel: +44 0191 334 6107
<http://etheses.dur.ac.uk>

Probing leptonic flavour with future long-baseline neutrino oscillation experiments

Peter A. Ballett

A thesis presented for the degree of
Doctor of Philosophy



Institute for Particle Physics Phenomenology
Department of Physics
Durham University
United Kingdom

July 2013

Probing leptonic flavour with future long-baseline neutrino oscillation experiments

Peter A. Ballett

Submitted for the degree of Doctor of Philosophy
July 2013

Abstract

Over the last 50 years, the study of the properties of neutrinos has unveiled a number of surprising facts that necessitate physics beyond the standard model. We now know that neutrinos are not only massive, but that there is a non-trivial alignment between the mass and flavour bases, inducing flavour changing transitions known as *neutrino oscillations*. Understanding the neutrino sector is a crucial first step in our attempts to extend our current theories of fundamental physics, and studies of neutrino oscillation provide us with a unique tool to probe these elusive particles. In this thesis, we assess the potential of the next generation of neutrino oscillation experiments to probe physical effects both within and beyond the current neutrino flavour paradigm: resolving existing unknowns, and constraining the correlations induced by theories of leptonic flavour.

Contents

Declaration	vii
1 Neutrinos in and beyond the Standard Model	2
1.1 Neutrinos and the Standard Model	2
1.2 Problems with the SM neutrino	8
1.3 Neutrino oscillation	10
1.3.1 Derivation	10
1.3.2 Matter effects	14
1.3.3 Approximate treatments of oscillation probability	17
1.4 Moving beyond the standard model	23
1.4.1 Dirac neutrinos	24
1.4.2 Majorana masses and see-saw mechanisms	25
1.5 Summary	27
2 Next-generation neutrino oscillation facilities	29
2.1 Conventional beams and superbeams	30
2.2 Neutrino Factories	34
3 Optimization of the Low-Energy Neutrino Factory	40
3.1 Low-energy neutrino factories	40
3.2 Simulation details	43
3.2.1 Technical details	47
3.3 Phenomenology of parameter sensitivity	51
3.3.1 θ_{13}	53
3.3.2 Mass hierarchy	58

3.3.3	CP violation	64
3.4	Optimal choices for the LENF	71
3.5	Summary	72
4	Discrete flavour symmetries and parameter correlations	74
4.1	Models with discrete flavour symmetries	75
4.2	Parameter correlations and sum rules	77
4.3	Hernandez-Smirnov symmetry building	80
4.3.1	Derivation of constraints	81
4.3.2	Impact of generator choice	83
4.3.3	Phenomenologically viable sum rules	85
4.3.4	Validity of linearization	87
4.3.5	Summary	89
4.4	Mass matrix decomposition	90
4.4.1	Construction of the parameterization	90
4.4.2	Computing s , r and a	93
4.4.3	Computing the CP phases	94
4.4.4	Simplified expressions for s , r and a ?	95
4.4.5	Summary	99
5	Constraining atmospheric sum rules at current and future facilities	100
5.1	Compatibility of sum rules with existing and projected data	101
5.2	Sum rules at next-generation facilities	106
5.2.1	Details of facilities and simulations	107
5.2.2	Impact of τ -contamination	109
5.2.3	Precision for a , r and $\cos \delta$	111
5.2.4	Joint determination of parameters	116
5.2.5	Measuring sum rule violations	117
5.2.6	Constraining λ	120
5.3	Conclusions	122
	Outlook	124
	Bibliography	126

List of Illustrations

1.1	Feynman diagrams for matter interactions	15
1.2	Oscillation probabilities and the effect of matter interactions	18
1.3	Oscillation probabilities for a range of values of δ	21
1.4	Effect of longer baselines on oscillation probability	22
1.5	Tree-level completions of the Weinberg dimension-5 operator	27
2.1	Spectrum of an off-axis neutrino beam	31
2.2	Schematic diagram of a muon storage ring	35
2.3	Feynman diagrams of the production channel for a Neutrino Factory .	36
3.1	Fluxes at a Neutrino Factory	51
3.2	θ_{13} discovery potential at $L = 2350$ km	53
3.3	θ_{13} discovery range as a function of baseline and stored-muon energy.	54
3.4	Hierarchy discovery potential at $L = 2350$ km	58
3.5	Hierarchy determination as a function of L and E_μ	59
3.6	Hierarchy discovery reach as a function of L for the TASD.	63
3.7	CP violating discovery potential at $L = 2350$ km	64
3.8	CP-violation discovery potential for large θ_{13}	65
3.9	CP-violation discovery potential for small θ_{13}	66
3.10	CP-violation discovery fraction as a function of exposure.	70
4.1	Solar angle predictions of the Hernandez-Smirnov approach	86
4.2	Impact of higher-order corrections to sum rules	87
4.3	Behaviour of $ M_{ij} $ as a function of absolute mass scale.	97
5.1	Predictions for $\cos \delta$ for the viable sum rules	101

5.2	$\lambda \approx 1$ sum rule predictions and current phenomenological intervals . .	103
5.3	$\lambda \approx -1/2$ sum rule predictions and current phenomenological intervals	104
5.4	Precision for a , r and $\cos \delta$	112
5.5	Impact of ν_τ on precision to a	115
5.6	Impact of ν_τ on precision to $\cos \delta$	116
5.7	Two-parameter determination of $\cos \delta$ and a	118
5.8	Exclusion regions for $\lambda \in \{1, -1/2\}$ for MIND and LAr	119
5.9	Exclusion regions for $\lambda \in \{1, -1/2\}$ for WBB	120
5.10	Ability to constrain λ	121

Declaration

The work in this thesis is based on research carried out at the Institute for Particle Physics Phenomenology, part of the Department of Physics of Durham University, United Kingdom. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

The work in Chapter 3 is based upon research undertaken by the author in collaboration with Silvia Pascoli. This work culminated in the publication below. The current author was responsible for all of the simulations, providing the majority of the phenomenological analysis, and writing the paper.

‘Understanding the performance of the low energy neutrino factory: the dependence on baseline distance and stored-muon energy’,

P.A. Ballett and S. Pascoli, Phys.Rev. **D86** (2012) 053002 [[arXiv:1201.6299](#)]

The work on the phenomenology of discrete symmetries which forms the basis of Chapters 4 and 5 is based on research undertaken by the author in collaboration with Stephen F. King, Christoph Luhn, Silvia Pascoli and Michael Schmidt. This work is to be published presently. The current author played a central role in the project, helping in the systematic exploration of viable sum rules, performing all simulations, and leading the writing of the paper.

Copyright © 2013 by Peter Ballett.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent, and information derived from it should be acknowledged.”

Chapter 1

Neutrinos in and beyond the Standard Model

In this thesis, we will explore the ways in which the study of neutrinos may further our understanding of the leptonic sector. Neutrinos must play a distinguished role in any modern theory of particle physics. Not only are they believed to be the only fundamental electrically neutral fermions, but their experimental investigation has uncovered a number of surprising facts over the last 50 years, and the method of incorporating these features into a new fundamental theory is not without ambiguities. In this chapter, we will explain the structure of the standard model, and present the role of neutrinos. We will then address the known facts of neutrino physics which necessitate a theory beyond the standard model. After an in depth discussion of the neutrino oscillation phenomenon, we shall end this chapter with a brief review of some minimal extensions of the standard model which may be able to incorporate these facts.

1.1 Neutrinos and the Standard Model

It is believed that all currently observed phenomena can be described by the actions of four forces: electromagnetism, the weak force, the strong force and gravity. Gravity is understood to be relevant only on the largest distances and at the highest energies, playing an integral role in the formation of astrophysical objects and

driving cosmological processes. Although these high energy effects are essential to the universe that we live in, at the scales accessible in the laboratory, including the laboratories of the predictable future, it is only the remaining three forces which act with sufficient strength to influence events. Our understanding of these three forces, those which dominate behaviour at the most accessible distances and energies, is contained in a theory called the Standard Model (SM). The SM adopted its current form around the 1970s, formed out of the union of a number of successful but more restricted models. The field of particle physics, and the path towards the SM, began with the field of nuclear physics. Attempts to understand the nuclei lead to the discovery of the proton and neutron. As novel experimental techniques were invented to probe these particles, namely the observation of both naturally and artificially accelerated particles, their scattering probabilities and their decay products, a plethora of additional particles were discovered. These were classified by their weight: the heaviest particles were termed *baryons*, whilst the lighter particles were termed *mesons* and *leptons*. Understanding the spectrum of these objects was a central goal of the early particle physics community, and great success was found in the organising structures of group theory.

It was noticed that the observed patterns of baryons and mesons were in strong agreement with the structure of the representations of $SU(3)$. These organised the mesons into octets and singlets, whilst the baryons formed octets and decuplets [1,2]. This line of thought resulted in the quark model [3–5] and the formulation of the strong nuclear force. This force provides the strong binding potential which holds nuclei together despite their large concentrations of positive electric charge. The strong force is an $SU(3)$ gauge theory acting on the quarks, which are fundamental fermions assigned to 3-dimensional representations of $SU(3)$ with each component labeled by a *colour* quantum number. A peculiar feature of the strong force is that it becomes stronger at low energies; the coupling constant which describes the strength of this force becomes large, and conventional perturbation expansions are no longer sufficient to compute quantities in this theory. This makes the low-energy spectrum of the strong force appear quite distinct to its high-energy field content. According to a property termed *confinement*, the lowest energy excitations in the

		SU(3)	SU(2) _L	U(1) _Y
Fermions	Q_L	3	2	$\frac{1}{3}$
	L_L	1	2	-1
	u_R	3	1	$\frac{4}{3}$
	d_R	1	1	$-\frac{2}{3}$
	e_R	1	1	-2
Bosons	G	8	1	0
	W	1	3	0
	B	1	1	0
	H	1	2	1

Table 1.1: The particle assignment of the standard model. Each particle is denoted by its irreducible representations of the sub-symmetries of the SM: the strong SU(3), weak SU(2) and hypercharge U(1) symmetries. For non-abelian groups the irreducible representation is denoted by its dimension, for the abelian hypercharge symmetry, the representation is denoted by the charge under that group. The fermions shown make up a single family of particles. There are three known families, these have identical quantum numbers and differ only in particle mass.

theory are composite objects with zero net colour. Modern studies of the strong force have revealed a remarkable agreement with experiment. In particular, great progress has been made in our understanding of the low-energy spectrum: although complicated by the non-perturbative nature of strong dynamics, the meson spectrum has been computed on the lattice (spatially and temporally discretized computer simulations) and shows a strong agreement with the spectrum of masses which have been measured experimentally [6].

The origins of the neutrino and our understanding of the weak force are intimately related. One of the earliest studied weak processes was the β -decay of atomic nuclei. Electrons emitted during these decays follow a continuous distribution of energies, in apparent violation with the law of conservation of energy. To explain this puzzling feature, Pauli proposed the existence of the neutrino [7], allowing the missing energy to be carried away by this hard to detect neutral par-

ticle. Our modern understanding of the weak force can be traced back to Fermi's four-point interaction [8]. In modern language, this interaction is an effective operator coupling fermions together at dimension-6. Studies of muon decay showed this to be a successful theory, but it was known to be non-renormalizable. Constructing a theory of the weak force based on a gauge symmetry was a desirable goal, but these attempts were met with an immediate complication: the gauge bosons corresponding to the new weak force needed to be massive. This not only allows them to escape detection, but also permits the reduction of the theory to the simplified model of Fermi for energies well below the gauge boson mass. However, introducing fundamental mass terms for the gauge bosons breaks gauge symmetry, rendering the physical model inconsistent. The solution to this problem was found by considering methods of dynamically generating gauge boson masses. This culminated in the electroweak symmetry breaking (EWSB) mechanism of Brout, Englert, Guralnik, Hagen, Higgs and Kibble [9–12]. This approach found a dynamical mechanism to break an $SU(2)_L \otimes U(1)_Y$ symmetry down to a single $U(1)_{EM}$ subgroup which is responsible for electromagnetism. The EWSB mechanism relied crucially on the introduction of a new particle: a fundamental scalar referred to as the Higgs boson, H . This process of discovery culminated in the Glashow–Weinberg–Salam model of unified electroweak interactions [13–15]. This theory describes fermions coupled to a renormalizable theory of gauge interactions under the symmetry group $SU(2)_L \otimes U(1)_Y$. This is a chiral theory: the $SU(2)_L$ subgroup couples only to the left-chiral components of the fermion fields: left-chiral leptons are unified in 2-dimensional representations of this subgroup *e.g.* $L_e = (\nu_e, e_L)^T$, whilst the right-chiral fields are assigned to singlets. The fields are charged under the remaining $U(1)_Y$ subgroup, called hypercharge, in such a way that the familiar electric charges emerge after EWSB. This model has four fundamental massless gauge boson degrees of freedom, but the process of EWSB leaves only a single gauge boson without mass (the photon), and the remaining gauge degrees of freedom combine to form three massive particles: two are identical up to a change of electric charge, W^+ and W^- , whilst the third is electrically neutral, Z^0 .

The SM finally takes shape as the trivial unification of the gauge theories of the

strong and electroweak interactions. The only connection between the two sectors is provided by the quarks, which are charged under both strong and electroweak groups: the left-chiral quark fields are given by an SU(2)-doublet $Q_L = (u_L, d_L)^T$ as with the leptons. The particle content of the SM is shown in Tab. 1.1 and can be described in terms of families. Each family contains two quarks, one with an electric charge of $+\frac{2}{3}e$ and the other with a charge of $-\frac{1}{3}e$, and two leptons, one with a charge of $-e$ and one uncharged, the neutrino. We can express the SM Lagrangian by a decomposition into a number of terms

$$\mathcal{L} = \mathcal{L}_G + \mathcal{L}_F + \mathcal{L}_H + \mathcal{L}_Y. \quad (1.1.1)$$

The first term contains the kinetic and self interaction terms for the gauge fields,

$$\mathcal{L}_G = -\frac{1}{4}G_{\mu\nu}^a G_{\mu\nu}^a - \frac{1}{4}W_{\mu\nu}^i W_{\mu\nu}^i - \frac{1}{4}B_{\mu\nu} B_{\mu\nu},$$

where the field strength tensors are given by $G_{\mu\nu}^a \equiv \partial_\mu G_\nu^a - \partial_\nu G_\mu^a - g_1 f^{abc} G_\mu^b G_\nu^c$, $W_{\mu\nu}^i \equiv \partial_\mu W_\nu^i - \partial_\nu W_\mu^i - g_2 \epsilon_{ijk} W_\mu^j W_\nu^k$, and $B_{\mu\nu} \equiv \partial_\mu B_\nu - \partial_\nu B_\mu$, where a is an SU(3) colour index and i is an SU(2) weak index. The quantities g_1 , g_2 and g_3 are the coupling constants of the three gauge groups, whilst f_{abc} and ϵ_{ijk} are the structure constants for the groups SU(3) and SU(2), respectively. The next term in the Lagrangian contains the kinetic terms for the fermions and their gauge interactions,

$$\mathcal{L}_F = \overline{Q}_L i \not{D} Q_L + \overline{u}_R i \not{D} u_R + \overline{d}_R i \not{D} d_R + \overline{L}_L i \not{D} L_L + \overline{e}_R i \not{D} e_R,$$

where the covariant derivative is given by $D_\mu = \partial_\mu + ig_1 G_\mu^a \frac{\lambda^a}{2} + ig_2 W_\mu^i \frac{\tau^i}{2} + ig_3 B_\mu \frac{Y}{2}$, denoting the generators of SU(3) by λ^a , the generators of SU(2) by τ^i , and Y the operator which returns the hypercharge of each field. If the field is a singlet of any of the gauge groups, the associated generator should be taken as 0. The next term $\mathcal{L}_H$ describes the Higgs kinetic and self-interaction terms

$$\mathcal{L}_H = (D^\mu H)^\dagger (D_\mu H) + \lambda H^\dagger H (v^2 - H^\dagger H),$$

where $D_\mu H = \left(\partial_\mu + ig_2 W_\mu^i \frac{\tau^i}{2} + ig_3 B_\mu \frac{Y}{2} \right) H$. The final part of the SM Lagrangian $\mathcal{L}_Y$ contains the interactions between the Higgs field and the fermions of the theory.

These are the only terms in the theory which couple fields in different families

$$\mathcal{L}_Y = -(\overline{Q_L}\tilde{H})Y_u u_R - (\overline{Q_L}H)Y_d d_R - (\overline{L_L}H)Y_e e_R + \text{h.c.}, \quad (1.1.2)$$

where Y_u , Y_d and Y_e are matrices in family space. The fact that the matrices Y_u and Y_d are not diagonal in the flavour basis leads to non-trivial flavour effects in the quark sector and generates CP violating phenomena arising from the complex phase of the Cabbibo-Kobayashi-Maskawa (CKM) matrix [16, 17].

The production of gauge boson masses by the EWSB mechanism can be understood by studying $\mathcal{L}_H$. The self-interaction terms of the Higgs contribute to the energy density, and must be minimized by the vacuum field configuration. This imposes a minimization condition

$$\langle H \rangle^\dagger \langle H \rangle (v^2 - 2\langle H \rangle^\dagger \langle H \rangle) = 0,$$

which leads to the Higgs VEV being given by $\langle H \rangle = (0, v/\sqrt{2})^T$, up to gauge equivalent choices. At leading order, the gauge couplings in the Higgs kinetic term will then produce mass terms for particular linear combinations of the W and B bosons,

$$\begin{aligned} |D_\mu \langle H \rangle|^2 &= \left| \frac{1}{2} \begin{pmatrix} g_3 B_\mu + g_2 W_\mu^3 & g_2(W_\mu^1 - iW_\mu^2) \\ g_2(W_\mu^1 + iW_\mu^2) & g_3 B_\mu - g_2 W_\mu^3 \end{pmatrix} \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix} \right|^2, \\ &= \frac{g_2^2 v^2}{4} W_\mu^+ W_\mu^- + \frac{v^2}{8} (g_2^2 + g_3^2) Z_\mu Z^\mu, \end{aligned}$$

where the new massive fields are defined in terms of the old by $W_\mu^\pm = (W_\mu^1 \mp iW_\mu^2)/\sqrt{2}$ and $Z_\mu = (g_3 B_\mu - g_2 W_\mu^3)/\sqrt{g_2^2 + g_3^2}$. The other orthogonal combination of fields, $A_\mu = (g_2 B_\mu + g_3 W_\mu^3)/\sqrt{g_2^2 + g_3^2}$, remains massless and describes the photon of quantum electrodynamics. The fermionic mass terms arising through the process of EWSB can be derived from the Yukawa terms in Eq. 1.1.2 by replacing the Higgs with its VEV

$$\mathcal{L}_Y = -\frac{v}{\sqrt{2}} Y_u^{\alpha\beta} \overline{u_L^\alpha} u_R^\beta - \frac{v}{\sqrt{2}} Y_d^{\alpha\beta} \overline{d_L^\alpha} d_R^\beta - \frac{v}{\sqrt{2}} Y_e^{\alpha\beta} \overline{e_L^\alpha} e_R^\beta + \text{h.c.},$$

where family indices have been reinstated as α and β . These terms correspond to mass terms for the fermions, given by the mass matrices

$$(M_u)^{\alpha\beta} = \frac{v}{\sqrt{2}} Y_u^{\alpha\beta}, \quad (M_d)^{\alpha\beta} = \frac{v}{\sqrt{2}} Y_d^{\alpha\beta} \quad \text{and} \quad (M_e)^{\alpha\beta} = \frac{v}{\sqrt{2}} Y_e^{\alpha\beta}.$$

As the charged lepton and neutrino in each family are two components of the same $SU(2)$ -doublet, each neutrino is naturally associated with the charged lepton in its family, and for this reason, they are referred to as the electron neutrino ν_e , muon neutrino ν_μ and tau neutrino ν_τ . Although the number of families may be larger than three, an important result was found at the LEP collider [18] which measured the number of neutrinos produced in the invisible decays of Z^0 -bosons (thereby requiring a mass of below around 45 GeV). These results have shown that there are three such neutrinos, and any fourth family would require either a neutrino with a mass a factor of 10^{12} higher than the other neutrinos, or a serious restructuring of the SM particle assignments.

The standard model has proved to be a very successful theory, being put through an impressive series of tests since its codification in the 1970s. Some of the most remarkable tests have been performed at collider facilities, confirming the SM predictions of the τ lepton [19], the W [20,21] and Z bosons [22,23], the top quark [24,25], and recently, a scalar resonance consistent with the Higgs boson [26,27]. Despite proving so valuable in the description of so many experimental searches, the SM is now known to be incomplete, and the most striking evidence of this fact has emerged from the neutrino sector. In the remainder of this chapter, we will discuss how this new evidence was discovered, how it can be explained by neutrino oscillation and finally, how we can extend the standard model to incorporate these phenomena in a consistent theoretical framework.

1.2 Problems with the SM neutrino

The direct experimental investigation of neutrinos began in 1956, when Cowan and Reines accomplished the first successful detection of electron antineutrinos emitted by a nuclear reactor [28]. This provided confirmation of the neutrino hypothesis, and the measured cross-sections were in excellent agreement with theoretical predictions [29]. The first suggestion that the neutrino sector may be other than that described in the standard model was given by the measurement of the solar neutrino deficit. In the 1960s the flux of electron neutrinos from the sun was measured and found to be

much smaller than expected from the known nuclear processes in the solar cycle [30]. Although other solutions were offered for this deficit, the problem could be explained by introducing a transition from electron neutrinos ν_e to some other flavour ν_α . These flavour-changed neutrinos would escape detection, as the experiments relied on the inverse β -decay reaction [31], which is only possible for ν_e . Neutral meson mixing was a known phenomenon, and it was noted by Pontecorvo that neutral fermions could oscillate too [32–35]. However, neutral fermion oscillation can only occur for massive particles, when the flavour states are misaligned from their mass states [32–36], and therefore, this mechanism could only be exploited if the SM was modified to incorporate non-zero neutrino masses.

Further anomalies came from experiments which studied the flux of neutrinos from the upper atmosphere. These are emitted as a consequence of cosmic ray interactions, when highly energetic protons collide with the particles making up the atmosphere and produce showers of secondary particles, including a number of neutrinos. The ratio of muon neutrinos to electron neutrinos was found to be significantly smaller than expected. The atmospheric neutrino deficit added to the significance of the known problems in the neutrino sector and, as with the solar deficit, could be explained by a flavour changing transition. Although both the solar and atmospheric deficits could be resolved by flavour transitions amongst the neutrinos, the very different environments under which these processes occur did not make the solution seem particularly clear. As we shall see in Section 1.3.2, the high-density solar medium significantly affects neutrino propagation, and the oscillation solution to the solar deficit only appeared likely once these effects were understood. The detailed knowledge of matter’s influence on flavour transitions made the presence of neutrino mass terms, and the subsequent flavour oscillations, far more appealing. However, the oscillatory explanation of the solar and atmospheric deficits was only widely accepted after Super-Kamiokande’s measurement of the atmospheric neutrino flux as a function of zenith angle (the angle from the normal to the Earth’s surface) [37]. This showed that the oscillatory hypothesis could consistently explain the whole distribution of observed events, and not just provide an explanation of a reduced count rate.

1.3 Neutrino oscillation

Neutrino mass terms have a kinematic effect on many processes which are known to occur in the SM. For example, the maximum of the β -decay electron energy spectrum is reduced by the rest energy of the associated neutrino [8, 38], and similar effects must arise in the weak decays of mesons and muons (see *e.g.* [39–43]). However, in experiments of this type, the effects of neutrino mass have been to date unobservable, suggesting that neutrino masses must be very small. Current bounds from cosmology place an upper bound on the sum of neutrino masses $\sum_i m_i \leq 0.23$ eV [44], whilst direct detection efforts suggest that $m(\nu_e) \lesssim 2.1$ eV [45, 46]. Nonetheless, to explain the anomalies discussed in Section 1.2, we are motivated to introduce small mass terms into an extension of the SM. In this section, we shall see how the introduction of neutrino masses leads to the neutrino oscillation phenomenon, which has allowed the inference of their existence despite their small absolute size. Understanding this mechanism, and probing the neutrino sector through its observation, has been a key step in the establishment of neutrino physics as a field and will be a major theme in the remainder of this thesis.

1.3.1 Derivation

There are a number of ways to describe the neutrino oscillation phenomenon mathematically. The simplest approach, and the first to be discovered, treated the neutrinos as plane waves [31, 47–49]. Plane wave treatments can produce the correct answer; however, these often rely on *ad hoc* kinematic assumptions which are known to be ultimately incorrect. A consistent quantum mechanical treatment which captures the spirit of the plane wave approach, whilst relying on no false assumptions, can be formulated using wave packets [50–52]. This treatment incorporates production and detection uncertainties, which are known to be essential to the oscillatory effect. Finally, a treatment in terms of quantum field theoretic language is possible, where the neutrino mass states exist virtually, connecting a vertex at production to another at detection [53–55]. This is known to produce essentially the same probabilities as the wave packet approach, and the parameters of these two formalisms

can be connected by explicit calculation [56]. In this section, we shall discuss a derivation based on wave packets, with an explicit form of the wave packet which will assist calculation. Our assumptions about the shape of the wave packets can be relaxed, and more general treatments are possible; however, the oscillation probabilities remain unaltered under these changes.

The oscillation phenomenon relies upon the existence of neutrino masses, and of a non-trivial relationship between the flavour and mass bases. We describe the relationship between bases with a matrix U , which is called the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [32, 33, 36] and, as a mapping between 3-dimensional bases, is assumed to be a 3×3 unitary matrix. If we denote the neutrino mass eigenstates by $|\nu_i\rangle$ for $i \in \{1, 2, 3\}$ and the neutrino flavour states $|\nu_\alpha\rangle$ for $\alpha \in \{e, \mu, \tau\}$, the PMNS matrix describes the linear transformation between these sets of states

$$|\nu_\alpha\rangle = U_{\alpha i} |\nu_i\rangle.$$

The PMNS matrix can be decomposed into a number of mixing angles and complex phases. A general $n \times n$ unitary matrix can be described by $\frac{n(n-1)}{2}$ mixing angles and $\frac{n(n+1)}{2}$ complex phases. However, not all of these complex phases are physical parameters: as the phase of a charged lepton is unphysical, we can absorb some of the phases of the PMNS matrix by rephasing these fields. Given this rephasing freedom, the PMNS matrix can be generally described by the factorisation below [6]

$$\begin{aligned} U_{\text{PMNS}} &= U_{23} U_{13} U_{12} P, \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha_1} & 0 & 0 \\ 0 & e^{i\alpha_2} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \\ &= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} P. \end{aligned}$$

The matrix P contains two complex phases referred to as the Majorana phases. If neutrinos are Dirac fermions (described by 2 Weyl spinor degrees of freedom) these phases will be unphysical and can be removed by rephasing, as with the charged

leptons. If instead neutrinos are Majorana fermions (described by a single Weyl spinor degree of freedom), these phases cannot be removed, and will be physical observables. Regardless of their physical status, we will show at the end of this section that these phases play no role in the phenomenon of neutrino oscillation.

To derive the oscillation probability consistently, we must take into account the nature of the superposition of mass eigenstates which is produced and subsequently detected. The oscillation effect occurs because of the coherent interference between eigenstates of different masses, and as these masses have different energies for a given momentum, this requires some uncertainty in the initial neutrino momentum and energy. In this derivation, we consider a neutrino produced with a definite flavour ν_α and described by a superposition of wave packets with characteristic momentum spread σ ,

$$|\nu_\alpha; p_0, \sigma\rangle = \frac{1}{(\sqrt{\pi}\sigma)^{\frac{3}{2}}} \int d^3p \, e^{-\frac{(p-p_0)^2}{2\sigma^2}} U_{\alpha i} |\nu_i; p\rangle,$$

where $|\nu_i; p\rangle$ is a momentum eigenstate with energy $E_i(p) = \sqrt{p^2 + m_i^2}$, normalised such that $\langle \nu_i; p | \nu_j; q \rangle = \delta_{ij} \delta^3(p - q)$. This state is allowed to propagate over a spacetime interval $x_\mu = (t, x)^T$,

$$|\nu_\alpha(x_\mu); p_0, \sigma\rangle = \frac{1}{(\sqrt{\pi}\sigma)^{\frac{3}{2}}} \int d^3p \, e^{-\frac{(p-p_0)^2}{2\sigma^2}} U_{\alpha i} e^{-iE_i(p)t + ip \cdot x} |\nu_i; p\rangle.$$

To find the probability for transition to another flavour, we introduce the final state neutrino ν_β . This must also be described by a superposition of wavepackets, as the detection process is necessarily spatially localised. For simplicity, we assume an identical form of the wavepacket to that at production

$$|\nu_\beta; p_0, \sigma\rangle = \frac{1}{(\sqrt{\pi}\sigma)^{\frac{3}{2}}} \int d^3q \, e^{-\frac{(q-p_0)^2}{2\sigma^2}} U_{\beta j} |\nu_j; q\rangle,$$

The amplitude for flavour transition $\nu_\alpha \rightarrow \nu_\beta$ is the projection of the propagated initial state onto the final state,

$$\begin{aligned} \mathcal{A}_{\alpha\beta} &= \langle \nu_\beta; p_0, \sigma | \nu_\alpha(x_\mu); p_0, \sigma \rangle, \\ &= \frac{U_{\beta j}^* U_{\alpha i}}{(\sqrt{\pi}\sigma)^3} \int d^3q \int d^3p \, e^{-\frac{(q-p_0)^2}{2\sigma^2} - \frac{(p-p_0)^2}{2\sigma^2}} e^{-iE_i(p)t + ip \cdot x} \langle \nu_j; q | \nu_i; p \rangle, \\ &= \frac{U_{\beta i}^* U_{\alpha i}}{(\sqrt{\pi}\sigma)^3} \int d^3p \, e^{-\frac{(p-p_0)^2}{\sigma^2}} e^{-iE_i(p)t + ip \cdot x}. \end{aligned}$$

If we assume that the uncertainty is small enough for the integrand to only have support for small value of $\delta p \equiv p - p_0$, we can Taylor expand the energy to first order

$$\mathcal{A}_{\alpha\beta} = \frac{U_{\beta i}^* U_{\alpha i}}{(\sqrt{\pi}\sigma)^3} e^{-iE_i(p_0)t + ip_0 \cdot x} \int d^3p e^{-\left(\frac{\delta p}{\sigma}\right)^2} e^{-i(v_i t - x) \cdot \delta p},$$

where $v_i \equiv \left. \frac{\partial E_i}{\partial p} \right|_{p_0}$ is the group velocity of the ν_i wave packet. The Gaußian integral can now be performed exactly,

$$\mathcal{A}_{\alpha\beta} = U_{\beta i}^* U_{\alpha i} e^{-iE_i(p_0)t + ip_0 \cdot x} e^{-\frac{\sigma^2}{4}(v_i t - x)^2}. \quad (1.3.3)$$

If we further assume that each propagating neutrino is highly relativistic, such that $m_i/|p_0| = \mathcal{O}(\varepsilon)$ with $\varepsilon \ll 1$, we can approximate E_i and v_i by

$$E_i(p_0) = |p_0| + \frac{m_i^2}{2|p_0|} + \mathcal{O}(\varepsilon^3) \quad \text{and} \quad v_i = 1 + \mathcal{O}(\varepsilon^2).$$

With these assumptions, the terminal exponents in Eq. 1.3.3 which contain dependence on the space time coordinates, $x - v_i t$, can be seen to heavily suppress the oscillation probability at any space-time point which does not closely approximate $x = t$. For this reason, we assume that our oscillation process is being observed at $x = t = L$. Finally, we find that the probability for $\nu_\alpha \rightarrow \nu_\beta$ transition is given by

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &\equiv |\mathcal{A}_{\alpha\beta}|^2, \\ &= \left| \sum_{j=1}^3 U_{\beta j}^* U_{\alpha j} e^{-i\Delta m_{j1}^2 \frac{L}{2E}} \right|^2, \end{aligned} \quad (1.3.4)$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$ and $E = |p_0|$.

We are now in a position to verify the claim that the Majorana phases, whether they are physical or not, will not influence the oscillation probability. As Eq. 1.3.4 depends upon the PMNS matrix only through the products $U_{\beta j}^* U_{\alpha j}$, the probabilities are invariant under the mapping $U \rightarrow UP$ for any diagonal unitary matrix P . This shows that there is no effect of the Majorana phases in the neutrino oscillation formulae. Therefore, there are four parameters from the PMNS matrix, $\{\theta_{12}, \theta_{13}, \theta_{23}\}$, and two mass-squared splittings, $\{\Delta m_{21}^2, \Delta m_{31}^2\}$, which influence the oscillation probability. Thanks to studies of neutrino oscillation, a large amount of

	best-fit	3σ interval
θ_{12} ($^\circ$)	33.36	(31.09, 35.89)
θ_{13} ($^\circ$)	8.66	(7.19, 9.96)
θ_{23} ($^\circ$)	$40.0 \oplus 50.4$	(35.8, 54.8)
Δm_{21}^2 (eV^2)	7.50×10^{-5}	$(7.00, 8.09) \times 10^{-5}$
$ \Delta m_{31}^2 $ (eV^2)	2.47×10^{-3}	$(2.28, 2.70) \times 10^{-3}$
δ ($^\circ$)	300	(0, 360)

Table 1.2: The current values of the oscillation parameters found by a global fit of the oscillation data [57]. This analysis includes the results of solar, atmospheric and reactor experiments.

information is now known regarding the values of these parameters. In Tab. 1.2, the current best-fit values are shown along with the 3σ allowed intervals. From this table, we can see that there are two significant unknowns among the oscillation parameters: the sign of Δm_{31}^2 (referred to as the mass hierarchy) and the value of δ . The search for these two parameters will be a crucial first step for any future experiment.

1.3.2 Matter effects

The neutrino oscillation phenomenon has so far only been considered in vacuo. However, the fact that most neutrinos will be propagating amidst a background of particles will lead to a non-negligible effect on the oscillation probability [58–60]. As neutrinos traverse any region of space occupied by matter, they will interact with a sea of electrons, protons and neutrons. Crucially, the type of interactions that they are subjected to will vary by neutrino flavour: although all neutrinos will interact via neutral current with these matter particles, the ν_e will have an additional charged current interaction with the e^- in the matter. The diagrams for these processes are shown in Fig. 1.1, and lead to an additional contribution to the neutrino Hamiltonian, referred to as the matter term, which is highly flavour asymmetric. This new term leads to an alteration of the probability for a neutrino to undergo




Figure 1.1: Two tree-level diagrams relevant to the propagation of neutrinos through matter. The charged-current process only exists for e -like neutrinos and leads to a flavour dependent effect.

flavour transition, particularly in scenarios of high matter density, as in the solar medium, or when propagation occurs over long distances inside the Earth, where the net effect of matter becomes significant. Exploiting the new parameter dependences introduced by the matter term is an essential element of modern neutrino oscillation experiments which hope to probe the underlying oscillation parameters. To illustrate the impact of the matter term, we will now derive the corrected oscillation probabilities for a simplified system of two neutrinos in the presence of matter.

The effect of the interactions in Fig. 1.1 is to modify the Hamiltonian by introducing an additional term dependent on a parameter $A = \sqrt{2}G_F n_e$ where n_e is the density of electrons in matter,

$$H_f = U \begin{pmatrix} 0 & 0 \\ 0 & \Delta m_{21}^2 \end{pmatrix} U^\dagger + \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix}.$$

This matrix can be re-diagonalized by a 2×2 unitary matrix $\tilde{U}$ to give

$$H_m = \tilde{U}^\dagger \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \tilde{U}.$$

The parameters λ_i and the mixing angle $\tilde{\theta}$ which are necessary to describe the 2×2 unitary matrix $\tilde{U}$ can be shown to be

$$\begin{aligned} \tan(2\tilde{\theta}) &= \frac{\sin(2\theta)A}{\Delta m_{21}^2 - \cos(2\theta)A}, \\ \lambda_2 &= \Delta m_{21}^2 + \sin^2 \theta A + \mathcal{O}(A^2), \\ \lambda_1 &= \cos^2 \theta A + \mathcal{O}(A^2). \end{aligned}$$

The derivation of the oscillation probability in matter now follows an identical procedure to the vacuum case, but instead of decomposing the flavour states into the mass basis, we use the new basis which diagonalise the full matter-corrected Hamiltonian. Aside from a reinterpretation of labels, the vacuum derivation is unmodified, and we can construct the matter-corrected oscillation probability by the substitution of $m_i \rightarrow \lambda_i$ and $\theta \rightarrow \tilde{\theta}$ in the two neutrino oscillation probability found by restricting Eq. 1.3.4 to two flavours

$$\begin{aligned} P(\nu_e \rightarrow \nu_\mu) &\approx \sin^2(2\tilde{\theta}) \sin^2\left(\frac{(\lambda_2 - \lambda_1)L}{4E}\right), \\ &\approx \sin^2(2\tilde{\theta}) \sin^2\left(\frac{\Delta m_{21}^2 L}{4E}\right) + \mathcal{O}(A^2). \end{aligned}$$

The presence of matter is seen here to lead to the observation of a different mixing angle than in vacuo. This has two particularly interesting consequences. First of all, although in vacuum the two neutrino oscillation probabilities are identical for neutrinos and antineutrinos, the presence of matter introduces an additional CP violating effect: the background's CP asymmetry leads to a separation of the two oscillation probabilities. It can be shown that the matter potential differs by a sign for neutrinos and antineutrinos, and we can map between the two by simply effecting the substitution $A \rightarrow -A$. In the two-neutrino case this leads to a separation in CP conjugate oscillation probabilities

$$P(\nu_e \rightarrow \nu_\mu) = \frac{[\Delta m_{21}^2 - \cos(2\theta)A]^2}{[\Delta m_{21}^2]^2 + A^2 - 2\Delta m_{21}^2 \cos(2\theta)A} \sin^2\left(\frac{\Delta m_{21}^2 L}{4E}\right), \quad (1.3.5)$$

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_\mu) = \frac{[\Delta m_{21}^2 + \cos(2\theta)A]^2}{[\Delta m_{21}^2]^2 + A^2 + 2\Delta m_{21}^2 \cos(2\theta)A} \sin^2\left(\frac{\Delta m_{21}^2 L}{4E}\right). \quad (1.3.6)$$

In general, therefore, we expect different behaviour for neutrinos and antineutrinos when observed propagating through matter, and from these two formulae, we see that one oscillation will be suppressed (its mixing angle decreasing), whilst the other enhanced (its mixing angle increasing). The relevant factor controlling which channel is suppressed and which is enhanced is the sign of Δm_{21}^2 . This discussion readily generalises to the full three neutrino case, and by observing this CP-asymmetric effect, the sign of the mass-squared splittings can be inferred. Indeed, this is the general strategy to uncover the currently unknown sign of Δm_{31}^2 by measurements of

the oscillation probability through the Earth [61, 62]. In Fig. 1.2, we plot the effect of the matter term on $P(\nu_\mu \rightarrow \nu_e)$ for a range of electron number densities n_e over a distance of $L = 2000$ km. The oscillation probability has been computed numerically and includes the full three-flavour effects. We have assumed $\Delta m_{31}^2 > 0$, which enhances the probability of oscillation for neutrinos (solid lines) whilst suppressing the probability for antineutrinos (dashed lines). The black central line shows the probability in the absence of matter.

It can be seen from Eqs. 1.3.5 and 1.3.6 that, for either neutrinos or antineutrinos, the mixing angle will achieve a maximum for $|A| = |\Delta m_{21}^2 \cos(2\theta)|$. Neutrinos produced in the high-density core of the Sun will pass through this resonant region as they propagate to its surface. The large mixing angle induces an efficient conversion of ν_e to ν_μ , significantly reducing the observed flux of ν_e at solar neutrino detection experiments. This effect, called the Mikheev–Smirnov–Wolfenstein (MSW) effect [58, 59], is now understood to be the solution to the solar neutrino problem, and lead to the determination of the correct ordering of the m_1 and m_2 eigenvalues, $\Delta m_{21}^2 > 0$ [63, 64].

1.3.3 Approximate treatments of oscillation probability

Although relatively compact to write down in Eq. 1.3.4, the full three-flavour oscillation probability is quite complicated to deal with without approximations. In this section we shall look at a few common approximations which will help us to understand the phenomenological properties of this formula.

Single dominant frequency approximation

Under the assumption of three active neutrino flavours, the formula given in Eq. 1.3.4 has two relevant phases of oscillation. If we assume that we are interested in distances L and energies E such that only one of these phases is large, we can approximate the full three flavour mixing with a two flavour model. For example, if we assume that $\Delta m_{31}^2 L/2E \approx 1$ whilst we suppress terms of order $\Delta m_{21}^2 L/2E \equiv \varepsilon$,




Figure 1.2: The oscillation probability for $\nu_\mu \rightarrow \nu_e$ transition at $L = 2000$ km. Solid (dashed) lines show the probabilities for neutrinos (antineutrinos). We have assumed normal hierarchy and $\delta = 0$, but the value of the matter density has been varied according to the legend (in g/cm^3). The black dashed line shows the vacuum probability for both neutrinos and antineutrinos. A Gaussian smear has been applied to regulate the fast low-energy oscillations.

Eq. 1.3.4 simplifies

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_\mu) &= \left| |U_{\mu 1}|^2 + |U_{\mu 2}|^2 + |U_{\mu 3}|^2 e^{i\Delta m_{31}^2 \frac{L}{2E}} + \mathcal{O}(\varepsilon) \right|^2, \\
 &= \left| 1 + |U_{\mu 3}|^2 \left(e^{i\Delta m_{31}^2 \frac{L}{2E}} - 1 \right) \right|^2 + \mathcal{O}(\varepsilon), \\
 &= 1 - 4|U_{\mu 3}|^2 (1 - |U_{\mu 3}|^2) \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right) + \mathcal{O}(\varepsilon). \tag{1.3.7}
 \end{aligned}$$

whilst for the flavour changing transition $\nu_\mu \rightarrow \nu_\tau$, we find

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_\tau) &= \left| U_{\tau 1}^* U_{\mu 1} + U_{\tau 2}^* U_{\mu 2} + U_{\tau 3}^* U_{\mu 3} e^{i\Delta m_{31}^2 \frac{L}{2E}} + \mathcal{O}(\varepsilon) \right|^2, \\
 &= \left| U_{\tau 3}^* U_{\mu 3} \left(e^{i\Delta m_{31}^2 \frac{L}{2E}} - 1 \right) \right|^2 + \mathcal{O}(\varepsilon), \\
 &= 4|U_{\tau 3}|^2 |U_{\mu 3}|^2 \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right) + \mathcal{O}(\varepsilon). \tag{1.3.8}
 \end{aligned}$$

These probabilities neglect matter effects, but the same formulae arise from the full matter-effect corrected probabilities for sufficiently small baseline distances $AL = \mathcal{O}(\varepsilon)$.

In a model with only two neutrino flavours, for example where only the ν_μ and ν_τ mix, the PMNS matrix satisfies $|U_{\mu 3}|^2 + |U_{\tau 3}|^2 = 1$, and we can express the

oscillation parameters by a new mixing angle $U_{\tau 3} = \sin \theta$,

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) + \mathcal{O}(\varepsilon).$$

Simplified formulae of this type are commonly used as approximate expressions for the full probability by the current generation of oscillation experiments [65,66]. This is generally a good approximation, as the full three-flavour effects are largely beyond their sensitivities. For $\nu_\mu \rightarrow \nu_\mu$ oscillations, the effective mixing angle satisfies $\sin(2\theta) = \sin(2\theta_{23})$ with around 2% accuracy for the known value of θ_{13} .

Three neutrino oscillation

In a true two flavour scheme, the PMNS matrix has no complex phases which can influence oscillation physics. Therefore, an understanding of the full three-flavour oscillation probability is necessary if we hope to probe CP violating effects. There are a number of three-flavour approximation schemes in the literature. In this section, we will present two of these approximations for $\nu_e \rightarrow \nu_\mu$ oscillations. Both methods rely on perturbation theory by expanding the oscillation probability in a set of small parameters. However, they differ in their exact implementation and the assumed set of small parameters.

First, we consider the approximation of Ref. [67]. This is an expansion in θ_{13} , Δ_{12}/Δ_{23} , Δ_{12}/A and $\Delta_{12}L$, where $\Delta_{ij} = \Delta m_{ij}^2/2E$. To derive the approximate formula, the full three-flavour Hamiltonian is decomposed into a zeroth order part,

$$H_0 = U \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \Delta_{31} \end{pmatrix} U^\dagger + \begin{pmatrix} A & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

and a perturbation,

$$H_1 = U \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{21} & 0 \\ 0 & 0 & 0 \end{pmatrix} U^\dagger.$$

Proceeding by perturbation theory, this leads to an expression for the probability of

$\nu_e \rightarrow \nu_\mu$ transition

$$\begin{aligned}
P(\nu_e \rightarrow \nu_\mu) = & \sin^2 \theta_{23} \sin^2(2\theta_{13}) \left(\frac{\Delta_{13}}{\Delta_{13} - A} \right)^2 \sin^2 \left(\frac{\Delta_{13}L}{2} - \frac{AL}{2} \right) \\
& + J \frac{\Delta_{12}}{A} \frac{\Delta_{13}}{\Delta_{13} - A} \sin \left(\frac{AL}{2} \right) \sin \left(\frac{\Delta_{13}L}{2} - \frac{AL}{2} \right) \cos \left(\delta - \frac{\Delta_{13}L}{2} \right) \\
& + \cos^2 \theta_{23} \sin^2(2\theta_{12}) \left(\frac{\Delta_{12}}{A} \right)^2 \sin^2 \left(\frac{AL}{2} \right)
\end{aligned} \tag{1.3.9}$$

where $J = \cos \theta_{13} \sin(2\theta_{12}) \sin(2\theta_{13}) \sin(2\theta_{23})$. This formula is very useful for understanding the relative roles of different parameters in the oscillation probability. The terms in this sum are ordered by their relative size, and the first term dominates the oscillation probability for most experimental set-ups. However, it is the second term which contains the dependence on δ — currently the most poorly constrained parameter. Novel experiments hoping to measure δ must, therefore, be sensitive to the sub-dominant effects introduced by this term. In Fig. 1.3, we show the impact that variations in δ have on the oscillation probabilities for two different baseline distances $L = 2300$ km and $L = 1160$ km. It is clear that these are subdominant effects, as the overall shape of the probability remains similar, and that the largest impact is found at lower energies, which can be understood from the formula by the factor of $\Delta_{12} \propto 1/E$ in the CP-dependent term, leading to a low-energy enhancement in sensitivity to δ . We also see a generic increase in the impact of δ at longer baselines, which can be seen as a consequence of the factor $\sin(AL/2)$. These two observations impact the measurement of the mass hierarchy by observation of the matter induced CP asymmetry. Indeed distinguishing between the effects of A and the effect of δ may be difficult. In the left-most panel of Fig. 1.4, we show the neutrino oscillation probabilities with $L = 1160$ km for the two different choices of mass hierarchy. The width of the bands indicates the impact of δ , and the presence of closely lying (or even overlapping) regions of these bands indicates the impossibility of differentiating between the hierarchy solutions at those points, if δ remains unconstrained. In the right-most panel of Fig. 1.4, we see the same quantities but for a longer baseline $L = 2300$ km. The longer baseline makes the matter effects more significant, and the two hierarchies now present a clear separation for all energies above 3 GeV. For this reason, longer baselines are desirable for the unambiguous




Figure 1.3: The effect of δ on the oscillation probability for $\nu_\mu \rightarrow \nu_e$ for neutrinos (left column) and antineutrinos (right column). We show results for two baseline distances: $L = 1160$ km (top row) and $L = 2300$ km (bottom row). A Gaussian smear has been applied to regulate the fast low-energy oscillations.

measurement of the mass hierarchy.

Eq. 1.3.9 was derived before the measurement of θ_{13} , and as it involves an expansion in this parameter, an improved formula may be able to be constructed by retaining the full θ_{13} dependence. Since the measurement of θ_{13} , alternate approximate formulae have been proposed which take this approach [68,69]. We will now focus on the method presented in Ref. [69], and show how this compares to the previous approximation and the full oscillation probability. The method in Ref. [69] relies on the fact that we mentioned in Section 1.3.2: the full matter-corrected oscillation probability can be expressed in the same form as the expressions in vacuum, provided a suitable substitution of the mixing angles and masses is made. As the vacuum oscillation probabilities are relatively simple, all that is required to compute the effect of matter is to compute these new angles and masses efficiently. This can




Figure 1.4: The variation in probability due to the value of δ for $\nu_\mu \rightarrow \nu_e$ for neutrinos. We show the behaviour for a shorter baseline on the left with $L = 1160$ km and a longer baseline on the right with $L = 2300$ km. A Gaussian smear has been applied to regulate the fast low-energy oscillations.

be done using the Jacobi method, a convergent recursive process where a matrix is diagonalised in the limit by diagonalising a series of submatrices. The derivation of a simple formula for the oscillation probability by this method is only made tractable by the observation that after two iterations, the Jacobi method finds matrices which are diagonal up to terms of order $0.2\Delta m_{21}^2$. This process produces the following matter-effect corrected mixing angles for the neutrino oscillation probability

$$\tan(2\theta'_{12}) = \frac{\Delta m_{21}^2 \sin(2\theta_{12})}{\Delta m_{31}^2 \cos(2\theta_{12}) - a \cos^2 \theta_{13}},$$

$$\tan(2\theta'_{13}) = \frac{(\Delta m_{31}^2 - \Delta m_{21}^2 \sin^2 \theta_{12}) \sin(2\theta_{13})}{(\Delta m_{31}^2 - \Delta m_{21}^2 \sin^2 \theta_{12}) \cos(2\theta_{13}) - a},$$

where $a = 2EA$. In this scheme θ_{23} and δ retain their vacuum values. The mass-squared splittings must be replaced with $\Delta m_{ij}^2 \rightarrow \lambda_i - \lambda_j$, where the new parameters are given by

$$\lambda_1 = \lambda'_-, \quad \lambda_2 = \lambda''_- \quad \text{and} \quad \lambda_3 = \lambda''_+,$$

assuming normal hierarchy, whilst for inverted hierarchy the λ_2 and λ_3 should be exchanged. The remaining unspecified parameters are given by

$$\lambda'_\pm = \frac{\Delta m_{21}^2 + a \cos^2 \theta_{13}}{2} \pm \frac{1}{2} \sqrt{(\Delta m_{21}^2 - a \cos^2 \theta_{13})^2 + 4a \cos^2 \theta_{13} \sin^2 \theta_{12} \Delta m_{21}^2},$$

$$\lambda''_\pm = \frac{1}{2} [\lambda'_+ + (\Delta m_{31}^2 + a \sin^2 \theta_{13})] \pm \frac{1}{2} \sqrt{[\lambda'_+ - (\Delta m_{31}^2 + a \sin^2 \theta_{13})]^2 + 4a^2 \sin^2 \theta_{12} \cos^2 \theta_{13} \sin^2 \theta_{13}},$$

It is debatable how much easier to manipulate these expressions are; however, these formulae have been shown to provide a very close approximation to the true oscillation probabilities over a wide range of experimental set-ups [69].

1.4 Moving beyond the standard model

To construct an extended standard model, one capable of accommodating the known facts of neutrino physics, a mechanism for the generation of neutrino masses must be found. In this section, we shall consider some of more popular proposals for the generation of neutrino masses which assume only minimal extensions of the SM.

Fundamental particle mass terms break the gauge group of the SM, and it is only through the process of electroweak symmetry breaking that mass terms may arise. At low energies, mass terms are generated effectively from the Yukawa interactions between fermions and the Higgs field: it is the expansion around the Higgs vacuum expectation value that produces the fermionic mass term. Focusing on a single lepton doublet $L_e = (\nu_e, e_L)^T$ for convenience, the relevant terms in the SM Lagrangian before electroweak symmetry breaking are given by

$$\mathcal{L} \supset -y_e \overline{L}_L H e_R + \text{h.c.} \quad (1.4.10)$$

At low energies, the dynamic Higgs field decouples from the theory, and is replaced by its vacuum expectation value $\langle H \rangle = (0, v/\sqrt{2})^T$. This leads to the generation of an effective mass for the electron $m = y_e v/\sqrt{2}$, and is referred to as a Dirac mass term. This mechanism allows masses to be generated by spontaneous symmetry breaking for all fermions which have left-chiral and right-chiral degrees of freedom, but without e_R , the Yukawa interaction in Eq. 1.4.10 could not be constructed. Due to the absence of right-handed neutrino degrees of freedom in the SM, neutrino masses cannot be written down in the same way. In fact, giving the neutrinos of the SM a mass term will necessitate the inclusion of new degrees of freedom.

There are a number of ways in which to construct a neutrino Standard Model (ν SM) capable of describing massive neutrinos. In the following sections, I will summarise a number of the most popular proposals.

1.4.1 Dirac neutrinos

The simplest way to construct a ν SM is to add additional degrees of freedom which play the role of the right-handed neutrinos. These fermions are gauge singlets and, accordingly, and have no tree-level interactions with other particles. With an additional right-handed neutrino ν_R , we can use the standard mechanism for mass generation in the SM for the neutrino fields

$$\mathcal{L} \supset -y_e \overline{L}_L H e_R - y_\nu \overline{L}_L \tilde{H} \nu_R + \text{h.c.},$$

producing a neutrino mass given by $m_\nu = y_\nu v / \sqrt{2}$.

Although this approach is perfectly consistent, it raises two issues. Firstly, we note that this solution introduces additional Yukawa couplings which must satisfy a quite staggering hierarchy of scales. In this model, all of the effective fermion masses are given by $m_f = y_f v / \sqrt{2}$. To satisfy the known experimental constraints on the particle masses, we must have a hierarchy of magnitudes in the leptonic Yukawa couplings given by their mass ratios: for example,

$$\frac{m_\tau}{m_\nu} = \frac{y_\tau}{y_\nu} \approx 10^9.$$

This large difference between the fundamental inputs of the theory would have no underlying explanation, and many physicists believe that such hierarchies are unnatural.

Another issue with the addition of gauge-singlet right-handed neutrinos to the SM is that there is no symmetry to forbid an additional term for these particles

$$\mathcal{L} \supset -\frac{M}{2} \overline{\nu}_R \nu_R^c + \text{h.c.},$$

which generates a mass M for the Weyl fermion degree of freedom ν_R , and is referred to as a Majorana mass term. Ideas of naturalness dictate that these terms should be included if they are allowed by the symmetries of the theory. Majorana mass terms are not invariant under $U(1)$ symmetries, and if this operator exists, the accidental lepton number symmetry of the SM will be broken, introducing new lepton number violating decay modes. This is not in itself a problem; however, we have been led to a situation where naturalness appears to disfavour purely Dirac mass terms. This

conclusion could be avoided entirely by imposing a global symmetry on the theory, for example $B - L$, which forbids a Majorana mass term, keeping neutrinos purely Dirac at the cost of an additional symmetry principle.

1.4.2 Majorana masses and see-saw mechanisms

We have seen that by considering natural implementations of Dirac neutrino masses, we have been prompted to consider Majorana masses for a set of right-chiral neutrinos. In fact, we can consider Majorana mass terms for the left-chiral neutrinos directly. This is complicated by the symmetries of the left-chiral neutrinos, which forbid Majorana mass terms at the renormalizable level. However, we can describe the mass terms at an effective level, and enumerate the possible simple extensions of the SM which could generate these terms. Working in an effective theory framework, there is a single dimension-5 operator which can be constructed out of the fields of the SM called Weinberg's dimension-5 operator [70, 71]

$$\mathcal{L}_5 \supset \frac{1}{2} c_{\alpha\beta} (\overline{L}_\alpha^c \tilde{H}^*) (\tilde{H}^\dagger L_\beta) + \text{h.c.},$$

where the new coefficients are expected to be small $c_{\alpha\beta} \sim 1/\Lambda$ for a new-physics scale Λ . This term generates a Majorana mass matrix M_ν for the left-handed neutrinos of the SM once the Higgs attains its VEV $\langle \tilde{H} \rangle = (v/\sqrt{2}, 0)^T$,

$$(M_\nu)_{\alpha\beta} = c_{\alpha\beta} \frac{v^2}{2}.$$

A nice feature of this approach is that the neutrino mass matrix comes with a natural suppression by the scale of new physics $M_\nu \sim 1/\Lambda$, which means that the smallness of neutrino masses can be naturally explained by the existence of high-energy new physics.

To understand the possible origins of neutrino mass, we are interested in how this operator may be UV-completed. To do this, we seek a non-effective model which can be reduced to the effective case via the process of integrating out some set of massive particles. There are many ways to do this, but if we focus on the case of tree-level completions which introduce a single new heavy field, we can show systematically that there are only three possibilities: known as the type I, type II

and type III see-saw mechanisms. In all of these cases, the scale of new physics Λ will be the mass of the new field.

The Weinberg operator contains four field operators, two lepton doublets and two Higgs doublets. To generate this at tree-level, we must divide these fields between two new gauge invariant and renormalizable interaction terms involving some new field φ , which will be ultimately integrated out. There are two possible ways to make this decomposition. The first option introduces a single new term into our Lagrangian, which we show without specifying the Lorentz or SU(2) structure

$$\mathcal{L}_1 \supset y(LH)\varphi + \text{h.c.}$$

The second option is to introduce two terms,

$$\mathcal{L}_2 \supset y(HH)\varphi + y'(LL)\varphi + \text{h.c.}$$

In $\mathcal{L}_1$, the Lorentz structure of our theory requires our new field to be a fermion, whilst the SU(2) structure affords two possibilities depending on how the L and H doublets are contracted, as $\mathbf{2} \otimes \mathbf{2} = \mathbf{1} \oplus \mathbf{3}$. If the lepton and Higgs doublets form a singlet, then φ must also be a singlet

$$(\bar{L}\tilde{H}^*) \sim \mathbf{1} \quad \text{and} \quad \varphi \sim \mathbf{1},$$

whereas if the lepton–Higgs term is contracted as an SU(2) triplet, φ must also be a triplet

$$(\bar{L}\tilde{H}^*) \sim \mathbf{3} \quad \text{and} \quad \varphi \sim \mathbf{3}.$$

Therefore there are two ways that we can complete the Weinberg operator at tree-level with the addition of an interaction of the type $(LH)\varphi$: either φ is an SU(2) singlet, or an SU(2) triplet. These two scenarios are commonly referred to as the Type I [72] and the Type III [73] see-saw mechanisms, respectively.

If instead we chose the decomposition of the Weinberg operator into two terms, as in $\mathcal{L}_2$, the Lorentz structure of the terms (LL) and (HH) tell us that φ must be a scalar field. Again, an analysis of the SU(2) structure allows us to deduce that φ must be an SU(2) singlet or a triplet. However, to ensure Lorentz invariance, the lepton doublets must be given by $(\bar{L}^c L)$, and this requires φ to have a hypercharge




Figure 1.5: The three tree-level completions of the Weinberg dimension-5 operator which can produce effective Majorana neutrino masses. The new fields are an $SU(2)$ singlet fermion ν_R , an $SU(2)$ triplet scalar Δ and an $SU(2)$ triplet fermion Σ . If these new fields have large masses, the smallness of the neutrino masses arises naturally.

of 2. Accordingly, the Higgs term must be $(\tilde{H}^\dagger \tilde{H}^*)\varphi$, which vanishes when H takes its VEV and produces no neutrino mass term [74]. We are left with a single choice for completing the Weinberg operator with a Lagrangian of the type $\mathcal{L}_2$: φ must be a scalar $SU(2)$ triplet. This is commonly referred to as the Type II seesaw mechanism [75, 76].

1.5 Summary

Neutrinos play a distinguished role in particle physics phenomenology. We know that they must be massive, that they have a non-trivial mixing between flavour and mass states, and that these properties require an extension of the standard model. We have shown how their new properties lead to the phenomenon of neutrino oscillation, which has been conclusively established by a number of dedicated oscillation experiments. The role of matter effects is crucial for the understanding of the oscillation probabilities, and provides an elegant explanation of the solar neutrino deficit, via the resonant flavour transitions of the MSW effect.

Although neutrino oscillation is now a widely accepted phenomenon, extending the SM to include neutrino masses is not straightforward. There are a number of ways in which to implement simple mechanisms which are responsible for the generation of neutrino mass; however, without further experimental data it is impossible

to say which is to be preferred. To further our knowledge of the neutrino sector, it is clearly necessary to perform detailed phenomenological studies. In this thesis, we will present work attempting to understand to what extent the properties of neutrinos can be explored with the next-generation neutrino oscillation experiments.

Chapter 2

Next-generation neutrino oscillation facilities

Determining the oscillation probabilities of neutrino flavour states, and therefore the misalignment between mass and flavour bases, is a very challenging task. As neutrinos only interact via the weak force, their interaction cross-section with matter is very small. This means that to build a neutrino detector with a significant number of interactions, we must look to very large facilities and very intense sources of neutrinos. Historically, the high fluxes of neutrinos from the sun and cosmic ray interactions in the upper atmosphere were some of the first to be exploited; however, as the field progressed into one of precision measurements, tunable man-made sources became desirable. This led the way to the construction of accelerator based neutrino beams.

The central goal of the next generation of oscillation experiments is to measure both the mass hierarchy, dictated by the sign of Δm_{31}^2 , and the value of the CP violating phase δ . As discussed in Section 1.3.2, a long baseline of the order 1000 km would allow the determination of the mass hierarchy by observation of the matter induced asymmetries, whilst also enhancing the sub-dominant terms in the oscillation probability which depend on δ . In this chapter, we will discuss two of the most promising candidates for a next-generation long-baseline neutrino oscillation experiment: wide-band superbeams and neutrino factories. These are both accelerator based, and offer a well controlled and tunable beam of neutrinos, which is essential

for precision measurements of the neutrino oscillation probabilities.

2.1 Conventional beams and superbeams

Conventional neutrino beams are produced as a product of the decay of mesons [77, 78], and all accelerator neutrino beams to date have relied on this method [79]. This process begins with the acceleration of protons towards a target, the subsequent collisions generating a large number of pions and kaons through strong interaction processes. Magnetic horns then focus the meson beam, selecting π^+ with a small number of K^+ (or π^- and K^- , by a change of horn polarity) and the decay of these particles generates the neutrino beam. The produced π^+ mesons decay predominately into ν_μ and μ^+ , whilst the subsequent decays of μ^+ amongst other sources lead to the production of ν_e and $\bar{\nu}_\mu$. The resultant flux of neutrinos is well approximated as almost purely composed of ν_μ , but with a contamination of ν_e at the 1% level. This flux permits the measurement of a number of oscillation channels. The two most accessible channels originate from the ν_μ particles which dominate the flux. The *disappearance channel* of a conventional beam is the observation of muons in the detector, arising from $\nu_\mu \rightarrow \nu_\mu$ oscillation. As discussed in Section 1.3.3, this channel has strong sensitivity to θ_{23} and $|\Delta m_{31}^2|$. The most important channel for testing the full three-flavour effects is the *appearance channel*, which is driven by $\nu_\mu \rightarrow \nu_e$ oscillations. Observing this channel is the main focus of many current and next-generation facilities, as it allows the mass hierarchy and CP phase to be probed. To measure this channel, the detectors must be sensitive to electrons and be able to consistently distinguish them from muons and Neutral Current (NC) events. NC interactions are flavour blind and provide no information on the oscillation probabilities. The subdominant ν_e component of the flux will constitute an intrinsic background to the appearance channel signal. This contamination presents an inherent limitation on the attainable sensitivity of experiments which rely on this method of neutrino production.

There are two main types of neutrino beam: *on-axis* and *off-axis*. The energy spectrum of neutrinos derived from pion decay is strongly dependent on the angle




Figure 2.1: The left-most plot shows the neutrino energy spectrum from the decay of pions as it is observed at different angles from the meson trajectory. The initial pion spectrum is given by $(E_p - E_\pi)^5$ with the proton energy E_p taken as 30 GeV. This narrowing of neutrino energy is due to the broad maximum neutrino energy produced at a certain angle as a function of pion energy (right). These figures have been reproduced from Ref. [80]

at which it is observed with respect to the trajectory of the parent meson. The most advanced experiments of the current generation of conventional neutrino beams exploit this dependence to produce a narrow-band beam, one which has a significant flux only over a small range of neutrino energies [81–83]. This is done by choosing the beamline to be off-axis, meaning that a neutrino collinear with its parent meson would miss the detector facility. As can be seen in the right-most panel of Fig. 2.1, the energy of a neutrino that is emitted at a small angle from the trajectory of its parent pion is almost independent of the energy of the pion [80, 84]. This kinematic effect leads to the narrowing of the neutrino energy spectrum when observed off-axis, as we can see in the simulated spectrum of the left-most panel in Fig. 2.1. The approximately monochromatic neutrino spectrum found using the off-axis technique provides a simple signal to analyse, whilst also leading to a kinematic suppression of the intrinsic ν_e background. The development of off-axis beams has been an important advance in technique, allowing T2K [81, 82] to measure the appearance

channel for the first time in an accelerator experiment [85], which paved the way towards a statistically significant exclusion of $\theta_{13} = 0$ by Daya Bay and RENO [66, 86]. The off-axis technique will be exploited further by the NO ν A experiment [83], which is expected to provide the most stringent constraints of the current generation of experiments [87]. The off-axis approach does, however, have some drawbacks: moving off-axis reduces the number of observed events, and crucially removes the possibility of measuring the oscillation spectrum over a range of energies. Off-axis experiments therefore behave more as total rate counting exercises, and lose valuable information about the full spectrum, making the problem of parameter degeneracies particularly acute and the full three-flavour effects, harder to conclusively observe.

Instead of an off-axis beam, the next generation of conventional neutrino beams is expected to be on-axis. These experiments are referred to as wide-band beams (WBBs) [88, 89], as they are designed to measure the oscillation probability over a wide range of neutrino energies. This allows for more information to be extracted from the observed event rates, and the impact of parameter degeneracies to be mitigated. Although the on-axis technique introduces a greater number of background events, with improved energy resolution and event reconstruction capabilities these can be addressed by other means. The term *superbeam* is sometimes used for a conventional beam experiment which has an increase in the power of its proton beam to energies of the order 1 MW (and, by consequence, increases in both the number and energies of the subsequent neutrinos) [90]. A superbeam is, therefore, just an extrapolation of the existing technology for neutrino production, resulting in a more intense and energetic beam, and studies have shown that this design could significantly extend our knowledge of the neutrino oscillation parameters [91–96].

There are a number of different designs for the detector of a conventional neutrino beam oscillation experiment. The main requirements are for a very large volume to help increase statistics, and a strong suppression of environmental backgrounds such as leptons from cosmic rays. This background suppression is usually obtained by placing the detectors underground, which provides extra shielding from external particles. One detector design which has proved very successful in neutrino oscillation experiments is the Water Čerenkov (WČ) detector [97], currently being

used in Kamioka as part of the T2K experiment [81,82]. This design measures the Čerenkov radiation emitted by charged leptons as they break the speed of light in water. Studying the topology of the Čerenkov cone allows the particle to be identified, its energy measured and its incident direction deduced. One of the main advantages of WČ detectors is that they are very scalable; a philosophy which is evident in the design of the IceCube experiment, which uses the WČ technique to study neutrinos impinging upon 1 km³ of Antarctic ice [98]. The WČ technology does, however, have certain limitations which limit its appeal for a next generation experiment. One particular problem is the lack of control over neutral current backgrounds involving the production of a π^0 . The resonant production of pions is especially common for low energy neutrino interactions [99,100], and the dominant pion decay $\pi^0 \rightarrow \gamma\gamma$ induces two electromagnetic showers, which can easily mimic the single shower expected of an electron. Therefore, alternative designs are desirable for the next-generation of experiments which can offer a better control over their NC backgrounds. A promising option is to use a detector filled with liquid scintillator (LSc) [101–103]. These detectors look similar to a small WČ detector but instead of water, they are filled with an organic compound known to emit scintillation light in the presence of a charged particle. By recording the energy and the time of arrival of observed photons, and LSc detector allows for both calorimetric and kinematic reconstruction of events. These detectors are known to have a very low threshold energy, ideal for the measurement of solar neutrinos, and good energy resolution properties. Due to the cost of liquid scintillator, LSc detectors can't be made at such large volumes as WČ, but their improved detection capabilities can make up for this shortcoming. Another detector technology which has been shown very promising performance is the liquid Argon time-projection chamber [104–107]. These detectors observe the ionisation track left as a charged particle passes through liquid Argon. The tracks are drifted in an electric field onto an electronic read-out, which provides an interaction topology and energy deposition information. These detectors have excellent particle identification and energy resolution capabilities, and are being considered in extrapolations to sizes of the order 100 kton [106,108–111]. This combination of high-statistics, low-threshold, and excellent event reconstruc-

tion make liquid Argon detectors (LArS) very effective detectors for long-baseline oscillation experiments [107].

After years of experimental work on conventional neutrino beams, the superbeam technology is very well understood and considerable expertise is to be found in the community. This makes the WBB a feasible medium-term project, and a number of facilities have been recently proposed. There are two leading designs for a next-generation WBB which will form the basis for the simulations in this thesis: a CERN to Pyhäsalmi superbeam (C2P) considered by the LAGUNA–LBNO group, and the LBNE experiment based in Fermilab. C2P has been developed, and recently recommended, by the LAGUNA–LBNO design study [112–115]. With a large value of θ_{13} , such a facility has been shown [90, 116–120] to provide a competitive physics reach compared to other designs, and for a significant fraction of parameter space, may be sensitive to δ and CP-violating effects [119, 120]. The baseline configuration of this facility consists of a new accelerator based at CERN and a large LAr detector [107] based in the Pyhäsalmi mine in Finland, producing a baseline distance of 2300 km. The detector is planned with an initial fiducial mass of 20 kt, but with the intention of an ultimate upgrade to 70 kt [115]. The LBNE proposal is for a long-baseline neutrino beam to be produced at Fermilab [121]. The current configuration is for a 10 kt surface LAr detector at the Homestake mine, producing a baseline distance of 1300 km [122]. This is a scaled down version of the initial proposal for an underground 33 kt LAr detector [123]; however, the current configuration is envisaged as the possible first step in a staged approach towards the original design goal [124]. For this reason, we will generally consider LAr detectors in the mass range 33 kt to 70 kt.

2.2 Neutrino Factories

Neutrinos, as the lightest SM fermions, can be produced as decay products in a large number of interactions. As mentioned in Section 2.1, the neutrinos which compose a conventional beam are produced by the decay of mesons generated in the collisions of high-energy protons with a target. Although this has proven to be a straightforward




Figure 2.2: The muon storage ring design of the International Design Study for the Neutrino Factory. The muons decay in the long straight sections to produce the neutrino flux. Figure taken from Ref. [125].

way to produce a controlled beam of neutrinos, there are some inherent uncertainties in this process. In particular, the systematic uncertainties in the beam's flavour composition lead to major limitations to the ultimate attainable sensitivity of any experiment using this technology. At a Neutrino Factory (NF) [126–128], to try and mitigate the uncertainties of the superbeams, neutrinos are produced by an alternative mechanism. Instead of studying neutrinos emitted during meson decay, the conventional production of mesons is used as a first stage to derive a secondary beam of muons (or anti-muons, depending on the choice of focusing). These leptons are then accelerated to a common energy E_μ (originally around 25 GeV, in latter designs 10 GeV) and are injected into a muon storage ring, an example of which is shown in Fig. 2.2. This acceleration stage consists of two parallel straight sections of beam pipe, around 600 m for $E_\mu = 25$ GeV and 360 m for $E_\mu = 10$ GeV, connected by semicircular end sections [125, 129]. Bunches of muons are allowed to circulate inside the ring, whilst maintaining their energy profile through the application of periodic acceleration and focusing. When passing along the straight sections of the decay ring, a significant number of high-gamma muons will decay, producing a well collimated beam of neutrinos, whose spectral content is known with high precision. The decay of μ^- (μ^+) in the storage ring produces an initial beam with two neutrino components: ν_μ and $\bar{\nu}_e$ ($\bar{\nu}_\mu$ and ν_e). This decay is well described by a tree-level charged-current process, shown in Fig. 2.3. Due to the simplicity of this process, the flux comes with very few theoretical uncertainties: higher-order corrections are known to introduce deviations to the tree-level flux of less than $\mathcal{O}(0.1\%)$ [130].

The detector for a NF is designed to be primarily sensitive to muons and anti-




Figure 2.3: The tree-level production channel for the beam of a neutrino factory. The leftmost (rightmost) panel shows the case of μ^- (μ^+) in the storage ring.

muons. For both μ^- and μ^+ in the storage ring, the observed particles can be divided into two categories, *right-sign* muons and *wrong-sign* muons, based on their agreement (or disagreement) with the charge of the muons in the storage ring. For instance with μ^- decays, the right-sign muon events derive from the observation of ν_μ coming from the disappearance channel, whilst the wrong-sign muon events are due to $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ oscillations. By determining the likelihood of these two classes of particles, the appearance channel oscillation probability can be reconstructed. A clean way to differentiate between particles with different charges is to use a magnetized detector. Detectors of this type are able to separate positively and negatively charged particles by the reconstruction of their trajectories, and most NF proposals employ this technique. The requirement of magnetization makes the optimal detector technology for a NF different to that of a WBB. The most popular design is the MIND [131,132], which has been adopted as the detector of choice of the International Design Study for the Neutrino Factory (IDS-NF) [129]. This is based on the successful MINOS detector [133], comprising of sheets of iron alternating with planes of scintillator, but scaled to a much larger mass. One of the advantages of the MIND is that it is relatively simple to magnetize at large volumes: the large quantity of iron in the detector helps to control the magnetic fields. Indeed, a 100 kton MIND has been proposed, supporting a toroidal magnetic field of around 1.7 T [134]. This detector has been shown to achieve an efficiency of greater than 50% for all energies above 2–4 GeV and has a threshold of around 1 GeV [129,134].

Although the MIND has been shown to be a strong candidate design for the NF detector technology, its detection efficiencies decrease sharply inside the low-energy region. This motivates the consideration of a number of alternative detector de-

signs which offer improved low-energy capabilities. These often use designs based on the same technologies as those discussed for the WBB, but are installed inside a magnetized cavern to permit charge identification. Magnetizing a large volume is very challenging and expensive, but could be done using large external superconducting magnets similar to those employed by the Compact Muon Solenoid (CMS) detector at the LHC [135]. Although providing an ideal combination of low-energy physics and charge identification, this does limit the feasible size of these necessarily large detectors. One detector of this type is the Totally-Active Scintillator Detector (TASD) [128], first proposed for a NF in Refs. [136] and [137]. The TASD concept has been successfully implemented in the MINER ν A experiment [138] and a larger scale device has been selected for the upcoming NO ν A [83, 139] design. The TASD for the NF is composed of a number of modules of plastic scintillator and measures the energy deposition and event timing in each module by collecting light in a dedicated photomultiplier tube. This technique is a direct extrapolation of the MINER ν A design [138], which itself builds upon the design of the DØ pre-shower detectors [140]. One possibility that has been suggested for cost-effective magnetization is to construct solenoids out of superconducting transmission lines as developed for the Very Large Hadron Collider [141]. The TASD detector has an improved low-energy efficiency compared to the MIND with a low-energy threshold of around 0.5 GeV and an efficiency of 94% above 1 GeV, providing good coverage of the second oscillation maximum over a wide range of low-energy configurations [136]. Another proposal which is known to have an excellent low-energy physics reach is the LAr [104–107]. Just as with the WBB, this technology allows for exceptional particle identification properties and detailed kinematic information which can aid the suppression of backgrounds. These detectors have a low threshold energy, similar to the TASD, and good low-energy efficiencies. These detectors can also be considered with very large volumes, of the order 50 to 100 kton, but as with the TASD, the magnetization of these volumes would require a magnetized cavern, which is technically challenging. The possibility of using a non-magnetized detector has also been studied [111]. In this case, although the event separation between wrong- and right-sign muons is impossible at an individual event level, the use of statistical

techniques allows the two channels to be partially distinguished and good sensitivity can be obtained thanks to the high event numbers which are associated with very large non-magnetized detectors. Additional studies are required in order to fully understand the capability of statistical separation of right- and wrong-sign muons and the impact of backgrounds.

Before the discovery of the large value of θ_{13} , the baseline configuration of the NF [129] was a High-Energy Neutrino Factory (HENF), using muons with an energy of 25 GeV and two MINDs. The first detector had a mass of 100 kton and was placed at a distance of 4000 km, and the second detector was 50 kton and placed at the *magic* baseline of 7500 km [142, 143]. The combination of two distinct baseline distances was designed to help resolve the problem of parameter degeneracies by providing complementary information of the oscillation probability at different points in parameter space. The special choice of 7500 km, was to exploit a particular cancellation in the oscillation probabilities that would make the determination of θ_{13} easier, even if it took very small values. This baseline distance is referred to as *magic* as it satisfies the condition

$$\sin\left(\frac{AL}{2}\right) = 0,$$

which, for average matter densities of 4 g/cm³ to 5 g/cm³, occurs at between 7000 km to 8500 km. If we look at the approximate oscillation probability of Ref. [67], given in Eq. 1.3.9. We see that at the magic baseline both the second and third terms vanish, leaving only the atmospheric term. This term does not depend upon δ , and therefore, the measurement of θ_{13} will be easier, as changes in the value of θ_{13} can no longer be accommodated for by shifts in δ [142]. The precise measurement of θ_{13} at the magic baseline helps to reduce the uncertainties on measurements made at the shorter baseline of 4000 km, which is sensitive to CP violation through the parameter δ . This dual baseline set-up would have been invaluable for small values of θ_{13} , and has been shown to have an excellent physics reach to θ_{13} , CP-violation and the mass hierarchy [129, 144, 145].

A modification of the standard design named the Low-Energy Neutrino Factory (LENF) was proposed as a more conservative option for large values of θ_{13} [136, 137, 146]. The first proposal for the LENS used a single baseline of 1300 km, correspond-

ing to the Fermilab to DUSEL distance, and consequently, a lower muon energy at around 4.5 GeV [136,137,146]. By using a detector with good energy resolution and a low threshold energy, the LENF can access the rich oscillatory pattern found at lower energies. As we have seen in Chapter 1, this low energy part of the spectrum is particularly sensitive to δ , and by focusing on this region, the LENF is expected to achieve significant improvements in sensitivity to CP violating effects. Central to the LENF concept is the role of the detector technology, where excellent low-energy physics performance is necessary. In the original proposal, the detector of choice was a TASD magnetized by means of a large magnetic cavern [136]. Other possible detectors include a magnetized liquid-argon time projection chamber, which would be ideal due to its large size and excellent detector performance [104–107]. Studies of the LENF using a MIND have also shown promising performance [145] and could serve as an intermediate step in the development of a future higher-energy facility. However, as the low-energy efficiency of the MIND detector does not allow the low-energy part of the neutrino spectrum to be exploited fully, the performance typically indicates a generic preference for high energies.

To maximize the physics potential of a low-energy facility we must ensure that the lowest parts of the neutrino energy spectrum are well measured. One possibility to increase our sensitivity is to consider detectors which are known to have strong low-energy properties. In the following chapter, we will present the results of an analysis of the LENF configured with detector options which are known to offer strong low-energy physics abilities.

Chapter 3

Optimization of the Low-Energy Neutrino Factory

The Neutrino Factory (NF) has, since its first proposal [126–128], shown exceptional ability to constrain the parameters which control the neutrino oscillation probability. In this chapter, we present an analysis of a modification on the standard NF design, the Low-Energy Neutrino Factory (LENF), which has been shown to be able to offer leading discovery reach in the case of a large value of θ_{13} . In particular, we focus on how two of the fundamental design parameters, the baseline distance L and the stored-muon energy E_μ , affect the ability of the LENF to establish two of its central goals: determining the sign of Δm_{31}^2 and discovering leptonic CP violation.

The chapter is structured as follows. After briefly reviewing the LENF in Section 3.1, we present the technical details of our simulations in Section 3.2. In Section 3.3, we present a phenomenological analysis of how the parameter sensitivity of the LENF changes as a function of L and E_μ . Finally, in Section 3.4, we discuss the optimization of such a facility in light of our simulations.

3.1 Low-energy neutrino factories

As touched upon in Chapter 2, the choice of detection technology affords an additional degree of freedom within the NF design. As the LENF is designed to focus on the lowest energy parts of the oscillation spectrum, where the effect of the CP phase

is largest, it is essential to choose a detector technology with strong low-energy capabilities. In general, there are a number of desirable characteristics for the detector of a LENF:

- Large volume: the number of events observed in a long-baseline neutrino experiment is directly proportional to the number of target molecules in the detector, and therefore the detector volume. To mitigate the small cross sections for neutrino–nucleon interaction, we require large detectors of the order 10–100 kton.
- Magnetization: the signal of oscillation in a NF is the appearance of wrong-sign muons. To measure this effect, the detector must have a means of distinguishing between positive and negatively charged leptons. Two means of affecting this separation have been proposed: magnetization, which allows for a event by event distinction to be made based on the dynamics of the observed lepton, and a statistical approach. The statistical approach introduces significantly higher backgrounds associated with the misidentification of lepton charges and is, therefore, not the optimal method.
- Low threshold energy: as mentioned previously, measuring the low energy spectrum is key to maximizing sensitivity to CP violating effects. To ensure that the proposed experiment has access to this region of parameter space, we require a detector which has a low threshold energy for detection. We aim to have a good sensitivity to neutrino events coming from the second oscillation maximum, where CP violating effects are most pronounced. The location of this region depends on the baseline distance, but for baselines between 1000 km and 4000 km, we require good sensitivity in the interval 1–4 GeV.
- Good energy resolution: as the neutrino energy decreases, the frequency of oscillation increases. To extract the maximum information possible from this region, we require strong energy resolution capabilities, allowing these oscillations to be observed.
- Good background suppression: although the NF beam has relatively low back-

grounds, the ultimate degree of background rejection depends upon the detector technology. The dominant backgrounds to a NF are due to charge misidentification, neutral current events which may produce signals which mimic a muon, and events which arise from τ -decay inside the detector. The event reconstruction and tracking abilities of each detector determines the ability to control these backgrounds.

With these factors in mind, there are a number of potential candidate technologies for the detector at a LENF. The MIND provides a well-developed and understood proposal and has received much interest from the community; however, its low-energy efficiencies are known to be poor in the region around the second maximum. To make the most of the LENF concept, we should instead focus on those detectors which are particularly sensitive to low-energy neutrinos. Despite the challenges that are posed by the magnetization of T ASD and LAr based detectors, their low-energy behaviour suggests that they are the candidates which are best suited to the requirements of the LENF. Therefore, to assess the optimal performance of the LENF, it is essential to understand the potential of a facility which employs these technologies, and the possible benefits to be found by fully exploiting the low-energy signal.

In light of the measurement of large θ_{13} , which has led to the LENF becoming the preferred experimental option, it is also crucial to study how we can optimize the LENF design. In the rest of this chapter, we perform a “green-field” study of the LENF, configured with T ASD and LAr detectors, in order to understand the dependence of the sensitivity to the oscillation parameters on the stored-muon energy E_μ and the choice of baseline distance L . In particular, we have conducted a finely-grained scan over E_μ and L and have identified performance indicators for two primary questions: the discovery of CP-violation and of the mass hierarchy. The analysis that we present has, for the most part, been performed without fixing the value of θ_{13} . For the measurements of CP-violation and mass hierarchy this allows us to present results for the full range of possible values of θ_{13} . For completeness, we have also computed the discovery potential for measurements of θ_{13} itself, as the discussion of this measurement allows the relevant physical effects governing

the performance and sensitivities of the neutrino factory to be explored and better understood.

3.2 Simulation details

We have performed simulations of the physics performance of the LENF, using T ASD and LAr detectors, over a range of experimental configurations. As our indicators of performance, we have computed the sensitivity of the LENF towards two of the most important potential discoveries for the future generation of neutrino oscillation experiments: the discovery of CP-violation arising from the Pontecorvo–Maki–Nakagawa–Sakata [36, 147, 148] matrix and the determination of the neutrino mass hierarchy. We have studied the potential for discovery of each of these fundamental phenomena over a range of stored-muon energies given by $4 \leq E_\mu \leq 25$ GeV and baseline distances by $1000 \leq L \leq 4000$ km. This range connects the regions of parameter space traditionally associated with the LENF design [136, 137, 146] to those of the conventional HENF set-up [129]. This allows us to understand these two designs not as distinct experiments but as part of a continuum.

It has been suggested that 1.4×10^{21} muon decays per year is an attainable goal for the lower-energy accelerator facility [149] and this estimate has been incorporated into previous studies of the performance of the LENF [146, 150]. In our simulations, so as to aid comparison with studies of the standard neutrino factory [144, 145], we have assumed 1.0×10^{21} useful muon decays per year per polarity and a run-time of 10 years which is divided evenly between the two polarities (10^{22} useful muons in total). This value is in accordance with the estimates for the conventional neutrino factory, assuming 10^7 operational seconds per year [129]. We have additionally performed simulations assuming the optimized value of 1.4×10^{21} muon decays per year per polarity, the data indicate a predictable uniform increase in performance and will not be presented separately. A similar survey of the performance of the LENF has been presented in Ref. [151]. In which the total number of useful muon decays was taken as 2.5×10^{22} . This leads to a further increase in statistics and consequently in the discovery potential of the LENF. To understand the performance of the LENF

in the context of similar experimental configurations it is important to ensure a fair comparison is made. For this reason, in the analysis that follows we have assumed the conventional 10^{22} total useful muon decays.

The detectors in our simulations measure muons arising through both the disappearance ($\nu_\mu \rightarrow \nu_\mu$ or $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$) and the appearance channels ($\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ or $\nu_e \rightarrow \nu_\mu$). We have chosen not to include the platinum channel (the observation of ν_e and $\bar{\nu}_e$) as previous work has shown that it offers only marginal improvement of the sensitivity to CP violation and the mass hierarchy [146]. This is due to the additional difficulties associated with the measurement of electrons: these particles are affected significantly by scattering events in the detector, and the reconstruction of their energy and charge by the observation of their tracks is subjected to large uncertainties. The production of $\tau^+(\tau^-)$ in the detector by the charged-current interactions of incident $\bar{\nu}_\tau(\nu_\tau)$ leads to the problem of *tau contamination* [152, 153]. τ -leptons have a lifetime at rest of 2.9×10^{-13} s and decay inside the detector into muons with a branching ratio of around 17% [154]. This effect leads to an increased number of both wrong- and right-sign muons. Due to the form of the oscillation probabilities, the number of additional muons in the right-sign channel (e.g. from $\nu_\mu \rightarrow \nu_\tau$ for μ^- in the storage ring) are considerably larger than in the wrong-sign channel (e.g. from $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ for μ^- in the storage ring). Subsequently, the effect of tau-contamination is most pronounced for measurements which rely on an accurate determination of the disappearance channels, for example in investigations of θ_{23} -maximality [152]. This contamination can lead to serious systematic uncertainties if unaccounted for. However, it has been shown that correctly incorporating this additional source of muons into the analysis of the golden channel, via migration matrices, can resolve the systematic deviations [155]. This has been confirmed in a recent study on the performance of the standard NF [145] where the change in sensitivity produced by correct incorporation of the contamination channel was found to be small. The effect of tau-contamination is also expected to be smaller for lower-energy facilities as the number of neutrinos with energies above the tau production threshold will be reduced. For these reasons, we have omitted the contamination channel in our study.

Our simulations have been performed numerically using the GLoBES package [156, 157] which incorporates the Preliminary Reference Earth Model [158, 159] for the computation of the matter density along the baseline. In this package, the oscillation probabilities are computed via numerical diagonalization of the full Hamiltonian assuming three neutrino flavors. However, it is convenient during the analysis of our results to use an approximate expression, and we will briefly recall the discussion in Chapter 1 to fix our notation. In Ref. [67], the $\nu_e \rightarrow \nu_\mu$ oscillation probability (often referred to as the *golden-channel*) is expressed as a first-order expansion in θ_{13} , $\alpha \equiv \Delta m_{21}^2 / \Delta m_{31}^2$, $\Delta m_{21}^2 L / E$ and $\Delta m_{21}^2 / EA$

$$\begin{aligned}
P_{e\mu} = & \sin^2 2\theta_{13} \sin^2 \theta_{23} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} - \frac{AL}{2} \right) \\
& + \alpha \sin 2\theta_{13} \sin 2\theta_{12} \sin 2\theta_{23} \frac{\Delta m_{31}^2}{2EA} \sin \left(\frac{AL}{2} \right) \sin \left(\frac{\Delta m_{31}^2 L}{4E} - \frac{AL}{2} \right) \cos \left(\frac{\Delta m_{31}^2 L}{4E} + \delta \right) \\
& + \alpha^2 \cos^2 \theta_{23} \sin^2 2\theta_{12} \left(\frac{\Delta m_{31}^2}{2EA} \right)^2 \sin^2 \left(\frac{AL}{2} \right). \tag{3.2.1}
\end{aligned}$$

The first summand in this expression is referred to as the *atmospheric term* and depends quadratically on θ_{13} . The *CP term* is second and introduces dependence on δ . The remaining part is called the *solar term*, which for small values of θ_{13} can dominate the oscillation probability. As the solar term is independent of θ_{13} , δ and $\text{sign}(\Delta m_{31}^2)$, this can lead to a significant loss of sensitivity in the measurement of these parameters.

To facilitate a comparison between the alternative detector options for the LENF, we have performed our simulations for three detectors simultaneously: a Totally-Active Scintillator Detector (TASD) and two different liquid-argon detectors which have optimistic and conservative performance estimates respectively. Our model of the TASD is based upon Ref. [136] and has a fiducial mass of 20 kton, 35 variable-width energy bins and a constant energy resolution of 10% for both quasi-elastic and non-quasi-elastic events. The efficiency rises linearly from 73% to 94% over a range of 0.5 GeV to 1 GeV and then remains constant for higher energies. The background on the golden channel is taken as a constant fraction of 1×10^{-3} of the events arising from neutral-current interactions and the same fraction of events from the disappearance channel which accounts for instances of charge misidentification. Both

of the liquid-argon detectors are based upon parameters first reported in Ref. [160] and elaborated on in subsequent optimization studies [146] by having a fiducial mass of 100 kton, an energy resolution on quasi-elastic events of 10% and a flat detection efficiency of 80%. The conservative (optimistic) model has 22 (35) variable-width energy bins and an energy resolution of 20% (10%) on non-quasi-elastic events. The backgrounds are taken as a fraction of 5×10^{-3} (1×10^{-3}) of events from both the neutral-current and disappearance channels.

Although a number of large mass-scale liquid-argon detectors have been proposed [107, 161–163], only a few designs [107, 162] discuss extensions to 100 kton. As mentioned previously, the magnetization of large-volume detectors at the scale considered in our simulations is a particular challenge and further research is needed to fully assess the feasibility of the design. Our choice of such large detector volumes (100 kton and 20 kton for liquid-argon and TASD respectively) is designed to provide an optimistic performance estimate which covers the full range of potentialities of the LENF if these technical difficulties can be overcome. However, it is worth noting that in our approximation the performance of the LENF will only depend upon the detector mass through its *exposure* (number of muon decays $\times$ fiducial detector mass). Using this equivalence, the performance of a 100 kton detector with 1.0×10^{21} muon decays per year is expected to be comparable to a 70 kton detector with a larger flux of 1.4×10^{21} decays.

We have assumed normal hierarchy to be true throughout our simulations and the known oscillation parameters were chosen to be $\sin^2 2\theta_{12} = 0.3$, $\theta_{23} = \pi/4$, $\Delta m_{21}^2 = 8.0 \times 10^{-5} \text{ eV}^2$ and $|\Delta m_{13}^2| = 2.5 \times 10^{-3} \text{ eV}^2$. The uncertainty on these values was accounted for by allowing the parameters to vary during the minimization procedure: we allowed an uncertainty of 4% and 10% for the solar and the atmospheric parameters, respectively. These parameter choices were made in accordance with previous optimization studies of the LENF [146] and are close to the best-fit values and uncertainties from recent global analyses of the existing neutrino oscillation data [164, 165].

3.2.1 Technical details

We will now provide a technical discussion of the method for computing parameter sensitivities at oscillation facilities. To compute the attainable parameter sensitivity, we perform a simulation of the experiment, which reproduces the spectral event rates expected at such a facility. We then fit these event rates, using an approach based on log-likelihood ratios. This allows us to assign a value to our estimator for each hypothesis and, following a standard procedure, this value can be converted into a statistical significance.

The main inputs into a simulation of a long-baseline oscillation facility are the flux of neutrinos produced at the source, the cross-sections for the neutrino–nucleon interaction at the detector and a description of the detector facility. Taking a spectral flux of neutrinos of a fixed flavour ν_α to be given by $f_\alpha(E_\nu)$ and the cross-section for interaction with a nucleon to be $\sigma_n(E_\nu)$, we can construct an expected number of interactions in the detector by their product

$$n_\beta(E_\nu) = f_\alpha(E_\nu) \times P_{\alpha \rightarrow \beta}(E_\nu, L) \times \sigma_n(E_\nu) \times \epsilon \times M_n \times T, \quad (3.2.2)$$

where $P_{\alpha \rightarrow \beta}$ is the oscillation probability from flavour α to flavour β , ϵ denotes the efficiency of the detector, M_n is the number of target nucleons in the detector and T is the total run-time of the experiment. It is often easier if we re-write M_n in terms of the total mass of the detector M as $M_n = \kappa M$. In this case κ is the number of nucleons per unit mass of our detector material. In this form we see the key dependences of the number of observed events: total run time, detector mass and incoming flux. To avoid statistical limitations, experiments are therefore driven to long runs of intense beams impinging upon very large detectors.

Our simulations are based on the above formula, and are implemented using the GLoBES package [156,157]. In GLoBES, the above formula is generalized to include different cross-sections for different detection channels, full oscillation probabilities computed for a realistic Earth matter profile, energy dependent efficiencies and detector reconstruction effects. GLoBES also allows for the inclusion of backgrounds and systematic uncertainties on the number of events in each channel independently. The expected rates produced by GLoBES form the basic data for our subsequent

analyses, and the question of parameter sensitivity is addressed via statistical fits of these rates.

Statistical method

Once we have specified the experimental assumptions in our simulations, GLoBES can be seen as a mapping from the true set of oscillation parameters to a set of expected event numbers, representing the energy-binned event spectrum. We view this as a typical data set of our experiment. For a given set of true oscillation parameters, we can perform statistical analyses on this data set and attempt to extract the oscillation parameters. In this section, we will overview how our test statistic is defined and illustrate our approach to parameter determination.

We assume that the number of events in a given bin y obeys a Poisson distribution with a mean given by the expected number of events n as calculated in GLoBES for the set of true oscillation parameters $y \sim \text{Pois}(n)$. The probability of measuring a given y is then given by the following formula

$$P(y|n) \equiv \frac{n^y e^{-n}}{y!}.$$

Our statistic is based on a log-likelihood ratio test [166], where the likelihood is defined by $L(n|y) \equiv P(y|n)$. We choose to compare the likelihood of the rate being n after an observed rate y , to the most likely possibility that the observed rate is the true rate $y = n$. We construct the log-likelihood ratio as follows

$$\begin{aligned} \Delta\chi^2 &= -2 \ln \left(\frac{L(n|y)}{L(y|y)} \right), \\ &= -2 \ln \left(\frac{P(y|n)}{P(y|y)} \right), \\ &= 2 \left[y - n + n \ln \left(\frac{n}{y} \right) \right]. \end{aligned}$$

We have chosen the notation $\Delta\chi^2$, as this log-likelihood ratio is known to approximate the $\Delta\chi^2$ distribution asymptotically [166], a result known as Wilks' theorem [167].

To incorporate the effects of uncertainties and backgrounds, the expected number of events computed by GLoBES is split into two components, $y = y_s + y_b$. The

term y_s is the number of events arising from the signal process in question: for a golden channel analysis, these are events from charged-current interactions following $\nu_e \rightarrow \nu_\mu$ oscillation. The term y_b contains events arising from background channels: for example, a small fraction of right-sign events would be added to simulate charge misidentification. We expect these event numbers to fluctuate due to uncertainties in our experiment, and we account for the effect of systematic uncertainties using the *pull method*. This introduces a new nuisance parameter for every systematic uncertainty, and the additional freedom that these introduce reflects the possible variation in event numbers arising from these errors. The dominant uncertainties are on the overall normalization of the number of events in signal and background, which are described by two new parameters ξ_s and ξ_b . These parameters are known as *pulls*, and they act multiplicatively on the total number of events

$$y(\xi_s, \xi_b) = (1 + \xi_s) y_s + (1 + \xi_b) y_b.$$

The range over which these systematic factors can vary must be estimated, and it is assumed that they fluctuate normally about zero with a standard deviation σ_i . To incorporate this information, additional penalty terms are added to the $\Delta\chi^2$ which disfavour pulls far from their expected range

$$\Delta\chi_{\text{pull}}^2 = \left(\frac{\xi_s}{\sigma_s}\right)^2 + \left(\frac{\xi_b}{\sigma_b}\right)^2.$$

So far, we have only considered the log-likelihood ratio for a single bin, but our approach can be extended to include $N_{\text{bins}} > 1$ different bins by forming the joint likelihood function and log-likelihood ratios as before. Ultimately, we find that the contributions to the $\Delta\chi^2$ from each bin combine additively

$$\Delta\chi_{\text{bins}}^2(\xi_s, \xi_b) = \sum_{i=1}^{N_{\text{bins}}} 2 \left[y_i(\xi_s, \xi_b) - n_i - n_i \ln \left(\frac{n_i}{y_i(\xi_s, \xi_b)} \right) \right].$$

The final piece of information that we would like to include in our analysis, is the data coming from other experimental searches. The LENF is not expected to improve the current measurements on all of the oscillation parameters (for example, θ_{12} is difficult to measure using the golden channel), and external information on these parameters can, therefore, help to improve the sensitivity of the LENF by

reducing the freedom of the fit parameters. To implement these constraints, we introduce Gaussian priors on a number of parameters θ_i which belong to the parameter set that is being tested,

$$\Delta\chi_{\text{priors}}^2 = \sum_{i=1}^{N_{\text{priors}}} \left(\frac{\theta_i - \theta_i^0}{\sigma_{\theta_i}} \right)^2,$$

which corresponds to the assumption that the likelihood of θ_i taking the value θ_i^0 is given by a Gaussian

$$L(\theta_i) \propto e^{-\frac{(\theta_i - \theta_i^0)^2}{2\sigma_{\theta_i}^2}}.$$

By combining all of these terms, we find our final test statistic

$$\Delta\chi^2 = \min_{\{\xi_a\}} [\Delta\chi_{\text{bins}}^2(\xi_s, \xi_b) + \Delta\chi_{\text{pulls}}^2] + \Delta\chi_{\text{priors}}^2. \quad (3.2.3)$$

The value of this statistic then tells us how likely a set of hypothesised parameters is, given the observed rates n generated by the true parameter set. Most of the time we are not interested in the values of all of these parameters, and we will remove the dependence on uninteresting parameters by a process of marginalization: minimizing over the set of uninteresting parameters holding the values of other parameters fixed. This tell us, for the fixed parameter combinations that we are interested in, what the best-fitting set of parameters is regardless of the values of any other factors. This method naturally takes into account all correlations and degeneracies between parameters.

As an example, will now illustrate the method to exclude $\theta_{13} = 0$. If we denote the true set of oscillation parameters, those which generate our data, by Θ^T and the hypothesised parameter set by Θ^H , then the $\Delta\chi^2$ can be seen as a function of these parameters, $\Delta\chi^2(\Theta^T, \Theta^H)$. We are interested in seeing how well the hypothesis of $\theta_{13}^H = 0$ can fit data produced with $\theta_{13}^T \neq 0$, where the superscripts denote the set to which these parameters belong. As we are only interested in the values of θ_{13}^H , we marginalize over all of the other hypothesised parameters,

$$\Delta\chi^2(\theta_{13}^H = 0) = \min_{\theta_i^H \in \Theta^H - \{\theta_{13}^H\}} \Delta\chi^2(\Theta^T, \Theta^H)|_{\theta_{13}^H=0}, \quad (3.2.4)$$

This expression still has functional dependence on the parameters of Θ^T , and we will be mostly interested in its dependence on the true value of θ_{13}^T . Seen in this




Figure 3.1: The unoscillated flux of $\bar{\nu}_e$ at a neutrino factory (with μ^- in the storage ring) observed over a range of distances L (left) and a range of energies (right). The leftmost plot assumes that the stored muon energy is held fixed at $E_\mu = 25$ GeV, whilst the rightmost plot assumes $L = 2000$ km.

way, Eq. 3.2.4 tells us the value of the best-fitting set of hypothesised parameters with $\theta_{13}^H = 0$ for a given θ_{13}^T . For very large values of θ_{13}^T , we would expect this to return a large value, indicating a poor fit, and that the hypothesis $\theta_{13}^H = 0$ can be excluded. The quality of the fit would increase as $\theta_{13}^H - \theta_{13}^T$ approaches zero.

3.3 Phenomenology of parameter sensitivity

In this section, we present the performance of the LENF by considering its ability to make three major discoveries. We start by computing the sensitivity to non-zero θ_{13} which, in spite of recent experimental results, is included to highlight a number of general physical effects which influence the performance of the the neutrino factory. After this, we focus on the discovery of CP violation and the mass hierarchy.

To aid our phenomenological analysis, we will briefly recall some of the generic factors which influence the variation in performance of the LENF over the L - E_μ parameter space. As seen in Eq. 3.2.2, maintaining a high flux of neutrinos is essential in long-baseline experiments, as this is directly proportional to the number of events that can be observed at the detector, and can limit the statistical significance of any observations. The flux of $\bar{\nu}_e$ particles produced at a NF is known to high precision [130] and is shown in Fig. 3.1. This is well approximated by its tree-level

expression

$$\frac{d^2 N_{\bar{\nu}_e}}{dE_{\bar{\nu}} d\cos\theta} = 48 \frac{E_\mu}{m_\mu^6} E_\nu^2 (1 - \beta \cos\theta) (m_\mu^2 - 2E_\nu^2(1 - \beta \cos\theta)),$$

where E_μ (m_μ) is the energy (mass) of the decaying muon, E_ν the energy of the emitted antineutrino, $\beta^2 = 1 - m_\mu^2/E_\mu^2$ and θ is the angle between the incoming muon and the emitted neutrino. This flux increases with the stored-muon energy as E_μ , leading to an enhanced number of observed events at higher stored-muon energies. The dependence on baseline distance enters our calculation as we consider the number of events from this flux which are incident on a detector separated from the decay by L . Integrating over the angular dependence, we pick up a factor which accounts for the solid angle occupied by the detector. This introduces a decrease in events as the baseline lengthens, behaving as L^{-2} .

The observed number of wrong-sign muons is heavily influenced by the oscillation probability, Eq. 3.2.1, which introduces an additional dependence on the baseline distance and on the energy of the individual neutrinos, E_ν . Observing events which come from the first oscillation maximum is important as it ensures a large signal in the appearance channel. Events from the low-energy part of the spectrum contain important information on CP-violation, as for these values of E_ν the oscillation probability exhibits a strong dependence on the CP-violating phase δ . Matter effects lead to an enhancement or suppression of the oscillation signal compared to the same process in vacuum. Observing the sign of this difference can provide us with information on the neutrino mass hierarchy. In Eq. 3.2.1, matter effects are present if $A \neq 0$ and their influence increases with baseline distance and neutrino energy. Sensitivity to a given measurement can also be affected by the presence of degeneracies, the locations of which generally depend on both L and E_ν . At the detector, the neutrino–nucleon interaction cross-sections depend on the energies of the incident neutrinos. Flavor-tagged detection is made possible for charged-current interactions by observing the emitted charged lepton. The total cross-section for this channel increases with energy [168, 169] and, assuming that all else is kept equal, leads to an improved number of events at higher-energy facilities.

To understand the variation in performance of the LENF as we vary L and E_μ , the individual dependences mentioned above must be considered in combination. In




Figure 3.2: θ_{13} discovery potential for a LENF with TASD and a baseline distance of 2350 km for a range of stored-muon energies.

the following sections, we will discuss how these effects can explain the sensitivity of the LENF to non-zero θ_{13} , CP violation and the mass hierarchy.

3.3.1 θ_{13}

Here we present the results of our simulation which address the ability of the LENF to discover non-zero θ_{13} . The ability to discover non-zero θ_{13} is dependent upon the true values of the oscillation parameters. As mentioned previously, the parameter settings in our simulations have been fixed at values close to the current global fits. However, this leaves an uncertainty in the value of δ , and the discovery potential of a new facility will depend crucially upon the true value of this parameter. To understand this dependence, we have computed the confidence at which each configuration of the LENF can exclude $\theta_{13} = 0$ at the 3σ confidence level. These plots will be referred to as *discovery contours*, and an example can be shown in Fig. 3.2. In this figure, we see the ranges of the true values of δ and θ_{13} for which we can make the discovery for the LENF with TASD experiment over a range of stored-muon energies at fixed baseline $L = 2300$ km. The discovery potential in these plots captures well the dependence of the other experiments and configurations. We see that for large values of θ_{13} , those given by $\sin^2(2\theta_{13}) \gtrsim -2$, all configurations of this




Figure 3.3: The θ_{13} discovery range as a function of baseline, L , and stored-muon energy, E_μ . The left-hand column shows the lowest value of $\sin^2 2\theta_{13}$ for which the discovery fraction is 100% for all higher values and the right-hand column shows the lowest value of $\sin^2 2\theta_{13}$ for which there is a non-zero discovery fraction. The intensity of each point is given by the value of $\log_{10}(\sin^2 2\theta_{13})$.

experiments can exclude $\theta_{13} = 0$. As the true value of θ_{13} gets smaller, the sensitivity to this parameter starts to weaken, and the experiment eventually becomes unable to make the distinction. To understand the dependence on L and E_μ in more detail, we compute two metrics of performance: the 100% discovery point and the 0% discovery point. The 100% discovery point provides a conservative estimate of the performance of the facility. It is defined to be the smallest value of θ_{13} for which all higher values can exclude $\theta_{13} = 0$ for all true values of δ at 3σ significance. The 0% discovery point provides a complementary optimistic performance estimate, the smallest value of θ_{13} which can exclude $\theta_{13} = 0$ for any true value of δ .

Physically, the first of these quantities gives the smallest value of θ_{13} above which we expect a discovery and provides a conservative estimator of performance. The second quantity is a complementary optimistic estimator and tells us the smallest value of θ_{13} at which we could possibly make a discovery. The range of these two parameters gives the region of intermediate performance where we find discovery fractions between zero and one; discovery in this region is dependent on the exact value of δ . These can also be seen as the end points of the θ_{13} discovery fraction curves seen in previous studies (see for example Fig. 2 in Ref. [145]).

A selection of these results are shown in Fig. 3.3. We see that the TASD expects θ_{13} -discovery to at least $\sin^2 2\theta_{13} \gtrsim 10^{-3}$ and has the possibility of extending this limit by an order of magnitude. In comparison, the optimistic liquid-argon detector can discover non-zero θ_{13} down to at least $\sin^2 2\theta_{13} \gtrsim 10^{-4}$ and possibly as low as $\sin^2 2\theta_{13} \gtrsim 3 \times 10^{-5}$. The TASD and the conservative liquid-argon detector generally offer comparable sensitivities, which are both worse than those of the optimistic liquid-argon detector. For each detector, we attribute this to different causes. The conservative liquid-argon detector sees a similar total number of events as the optimistic liquid-argon detector however its higher backgrounds lead to a poor signal to background ratio which reduces its sensitivity. In contrast, the TASD has a superior control of the dominant backgrounds to the conservative liquid-argon detector but sees far fewer events due to its smaller size. These two effects reduce the attainable experimental sensitivity by a similar degree.

To further the analysis of our results, we will discuss the four quarters of this

parameter space separately. In the short-baseline and low-energy region (SB-LE) with $L \lesssim 2500$ km and $E_\mu \lesssim 14$ GeV, we see 100% discovery fractions for $10^{-4} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-2.5}$ depending on the choice of detector technology. In this region, the atmospheric term is relatively suppressed in the $P_{e\mu}$ oscillation probability. This suppression is a result of small θ_{13} and the relative enhancement of the CP and solar terms at lower neutrino energies. This leads to a poor performance as the signal becomes decreasingly sensitive to θ_{13} and the CP term introduces a more complicated dependence on the oscillation parameters. The decline in performance found towards the very shortest baselines is a consequence of the reduction in oscillation probability at small L . This leads to a poor signal-to-background ratio as the number of wrong-sign muons decreases. In the region of short baselines and high energies (SB-HE) with $L \lesssim 2500$ km and $E_\mu \gtrsim 14$ GeV, the larger energies increase the relative importance of the atmospheric term and we see that the situation is marginally improved with respect to the SB-LE region with 100% discovery fractions for $10^{-4} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-3}$ depending on detector technology. The overall improvement in this region can be seen as simply a leveling-off of the higher energy improvements of the SB-LE regions: for a fixed baseline distance, there is negligible improvement to be found when moving from $E_\mu = 15$ GeV to $E_\mu = 25$ GeV. This plateauing effect is associated with the energy spectrum of the neutrino factory, which rises almost linearly from $E_\nu = 0$ to its peak at $E_\nu \approx 2E_\mu/3$ and then drops sharply at the kinematic cut-off $E_\nu = E_\mu$. As we increase the stored-muon energy, the approximately linear tail of this spectrum only decreases slightly and therefore increasing the stored-muon energy at a LENF can be thought of as simply adding higher-energy neutrino events on top of the previous low-energy spectrum. Consequently, the difference in performance between an experiment with a low and a high stored-muon energy can be estimated by determining the importance of the neutrino events occurring with energies between the two stored-muon energies. Moving up through the SB-HE region, the information provided by the additional high-energy neutrinos is decreasingly useful because for higher-energy neutrinos, the L/E ratio is smaller and the oscillation probability decreases. This effect leads to a law of diminishing returns, where the discovery reach remains approximately constant.

However, as we move to higher energies, the signal to background ratio decreases and with the necessary introduction of additional backgrounds and the slight decrease in the low-energy part of the spectrum, the sensitivity in this region is expected to ultimately be reduced.

For the regions of parameter space with baselines of $L \gtrsim 2500$ km, we again divide the parameter space along $E_\mu = 14$ GeV into two quarters: the long-baseline, low-energy region (LB-LE) and the long-baseline, high-energy region (LB-HE). The LB-HE region has in general the best 100% discovery fractions of all of the parameter space. In this region, the design is approaching the HENF configuration where the signal to background ratio is increased and, thanks to the energy-dependent relative suppression of the solar and CP terms, the atmospheric contribution to the oscillation probability has a significant influence on the appearance channel. This permits the true value of θ_{13} to be smaller than in the SB-LE region whilst still providing an appreciable signal and therefore furthering the discovery reach. In contrast, the LB-LE region displays the poorest performance across the parameter space. We find 100% discovery fractions at around $\sin^2 2\theta_{13} \approx 10^{-2.5}$ for the T ASD and conservative liquid-argon detectors, whilst the optimistic liquid-argon detector has a 100% discovery fraction for $10^{-3} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-4}$. The poor performance in this region can be explained as an effect of low statistics: the neutrino flux is reduced as the baseline increases due to the effects of dispersion on the beam. In the LB-LE region this effect is compounded with a small neutrino flux at production. Consequently, the number of wrong-sign muons incident on the detector becomes increasingly restricted for these parameter choices.

In general, the optimal configuration depends upon the magnitude of the true value of θ_{13} . For large values of θ_{13} , as recently confirmed by Daya Bay and RENO [66, 86], all configurations are equally capable of confirming this effect.

The θ_{13} discovery reaches reported here are consistent with studies performed on the conventional NF [145] and, although making a direct quantitative comparison is difficult, the discovery reaches are of a similar magnitude. This behavior is also reported by Ref. [151] although, due to a difference in exposure, once again the results differ quantitatively.




Figure 3.4: Hierarchy discovery potential for a LENF with TASD and a baseline distance of 2350 km for a range of stored-muon energies.

3.3.2 Mass hierarchy

In order to present the results of our simulations regarding the determination of the mass hierarchy, we have computed analogous quantities to those used in Section 3.3.1. We have computed, for each detector type and choice of L and E_μ , the region of parameter space spanned by the true values of δ and θ_{13} for which the incorrect hierarchy can be excluded at the 3σ confidence level. Examples of these plots are shown in Fig. 3.4 for the TASD detector and a baseline of $L = 2350$ km. For large values of θ_{13} , $\sin^2(2\theta_{13}) \gtrsim 10^{-2}$, the determination of the mass hierarchy is possible for all configurations in these plots.

As in Section 3.3.1, we introduce 100% and 0% discovery points to allow us to display this data and extract the dependence on L and E_μ . The 100% discovery point is defined as the smallest value of θ_{13} for which all higher values allow the determination of the mass hierarchy at 3σ regardless of the true value of δ , whilst the 0% discovery point is defined as the smallest value of θ_{13} for which there is any value of δ which allows the determination of the mass hierarchy at 3σ . These quantities are plotted for the TASD and both variants of the LAr design in Fig. 3.5. We see that, for most of the parameter space, the TASD has a 100% discovery point around $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ and 0% discovery points that are smaller by an order of




Figure 3.5: Hierarchy determination as a function of baseline, L , and stored muon energy, E_μ , for the TASD (top row) and the liquid-argon detector with conservative and optimistic performance estimates (middle and bottom row, respectively). The left column shows the 100% discovery point, whilst the right column shows the 0% discovery point. In each plot, the black lines describe regions for which discovery is possible up to the value of $\sin^2 2\theta_{13}$ as labeled. The intensity at each point is given by $\log_{10}(\sin^2 2\theta_{13})$ of the relevant discovery limit.

magnitude. The conservative liquid argon detector performs very similarly to the T ASD, whilst the optimistic liquid-argon detector offers a similar 100% discovery point for baselines below $L = 2500$ km but, for baselines greater than this, can produce a significantly lower limit between $\sin^2 2\theta_{13} \gtrsim 10^{-3}$ and $\sin^2 2\theta_{13} \gtrsim 10^{-4}$. The 0% discovery point in this case is generally $\sin^2 2\theta_{13} \gtrsim 3.2 \times 10^{-4}$ going down to $\sin^2 2\theta_{13} \gtrsim 3.2 \times 10^{-5}$ for the longest baselines at around $L = 3500$ km. It is clear that these plots exhibit a stronger dependence on baseline distance than has been seen in previous plots. This was to be expected, as it is well known that matter effects are crucial in lifting the hierarchy degeneracy and that these effects increase with longer baselines. This can be seen in the golden-channel oscillation probability by considering the difference between the appearance probability for wrong- and right-sign muons. Depending on the neutrino mass hierarchy, one of these channels is suppressed and the other enhanced: this discrepancy grows with longer baselines. An exception to this pattern is found in the LB-LE region where there is a notably poor performance compared to the other regions: the particularly low neutrino flux arising from the combination of long baselines and low stored-muon energies leads to this decreased sensitivity. These results are in qualitative agreement with predictions for the standard NF [145] and with previous studies of the LENF [146, 151].

It has been suggested [170] that a low-energy neutrino factory with a “bimagic” baseline of around $L = 2540$ km and stored-muon energy of $E_\mu = 5$ GeV would offer a pronounced sensitivity to the neutrino mass hierarchy. This claim was motivated by studies of superbeams [116, 117] which looked for points of L – E_μ parameter space which show a clean separation between the hierarchy-conjugate oscillation probabilities. Starting from Eq. 3.2.1, some terms in the oscillation probability vanish at special choices of L and E_ν . This reduces the number of free parameters influencing the oscillation probability, mitigating the effect of parameter correlations. For the determination of the mass hierarchy, the uncertainty on the parameter δ introduces has a significant influence, and the idea of the bimagic baseline is to choose L and E_μ to reduce this effect. The dependence on δ in the oscillation probability can be removed for IH, and maximised for NH, by choosing L and E_ν

to satisfy the following relations

$$\sin\left(\frac{\Delta m_{31}^2 L}{4E_\nu} - \frac{AL}{2}\right) = \begin{cases} 0 & \text{for IH,} \\ \pm 1 & \text{for NH.} \end{cases}$$

These relations admit solutions given by

$$\begin{cases} \left(\frac{|\Delta m_{31}^2|}{2E_\nu} + A\right) \frac{L}{2} = n\pi & \forall n \in \mathbb{N}_1, \\ \left(\frac{|\Delta m_{31}^2|}{2E_\nu} - A\right) \frac{L}{2} = \frac{2m-1}{2}\pi & \forall m \in \mathbb{N}_1. \end{cases}$$

Solving these equations for L and E_ν , we find the expressions

$$\begin{aligned} L &= [2(n-m) + 1] \frac{\pi}{2A}, \\ E_\nu &= \frac{|\Delta m_{31}^2|}{2A} \frac{2(n-m)+1}{2(n+m)-1}. \end{aligned}$$

If we assume a constant matter density of $\rho = 3.2 \text{ g/cm}^3$, these formulae are simplified to

$$L \approx [2(n-m) + 1] 2700 \text{ km} \quad \text{and} \quad E_\nu \approx 10 \frac{2(n-m) + 1}{2(m+n) - 1} \text{ GeV.} \quad (3.3.5)$$

We see that, although there are a large set of solutions, for realistic baseline distances we should only consider solutions of the type $n = m$. For $m = n = 1$, Eqs. 3.3.5 produce $L \approx 2700 \text{ km}$ and $E_\nu \approx 3.3 \text{ GeV}$. (For the larger values of $m = n$, the energy becomes small and hard to access at a feasible facility.) We can follow an analogous process to look for L and E_ν combinations which remove the δ dependence for NH whilst maximizing it for IH. This leads to different set of L and E_ν combinations

$$L \approx [2(m-n) - 1] 2700 \text{ km} \quad \text{and} \quad E_\nu \approx 10 \frac{2(m-n) - 1}{2(m+n) - 1} \text{ GeV.} \quad (3.3.6)$$

We see that the only physical solutions will have $m > n$, and $m = n + 1$ is required for a reasonable baseline distance. For $m = 2$ and $n = 1$, Eqs. 3.3.6 is given by $L \approx 2700 \text{ km}$ and $E_\nu \approx 2 \text{ GeV}$. As before, the higher $m = n + 1$ solutions have harder to access energies.

As the flux of the neutrino factory covers a wide range of E_ν for any given E_μ , it is possible to choose a baseline and stored-muon energy combination which will cover both of these (L, E_ν) points: the first satisfying Eqs. 3.3.5 and the second satisfying Eqs. 3.3.6. It was the existence of these two “magic” energies for the same value of

L that lead to the idea of a “bimagic” baseline. A neutrino factory with a stored-muon energy of 5 GeV and a baseline of $L = 2540$ km would have a significant flux of neutrinos at both of the energies $E_\nu = 3.3$ GeV and $E_\nu = 2$ GeV which we have identified as the only practical solutions to Eqs. 3.3.6 and Eqs. 3.3.5. This configuration has been shown to lead to a strong sensitivity to the neutrino hierarchy [170]; however, it was not shown that this is the optimal configuration for hierarchy determination or that the bimagic baseline was responsible for the sensitivity. Using our results, we can approach these questions, and in Fig. 3.6 we show how the two discovery points depend on baseline distance for a selection of muon energies similar to the bimagic set-up. At a muon storage energy of $E_\mu = 5$ GeV we see evidence for a minimum in the 0% discovery limit at baseline distances around 2500 to 2800 km. This lower limit corresponds to the optimistic performance of the LENF and it is important to note that there is no corresponding minimum in the conservative estimate shown by the 100% discovery limit. We see from the other bands in Fig. 3.6 how this feature changes as the stored-muon energy increases: the minimum flattens out and drifts to higher baselines. We see that, for all stored-muon energies, a facility with a baseline distance below $L = 2500$ km can improve its discovery reach notably by increasing its baseline to at least 2500 km. Beyond this, it appears that if muon storage energies higher than 5 GeV are available then the bimagic choice is not the optimal configuration as the discovery reach can be additionally furthered by increasing both the energy and baseline. Furthermore, it is important to remember that even if stored muon energies are fixed at 5 GeV, a baseline of $L \approx 2500$ km only maximizes the optimistic performance of the design and any potential sensitivity advantage would crucially depend on the value of δ .

The behavior in Fig. 3.6 can be understood in light of our previous analysis. Matter effects are necessary to lift the hierarchy degeneracy and these are increased significantly by the use of long baselines. The reason that this trend does not appear for the $E_\mu = 5$ GeV case is because, at such a low energy, the experiment becomes statistically limited in the LB-LE region and necessitates the move to higher energies. The fact that the configurations near the bimagic baseline do not confer equal improvements for the conservative, 100% discovery point is a result of the construction




Figure 3.6: The hierarchy discovery reach as a function of baseline distance L for the T ASD design. For each band the upper (lower) boundary shows the 100% (0%) discovery reach. These correspond to the smallest value of $\log_{10}(\sin^2 2\theta_{13})$ for which the hierarchy can be resolved, independently of δ , for all higher values and the smallest value for which the discovery can be resolved for at least one value of δ , respectively.




Figure 3.7: CP violation discovery potential for a LENF with TASD and a baseline distance of 2350 km for a range of stored-muon energies.

of the bimagic baseline criteria: the degree of contrast between hierarchy-conjugate probabilities is not ensured to be large and its magnitude depends on the exact value of δ . Although for some values the contrast is sizable, for others it is greatly reduced. For example, at the bimagic configuration, when $\delta = 0$ the lowest magic energy around $E_\nu = 2$ GeV predicts rates which are almost identical for both hierarchies and can offer little discriminatory information. The presence of values of δ for which the hierarchy distinction is less marked leads to a larger 100% discovery point, which is designed to measure exactly this worst case scenario.

Provided the LB-LE region is avoided, identifying the optimal configuration for measurements of the mass hierarchy reduces to the observation of a simple correlation between the expected magnitude of $\sin^2 2\theta_{13}$ and the baseline distance. For the smallest values of θ_{13} , configurations in the LB-HE region are necessary. However, for the larger values as measured by Daya Bay and RENO, the physics reach is quite stable and the exact configuration in our parameter space is largely unimportant.

3.3.3 CP violation

In this section we analyse the ability of the LENF to discover leptonic CP violation. We define discovery to be the exclusion at 3σ significance of $\delta \in \{0, \pi\}$, the CP conserving values of δ . As for the discovery of $\theta_{13} \neq 0$ and the mass hierarchy,




Figure 3.8: CP-violation discovery fractions as a function of baseline, L , and stored-muon energy, E_μ for large θ_{13} . Each column (row) shows the discovery fraction for a different ‘true’ value of $\sin^2 2\theta_{13}$ (detector option).




Figure 3.9: CP-violation discovery fractions as a function of baseline, L , and stored-muon energy, E_μ for small θ_{13} . Each column (row) shows the discovery fraction for a different ‘true’ value of $\sin^2 2\theta_{13}$ (detector option) which is given at the top-left (top-right) of each plot.

the ability to make this exclusion depends upon the true values of the oscillation parameters. One of the most important unknowns influencing the difficulty of this search will be the true value of δ itself. This dependence becomes manifest in Fig. 3.7, where we show the regions of the parameter space spanned by the true values of θ_{13} and δ for which the discovery of CP violation can be made. The characteristic shape of these contours, two lobes which peak around $\delta = \pm 90^\circ$, is due to the nature of our discovery criteria: we are looking to exclude $\delta \in \{0, \pi\}$ and this becomes increasingly difficult as the true value of δ approaches 0 or π . As such, a facility will never be able to discover CP violation for all of the unknown parameter space, and the sensitivity to δ must always be expressed by degree. For this reason, in order to show the sensitivity to CP violation as a function of L and E_μ , we compute a quantity called the *CP-violation discovery fraction*. This is defined as the fraction of true values of δ for which we can exclude $\delta \in \{0, \pi\}$.

In Figs. 3.8 and 3.9 we present the CP-violation discovery fraction as a function of baseline and stored-muon energy at different true values of $\sin^2 2\theta_{13}$ for the TASD and both models of the liquid-argon detector. In the SB-LE region, for $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ we generally see very strong performance with discovery fractions of 70% to 90% depending on the choice of detector technology. These sensitivities confirm the expectations behind the original motivation for the LENF: a facility with the ability to measure low-energy neutrino events has access to the oscillation spectrum near the second maximum where CP-violating effects are most pronounced. In Fig. 3.9, we show the results of our simulations for smaller values of $\sin^2 2\theta_{13}$, where the δ -dependent terms in the probability are increasingly suppressed and the performance of the SB-LE region starts to worsen significantly. At the lowest energies, this suppression is compounded by the enhancement of the θ_{13} - and δ -independent solar term. For $\sin^2 2\theta_{13} \approx 10^{-4}$ there is negligible coverage for the TASD, but the optimistic liquid-argon detector still maintains a discovery fraction of around 60% due to its combination of high statistics and strong signal to background ratio. Considering experiments with higher stored-muon energies, we see scant improvement in the SB-HE region as the additional events at high energy provide little information on the parts of the oscillation spectrum which exhibit the most sensitivity to the

CP-violating phase. Generally we see discovery fractions of around 70% to 90% for $\sin^2 2\theta_{13} \gtrsim 10^{-2}$, which drops to 60% to 70% for $\sin^2 2\theta_{13} \approx 10^{-3}$. For the smallest values of θ_{13} , the T ASD has once again negligible sensitivity whilst the optimistic liquid-argon detector can still determine the effect of CP-violation in 60% to 70% of cases.

As noted in Section 3.3.1, while the baseline distance is increased, there is a decrease in event numbers due to a weakening of the neutrino flux arising from long baselines, and this may be compounded by an additional weakening of the flux at low energies. We generally find the lowest discovery fractions in the LB-LE region and this contrast is especially marked in the case of small θ_{13} . However, for long baselines but high energies (LB-HE), we see good sensitivity to CP-violation, especially when θ_{13} is small. In this region, both the neutrino flux at production and the neutrino–nucleon cross-sections are increased and this helps to mitigate the effect of baseline distance on the event numbers. The additional influence of appreciable matter effects over longer baselines and the inclusion of neutrinos which probe the most CP-sensitive parts of the oscillation spectrum further improve the sensitivity. For $\sin^2 2\theta_{13} \gtrsim 10^{-3}$, the LB-HE region has comparable discovery fractions to those of the SB-LE region, whilst for $\sin^2 2\theta_{13} \approx 10^{-4}$ the only significant sensitivity is to be found in the LB-HE region, with discovery fractions of 40% for the T ASD and conservative liquid-argon detectors and 70% for the optimistic liquid argon detector.

The true dependence of our simulations on detector mass, runtime and the number of useful muon decays per year is through their product, referred to as the *exposure*. In Fig. 3.10, we can further our understanding of the CP-violation discovery fraction by considering its dependence on this parameter. In general, for large values of θ_{13} the discovery fraction reaches a plateau for each experimental configuration. This limiting value represents the inherent limitations of the experiment: exposure directly determines the statistics of our experiment and at some point the measurement uncertainties will be dominated by systematic effects for which an increase in statistics can confer only slight improvement to the experiment’s sensitivity. The plots in Fig. 3.10 show the difference in performance for a modest (10%, left plot) and a more optimistic (1%, right plot) energy resolution

over a range of values of a uniform systematic error on the signal and backgrounds. We see that variations in magnitude of the systematic errors induce the greatest change in the attainable CP-violation discovery fraction. These effects can lead to a significant decline in performance; for 10% (1%) energy resolution, there is a decrease in the discovery fraction of 8% (6%). This limiting influence on the discovery fraction, which arises through systematic uncertainties, is quite stable to variations under baseline and energy, assuming that these choices do not generate a significant probabilistic suppression of the number of events. This effect helps to explain the observed uniformity in parts of Fig. 3.8, for example in the SB-HE region. We have also investigated the impact of alternative sources of systematic limitations. For instance, reducing the prior uncertainty associated with the matter density can lead to modest improvements in sensitivity; however, we find that the unilateral improvement of any one systematic factor leads to little impact on the sensitivity obtained in Fig. 3.10: it is necessary to reduce all systematics uniformly to significantly further the physics reach. Generally, for the parameter ranges that we have studied, it is the energy resolution and overall systematic error that are responsible for the greatest variation in the attainable CP-discovery fraction.

In common with the discovery of $\theta_{13} \neq 0$, the optimal configuration of the LENF for CP-violation discovery divides into two scenarios depending on the size of θ_{13} . For $\sin^2 2\theta_{13} \gtrsim 10^{-3}$, provided the LB-LE region is avoided, the CP-violation discovery fractions are almost uniform and the exact configuration is unimportant. As θ_{13} decreases, the δ -dependent signal is suppressed and appreciable sensitivities can only be found in the LB-HE region where the advantages of the HENF design start to become relevant.

The performance of the single-baseline HENF has been shown [145] to share generic features with our data: the SB-LE region suffers from lower discovery fractions and once θ_{13} has decreased to around $\sin^2 2\theta_{13} = 10^{-4}$, the only sensitivity can be found in the LB-HE region. The CP-violation discovery fractions of the single-baseline HENF are very similar to those of the LENF with the T ASD whereas, for all values of $\sin^2 2\theta_{13}$ that we have studied, the liquid-argon detector has discovery fractions higher by around 10%. For a two-baseline HENF, with the longer baseline




Figure 3.10: The CP-violation discovery fraction as a function of exposure for a range of systematic errors. The plot on the left (right) shows the discovery fraction assuming an energy resolution of 10% (1%). The scale of the horizontal axis is chosen to coincide with runtime assuming a 20 kton detector and 10^{21} useful muon decays per year per polarity. These plots use the TASD and have θ_{13} set to the Daya Bay best-fit value of $\sin^2 2\theta_{13} = 0.092$ [66]. All otherwise unspecified parameters are as described in Section 3.2.

chosen to be at the magic baseline [145], the comparison changes depending on the size of $\sin^2 2\theta_{13}$. For $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ the HENF performs similarly to the LENF with TASD whilst the optimistic liquid-argon detector has higher discovery fractions, once again by around 10%. For values of θ_{13} in the range $10^{-4} \lesssim \sin^2 2\theta_{13} \lesssim 10^{-3}$, the HENF starts to out-perform the LENF with TASD. However, the performance of the LENF with an optimistic liquid-argon detector remains comparable. Our results in the SB-LE region agree qualitatively with those computed in Ref. [146] for a LENF with $L = 1300$ km and $E_\mu = 4.5$ GeV and also agree qualitatively with the recent simulations of the LENF [151].

3.4 Optimal choices for the LENF

In light of the latest experimental results [66, 85, 86, 171, 172], the optimization of the LENF for large values of θ_{13} has become essential. The Daya Bay and RENO experiments can both independently exclude $\theta_{13} = 0$ at around 5σ and suggest an allowed region of around $0.07 \leq \sin^2 2\theta_{13} \leq 0.14$ at 1σ [66, 86]. Our simulations show that these values lie in the optimal region for the performance of the LENF: CP-violation discovery fractions of 80–90% are attainable in most of the parameter space with an optimistic liquid-argon detector or TASD. This sensitivity can be understood by the low-energy enhancement of the CP term in the golden-channel oscillation probability combined with a relative suppression of the solar term because of the large value of θ_{13} . Additionally, the mass hierarchy will be measurable independently of the choice of δ for all of the detectors, baselines and stored-muon energies that we have considered in this study. This capability arises because of the effect of the matter potential which generates an enhancement or suppression in the expected number of wrong-sign muons depending on the neutrino hierarchy. For a large value of θ_{13} , a significant number of oscillated events should occur and it is unlikely that the parameter sensitivity will be statistically limited. As the motivation for detectors with very large masses is primarily one of statistics, a large value of θ_{13} allows a reduction in detector mass whilst maintaining a similar level of performance. Although statistical errors are less relevant with a larger mixing angle, systematic

errors become increasingly important. A careful consideration of the systematic uncertainties is a vital next step in the analysis of the LENF; work has been done on the role of systematics for the LENF with MIND [173], but alternative detectors are yet to be studied.

For a related discussion of the effect of large values of θ_{13} on the sensitivities of the LENF, see also Ref. [151].

3.5 Summary

The Neutrino Factory has previously been studied in two distinct configurations: the conventional HENF uses 25 GeV muons and has two baselines at 4000 km and 7000 km. More recently, the idea of the LENF has been presented: this uses low-energy muons, typically around 5–10 GeV, and a unique baseline at 1500–2000 km. This set-up exploits the rich oscillatory pattern of the appearance probability and, by mitigating the effect of degeneracies among the unknown neutrino parameters, can provide an excellent physics reach for larger values of θ_{13} . Motivated by recent experimental results, we have performed a “green-field” study of the dependence of the performance of the LENF on the choice of stored-muon energy, $4 \leq E_\mu \leq 25$ GeV, and baseline distance, $1000 \leq L \leq 4000$ km, in order to ultimately identify the configuration of the optimal LENF.

In this chapter we have presented the results of numerical simulations on the ability of the LENF to answer three questions: is θ_{13} non-zero, does the neutrino mixing matrix give rise to CP-violating phenomena and what is the correct neutrino mass hierarchy? We find that non-zero θ_{13} can be discovered for almost all values of baseline and stored-muon energy down to $\sin^2 2\theta_{13} \gtrsim 10^{-3}$ and we expect limited sensitivity for a further order of magnitude. For CP-violation, discovery fractions of 70% to 90% are expected to be attainable for $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ and of 50% to 70% for $\sin^2 2\theta_{13} \gtrsim 10^{-3}$. We have shown that, provided that the extremal configurations of our L – E_μ parameter space are avoided, specifically very short baselines or long baselines paired with low energies, these performance estimates are quite general and do not require significant fine-tuning of the baseline or the stored-muon energy. The

neutrino mass hierarchy is expected to be accessible to the LENF for $\sin^2 2\theta_{13} \gtrsim 10^{-2}$ with the possibility of discovery generally extending to $\sin^2 2\theta_{13} \gtrsim 4 \times 10^{-4}$. For the hierarchy determination, the dependence on baseline is clearly seen and if baselines of $L > 3000$ km are selected, discovery reach could be extended even further to $\sin^2 2\theta_{13} \gtrsim 4 \times 10^{-5}$.

We have also considered the potential advantages of the bimagic baseline configuration given by $L = 2540$ km and $E_\mu = 5$ GeV. We have shown that for this choice of parameters, although a performance maximum is present in the 0% discovery reach, there is no corresponding maximum in the 100% discovery reach. This means that, if considered conservatively, such a configuration is of limited benefit for a NF. We have also shown that higher energy configurations can provide an improved discovery reach and we see a drift in the optimal configuration towards higher energies and baselines.

Now that θ_{13} has been measured to be relatively large and $\sin^2 2\theta_{13} \approx 0.09$, we see that the sensitivity of the LENF is comparable for the majority of baseline distance and stored-muon energy arrangements. Furthermore, we expect the results not to be statistically limited, implying that smaller detectors might be considered to provide the required sensitivities and that a detailed study of the impact of systematic errors is required. For these values of θ_{13} , the optimal LENF for the resolution of the most compelling open questions of long-baseline physics has a very broad design. This stability of performance shows that the LENF concept is versatile, with flexibility to accommodate additional design criteria, whilst also providing a strong sensitivity to the most important physical quantities.

Chapter 4

Discrete flavour symmetries and parameter correlations

The flavour structure of the standard model (SM) can be seen as one of its most puzzling features. The Yukawa couplings which govern the quark and charged lepton mass terms account for 13 out of the 19 free SM parameters, and if we naively extended this to include neutrino masses, another 7 to 9 parameters are introduced. However, the enlarged neutrino sector should not be seen as simply an extension of the problem of flavour, but as offering another path towards a solution through novel experimental inquiry. In this chapter, we will discuss a possible way for the neutrino flavour sector to be explored experimentally. In particular, we will analyse the testable predictions of theories which attempt to explain the leptonic flavour sector through the application of a new symmetry principle. After a recap of the structure and successes of these models in Section 4.1, we will discuss potential signatures of these models in the form of certain parameter correlations which could be tested in the near future at a precision facility. In Section 4.4 and Section 4.3, we will then present two phenomenological frameworks in which we can start to understand how these correlations arise, and what we can learn from their investigation.

4.1 Models with discrete flavour symmetries

The combination of independent sectors of a theory into a unifying structure is a very popular technique in the community of model builders who are interested in extending the standard model (SM). The three generations of particles in the SM present an obvious candidate for such a unification, and many models have been proposed which organise these fields into larger structures. One way in which the three generations could be unified is by applying a flavour symmetry. This extends the symmetry group of the standard model by a symmetry G_f , and in some limit, all three generations can be described by the same object. To produce the known leptonic mixing pattern, this must clearly be a broken symmetry, as the flavour symmetries of the known leptonic mass terms admit no non-trivial symmetries, and in a complete theory this breaking would be expected to occur spontaneously. The additional group may be either continuous or discrete. Although models of continuous flavour symmetries are possible, the large mixing angles of the PMNS matrix have suggested simple permutation relations to many authors. These are prototypical examples of discrete symmetries, and models based on these ideas have proven to be a successful way of implementing flavour symmetric principles.

Many models have been proposed which invoke a discrete symmetry to help resolve the problem of leptonic flavour [174–176]. These models vary greatly in their complexity and naturalness; however, most of them follow the same essential steps. At some scale the discrete group, G_f , must be specified; common choices for this group include A_4 , S_4 and $\Delta(27)$. These are small finite groups with three-dimensional representations, and frequently, the three generations of leptonic SU(2) doublets are assigned to a triplet representation ensuring that their mixing is highly constrained. New scalar fields are then introduced, called *flavons*, which are also assigned to representations of G_f , but are typically neutral under the standard model gauge group. The lagrangian can then be written down in the conventional fashion, with all terms included that are consistent with the symmetries of the theory. The terms which constitute the flavon-flavon interactions are referred to as the flavon potential; in successful models the minimum of this potential will require non-zero vacuum expectation values (VEVs) for a subset of the flavon fields, a feature which

will spontaneously break G_f . A pattern of masses and mixings should then emerge, shaped by the breaking mechanism and the possible presence of residual symmetries.

Approaches of this type typically generate first-order expressions for the PMNS matrix which are populated by simple algebraic values, and a number of such patterns have been proposed, see *e.g.* [174–176] for reviews with extensive lists of references. A noteworthy example is the tribimaximal (TB) mixing matrix [177]:

$$U_{\text{TB}} = \begin{pmatrix} \sqrt{\frac{2}{3}} & \frac{1}{\sqrt{3}} & 0 \\ -\frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{3}} & \frac{1}{\sqrt{2}} \end{pmatrix} P, \quad (4.1.1)$$

which has received much attention by model builders. Another option, referred to as golden ratio (GR) mixing, is given by the following matrix [178, 179]:

$$U_{\text{GR}} = \begin{pmatrix} \cos \vartheta & \sin \vartheta & 0 \\ -\frac{1}{\sqrt{2}} \sin \vartheta & \frac{1}{\sqrt{2}} \cos \vartheta & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \sin \vartheta & -\frac{1}{\sqrt{2}} \cos \vartheta & \frac{1}{\sqrt{2}} \end{pmatrix} P, \quad (4.1.2)$$

where $\tan \vartheta = 1/\varphi$ with φ given by the golden ratio, $\varphi = (1 + \sqrt{5})/2$. It is clear that these patterns must receive corrections to be consistent with the known phenomenological data: in particular, both patterns predict $\theta_{13} = 0$. However, as we shall show later, these patterns provide a helpful framework for classifying a large class of models with phenomenologically viable mixing angles.

The knowledge of non-zero θ_{13} has excluded many previously compelling models based on discrete symmetries. However, there is an ever expanding list of models which have adapted to deal with this new information, and many are currently viable and predictive. The impact that the measurement of θ_{13} has had on the field of discrete leptonic flavour symmetries shows how essential it is to establish phenomenological connections between the model building and experimental aspects of the discipline. In the remainder of this chapter, we present our work along these lines: developing an understanding of how we can further test models of discrete flavour symmetries at current and future experimental facilities.

4.2 Parameter correlations and sum rules

The incorporation of discrete flavour symmetries into any extension of the standard model will only further our understanding of leptonic flavour if it manages to reduce the number of free parameters controlling this sector. Therefore, we generally expect for these models to generate correlations amongst the physical parameters governing the leptonic Yukawa matrices. Correlations amongst the Yukawa terms lead to constraints on the possible mass matrix textures that can be achieved in that theory. In turn, these are diagonalised by matrices with a certain pattern to their entries, and this loss of generality is ultimately inherited by the PMNS matrix. For a general model based on discrete flavour symmetries, the correlations between the PMNS matrix elements will correspond to a number of complicated relations amongst the mixing angles and phases. We refer to these relations as *sum rules*, and each sum rule provides a constraint which reduces the number of degrees of freedom in the leptonic mixing sector by one. For example, in a recent model presented in Ref. [180], an A_4 symmetry is imposed which is broken spontaneously by a set of flavons. This leads to the second column of the PMNS mixing matrix being fixed at its tribimaximal value,

$$|U_{e2}| = |U_{\mu 2}| = |U_{\tau 2}| = \frac{1}{\sqrt{3}}.$$

These correlations amongst the PMNS matrix elements can be expressed in terms of the mixing parameters by two sum rules

$$\sin^2 \theta_{12} = \frac{1}{3(1 - \sin^2 \theta_{13})}, \quad (4.2.3)$$

$$\tan(2\theta_{23}) \cos \delta = \frac{\cos^2 \theta_{12} - \sin^2 \theta_{12} \sin^2 \theta_{13}}{\sin(2\theta_{12}) \sin \theta_{13}}. \quad (4.2.4)$$

The sum rules associated with a given model present a set of predicted relationships which can, in principle, be tested experimentally. In this chapter we aim to discuss the types of sum rules that we might expect to see at a facility capable of making precision measurements of the neutrino mixing parameters. We will address this topic using two different phenomenological frameworks. The first, in Section 4.3, provides a general class of models which predict a number of well defined sum rules,

whilst the second, in Section 4.4, discusses what a given sum rule could tell us about the structure of the leptonic mass matrix.

Before we begin a more detailed analysis, it is clear from Eqs. 4.2.3 and 4.2.4 that the sum rules arising from discrete symmetries are not necessarily simple nor easy to manipulate. For this reason we will often work with approximate forms of the sum rules, found by linearizing the full relations in a set of small parameters. We use the notation of Ref. [181], which introduces the parameters s , r and a defined by

$$\sin \theta_{12} \equiv \frac{1+s}{\sqrt{3}}, \quad \sin \theta_{13} \equiv \frac{r}{\sqrt{2}}, \quad \sin \theta_{23} \equiv \frac{1+a}{\sqrt{2}}.$$

These parameters were originally introduced to describe deviations from tribimaximality and are referred to as the *solar*, *reactor* and *atmospheric* parameters, respectively. They provide a close phenomenological fit to the known mixing angles; the recent global fit in Ref. [57] provides the following 1σ intervals (for normal neutrino mass ordering)

$$\begin{aligned} -0.07 &\leq s \leq -0.02, \\ 0.20 &\leq r \leq 0.23, \\ -0.12 &\leq a \leq -0.05. \end{aligned} \tag{4.2.5}$$

Our analysis will lead us to focus on a specific set of correlations which are primarily dependent on the atmospheric mixing angle θ_{23} , reactor mixing angle θ_{13} and the cosine of the Dirac CP phase, $\cos \delta$. These correlations are referred to as *atmospheric sum rules*, as they do not depend on the solar parameter. Using the parameters above, we will often work with the linearized atmospheric sum rule which, for the models that we are interested in, will take the form

$$a = a_0 + \lambda r \cos \delta + \mathcal{O}(r^2, a^2), \tag{4.2.6}$$

where we will treat a_0 and λ as new model-dependent constants.

We will now illustrate the preceding discussion with two examples from the literature which will highlight our terminology, how sum rules arise, and will also introduce two special values of λ . The sum rules given previously in Eqs. 4.2.3 and 4.2.4,

from the model presented in Ref. [180], can be expressed in terms of s , r and a as

$$\begin{aligned} s &= \sqrt{\frac{2}{2-r^2}} - 1, \\ a &= -\frac{1}{2}r \cos \delta + \mathcal{O}(a^2, r^2), \end{aligned} \quad (4.2.7)$$

where the atmospheric sum rule is shown in its linearized form. This is a specific realisation of our general rule, Eq. 4.2.6, with $a_0 = 0$ and $\lambda = -1/2$. A different pair of sum rules has been found in Ref. [182], once again by spontaneously breaking the group A_4 ; however, in this model the first column of the PMNS matrix is fixed at its tribimaximal value. This imposes the relations,

$$|U_{e1}| = \sqrt{\frac{2}{3}} \quad \text{and} \quad |U_{\mu 1}| = |U_{\tau 1}| = \frac{1}{\sqrt{6}}.$$

Using these relations to compute the solar sum rule and the associated linearized atmospheric sum rule, we find

$$\begin{aligned} s &= \sqrt{1 - \frac{2r^2}{2-r^2}} - 1, \\ a &= r \cos \delta + \mathcal{O}(a^2, r^2), \end{aligned} \quad (4.2.8)$$

which corresponds to $a_0 = 0$ and $\lambda = 1$.

Although we will show that there exist a number of values of interest for λ , the two sum rules which we have found in Eqs. 4.2.7 and 4.2.8, with $\lambda = -1/2$ and $\lambda = 1$, will play a distinguished role in our analysis. These two relations have a degree of universality, having arisen in the literature from the fully consistent models discussed above, whilst also well describing many of the simple rules that we have found in our more phenomenological treatments. These will appear in the analysis of Section 4.3 and also in the simplified relations of Section 4.4. Their appearance can be understood quite generally by considering the form of the PMNS matrix expanded to first order in s , r and a [181]

$$U_{\text{PMNS}} = \begin{pmatrix} \sqrt{\frac{2}{3}}(1 - \frac{1}{2}s) & \frac{1}{\sqrt{3}}(1 + s) & \frac{1}{\sqrt{2}}re^{-i\delta} \\ -\frac{1}{\sqrt{6}}(1 + s - a + re^{i\delta}) & \frac{1}{\sqrt{3}}(1 - \frac{1}{2}s - a - \frac{1}{2}re^{i\delta}) & \frac{1}{\sqrt{2}}(1 + a) \\ \frac{1}{\sqrt{6}}(1 + s + a - re^{i\delta}) & -\frac{1}{\sqrt{3}}(1 - \frac{1}{2}s + a + \frac{1}{2}re^{i\delta}) & \frac{1}{\sqrt{2}}(1 - a) \end{pmatrix} P.$$

The first of our distinguished relations, $a = -\frac{1}{2}r \cos \delta$, appears when we combine the phenomenological requirement that $s \approx 0$ with the constraint $|U_{\mu 2}| = |U_{\tau 2}|$. This leads to the following first order relation

$$\Re \left[a + \frac{1}{2} r e^{i\delta} \right] = 0 + \mathcal{O}(a^2, r^2),$$

where $\Re[z]$ denotes the real part of $z \in \mathbb{C}$. This is equivalent to the sum rule Eq. 4.2.7. The second relation, $a = r \cos \delta$, follows from a parallel procedure applied to the first column of the PMNS matrix. To satisfy the experimental constraints, s must be small and this enforces the distinguished sum rules that we have identified. However, s need not be zero, and this allows the sum rule to deviate from these two forms. In the next two sections we will see that a number of relations which are numerically close to these sum rules can arise consistently with the known data.

4.3 Hernandez-Smirnov symmetry building

In this section, we present an analysis of the sum rules which arise in a framework first introduced in Ref. [183], following earlier work by Refs. [184–187]. This framework allows us to derive flavour-symmetric correlations amongst the PMNS mixing matrix elements, whilst making minimal assumptions about the details of the model. This approach was built around the assumption that there exists a discrete flavour group which is broken spontaneously into two subgroups. These subgroups act independently on the charged lepton and neutrino sectors of the theory, and their misalignment leads to a non-trivial PMNS matrix. If we additionally assume that some of the known symmetries of the leptonic mass terms are in fact residual symmetries arising from this larger broken group, constraints can be placed on the PMNS matrix in a general manner, regardless of the precise implementation of the symmetry breaking. These assumptions are common to many models, so called *direct models* in the terminology of a recent review [176], but are not the most general scenario, and alternative scenarios have also been shown to produce viable mixing patterns.

The method assumes the breaking of a discrete flavour group into two distinct $\mathbb{Z}_n$ subgroups which remain unbroken in either the charged lepton or neutrino sector,

whilst broken in the other. Based on this construction, the authors of Ref. [183] reported a number of parameter correlations; however, these correlations led to linearized sum rules identical to those reported in previous studies. In this section, we review the approach and weaken some of the assumptions made in the derivations of these relations, this generates a number of additional sum rules with distinct linearized relations to those reported previously.

4.3.1 Derivation of constraints

The approach in Ref. [183] assumes that the flavour group is a von Dyck group, $D(n, m, p)$. These are defined by the presentation

$$\langle S, T, W | S^n = T^m = W^p = STW = 1 \rangle.$$

The generators S and T are assumed to describe residual symmetries of the Majorana neutrino and charged lepton mass terms, respectively, whilst W is defined to be the inverse of the product ST . The symmetry of the Majorana neutrino mass term is the Klein group $\mathbb{Z}_2 \times \mathbb{Z}_2$. Every member of this group has order 2, and this fixes n to be given by $n = 2$. Only one of the $\mathbb{Z}_2$ factors originates from the flavour symmetry and is generated by S , while the other one arises accidentally. If the second $\mathbb{Z}_2$ would be embedded in the group as well, another parameter relation would appear, which fixes the mixing angles as it has been discussed in Ref. [184, 185]. The symmetry of the charged lepton mass term is $U(1)^3$, and we consider only the discrete sub-groups given by $\mathbb{Z}_m$. The choice of m and p remains free; however, the assumption that the unbroken group is finite restricts these to specific values.¹ Representing each choice by the ordered pair (m, p) , the choices which lead to finite groups are exhausted by five special pairs

$$(3, 3), \quad (3, 4), \quad (3, 5), \quad (4, 3), \quad (5, 3),$$

and two infinite sequences

$$(2, N) \quad \text{and} \quad (N, 2) \quad \forall N \in \{n \in \mathbb{N} | n \geq 2\}.$$

¹For a generalisation to infinite von Dyck groups see Ref. [188].

The former are isomorphic to the groups A_4 , S_4 , A_5 , S_4 and A_5 , respectively. The two infinite sequences lead to dihedral symmetry groups which do not have irreducible triplet representations and are therefore not considered any further.

For a given (m, p) , the two generators S and T must be chosen from the symmetries of the leptonic mass terms, assuming that they are residual symmetries following the spontaneous breakdown of G_f . For this to be the case, the generators S and T must have at least one unit eigenvalue. This is necessary for there to exist a VEV alignment that remains invariant under their action. Under the further assumption that the discrete groups are subgroups of $SU(3)$, we find that the generator of the symmetry of the diagonalised neutrino mass matrix must be given by either

$$S'_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad S'_2 = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \text{or} \quad S'_3 = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Similarly, these constraints imply that the symmetry of the diagonalised charged lepton mass matrix is given by one of the three order- m generators

$$T'_e = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\frac{2\pi k}{m}} & 0 \\ 0 & 0 & e^{-i\frac{2\pi k}{m}} \end{pmatrix}, \quad T'_\mu = \begin{pmatrix} e^{i\frac{2\pi k}{m}} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{-i\frac{2\pi k}{m}} \end{pmatrix}, \quad \text{or} \quad T'_\tau = \begin{pmatrix} e^{i\frac{2\pi k}{m}} & 0 & 0 \\ 0 & e^{-i\frac{2\pi k}{m}} & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

where $k \in \{n \in \mathbb{Z}_m \mid n \text{ and } m \text{ are coprime}\}$. We choose to work in the basis of diagonal charged leptons such that $T_\alpha = T'_\alpha$ and $S_i = U_{\text{PMNS}} S'_i U_{\text{PMNS}}^\dagger$.

With a choice of generators T_α - S_i , we can construct W

$$W^{-1} = S_i T_\alpha = U_{\text{PMNS}} S'_i U_{\text{PMNS}}^\dagger T'_\alpha. \quad (4.3.9)$$

In Ref. [183], it is assumed that W has an eigenvalue 1. This constrains $\text{Tr}[W]$ to be real: from Eq. 4.3.9 we see that $\det W = 1$, if W has eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$, this requires $\lambda_1 \lambda_2 \lambda_3 = 1$. Taking $\lambda_1 = 1$ enforces the remaining eigenvalues to be complex conjugates $\lambda_2 = \lambda_3^*$. The trace is then given by

$$\text{Tr}[W] = 1 + \lambda_2 + \lambda_3 = 1 + 2\Re[\lambda_2] \in \mathbb{R}.$$

For the three finite von Dyck groups with a 3-dimensional irreducible representation, it can be shown by considerations of the group character tables that the existence of

an eigenvalue $\lambda_1 = 1$ is in fact a necessary property. From the group presentation, we see that the remaining eigenvalues must be p -th roots of unity, and therefore, we can express

$$\text{Tr}[W] = 1 + 2 \cos\left(\frac{2\pi d}{p}\right) \quad \text{s.t. } d \in \{n \in \mathbb{Z}_p \mid d \text{ and } p \text{ coprime}\}. \quad (4.3.10)$$

The trace of W is therefore real and given by a discrete set of values dictated by the choice of d and p . Computing the real and imaginary parts of this trace in terms of our chosen expressions for the generators S'_i and T'_α ,

$$\text{Tr}[W] = \text{Tr}[W^{-1}] = \text{Tr}[U_{\text{PMNS}} S'_i U_{\text{PMNS}}^\dagger T'_\alpha],$$

produces constraints on the PMNS matrix elements. These fix the values of one of the columns of the PMNS matrix, where the column fixed corresponds to the choice of generator S_i , and the order of the rows to the choice of T_α . In general, these constraints are given by

$$|U_{\beta i}|^2 = |U_{\gamma i}|^2 = \frac{1 - \eta}{2} \quad \text{and} \quad |U_{\alpha i}|^2 = \eta, \quad (4.3.11)$$

where $\{\alpha, \beta, \gamma\} = \{e, \mu, \tau\}$, and η is defined by

$$\eta = \frac{1 + \text{Tr}[W]}{4 \sin^2\left(\frac{\pi k}{m}\right)}.$$

Combined with Eq. 4.3.10, this produces an expression for η in terms of k and d

$$\eta = \frac{\cos^2\left(\frac{\pi d}{p}\right)}{\sin^2\left(\frac{\pi k}{m}\right)}. \quad (4.3.12)$$

Using Eq. 4.3.12, we can scan systematically over the parameters k and d to find all possible values of η , and these are given in Tab. 4.1.

4.3.2 Impact of generator choice

The constraints on the PMNS matrix given in Eq. 4.3.11 lead to different sum rules depending upon the choice of generators S_i and T_α . There are nine possible choices of generator pairs and each pair produces a different set of sum rules. In this section, we will compute the sum rules imposed by each choice.

	(m, p)	(k, d)	η
A_4	(3, 3)	(1, 1)	$\frac{1}{3}$
S_4	(3, 4)	(1, 1)	$\frac{2}{3}$
	(4, 3)	(1, 1)	$\frac{1}{2}$
A_5	(3, 5)	(1, 1)	$\frac{1+\varphi}{3}$
	(3, 5)	(1, 2)	$\frac{2-\varphi}{3}$
	(5, 3)	(1, 1)	$\frac{2+\varphi}{5}$
	(5, 3)	(2, 1)	$\frac{3-\varphi}{5}$

Table 4.1: The unique values of η for finite von Dyck groups following the construction of Ref. [183]. The fractions k/m and d/p must be less than 1 and irreducible. Respecting these constraints, we are still left with some pairs of solutions with the same η , in which case the smallest (k, d) are shown.

Generators $T_\alpha-S_1$ or $T_\alpha-S_2$

For the six choices of generators $T_\alpha-S_1$ or $T_\alpha-S_2$, we can express the constraints as an exact solar sum rule and a linearized atmospheric sum rule. The pairs composed of S_1 are given by

$$\begin{aligned}
T_e-S_1 \quad & \begin{aligned} s &= \sqrt{3 \left(1 - \frac{2\eta}{2-r^2}\right)} - 1, \\ a &= \sqrt{\frac{\eta}{2(1-\eta)}} r \cos \delta. \end{aligned} \\
T_\mu-S_1 \quad & \begin{aligned} s &= \sqrt{3 \left(1 - \frac{1-\eta}{2-r^2}\right)} - 1, \\ a &= \frac{1-3\eta}{2(1+\eta)} + \sqrt{\frac{1-\eta}{2(1+\eta)}} r \cos \delta. \end{aligned} \\
T_\tau-S_1 \quad & \begin{aligned} s &= \sqrt{3 \left(1 - \frac{1-\eta}{2-r^2}\right)} - 1, \\ a &= -\frac{1-3\eta}{2(1+\eta)} + \sqrt{\frac{1-\eta}{2(1+\eta)}} r \cos \delta. \end{aligned}
\end{aligned}$$

The relations imposed by assuming the choice $T_\alpha-S_2$ are given by

$$\begin{aligned}
T_e-S_2 \quad & s = \sqrt{\frac{6\eta}{2-r^2}} - 1, \\
& a = -\sqrt{\frac{\eta}{2(1-\eta)}} r \cos \delta. \\
\\
T_\mu-S_2 \quad & s = \sqrt{\frac{3(1-\eta)}{2-r^2}} - 1, \\
& a = \frac{1-3\eta}{2(1+\eta)} - \sqrt{\frac{1-\eta}{2(1+\eta)}} r \cos \delta. \\
\\
T_\tau-S_2 \quad & s = \sqrt{\frac{3(1-\eta)}{2-r^2}} - 1, \\
& a = -\frac{1-3\eta}{2(1+\eta)} - \sqrt{\frac{1-\eta}{2(1+\eta)}} r \cos \delta.
\end{aligned}$$

Generators $T_\alpha-S_3$

These choices enforce either the relation $|U_{e3}|^2 = \eta$ or $|U_{e3}|^2 = \frac{1-\eta}{2}$. These lead to predictions for the reactor parameter given by

$$r = \sqrt{2\eta} \quad \text{or} \quad r = \sqrt{1-\eta}.$$

Combining these with the possible values of η in Tab. 4.1, it can be quickly seen that they violate the known experimental data on θ_{13} . We will not consider these further.

4.3.3 Phenomenologically viable sum rules

As we have shown, the constraints imposed by the Hernandez-Smirnov approach fix the i -th column of the PMNS matrix by symmetry alone. Fixing a column of the PMNS matrix introduces two independent constraints on the mixing angles. Some of these relations are quickly seen to be excluded by the phenomenological data, for example if we choose the generators $T_\alpha-S_3$; however, the remaining six generator choices must be compared more carefully to the known experimental data on the mixing angles to assess their viability.

We see that for all of the cases of interest in Section 4.3.2, the constraints can be expressed as a solar sum rule and a linearized atmospheric sum rule. The next generation of long baseline oscillation experiments is expected to increase our precision on the atmospheric parameter and also provide, for the first time, experimental constraints on δ . Therefore, in this section we aim to identify which atmospheric sum rules are consistent with the experimental data to date.




Figure 4.1: The predictions for the solar parameter s in the Hernandez-Smirnov framework. On the left we show all the possible predictions, whilst the right panel shows an enlarged version which focuses on those that intersect the allowed regions for s . The grey bands show best-fit regions for s : 5σ (left and right panel), 3σ (right only) and 1σ (right only).

We consider an atmospheric sum rule to be *viable* if its corresponding solar sum rule is consistent with the experimental data. To determine this, we systematically combine the predictions of all of the solar sum rules generated by the six pairs of generators in Section 4.3.2 with the values of η consistent with the Hernandez-Smirnov framework, as listed in Tab. 4.1. We show the results of this process in Fig. 4.1. In the left panel we have shown the prediction of the solar parameter as a function of r over its 3σ range [189]. The right panel shows an enlarged version over the region of phenomenological interest: the best-fit regions for s shown at 1σ , 3σ and 5σ confidence levels. We see that only six distinct predictions for the solar angle are compatible with the global data, although some of these are produced by more than one pair of generators: for A_4 , all choices $T_\alpha-S_2$ agree with T_e-S_2 , and for S_4 and A_5 , $T_\tau-S_2$ has an identical prediction to $T_\mu-S_2$.

In Tab. 4.2, we have shown all of the viable sets of sum rules, along with the choices of (m, p) and (k, d) which generate them. The atmospheric sum rule is in each case given by its linearized form, described by the parameters a_0 and λ introduced in Eq. 4.2.6. A number of the scenarios that we have identified in Tab. 4.2 can be explained in terms of the TB and GR matrices given in Eqs. 4.1.1 and 4.1.2. The three scenarios based on an A_4 symmetry all lead to a value of the second column of




Figure 4.2: A comparison between linearized and exact sum rules for $\lambda = 1$ and the model presented in Ref. [182], which fixes the elements of the first column of the PMNS matrix to their tribimaximal values. The solid (empty) region denotes the exact (linearized) prediction for $\cos \delta$ which is produced by varying r over its current 3σ allowed interval.

the PMNS matrix fixed at its tribimaximal value; similarly, the S_4 scenario with the generator choice T_e-S_1 fixes the prediction of the first column to be tribimaximal. The scenario based on A_5 with unbroken generators T_e-S_1 (T_e-S_2) fixes the first (second) column of the PMNS matrix to the equivalent values of the GR mixing matrix. As we see in the final three columns of Tab. 4.2, all of these models predict sum rules with values of λ numerically close to 1 or $-1/2$.

4.3.4 Validity of linearization

So far we have focused only on the first-order expression of these correlations in the parameters r and a ; we will now address the impact of higher-order terms. As we have shown, in general correlations predicted by flavour symmetric models are non-linear relations between the oscillation parameters. In this chapter, we have computed all viable atmospheric sum rules in a linearized form, and the validity of this approach can be determined. To do so requires some knowledge of the true correlations and, as an example, we will consider the model presented in Ref. [182],

G_f	m	$T_\alpha-S_i$	(k, d)	s	a_0	λ	s (3 d.p.)	a_0 (3 d.p.)	λ (3 d.p.)
A_4	3	T_e-S_2	(1, 1)	$\frac{1}{\sqrt{1-r^2/2}} - 1$	0	$-\frac{1}{2}$	0.012	0	-0.5
	3	$T_\mu-S_2$	(1, 1)	$\frac{1}{\sqrt{1-r^2/2}} - 1$	0	$-\frac{1}{2}$	0.012	0	-0.5
	3	$T_\tau-S_2$	(1, 1)	$\frac{1}{\sqrt{1-r^2/2}} - 1$	0	$-\frac{1}{2}$	0.012	0	-0.5
S_4	3	T_e-S_1	(1, 1)	$\sqrt{1 - \frac{2r^2}{2-r^2}} - 1$	0	1	-0.024	0	1
	4	$T_\mu-S_2$	(1, 1)	$\sqrt{\frac{3}{2(2-r^2)}} - 1$	$\frac{1}{6}$	$-\frac{1}{\sqrt{6}}$	-0.124	0.167	-0.408
	4	$T_\tau-S_2$	(1, 1)	$\sqrt{\frac{3}{2(2-r^2)}} - 1$	$-\frac{1}{6}$	$-\frac{1}{\sqrt{6}}$	-0.124	-0.167	-0.408
A_5	5	T_e-S_1	(1, 1)	$\sqrt{3 + \frac{6}{(3-\varphi)(r^2-2)}} - 1$	0	$\frac{\varphi}{\sqrt{2}}$	-0.118	0	1.144
	5	T_e-S_2	(2, 1)	$\sqrt{\frac{6}{(2+\varphi)(2-r^2)}} - 1$	0	$\frac{1-\varphi}{\sqrt{2}}$	-0.079	0	-0.437
	5	$T_\mu-S_2$	(2, 1)	$\sqrt{\frac{3\varphi}{(2\varphi-1)(2-r^2)}} - 1$	$-\frac{5-4\varphi}{22}$	$-\sqrt{\frac{3+2\varphi}{22}}$	0.054	0.067	-0.532
	5	$T_\tau-S_2$	(2, 1)	$\sqrt{\frac{3\varphi}{(2\varphi-1)(2-r^2)}} - 1$	$\frac{5-4\varphi}{22}$	$-\sqrt{\frac{3+2\varphi}{22}}$	0.054	-0.067	-0.532

Table 4.2: Analytical and numerical expressions for the phenomenologically viable linearized sum rules arising in the Hernandez-Smirnov framework for finite von Dyck groups. In this table, m gives the order of the generator which controls the charge lepton mass matrix, $T_\alpha^m = 1$, and $\varphi = (1 + \sqrt{5})/2$ is the golden ratio.

which fixes the elements of the first column of the PMNS to their tribimaximal values. As a function of r and a , this model predicts that $\cos \delta$ is given by the composition of the following functions:

$$\cos \delta = \frac{\sqrt{2}}{3} \frac{1 - 6 \sin^2 \theta_{12} \cos^2 \theta_{23} - 3r^2 \cos^2 \theta_{12} \sin^2 \theta_{23}}{r \sin(2\theta_{12}) \sin(2\theta_{23})},$$

$$\cos \theta_{12} = \frac{2}{\sqrt{3(2-r^2)}}, \quad \text{and} \quad \sin \theta_{23} = \frac{1+a}{\sqrt{2}}.$$

When linearized, these relations lead to the simpler expression $\cos \delta = a/r$. In Fig. 4.2 we have computed the predictions of $\cos \delta$ as a function of a for both the linearized and exact sum rules. We see that for this model the difference between the two treatments is small. The impact of higher order corrections can only be assessed on a case by case basis once the exact sum rules are known; however, due to the smallness of the s , r and a parameters, we expect the linear approximation to be a good one. This is confirmed by our simulations for the known exact sum rules, and therefore we will focus our later analysis only on the linearized relations.

4.3.5 Summary

For a von Dyck group $D(2, m, p)$ broken into two cyclic subgroups generated by S and T , which respectively act on the neutrino and charged lepton mass terms, a column of the PMNS matrix is fixed by the symmetry alone. The values of the elements of this column are given by the choice of (m, p) and the choice of two integers k and d . Which column is fixed, and the pattern of values that are imposed, is governed by a choice of one of nine possible pairs of generators. Only 4 of these choices appear interesting phenomenologically: T_e-S_1 , T_e-S_2 , $T_\mu-S_2$ and $T_\tau-S_2$. The constraints on the mixing angles imposed for these generator choices can be expressed as an exact expression for s as a function of r and a linearized sum rule for the atmospheric mixing angle. By comparing the predictions of s for all of the choices of (m, p) , k and d with the known phenomenological interval, we identify 8 viable scenarios, which are listed in Tab. 4.2.

4.4 Mass matrix decomposition

In this section, we will approach the question of parameter correlations in a different way, by investigating the relationship between the s , r and a parameters and the structure of the neutrino mass matrix. This will allow us to connect the observation of deviations from tribimaximality to features of the theory which, in theories with discrete non-Abelian symmetries, are ultimately connected to the choice of group and the allocation of particles to representations. We will start at the level of the neutrino (Majorana) mass matrix and decompose it into a basis of simple matrices. This basis is chosen such that there are three terms which lead to deviations from tribimaximality, and therefore non-zero s , r and a terms. We will then compute the mixing angles for this mass matrix by perturbation theory, and derive expressions for the deviation parameters in terms of the parameters in our mass matrix decomposition. This will highlight how the structure of the mass matrix influences the correlations between PMNS parameters, and elucidate what the observation of certain patterns of deviations can tell us about the underlying mass matrix. We will then show that in certain simplified cases, familiar sum rules can be re-derived.

4.4.1 Construction of the parameterization

In general, the Majorana mass matrix in the conventional three neutrino scenario is a complex-valued 3×3 symmetric matrix, M . We can identify the form of a general tribimaximal mass matrix by reversing the diagonalisation process. Working with a diagonal charged lepton mass matrix, the neutrino mass matrix in the flavour basis, M , satisfies

$$U_{\text{TB}}^\dagger M U_{\text{TB}}^* = M_{\text{D}},$$

where M_{D} is the diagonalised mass matrix $M_{\text{D}} = \text{diag}(m_1, m_2, m_3)$, and U_{TB} is the tribimaximal matrix given in Eq. 4.1.1. This can be decomposed into three matrices,

$$M = \sum_{i=1}^3 \tau_i T_i,$$

where $\tau_i \in \mathbb{C}$, and the T -matrices are given in the flavour basis by

$$T_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad T_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \quad T_3 = \frac{1}{\sqrt{6}} \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}.$$

A general complex symmetric 3×3 matrix has 6 complex degrees of freedom. Therefore, to describe a general mass matrix we must introduce three additional matrices which will not preserve tribimaximality. We extend the decomposition to include three D_i matrices,

$$M = \sum_{i=1}^3 \tau_i T_i + \sum_{i=1}^3 \delta_i D_i,$$

with $\delta_i \in \mathbb{C}$, and the D -matrices are chosen to be given in the flavour basis by

$$D_1 = \frac{1}{2} \begin{pmatrix} 0 & 1 & -1 \\ 1 & 2 & 0 \\ -1 & 0 & -2 \end{pmatrix}, \quad D_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & -1 \\ 1 & -1 & 0 \\ -1 & 0 & 1 \end{pmatrix}, \quad D_3 = \frac{1}{2\sqrt{3}} \begin{pmatrix} 4 & 1 & 1 \\ 1 & -2 & -2 \\ 1 & -2 & -2 \end{pmatrix}.$$

Using this decomposition, if we make a transformation into the tribimaximal basis the τ -parameters generate the diagonal terms and the δ -parameters, the off-diagonals

$$U_{\text{TB}}^T M_F U_{\text{TB}} = \begin{pmatrix} \mu_1 & 0 & 0 \\ 0 & \mu_2 & 0 \\ 0 & 0 & \mu_3 \end{pmatrix} + \sqrt{\frac{3}{2}} \delta_1 \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} + \sqrt{\frac{3}{2}} \delta_2 \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} + \sqrt{\frac{3}{2}} \delta_3 \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

where the diagonal terms are given by

$$\mu_1 = \tau_1 - \frac{\tau_2 - \sqrt{3}\tau_3}{\sqrt{2}}, \quad \mu_2 = \tau_1 + \sqrt{2}\tau_2, \quad \mu_3 = -\tau_1 + \frac{\tau_2 + \sqrt{3}\tau_3}{\sqrt{2}}.$$

The basis introduced above is entirely general; however, we are interested in mass matrices which can produce phenomenologically viable mixing patterns. Our scheme of parameterizing sum rules in terms of the s , r and a parameters, introduced in Section 4.2, means that these parameters are known to vanish when TB mixing is satisfied. It may be constructive therefore to consider the form of matrices which have approximately TB mixing angles. If we start with the diagonal mass matrix,

and assume that a small rotation can be made to put it into the TB basis, we see that we generate symmetric off diagonal terms

$$(M)_{\text{TB}} = \begin{pmatrix} m_1 & 0 & 0 \\ 0 & m_2 & 0 \\ 0 & 0 & m_3 \end{pmatrix} + \begin{pmatrix} 0 & (m_1 - m_2)x_{12} & (m_1 - m_3)x_{13} \\ (m_1 - m_2)x_{12} & 0 & (m_2 - m_3)x_{23} \\ (m_1 - m_3)x_{13} & (m_2 - m_3)x_{23} & 0 \end{pmatrix}.$$

For the corrections to the TB angles to be small, we require $|x_{ij}| \ll 1$. If we make the identification of δ_k with x_{ij}

$$\delta_k = x_{ij}(m_i - m_j),$$

then we see that for $|x_{ij}|$ to be small we require that $|\delta_k| \ll |m_i - m_j|$. This gives us an idea of the conditions required to produce an approximate TB mixing scheme. With this fact in mind, we revise our notation to make it clear when we expect approximate TB mixing. We choose the coefficients in front of the D-matrices to include a factor of the relevant difference in diagonal terms. We define the parameters ζ_1 , ζ_2 and ζ_3 by

$$\delta_1 e^{i(\beta_2 + \beta_3)} = \zeta_1 (|\mu_3| - |\mu_2|),$$

$$\delta_2 e^{i(\beta_1 + \beta_3)} = \zeta_2 (|\mu_1| - |\mu_3|),$$

$$\delta_3 e^{i(\beta_1 + \beta_2)} = \zeta_3 (|\mu_2| - |\mu_1|),$$

with the choice $2\beta_i = -\arg(\mu_i)$. Substituting these expressions into the mass matrix decomposition, we find

$$\begin{aligned} P^T U_{\text{TB}}^T M_\nu U_{\text{TB}} P = & \begin{pmatrix} |\mu_1| & 0 & 0 \\ 0 & |\mu_2| & 0 \\ 0 & 0 & |\mu_3| \end{pmatrix} + \sqrt{\frac{3}{2}} \zeta_2 (|\mu_1| - |\mu_3|) \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \\ & + \sqrt{\frac{3}{2}} \zeta_1 (|\mu_3| - |\mu_2|) \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} + \sqrt{\frac{3}{2}} \zeta_3 (|\mu_2| - |\mu_1|) \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \end{aligned} \quad (4.4.13)$$

where, in addition, we have multiplied by a diagonal unitary matrix P to make the diagonal terms real, $P_{ii} = e^{i\alpha_i}$. This decomposition is again entirely general, but we expect the mixing angles to be close to tribimaximal for $|\zeta_i| \ll 1 \forall i$.

4.4.2 Computing s , r and a

For a matrix of the form given in Eq. (4.4.13) (and assuming a diagonal charged lepton mass matrix) the PMNS matrix is given by $U = U_{\text{TB}}PV$, where U_{TB} is the tribimaximal matrix, P is a diagonal phase matrix and V is the unitary matrix necessary to diagonalise the TB-basis mass matrix. From a knowledge of the PMNS matrix, the mixing angles are found by comparing it to the PDG form,

$$\sin^2 \theta_{13} = |U_{e3}|^2, \quad \sin^2 \theta_{12} = \frac{|U_{e2}|^2}{1 - |U_{e3}|^2}, \quad \sin^2 \theta_{23} = \frac{|U_{\mu 3}|^2}{1 - |U_{e3}|^2}. \quad (4.4.14)$$

To compute these quantities for a general matrix, it is necessary to calculate V . This will be performed by perturbation theory, assuming small parameters $|\zeta_i|$. As M is a complex symmetric matrix, computing its diagonalising matrix, and therefore computing V , can be done using a special case of the singular value decomposition theorem, which states that there exists a unitary matrix W such that $W^\dagger MW^* = D$, such that D is a diagonal matrix with real positive entries. The matrix W can be computed by considering the Hermitian matrix $H = MM^\dagger$, which satisfies the following relation

$$W^\dagger HW = W^\dagger MM^\dagger W = W^\dagger MW^*(W^\dagger MW^*)^\dagger = D^2,$$

which is diagonal, and so W can be identified with the basis of eigenvectors of H .

The perturbative analysis leads to the following result for U

$$U = \sqrt{\frac{3}{2}} U_{\text{TB}} \begin{pmatrix} e^{i\beta_1} & 0 & 0 \\ 0 & e^{i\beta_2} & 0 \\ 0 & 0 & e^{i\beta_3} \end{pmatrix} \begin{pmatrix} \sqrt{\frac{2}{3}} & f(\zeta_3) & -f(\zeta_2) \\ -f(\zeta_3^*) & \sqrt{\frac{2}{3}} & f(\zeta_1) \\ f(\zeta_2^*) & -f(\zeta_1^*) & \sqrt{\frac{2}{3}} \end{pmatrix}, \quad (4.4.15)$$

where the $f(\zeta_i)$ functions are defined by

$$\begin{aligned} f(\zeta_1) &= \Re[\zeta_1] + i\Im[\zeta_1] \frac{|\mu_2| - |\mu_3|}{|\mu_2| + |\mu_3|}, \\ f(\zeta_2) &= \Re[\zeta_2] + i\Im[\zeta_2] \frac{|\mu_1| - |\mu_3|}{|\mu_1| + |\mu_3|}, \\ f(\zeta_3) &= \Re[\zeta_3] + i\Im[\zeta_3] \frac{|\mu_1| - |\mu_2|}{|\mu_1| + |\mu_2|}. \end{aligned} \quad (4.4.16)$$

Using the relations in Eq. 4.4.14, we can compute the expressions for s , r and a to first order in $|\zeta_i|$

$$\begin{aligned} s &= \sqrt{3} \Re [f(\zeta_3) e^{i\beta_{12}}], \\ r &= \left| f(\zeta_1) - \sqrt{2} f(\zeta_2) e^{i\beta_{12}} \right|, \\ a &= \frac{1}{\sqrt{2}} \Re [f(\zeta_2) e^{i\beta_{13}}] + \Re [f(\zeta_1) e^{i\beta_{23}}], \end{aligned} \quad (4.4.17)$$

where $\beta_{ij} \equiv \beta_i - \beta_j$.

4.4.3 Computing the CP phases

The computation of the CP phases of the PMNS matrix is complicated by the presence of unphysical phases. The rephasing of the charged leptons transforms the PMNS matrix as

$$(U)_{ij} \rightarrow e^{i\phi_i} (U)_{ij}. \quad (4.4.18)$$

As the phase of these leptons is unphysical, so is the phase on any matrix element of the PMNS matrix. To circumvent this problem, we compute the Majorana and Dirac phases by the use of rephasing invariants: functions of U invariant under Eq. 4.4.18. We can construct bilinear rephasing invariants of the form

$$s_{\alpha ij} = U_{\alpha i} U_{\alpha j}^*,$$

which are seen to be invariant under the transformation of Eq. 4.4.18. This ensures that the unphysical degrees of freedom surrounding the choice of phases will not enter our final expressions. Using the standard PDG parameterization of the mixing matrix, we find that the arguments of these invariants give us sufficient information to reconstruct the CP phases of the PMNS matrix

$$e^{i\alpha_{21}} = \frac{U_{12} U_{11}^*}{|U_{12} U_{11}|}, \quad e^{i\alpha_{32}} = \frac{U_{23} U_{22}^*}{|U_{23} U_{22}|}, \quad e^{i(\alpha_{31}-\delta)} = \frac{U_{13} U_{11}^*}{|U_{13} U_{11}|},$$

such that $\alpha_{ij} \equiv \alpha_i - \alpha_j$. Thanks to the linear dependence of the phase differences, $\alpha_{21} + \alpha_{13} + \alpha_{32} = 0$, the knowledge of these three invariants is sufficient to uniquely determine all of the physical CP phases.

After a straightforward calculation, we find the following expressions for α_{21} and α_{32} ,

$$\begin{aligned}\alpha_{21} &= \beta_{21} + \frac{\sqrt{3}}{2} \Im [f_3 e^{-i\beta_{21}}] + \mathcal{O}(\zeta_i^2), \\ \alpha_{32} &= \beta_{32} - \frac{1}{2} \Im [f_1 e^{-i\beta_{32}}] + \frac{1}{\sqrt{2}} \Im [f_2 e^{-i\beta_{31}}] + \frac{\sqrt{3}}{2} \Im [f_3 e^{-i\beta_{21}}] + \mathcal{O}(\zeta_i^2),\end{aligned}$$

and α_{31} can be found by exploiting the linear dependence of the phase differences

$$\alpha_{31} = \alpha_{32} + \alpha_{21} = \beta_{31} - \frac{1}{2} \Im [f_1 e^{-i\beta_{32}}] + \frac{1}{\sqrt{2}} \Im [f_2 e^{-i\beta_{31}}] + \sqrt{3} \Im [f_3 e^{-i\beta_{21}}] + \mathcal{O}(\zeta_i^2).$$

Now to use the formulae for the bilinear rephasing invariants to find δ , we must compute the phase of the matrix element U_{13} . As this element is zero to leading order, its phase is very sensitive to the perturbations and difficult to compute in general. However, if we denote its phase by Φ , defined as

$$\frac{U_{13}}{|U_{13}|} = e^{i\Phi},$$

then we can find a formula for the Dirac CP phase in terms of this unphysical parameter, which can be found explicitly in a number of relevant regimes. This process leads us to the following first-order expression for δ ,

$$\delta = \beta_3 - \Phi - \frac{1}{2} \Im [f_1 e^{-i\beta_{32}}] + \frac{1}{\sqrt{2}} \Im [f_2 e^{-i\beta_{31}}] + \frac{3\sqrt{3}}{2} \Im [f_3 e^{-i\beta_{21}}] + \mathcal{O}(\zeta_i^2).$$

4.4.4 Simplified expressions for s , r and a ?

Although they elucidate how the mass matrix generates deviations from tribimaximality, the expressions in Eq. 4.4.17 are too complicated to derive simple relationships amenable to phenomenological analysis. However, in certain special cases, Eq. 4.4.17 might be expected to simplify. In this section, we will consider two such scenarios. Firstly, we will consider the effect of the mass hierarchy. Then we will study some restricted scenarios when only one ζ_i parameter is non-zero.

Hierarchical spectra

Let us define M_{ij} by

$$M_{ij} = \frac{|\mu_i| - |\mu_j|}{|\mu_i| + |\mu_j|},$$

which is related to the $f(\zeta_k)$ functions introduced in Eq. 4.4.16. The bound $|M_{ij}| \leq 1$ is clearly satisfied, implying bounds for $f(\zeta_k)$

$$|\Re[\zeta_k]| \leq |f(\zeta_k)| \leq |\zeta_k|.$$

These bounds are saturated in extreme mass configurations: the upper limit is met only when one of the μ_i terms approaches zero and the lower limit, when the μ_i terms approach degeneracy. The two bounds converge when ζ_k is itself real.

We consider mass hierarchies which satisfy the known values of the mass squared differences

$$\Delta m_{21}^2 = 8.0 \times 10^{-5} \quad \text{and} \quad |\Delta m_{31}^2| = 2.5 \times 10^{-3}.$$

The preceding discussion has shown that simplifications will occur for sufficiently extreme hierarchies; however, we are only interested in such configurations if they are attainable at reasonable neutrino mass scales. In Fig. 4.3, we have plotted the values of $|M_{ij}|$ as a function of the lowest neutrino mass in each scenario. Assuming that $|\zeta_k| = \mathcal{O}(10^{-2})$ then, regardless of hierarchy, $|M_{ij}| = \mathcal{O}(|\zeta_k|) \forall i, j$ holds if the lightest mass is greater than around 0.15 eV. This leads to a constraint on $f(\zeta_k)$

$$f(\zeta_k) = \Re[\zeta_k] + \mathcal{O}(|\zeta_k|^2).$$

For smaller, yet still non-zero, values of the lowest neutrino mass, the only simplification we see is that $M_{12} = \mathcal{O}(|\zeta_k|)$ remains valid down to lowest masses of around 2.5×10^{-2} eV. As the lowest mass approaches zero, we see a new hierarchy dependent simplification: for normal hierarchy, $M_{13} = M_{12} = -1$ and $M_{23} \approx -0.7$, whereas for inverted hierarchy $M_{13} = M_{23} = 1$ and $M_{12} = \mathcal{O}(|\zeta_k|)$. This leaves us with the following matrix of possibilities

	NH, $m_1 = 0$	IH, $m_3 = 0$
$f(\zeta_1) =$	$\Re[\zeta_1] - i0.7\Im[\zeta_1]$	ζ_1
$f(\zeta_2) =$	ζ_2^*	ζ_2
$f(\zeta_3) =$	ζ_3^*	$\Re[\zeta_3] + \mathcal{O}(\zeta_3 ^2)$

These relations clearly simplify the general expressions found in Eq. 4.4.17. However, these simplifications are not sufficient to significantly alter the pattern of correlations observed in the general case.




Figure 4.3: $|M_{ij}|$ as a function of the smallest mass in the case of a normal (left) and inverted (right) hierarchy. In each plot, the value 5×10^{-2} is also plotted for comparison.

Single matrix contributions

Although we have seen that the general expressions for s , r and a in terms of a perturbation series in ζ_i are quite complicated. We can find significant simplifications if we consider restricted scenarios, where some of the ζ_i parameters vanish. In this section, we will consider the simplest of these configurations, when there is a single non-zero ζ -parameter, and show that the correlations predicted in these scenarios are closely related to those already discussed in Section 4.3.

$\zeta_1 = \zeta_2 = 0$ **and** $\zeta_3 \neq 0$

As we can see in Eq. 4.4.17, at first order the parameter ζ_3 decouples from the other parameters and only influences the solar parameter s . In the case that the ζ -parameters satisfy $\zeta_1 = \zeta_2 = 0$, the prediction for the solar parameter is given by

$$s = \sqrt{3}|f(\zeta_3)| \cos(\beta_{21} - \phi_3),$$

for $\phi_3 = \arg[f(\zeta_3)]$, whilst $a = r = 0$ to first order. This latter prediction is clearly in violation of the known value of θ_{13} , and we can therefore conclude that the true neutrino mass matrix must have at least one contribution from ζ_1 and ζ_2 .

$\zeta_1 = \zeta_3 = 0$ **and** $\zeta_2 \neq 0$

As $\zeta_3 = 0$ in this configuration, the solar parameter is fixed at leading order to be zero. However, the non-zero value of ζ_2 introduces deviations from tribimaximality for both a and r . These are given by

$$\begin{aligned} r &= \sqrt{2} |f(\zeta_2)| + \mathcal{O}(\zeta_i^2), \\ a &= \frac{1}{\sqrt{2}} |f(\zeta_2)| \cos(\beta_{31} - \phi_2) + \mathcal{O}(\zeta_i^2), \end{aligned}$$

for $\phi_2 = \arg[f(\zeta_2)]$. These expressions can be used to derive a sum rule at first-order

$$a = \frac{1}{2} r \cos(\beta_{31} - \phi_2). \quad (4.4.19)$$

In this simplified setting we can compute the phase in Eq. 4.4.19 in terms of δ by calculating Φ . As $\zeta_1 = 0$ and $\zeta_2 \neq 0$, U_{13} is given by $U_{13} = -f(\zeta_2)e^{i\beta_1}$, and therefore $\Phi = \phi_2 + \beta_1 + \pi$. To first order, this implies the relation

$$\delta = \beta_3 - \Phi = \beta_{31} - \phi_2 - \pi.$$

Substituting this into Eq. 4.4.19, we find a familiar sum rule

$$a = -\frac{1}{2} r \cos \delta.$$

$\zeta_2 = \zeta_3 = 0$ **and** $\zeta_1 \neq 0$

As with the previous case, having ζ_3 equal to zero fixes $s = 0$ to first order. The remaining parameters are given by

$$\begin{aligned} r &= |f(\zeta_1)| + \mathcal{O}(\zeta_i^2), \\ a &= |f(\zeta_1)| \cos(\beta_{32} - \phi_1) + \mathcal{O}(\zeta_i^2), \end{aligned}$$

where $\phi_1 = \arg[f(\zeta_1)]$. As before, these relations can be combined into a sum rule given by

$$a = r \cos(\alpha_{32} - \phi_1). \quad (4.4.20)$$

For these parameters the PMNS matrix element U_{13} is given by $U_{13} = f(\zeta_1)e^{i\beta_2}/\sqrt{2}$, which implies that $\Phi = \beta_2 + \phi_1$, and to first order δ is given by

$$\delta = \beta_3 - \Phi = \beta_{32} - \phi_1 = \alpha_{32} - \phi_1.$$

Combining this relation with Eq. 4.4.20, we find the sum rule

$$a = r \cos \delta.$$

4.4.5 Summary

In this section, we have shown that certain structures in the mass matrix can be connected to the correlations between the mixing parameters of the PMNS matrix. To satisfy the experimental data on the mixing angles, we consider a tribimaximal mass matrix subject to three distinct perturbations governed by the parameters ζ_1 , ζ_2 and ζ_3 . In general, the parameters ζ_1 and ζ_2 govern the linearized atmospheric sum rule, whilst ζ_3 dictates the solar sum rule to first order. If we consider simplified scenarios, with only a single non-zero ζ_i parameter, we can re-derive sum rules previously seen in the Hernandez-Smirnov approach of Section 4.3. We have shown that $\zeta_1 \neq 0$ leads to the sum rule $a = r \cos \delta$, and $\zeta_2 \neq 0$ leads to $a = -\frac{1}{2}r \cos \delta$. This suggests that understanding the other sum rules found in Section 4.3 may be possible if we consider a hierarchy of perturbations about these two configurations. Measuring deviations from $\lambda = 1$ and $\lambda = -1/2$ could, therefore, tell us important information about the terms present in M .

Chapter 5

Constraining atmospheric sum rules at current and future facilities

In Chapter 4, we derived a number of constraints on the parameters of the PMNS matrix, and investigated how these constraints were related to the flavour structure of the leptonic sector. We saw, by working in the frameworks of Section 4.3 and Section 4.4, that a study of correlations between a , r and $\cos \delta$ could distinguish between a number of viable models. These correlations, referred to as atmospheric sum rules, can be described by the general linearised expression

$$a = a_0 + \lambda r \cos \delta + \mathcal{O}(a^2, r^2). \quad (5.0.1)$$

In this chapter, we study how these correlations are constrained by current data, and how well they can be measured at two of the leading designs for a next-generation oscillation facility: a low-energy neutrino factory and a wide-band superbeam. In Section 5.1, we consider what the current data on the neutrino mixing parameters can tell us about atmospheric sum rules, and in Section 5.2 we present the results of our simulations of the superbeam and neutrino factory, assessing the ability of these experiments to constrain atmospheric sum rules.




Figure 5.1: The predictions for $\cos \delta$ made by the linearized sum rules listed in Tab. 4.2 with a and r varying over their current 3σ allowed intervals [189]. The region of physical $\cos \delta$ is shown in blue. The prediction for $S_4 (T_\tau - S_2)$ actually extends further than plotted into the unphysical region.


5.1 Compatibility of sum rules with existing and projected data

Sum rules make definite predictions about the relationships between the parameters of the PMNS matrix. These parameters have been investigated by a number of different experiments over the years, and our knowledge of them is quite advanced [57, 189]. Of the three parameters of the PMNS matrix which appear in the general linearised atmospheric sum rule, shown in Eq. 5.0.1, the values of the two mixing angles θ_{13} and θ_{23} are known to within 5% and 10% respectively, whilst the CP violating phase δ is still unconstrained at a significant level [189]. Given these constraints, the linearized atmospheric sum rule for a model can be seen as a prediction for the value of $\cos \delta$. We define the prediction range for $\cos \delta$ to be given by the mapping from r and a which is found by inverting Eq. 5.0.1,


$$\cos \delta = \frac{a - a_0}{\lambda r}.$$

In Fig. 5.1, we allow the parameters a and r to vary across their 3σ intervals and the image of this mapping is taken to be the range of potential values for $\cos \delta$. This has been done for all of the viable sum rules found in Section 4.3, with each prediction labeled by a representative choice of group and generators. All of these models have some overlap between their predictions and the physical region $|\cos \delta| \leq 1$, but for some models this is a prediction at the extremes of their allowed intervals. In models of this type, the sum rules are starting to be constrained by the requirement of their own consistency: an unphysical prediction for $\cos \delta$ means that the sum rule does not fit comfortably with the known data. To further understand the relationship between a and the predictions for $\cos \delta$, we have plotted the predictions of the viable sum rules as a function of the true value of a whilst r varies over its 1σ interval. This can be seen in Figs. 5.2 and 5.3, where we also show the current global data on a (the grey regions) and the projected sensitivity to the a parameter as reported in Ref. [95] (the red bands). These projections are for the global parameter sensitivity in 2025 assuming only the current experimental program: 5 years of data from T2K, 6 from NO ν A, and 3 years each for Double Chooz, RENO and Daya Bay. As we cannot predict the future best-fit value, the horizontal location of the predicted regions is largely irrelevant, and in Figs. 5.2 and 5.3 they have been arbitrarily centred around the current best-fit value to indicate the expected increase in precision.

For the case of the $\lambda \approx 1$ sum rules, shown in Fig. 5.2, both sum rules make very similar predictions, and are largely consistent with the known data on the mixing angles: only the most extreme values of $\cos \delta$ are inconsistent with the 2σ allowed regions for a . If the current global best fit values remain fixed as more data is collected, reducing the uncertainty on a , then both sum rules will predict values of $\cos \delta$ around $\cos \delta = -\frac{1}{2}$, regardless of the true mass hierarchy. If the IH solution becomes favoured, and a is found to be positive, these relations would instead predict $\cos \delta \approx \frac{1}{2}$. There is a wider range of predictions for $\cos \delta$ amongst the models with $\lambda \approx -\frac{1}{2}$, as shown in Fig. 5.3. This is due to the non-zero values of a_0 , which produce linear relations that don't pass through the origin at $a = \cos \delta = 0$, making predictions less tightly clustered. For the two sum rules with $a_0 = 0$, the situation is similar to that of the $\lambda \approx 1$ sum rules: two very closely lying sum rules which are




a) Normal Hierarchy




b) Inverted Hierarchy

Figure 5.2: The current experimental status of the sum rules given by $\lambda \approx 1$. The diagonal lines show the regions predicted for a and $\cos \delta$ given the 3σ bounds on r , assuming both normal ordering (top) and inverted ordering (bottom). The vertical line shows the current best-fit for a whilst the dark (light) grey regions show the 1σ (2σ) allowed intervals [57].



a) Normal Hierarchy



b) Inverted Hierarchy

Figure 5.3: The current experimental status of the sum rules given by $\lambda \approx -0.5$. The diagonal lines show the regions predicted for a and $\cos \delta$ given the 3σ bounds on r , assuming both normal ordering (top) and inverted ordering (bottom). The vertical line shows the current best-fit for a whilst the dark (light) grey regions show the 1σ (2σ) allowed intervals [57].

largely consistent with all values of $\cos \delta$. The values predicted by these relations for the current best fit value of a , are close to maximal $\cos \delta \approx 1$. For sum rules with $a_0 \neq 0$, the values of $\cos \delta$, assuming values of a inside the current 2σ interval, become more predictive. For example, the sum rule found with the group S_4 and the generators $T_\mu-S_2$ has only a small range of predictions consistent with current bounds: requiring IH and $\cos \delta \approx 1$. Although the range of predictions for $\lambda \approx -\frac{1}{2}$ is quite large, there are regions of parameter space which are only consistent with a few sum rules, and some broad additional constraints could significantly improve our knowledge of the viability of these relations. For example, recent global fits have provided some weak constraints on δ : Ref. [189] found a 1σ region roughly given by $-1 \leq \cos \delta \leq 0$, which broadens to the full $[-1, 1]$ range at 2σ . If this behaviour persists, it may be possible for the current generation of experiments to globally hint towards the exclusion of certain models. With NH and $\lambda \approx -\frac{1}{2}$, most models predict $\cos \delta > 0$ for the current best fit values of a : the only sum rule consistent with $\cos \delta < 0$ is given by the group S_4 and the generators $T_\tau-S_2$. To accommodate $\cos \delta < 0$. Another scenario when we could expect interesting constraints would be if no signs of CP violation were found in the near future, suggesting approximate CP conservation, $|\cos \delta| \approx 1$. In this case, most models would be in severe tension with the data; although, some of the $\lambda \approx -\frac{1}{2}$ could accommodate this very well. Broad constraints of this type could emerge in the current generation of experiments at the global level, or early on in the next generation. These constraints would suggest interesting patterns amongst the oscillation parameters, and restrict the likelihood of some of the models that we are considering; however, the possibility of constraining any of these models with sufficient statistical significance is small even at the global level.

All of the sum rules that we have considered in this section offer some range of predictions for $\cos \delta$ which is consistent with the known constraints from global fits. Therefore, to separate between competing predictions, we require improved knowledge of the parameter $\cos \delta$. Measuring the parameter $\cos \delta$ is one of the primary goals of the next generation of oscillation experiments, and for this reason, they are ideal experimental tools for investigating the presence of flavour symmetric

relations amongst the PMNS parameters. In the remainder of this chapter, we shall study how the next generation of oscillation experiments will be able to improve constraints on atmospheric sum rules.

5.2 Sum rules at next-generation facilities

With the measurement of the value of θ_{13} , the campaign for a next-generation facility, designed to make precision measurements of the neutrino mixing parameters, has been greatly strengthened. It is likely that within the extant experimental neutrino physics program, we will see hints towards the measurement of two of the most important unknowns in the conventional neutrino flavour-mixing paradigm: the neutrino mass hierarchy and the value of the CP-violating phase, δ . It is, however, unlikely that these questions will be resolved at an acceptable statistical confidence level: the projected 3σ CP-violation discovery fraction with the current experimental program only reaches around 20% of the parameter space [95] and it is only modestly higher for the determination of the mass hierarchy at around 40%. The desire for a definitive 5σ answer to these questions provides the first motivation for the construction of a next-generation neutrino oscillation facility. However, after these questions of discovery have been addressed, any new facility should have the ability to approach questions of precision. In this chapter, we will focus on two designs capable of making precise determinations of the oscillation parameters: the Low-Energy Neutrino Factory (LENF) and a wide-band superbeam (WBB), as discussed in Chapter 2.

In Section 5.2.1, we will present the details of our simulations and the facilities under consideration. Following this will be a brief discussion, in Section 5.2.2, on the effect of τ -contamination on precision measurements, before moving on to the results of our simulations: single parameter determination in Section 5.2.3, joint parameter determination in Section 5.2.4 and, finally, in Sections 5.2.5 and 5.2.6, the full simulation of the exclusion of sum rules, and possible constraints on the parameter λ .

5.2.1 Details of facilities and simulations

We are interested in knowing all of the physical parameters of the SM to a reasonable precision, one which would allow us to test our paradigm and rule out alternative explanations, but there is no limit to how small the influence of an interesting effect could be, and no clear cut definition of reasonable. A sensible goal is to aim for measurements of the parameters of the PMNS matrix at the same level of precision as the current measurements of the CKM matrix in the quark sector, around 5% [6]. However, the search for precision is seen as a secondary goal of the next-generation of experiments, which have been designed with discovery questions in mind. As such, we restrict our attention to the WBB and LENF, and how well they can address questions of precision.

The first facility we consider is the WBB with a baseline of around 2300 km. This is a technically well-understood medium-term project, which has been shown to offer a competitive discovery potential for CP violation and the mass hierarchy [119]. Its ability to make precision measurements has been shown to be good [119, 120, 173, 190], achieving a precision of δ of slightly higher than 12% [173]. Although this is not as precise as the current knowledge of the CKM matrix, it would be a significant step forward in our knowledge of the oscillation parameters. The standard detector for the WBB is taken to be a large underground LAr detector, and we shall consider this design at a mass scale of 35 kton and 70 kton, the lower bound given by the projected mass for an upgraded LBNE facility [121–123] and the upper bound given by the final design goal of the LBNO detector [112–115].

The neutrino factory is a more technically complicated proposal, but with this complexity comes a significantly improved discovery reach and attainable precision. It has been shown that a high-energy neutrino factory has excellent sensitivity to small non-standard effects in neutrino propagation [191–194], and this is expected to be an ideal facility for precision measurements in general. However, the discovery-based targets of the next-generation have favoured a lower energy facility: the LENF, which has excellent discovery potential as discussed in Chapter 3. The very well understood beam and low backgrounds of the LENF can achieve precisions on δ of around 5% [173] using a MIND. This marks the LENF as an excellent

choice for precision studies of the PMNS parameters. The possibility of using alternative detectors to the MIND, ones that offer improved energy resolution and event reconstruction capabilities, has not been thoroughly investigated in the context of precision, although some discussion of θ_{23} sensitivity can be found in Ref. [146]. In this chapter, we are not looking to make a detailed comparison of different designs, but instead to show the feasibility of constraining sum rules at next-generation facilities. As such, we have restricted our attention to two general variants of the LENF detector design: a 100 kton MIND and a magnetized 50 kton LAr detector. The first of these technologies provides us with a state of the art estimate of performance for the current proposals, whilst the second offers a more optimistic assessment of the potential of a LENF.

We have used the GLoBES package [195, 196] to perform our simulations of the LENF and WBB experiments. In our model of the LENF with MIND, we have incorporated all detector effects by the use of migration matrices. These have an absolute energy threshold of 100 MeV, and account for the dependence of the detection efficiency on the energy as calculated in dedicated detector simulations [197]. The backgrounds to the golden channel are from charge misidentified muons, neutral current events and incorrectly identified electrons. A 1% (10%) systematic error has been associated with the signal (background) normalisation. For the simulation of the LENF with LAr, we use similar parameters to our model in Chapter 3. We assume a 0.5 GeV threshold, with a flat 80% efficiency for higher energies. The energy resolution is assumed to be 5% for quasi-elastic events and 10% for non-quasi-elastic events. The backgrounds to the appearance channel are taken as 0.1% of the total charge misidentified muons and neutral current events. A 2% uncertainty is attached to both the signal and background normalisations. Both models of the LENF have a stored-muon energy of 10 GeV and a baseline distance of 2000 km. These have been shown to be near optimal choices for large θ_{13} [145, 151, 198]. Similar parameter choices have recently been recommended by the EUROnu Design Study [199], and coincide with the expected specifications of the International Design Study for the Neutrino Factory [200]. All of our simulations of the LENF design assume 10^{22} total useful muon decays divided equally between μ^- and μ^+ . Our model of the WBB de-

sign is based on Ref. [190], and assumes 10^{21} protons on target per year at 50 GeV, a baseline distance of 2300 km and a 70 kton or 35 kton liquid Argon detector similar to the GLACIER [107] design. The fluxes for this set-up are taken from Ref. [201] (for discussion see Ref. [202]). We have assumed a 90% detection efficiency and the backgrounds are taken as arising from a combination of the contamination of the beam and 0.5% of neutral-current events at the detector. The detector has a low-energy threshold of 100 MeV with an energy resolution taken to be a flat 150 MeV for electrons and $0.2\sqrt{E}$ for muons. An uncertainty of 5% has been imposed on the signal and background, and a 2% uncertainty on the matter density. In most of the subsequent simulations, we have neglected to include events arising from the τ -contamination channel. This channel is known to influence measurements at long-baseline experiments, and in the following section we shall consider the impact of this omission.

5.2.2 Impact of τ -contamination

Regardless of the flavour composition of the beam in a long-baseline oscillation experiment, thanks to the effects of neutrino oscillation, we can expect to observe neutrinos of all flavours impinging upon the detector. For the next generation of long-baseline facilities, the channel of most interest is the appearance channel (*e.g.* $\nu_e \rightarrow \nu_\mu$) or its time-conjugate (*e.g.* $\nu_\mu \rightarrow \nu_e$). NFs use detectors looking for the wrong-sign muon associated with a golden channel event, whilst WBBs use a beam which is mainly ν_μ and seek to record these events by measuring incident electrons. As with all experimental searches, these signals are confused by the presence of backgrounds: channels which lead to the production of muons and electrons which the detector cannot separate from the desired signal. One particularly interesting background is referred to as τ -contamination [152, 153]. Whilst most backgrounds are largely independent of the oscillation phenomena, τ -contamination relies crucially upon it. In this section, we consider the impact of the τ -background on precision measurements of the oscillation parameters.

τ -contamination is caused by the flux of ν_τ or $\bar{\nu}_\tau$ impinging upon the detector. Some of these particles will undergo charged current interactions and produce a τ -

lepton. This lepton is very unstable and decays with a lifetime of 10^{-17} s, frequently producing electrons or muons (with branching ratios of around 19% and 17%, respectively [6]). Left untreated, these additional leptons may be counted as appearance channel events and, when reconstructed by the detector algorithms, will be marked as coming from ν_μ neutrinos with spuriously low energies [152, 153, 155]. The extraction of the oscillation probabilities will be subsequently distorted, leading to serious systematic errors in the inferred oscillation parameters [155]. The magnitude of the effect of τ -contamination depends upon a number of factors, most importantly the size of the ν_τ component in the oscillated flux. Oscillations from ν_e into ν_τ are very rare and we will neglect this channel. The oscillations from ν_μ to ν_τ are, in contrast, almost maximal being driven by θ_{23} , and we expect a significant number of τ -leptons to arise from this channel. For the NF, the effect of τ -contamination is expected to be most pronounced for measurements of the disappearance channel ($\nu_\mu \rightarrow \nu_\mu$ or $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$), such as the determination of θ_{23} and Δm_{32}^2 . For measurements which rely upon the appearance channel ($\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ or $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$), the τ -background plays a far less significant role due to the additional CP conjugation: the muons arising from τ -decay are predominately right-sign, and only those τ -decay derived muons which, in addition, have their charges misidentified will contribute to the appearance channel background. For the WBB, the flux is predominately composed of a single flavour ν_μ (or $\bar{\nu}_\mu$, depending on horn polarity). The signal of the appearance channel is the presence of electrons in the detector, arising from the oscillation of the dominant ν_μ contribution to the flux into ν_e . This signal will be overlaid with a background arising from τ -contamination: the oscillation of the same ν_μ flux component into ν_τ with the subsequent τ -decay producing electrons. These τ -decays will also produce muons and, therefore, τ -contamination is expected to affect both appearance and disappearance measurements at the WBB.

The degree to which the τ -background can be controlled at a facility depends upon the detector technology being used. As leptons arising from τ -decay generally have a smaller energy than those from events with fewer decay products, they predominately affect the lowest-energy bins, and facilities with a higher threshold expect a smaller impact from this channel. Another factor is the degree of kine-

matic information available at the detector. The kinematics of a lepton which is produced by τ -decay are quite different compared to one coming directly from the charged-current neutrino–nucleon interaction, and this information can be used in a cut-based analysis to heavily reduce the τ -background. Experimental designs which have strong event reconstruction abilities, for example those based on liquid-argon time-projection chamber technology, will have a much better ability to control this background. For these reasons, estimating the impact of the τ -background requires detailed input from the experimental groups which study these technologies. Simulations of this kind have been performed for the MIND [197], but no such information is currently generally available for LAr detectors. In the following analysis, when making a comparison between facilities, we have omitted the τ -backgrounds as they cannot be consistently implemented. At the end of Section 5.2.5, after discussing precision measurements of the oscillation parameters, we will address this omission using the simulations of the MIND detector, and review the validity of this approach.

5.2.3 Precision for a , r and $\cos \delta$

We start our study by computing the precision with which the next-generation facilities can individually measure the parameters a , r and $\cos \delta$. An understanding of this precision should give us an indication of the potential precision towards generic sum rules in these variables and help us to identify the dominant uncertainties and functional dependence of such a measurement. In the following analysis, we will refer to the parameter values which are used to generate the simulated data as the *true values* and the parameters which are extracted by fitting our models to the data as the *fitted values*. When necessary, true and fitted values will be distinguished by subscripts *i.e.* $\Delta a \equiv a_F - a_T$. For each parameter of interest, we have scanned over a range of true values and then computed the allowed region (at 1, 3 and 5σ) in the fitted value of this parameter for both experimental set-ups. We marginalize over all of the otherwise unspecified oscillation parameters in each case. We then subtract off the true value from the fitted values to produce the allowed region expressed in terms of the permitted deviation from the true parameter value as a function of the true value itself.




Figure 5.4: The sensitivity of the next-generation facilities to the a , r and $\cos \delta$ parameters. In the leftmost plots, the empty (shaded) regions are for the LAr with MIND (LAr), whilst in the rightmost plots, the empty (shaded) regions are for the WBB with a 35 kton (70 kton) mass. In all of the plots the progressive regions show the 1σ , 3σ and 5σ confidence intervals.

The top row in Fig. 5.4 shows the sensitivity to a for both the LENF (leftmost panel, empty regions for MIND and solid for LAr) and the WBB (rightmost panel, empty regions for 35 kton and solid regions for 70 kton detectors). For large values of a_T , we find the magnitude of $\Delta a \equiv a_F - a_T$ to be between 0.005 and 0.015 at 3σ for the LENF, whilst the WBB has worse performance with a range of between 0.014 and 0.021. The attainable precision increases notably for both experiments around $|a_T| \lesssim 0.05$, where Δa can become potentially as high as 0.041 (0.112) for the LENF with LAr (LENF with MIND) and 0.095 (0.122) for the WBB with 35 kton LAr (70 kton LAr). This increase is due to the presence of a degeneracy; For a given value of a_T , we get two reasonably good solutions for the fit $a_F \approx \pm a_T$: a manifestation of the θ_{23} octant degeneracy [203]. This is not an exact degeneracy of the 3-neutrino oscillation probability, and the ambiguity only appears for the smallest deviations from θ_{23} -maximality. For all values of a_T , WBB performs worse than the LENF, and for both facilities, the optimistic detectors perform better than the more conservative ones. However, if we focus on the best-fit values for a given by recent global fits, at around $a = -0.12$ [189], the discrepancy between the four experimental designs considered here is small, with a difference of around ± 0.003 at 1σ , less than 3% of the best-fit value of a .

In the middle row of Fig. 5.4, we have computed the sensitivity of the LENF and WBB to the parameter r . Over the region of r_T that is phenomenologically interesting, this sensitivity is relatively constant at about 0.007 (0.025) for the LENF (WBB) at 3σ . There is a slight broadening of the allowed region towards larger values of r ; an effect which is less marked for weaker confidence levels. Once again, we see that LENF uniformly out-performs WBB. The discrepancy is particularly marked at 5σ where the WBB allowed region is around 3.5 times broader than the corresponding region for the LENF. In recent work on the precision of next-generation facilities, it has been shown [190] that only the LENF will be able to surpass the precision on θ_{13} that is expected to be attained by the current generation of reactor experiments. However, the improvement in precision possible with the LENF is rather small, at around 1%, and effectively the constraints on θ_{13} will be set by the reactor experiments alone. For this reason, the observed discrepancy

in precision for r between the LENF and WBB is only expected to influence the ability of the experiments to place individual constraints on sum rules, and should not influence constraints extracted from global analyses of the oscillation data.

The bottom row of Fig. 5.4 shows the expected sensitivity to $\cos \delta$ for the LENF and WBB. This measurement has a 3σ precision at its widest point of 0.28 (0.53) for the LENF with LAr (MIND) and 0.65 (0.89) for the WBB with 35 kton (70 kton) LAr. This decreases dramatically for the extreme points of the spectrum where the true value of $\cos \delta$ approaches ± 1 and the uncertainty becomes very small for the LENF, and is reduced but remains sizable at higher significances for WBB. We see that the LENF performs significantly better at this measurement than WBB: at 5σ , even the WBB with 70 kton LAr offers little discriminatory power, with a region that almost covers the whole parameter space, while the LENF offers a reasonable precision which becomes excellent for large values of $|\cos \delta|$. The boundaries of the allowed regions at low significance can be approximated analytically as ellipses: this can be seen by considering a uniform precision on δ itself, $\Delta\delta = \epsilon$, which implies $\Delta(\cos \delta) \equiv \cos \delta_F - \cos \delta_T = -\epsilon \sin \delta_T + \mathcal{O}(\epsilon^2)$. The coordinates $(-\epsilon \sin \delta, \cos \delta)$ provide a parametric description of the ellipse. The assumption of approximately uniform precision in δ is consistent with the simulations performed in Ref. [190] where $\Delta\delta \approx 5^\circ \pm 2^\circ$ for all δ_T . The deviations from ellipticity can be explained by assuming a variable precision on δ as shown in Ref. [190]. Generally, $\cos \delta$ is considerably harder to constrain than r and a . As such, it is expected to introduce a significant uncertainty and should be the dominant limiting factor in the possible constraints on the sum rules. However, we must remember that the measurements in this section have focused on a single parameter at a time, and therefore their results can not be simply combined to understand the precision on a sum rule. Measurements of parameter combinations will in general introduce correlations which may strongly influence the precision as we will see in the next section.

We can now return to address the impact of τ -contamination as discussed in Section 5.2.3. We have computed the sensitivities to a and $\cos \delta$ for the LENF with MIND, with and without the τ -background. The events arising from τ -decay have been implemented into our GLoBES simulation using a migration matrix provided




Figure 5.5: The precision on θ_{23} , shown via the combination $a \equiv \sqrt{2} \sin \theta_{23} - 1$, as a function of its true value for the LENF with MIND with τ -background taken into account (empty regions) and neglecting the τ -background (solid regions).

by Ref. [197]. In Fig. 5.5, we show the impact of the τ -contamination effect on the determination of a . This measurement is dominated by the disappearance channel, and for this reason, is expected to be one of the quantities that is most affected by the events arising from τ -decay, as discussed in Section 5.2.2. The impact of this background is significant, with the precision at large values of a dropping from around 0.015 to 0.010. In Fig. 5.6, we see the equivalent plots for sensitivity to $\cos \delta$. This measurement relies heavily on the measurement of the appearance channel, where the effect of τ -decays is less pronounced, and accordingly, we see very little change in the allowed regions, except at the highest significance.

Correctly incorporating the τ -background requires knowledge of the detector capabilities, aiding by dedicated simulations. However, if we extrapolate the results found for the LENF with MIND to the LENF in general: as the dominant uncertainty in the single parameter determination is given by the measurements of $\cos \delta$, we expect the τ -background to only marginally influence the possible constraints on any given sum rule at a neutrino factory. For the WBB, the τ -decays can also influence the appearance channel, and extrapolating our results is less reliable, but we would expect a worsening of the performance of $\cos \delta$ as well as a . However, the




Figure 5.6: The precision on $\cos \delta$ as a function of its true value for the LENF with MIND with τ -background taken into account (empty regions) and neglecting the τ -background (solid regions). We see little impact from the τ -backgrounds on this measurement as it is based on the appearance channel, which is less affected than the disappearance channel.

additional kinematic information accessible with a LAr detector can help identify τ -events, and may be able to mitigate some of this loss of sensitivity. As we don't have access to information on the abilities of these detectors, for our study of sum rules in Section 5.2.5, we will restrict our attention to detector models without the τ -background, which should allow for a fair comparison between facilities, and a good estimate of their performance.

5.2.4 Joint determination of parameters

To determine more than one parameter from a data set, the effect of correlations must be taken into account. In general, we expect the ability to determine a given parameter to depend upon the values of the other parameters. In the analysis of the previous section, this correlation was accounted for by marginalisation: to calculate how well data generated with the parameter a_T can be explained by a hypothesised parameter set with fixed a_F , all hypothesised parameter sets are considered which hold a_F constant, and only the minimum $\Delta\chi^2$ is returned. This guarantees the

existence of a parameter set with $a = a_F$ that fits the data well, but ignores the values of the other parameters in that set. However, if we choose to fix another hypothesised parameter, $\cos \delta_F$, the possible minima will change. To compute how the precision for one parameter can vary as a function of another, we must compute joint parameter determinations, which incorporate these correlations. In Fig. 5.7, we have computed the joint determination of the parameters $\cos \delta$ and a for a selection of sets of true parameters which obey the sum rule $a = r \cos \delta$. This simulation uses the LENF with MIND experiment, and incorporates the τ -background which, as we have shown in Section 5.2.2, is known to impact the attainable precision on a . This plot gives us an indication of the severity of correlations between these two parameters. We see that there is some correlation: the allowed intervals for $\cos \delta$ have some dependence on the values of a . The width of the allowed regions in both parameters decreases for large values, and this behaviour can be understood by comparing it with the results of Section 5.2.3, where the precision to both a and $\cos \delta$ becomes worse near the origin.

In these plots we have assumed that the true parameters obey the sum rule $a = r \cos \delta$, and we have marginalised over all parameters other than a_F and $\cos \delta_F$. These plots, therefore, show how well the parameters could be constrained if the sum rule were true. However, the solutions found in the allowed regions are not required to obey the sum rule. For example, although there are plenty of solutions around the origin for $a_T = \cos \delta_T = 0$, the parameter r_F is allowed to vary in the marginalisation and can take any reasonable value. If we are interested in excluding the sum rule without assuming its validity, we must ask a slightly different question: for a general set of true parameter values, which sets of parameters obeying a hypothesised sum rule can be excluded. We will address this question in the next section.

5.2.5 Measuring sum rule violations

In this section, we compute the ability of the LENF and WBB experiments to directly constrain and exclude the linearized sum rules discussed in Section 5.1, whilst fully incorporating the effects of parameter correlations.

To do this, we have scanned over a parameter space spanned by the true value




Figure 5.7: The allowed regions for the joint parameter determination of $\cos \delta$ and a for LENF the MIND detector, incorporating τ -contamination. The true set of parameters for each set of concentric contours lies at their centre on the dashed line $a = r \cos \delta$.

of $\cos \delta$ and the true value of a . At each point in this parameter space, we have found the best fitting set of oscillation parameters which obey a given sum rule, and plotted the corresponding value of $\Delta\chi^2$. Once this value exceeds a chosen significance threshold, we can consider that sum rule excluded: there are no sets of parameters which obey that sum rule and provide a reasonable fit to the data. When the true parameter set approximately obeys the sum rule in question, we get a good fit, and the width of the surrounding allowed region gives an indication of how sensitive the experiment is to deviations from the sum rule. Technically, this search has been implemented by using a modified form of the $\Delta\chi^2$ statistic. We have extended the $\Delta\chi^2$ to include an additional prior which enforces the sum rule on our set of hypothesised parameters,

$$\Delta\chi^2 \supset \left(\frac{a - a_0 - \lambda r \cos \delta}{\sigma} \right)^2,$$

where σ is a parameter chosen to be small, ensuring that the sum rule is held to high precision. This term forces the minimal parameter set to obey the sum rule, whilst not dictating any of the values of the parameters themselves.

We have focused our analysis on the two simplest sum rules $\lambda = 1$ and $\lambda = -\frac{1}{2}$




Figure 5.8: The left (right) plot shows the ability to exclude models with $\lambda = 1$ ($\lambda = -0.5$) as a function of the true parameters. The blue solid regions show the 2σ and 3σ regions for the LENF with LAr, whilst the empty regions show the equivalent regions for the LENF with MIND.

both with $a_0 = 0$. This is to illustrate the type of constraints that can be placed on parameter correlations in the PMNS matrix, and our approach can be easily generalised to include other types of correlations, beyond the atmospheric sum rules discussed so far. The coloured regions of the left-hand (right-hand) panel of Fig. 5.8 show the allowed regions for $\lambda = 1$ ($\lambda = -0.5$) for the LENF with MIND and LAr detectors. We see that the largest allowed region, and therefore the hardest point to exclude the sum rule, is when $\cos \delta \approx a \approx 0$. Whilst the best sensitivity is generally found at large values of $|\cos \delta|$. As expected, this behaviour is largely inherited from the sensitivity to $\cos \delta$; however, around the origin we see a novel feature associated with solutions of the type $a = 0$ and $\cos \delta = 0$. For any hypothetical sum rule of the type $a = \lambda r \cos \delta$, a trivial solution can be found for small $a = \cos \delta = 0$. At this point, the ability to constrain both a and $\cos \delta$ weaken, and we find that regardless of the relationship between the true parameters, provided they are sufficiently close to the origin, we can use this solution to describe the data and satisfy the sum rule. The LAr detector allows for the sum rule to be excluded over a larger region of parameter space: the 2σ allowed region for the LAr is contained completely inside the 2σ region for the MIND detector. At the widest points, the allowed regions for $\cos \delta$ cover around 24% (42%) of the parameter space for $\cos \delta$ for the LENF with




Figure 5.9: The left (right) plot shows the ability to exclude models with $\lambda = 1$ ($\lambda = -0.5$) as a function of the true parameters. The blue solid regions show the 2σ and 3σ regions for the WBB with 70 kt detector, whilst the empty regions show the equivalent regions for the WBB with 35 kt detector.

LAr (MIND) at 3σ . In Fig. 5.9, we show the equivalent regions for the WBB. These follow the same shape, inherited from the uncertainties in measurement of $\cos \delta$. The WBB is unable to constrain the parameter $\cos \delta$ to the same extent as the LENF, and we see that the allowed region for the sum rules are correspondingly much larger. At its widest point, the WBB with 70 kton (35 kton) LAr has an allowed region for $\cos \delta$ which covers 56% (81%) of the parameter space at 3σ . For both LENF and WBB, excluding models over even 50% of the parameter space would be an interesting result; however, we have seen that these measurements are challenging, and the more optimistic facilities are required to make significant advances.

5.2.6 Constraining λ

Both LENF and WBB will be able to observe violations of a given sum rule for a significant fraction of parameter space, especially if $|\cos \delta|$ is large. In the scenario that the true parameter set appears to agree with some sum rule, it is interesting to see what constraints we can put on the parameters describing such a rule. In this section, we consider the ability of the next-generation oscillation experiments to distinguish between models with similar λ and a_0 parameters, introduced in Eq. 5.0.1. Our interest here is in illustrating the possible constraints that can be




Figure 5.10: The fraction of values of δ for which the hypothesised value of λ_F can be excluded assuming different true values of λ_T . In these plots $\Delta\lambda = \lambda_F - \lambda_T$.

posed by a next-generation oscillation experiment, and as such we will restrict our attention to some specific cases; however, the analysis of this section could be simply extended to address other classes of models.

In the models that have been either reported in the literature or identified in Chapter 4, most values of λ are close to $\lambda = 1$ or $\lambda = -0.5$. In this section, we consider general relations of the type $a = \lambda r \cos \delta$, with continuous ranges of λ in the neighbourhoods of these special values. In the left panel of Fig. 5.10 we show how well the hypothesis $\lambda_F = -1/2$ can be excluded as a function of $\Delta\lambda \equiv \lambda_T - \lambda_F$ for sum rules with $a_0 = 0$. The lines of parameters which obey sum rules of this type intersect at the origin, and we will always be able to find true parameter values close to the origin which satisfy any pair of sum rules. Therefore, it is impossible to distinguish two similar models in all possible cases, and instead we must assess this ability by degree. In Fig. 5.10, to measure the degree of distinguishability at different facilities, we have plotted a continuous parameter which gives the fraction of values of δ_T for which we can exclude the hypothesis $\lambda = \lambda_F$. The corresponding fraction of distinguishability for the hypothesis $\lambda_F = 1$ as a function of $\Delta\lambda$ is shown in the right panel of Fig. 5.10. If we choose our threshold to be 50%, the LENF with LAr can distinguish between sum rules of the type $\lambda \approx -1/2$ which deviate by $|\Delta\lambda| \approx 0.2$. If we instead use a MIND, this region increases to $|\Delta\lambda| \approx 0.3$, whilst the WBB superbeam with a detector of 100 kton is closer to $|\Delta\lambda| \approx 0.5$. For sum rules with $\lambda \approx 1$ the size of these deviations approximately doubles.

For the models presented in Chapter 4, which cluster around $\lambda = 1$ or $\lambda = -0.5$, the values of λ differ by around ± 0.1 . The ability to separate these candidate models experimentally is clearly dependent on the true value of $\cos \delta$; however, the LENF with LAr can make this discrimination for about 25% of the values δ . This will be a very challenging measurement and is unlikely to be feasible in the next-generation of oscillation experiments unless an aggressive strategy is adopted.

5.3 Conclusions

Next-generation neutrino oscillation facilities are not only necessary to resolve the traditional questions about the PMNS matrix, but will also lead the way in a new programme of precision neutrino flavour physics. Over the years, many attempts have been made to understand the origin of flavour in the Standard Model. One popular approach is to invoke a symmetry to explain the pattern of mixing angles that have been discovered experimentally in the PMNS matrix: an idea which has met with great success and generated a large number of candidate models. Thanks to the precision that is expected at the next-generation oscillation facilities, it will soon be possible to put these theories to the test.

A predictive model of flavour will generally introduce correlations amongst the parameters of the Yukawa sector. Expressions of these correlations are called sum rules, and testing them is a direct way to confirm or exclude a given model. In this chapter, we have studied how these correlations will be constrained by current and future oscillation experiments. We have seen that, when viewed as predictions for $\cos \delta$, these sum rules are constrained by their consistency with the current data, and although all of the models that we have investigated have some region of applicability, some models require deviations from the current best-fits which may be excluded in the near future. The major difficulty in constraining the sum rules found in Chapter 4, is the absence of information on the parameter $\cos \delta$, and we must look to the next generation of oscillation experiments to provide this. We have studied the ability of two candidate next-generation neutrino oscillation experiments, a low-energy neutrino factory and a wide-band superbeam, to constrain

these correlations. To illustrate the general constraints that these experiments can place on flavour effects, we have chosen to focus our attention on sum rules with the linearized form $a = a_0 + \lambda r \cos \delta$, and specifically on the choices $\lambda = 1$ and $\lambda = -0.5$. These have arisen previously in the literature, and we have shown in Chapter 4 that these two special values appear to well characterize a large class of models. We have seen that violations of these sum rules will be readily testable at the LBNF and WBB: for example, a LBNF with LAr is expected to be able to exclude the relation $a = r \cos \delta$ for over 75% of the parameter space. We have also considered the ability to distinguish between models which predict similar sum rules with separations in λ of only around 0.1. We have found that this ability is dependent on the exact value of $\cos \delta$; however, it is likely that only the LBNF with LAr is precise enough to make such a distinction at a reasonable statistical significance for 25% of the parameter space.

We have shown that correlations amongst the parameters of the PMNS matrix present a well motivated goal for next-generation oscillation facilities. These correlations can be measured, and constraints can be inferred on the underlying models which generate them. However, these are difficult measurements, requiring optimistic experimental parameters, and must be pursued aggressively if we hope to make progress.

Summary and outlook

From their very first measurement, neutrinos have displayed a number of unexpected properties which have made neutrino physics a fast-paced and exciting field. The flavour transitions which neutrinos are believed to undergo during propagation are evidence of small mass terms and non-trivial mixing between the mass and flavour bases. Understanding the mechanisms for neutrino oscillation, in vacuo and in background media, has allowed these mass terms to be explored experimentally, and much is now known. However, the incorporation of neutrino masses into a new standard model (ν SM) and the explanation of their flavour structure are problems yet to be resolved. At a theoretical level, many mechanisms are known to address these issues; however, discovering the kind of solution which manifests itself in nature will require the continued experimental investigation of the neutrino sector.

In this thesis we have studied how the future programme of neutrino oscillation experiments can be used to further our understanding of the leptonic sector. These experiments will be necessarily over long distances, aiming to resolve the mass hierarchy and constrain the CP-violating phase δ . As we have shown in Chapter 3, by exploiting the hierarchy-discriminating matter effects and the CP sensitive low-energy oscillation spectrum, there are excellent prospects for a Low-Energy Neutrino Factory (LENF) to make these measurements; indeed, the LENS has since been adopted as the standard neutrino factory by the International Design Study for the Neutrino Factory [204].

Once the questions of discovery are settled, measurements of the neutrino oscillation probabilities can be used as a unique tool in probing the leptonic mixing parameters at high precision. This will be able to provide essential phenomenological data on the flavour properties of leptons, and by extension, the flavour structure

of the SM. As an example of the kind of questions that can be addressed by a precision neutrino oscillation facility, we have studied a class of models which attempt to explain neutrino flavour by the application of leptonic non-Abelian discrete symmetries. In Chapter 4 we have shown that models of this type introduce parameter correlations which can be condensed into sum rules, and our simulations in Chapter 5 have shown that the investigation of these at next generation facilities is an achievable goal. One which provides a solid motivation for increased precision in oscillation experiments.

The next generation of experiments probing leptonic flavour will be the first to really start testing the paradigm of the ν SM. Progress will only be possible with the aid of a detailed phenomenological understanding of the models which might upset these predictions. Hopefully, the work presented in this thesis helps us to advance in this direction.

Bibliography

- [1] M. Gell-Mann, Phys.Rev. **125**, 1067 (1962)
- [2] Y. Ne'eman, Nucl.Phys. **26**, 222 (1961)
- [3] M. Gell-Mann, Phys.Lett. **8**, 214 (1964)
- [4] G. Zweig, CERN-TH-401 (1964)
- [5] G. Zweig, CERN-TH-412, NP-14146 (1964)
- [6] J. Beringer *et al.* (Particle Data Group), Phys.Rev. **D86**, 010001 (2012)
- [7] W. Pauli, Phys.Today **31N9**, 27 (1978)
- [8] E. Fermi, Nuovo Cim. **11**, 1 (1934)
- [9] P. W. Higgs, Phys.Lett. **12**, 132 (1964)
- [10] P. W. Higgs, Phys.Rev.Lett. **13**, 508 (1964)
- [11] G. Guralnik, C. Hagen, and T. Kibble, Phys.Rev.Lett. **13**, 585 (1964)
- [12] F. Englert and R. Brout, Phys.Rev.Lett. **13**, 321 (1964)
- [13] S. Glashow, Nucl.Phys. **22**, 579 (1961)
- [14] S. Weinberg, Phys.Rev.Lett. **19**, 1264 (1967)
- [15] A. Salam, Conf.Proc. **C680519**, 367 (1968)
- [16] N. Cabibbo, Phys.Rev.Lett. **10**, 531 (1963)
- [17] M. Kobayashi and T. Maskawa, Prog.Theor.Phys. **49**, 652 (1973)

- [18] S. Schael *et al.* (ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SLD Collaboration, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group), Phys.Rept. **427**, 257 (2006), arXiv:hep-ex/0509008
- [19] M. L. Perl *et al.*, Phys.Rev.Lett. **35**, 1489 (1975)
- [20] G. Arnison *et al.* (UA1 Collaboration), Phys.Lett. **B122**, 103 (1983)
- [21] M. Banner *et al.* (UA2 Collaboration), Phys.Lett. **B122**, 476 (1983)
- [22] G. Arnison *et al.* (UA1 Collaboration), Phys.Lett. **B126**, 398 (1983)
- [23] P. Bagnaia *et al.* (UA2 Collaboration), Phys.Lett. **B129**, 130 (1983)
- [24] F. Abe *et al.* (CDF Collaboration), Phys.Rev.Lett. **74**, 2626 (1995), arXiv:hep-ex/9503002
- [25] S. Abachi *et al.* (D0 Collaboration), Phys.Rev.Lett. **74**, 2632 (1995), arXiv:hep-ex/9503003
- [26] G. Aad *et al.* (ATLAS Collaboration), Phys.Lett. **B716**, 1 (2012), arXiv:1207.7214
- [27] S. Chatrchyan *et al.* (CMS Collaboration), Phys.Lett. **B710**, 26 (2012), arXiv:1202.1488
- [28] C. Cowan, F. Reines, F. Harrison, H. Kruse, and A. McGuire, Science **124**, 103 (1956)
- [29] F. Reines, C. Cowan, F. Harrison, A. McGuire, and H. Kruse, Phys.Rev. **117**, 159 (1960)
- [30] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, Phys.Rev.Lett. **20**, 1205 (1968)
- [31] H. Fritzsch and P. Minkowski, Phys.Lett. **B62**, 72 (1976)
- [32] B. Pontecorvo, Sov.Phys.JETP **6**, 429 (1957)

- [33] B. Pontecorvo, Sov.Phys.JETP **7**, 172 (1958)
- [34] B. Pontecorvo, Sov.Phys.JETP **26**, 984 (1968)
- [35] V. Gribov and B. Pontecorvo, Phys.Lett. **B28**, 493 (1969)
- [36] Z. Maki, M. Nakagawa, and S. Sakata, Prog.Theor.Phys. **28**, 870 (1962)
- [37] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), Phys.Rev.Lett. **82**, 2644 (1999), arXiv:hep-ex/9812014
- [38] E. Fermi, Z.Phys. **88**, 161 (1934)
- [39] R. Robertson and D. Knapp, Ann.Rev.Nucl.Part.Sci. **38**, 185 (1988)
- [40] C. Weinheimer, Nucl.Phys.Proc.Suppl. **118**, 279 (2003)
- [41] G. Drexlin, V. Hannen, S. Mertens, and C. Weinheimer, Adv.High Energy Phys. **2013**, 293986 (2013), arXiv:1307.0101
- [42] C. Aalseth *et al.*(2004), arXiv:hep-ph/0412300
- [43] N. Angelov *et al.*, Nucl.Phys. **A780**, 78 (2006), arXiv:nucl-ex/0605002
- [44] P. Ade *et al.* (Planck Collaboration)(2013), arXiv:1303.5076
- [45] C. Kraus *et al.*, Eur.Phys.J. **C40**, 447 (2005), arXiv:hep-ex/0412056
- [46] V. Aseev *et al.* (Troitsk Collaboration), Phys.Rev. **D84**, 112003 (2011), arXiv:1108.5034
- [47] S. Eliezer and A. R. Swift, Nucl.Phys. **B105**, 45 (1976)
- [48] S. M. Bilenky and B. Pontecorvo, Phys.Rept. **41**, 225 (1978)
- [49] S. M. Bilenky and B. Pontecorvo, Lett.Nuovo Cim. **17**, 569 (1976)
- [50] S. Nussinov, Phys.Lett. **B63**, 201 (1976)
- [51] B. Kayser, Phys.Rev. **D24**, 110 (1981)
- [52] C. Giunti, C. Kim, and U. Lee, Phys.Rev. **D44**, 3635 (1991)

- [53] J. Rich, Phys.Rev. **D48**, 4318 (1993)
- [54] C. Y. Cardall, Phys.Rev. **D61**, 073006 (2000), arXiv:hep-ph/9909332
- [55] A. Dolgov, Phys.Rept. **370**, 333 (2002), arXiv:hep-ph/0202122
- [56] E. K. Akhmedov and J. Kopp, JHEP **1004**, 008 (2010), arXiv:1001.4815
- [57] M. C. Gonzalez-Garcia, M. Maltoni, J. Salvado, and T. Schwetz, JHEP **1212**, 123 (2012), arXiv:1209.3023
- [58] L. Wolfenstein, Phys.Rev. **D17**, 2369 (1978)
- [59] S. Mikheev and A. Y. Smirnov, Sov.J.Nucl.Phys. **42**, 913 (1985)
- [60] S. J. Parke, Phys.Rev.Lett. **57**, 1275 (1986)
- [61] V. D. Barger, K. Whisnant, S. Pakvasa, and R. Phillips, Phys.Rev. **D22**, 2718 (1980)
- [62] M. Freund, M. Lindner, S. Petcov, and A. Romanino, Nucl.Phys. **B578**, 27 (2000)
- [63] H. Bethe, Phys.Rev.Lett. **56**, 1305 (1986)
- [64] S. J. Parke and T. P. Walker, Phys.Rev.Lett. **57**, 2322 (1986)
- [65] K. Abe *et al.* (T2K Collaboration), Phys.Rev. **D85**, 031103 (2012), arXiv:1201.1386
- [66] F. P. An *et al.* (Daya-Bay Collaboration), Phys.Rev.Lett. **108**, 171803 (2012), arXiv:1203.1669
- [67] A. Cervera *et al.*, Nucl.Phys. **B579**, 17 (2000)
- [68] K. Asano and H. Minakata, JHEP **1106**, 022 (2011), arXiv:1103.4387
- [69] S. K. Agarwalla, Y. Kao, and T. Takeuchi(2013), arXiv:1302.6773
- [70] S. Weinberg, Phys.Rev.Lett. **43**, 1566 (1979)

- [71] S. Weinberg, Phys.Rev. **D22**, 1694 (1980)
- [72] P. Minkowski, Phys.Lett. **B67**, 421 (1977)
- [73] R. Foot, H. Lew, X. He, and G. C. Joshi, Z.Phys. **C44**, 441 (1989)
- [74] P. W. Angel, N. L. Rodd, and R. R. Volkas, Phys.Rev. **D87**, 073007 (2012), arXiv:1212.6111
- [75] M. Magg and C. Wetterich, Phys.Lett. **B94**, 61 (1980)
- [76] J. Schechter and J. Valle, Phys.Rev. **D22**, 2227 (1980)
- [77] M. Schwartz, Phys.Rev.Lett. **4**, 306 (1960)
- [78] B. Pontecorvo, Zh.Eksp.Teor.Fiz. **39**, 1166 (1960)
- [79] S. E. Kopp, Phys.Rept. **439**, 101 (2007), arXiv:physics/0609129
- [80] J.-M. Levy(2010), arXiv:1005.0574
- [81] Y. Itow *et al.* (T2K Collaboration), 239(2001), arXiv:hep-ex/0106019
- [82] K. Abe *et al.* (T2K Collaboration), Nucl.Instrum.Meth. **A659**, 106 (2011), arXiv:1106.1238
- [83] D. Ayres *et al.* (NO ν A Collaboration)(2004), arXiv:hep-ex/0503053
- [84] K. T. McDonald(2001), arXiv:hep-ex/0111033
- [85] K. Abe *et al.* (T2K Collaboration), Phys.Rev.Lett. **107**, 041801 (2011), arXiv:1106.2822
- [86] J. K. Ahn *et al.* (RENO collaboration), Phys.Rev.Lett. **108**, 191802 (2012), arXiv:1204.0626
- [87] S. K. Agarwalla, S. Prakash, S. K. Raut, and S. U. Sankar, JHEP **1212**, 075 (2012), arXiv:1208.3644
- [88] M. Diwan *et al.*, Phys.Rev. **D68**, 012002 (2003)

- [89] P. Huber and J. Kopp, JHEP **1103**, 013 (2011)
- [90] B. Richter(2000), hep-ph/0008222
- [91] V. D. Barger, S. Geer, R. Raja, and K. Whisnant, Phys.Rev. **D63**, 113011 (2001), arXiv:hep-ph/0012017
- [92] J. J. Gomez-Cadenas *et al.* (CERN working group on Super Beams), 463(2001), arXiv:hep-ph/0105297
- [93] H. Minakata and H. Nunokawa, JHEP **0110**, 001 (2001)
- [94] P. Huber, M. Lindner, and W. Winter, Nucl.Phys. **B645**, 3 (2002)
- [95] P. Huber, M. Lindner, T. Schwetz, and W. Winter, JHEP **0911**, 044 (2009), arXiv:0907.1896
- [96] H. Minakata and H. Sugiyama, Phys.Lett. **B580**, 216 (2004), arXiv:hep-ph/0309323
- [97] A. de Bellefon *et al.*(2006), arXiv:hep-ex/0607026
- [98] J. Ahrens *et al.* (IceCube Collaboration), Nucl.Phys.Proc.Suppl. **118**, 388 (2003), arXiv:astro-ph/0209556
- [99] S. L. Adler, Annals Phys. **50**, 189 (1968)
- [100] D. Rein and L. M. Sehgal, Annals Phys. **133**, 79 (1981)
- [101] J. Peltoniemi(2009), arXiv:0911.4876
- [102] T. Marrodan Undagoitia *et al.*, J.Phys.Conf.Ser. **120**, 052018 (2008)
- [103] M. Wurm *et al.* (LENA Collaboration), Astropart.Phys. **35**, 685 (2012), arXiv:1104.5620
- [104] C. Rubbia, CERN-EP-INT-77-08 (1977)
- [105] P. Cennini *et al.* (ICARUS Collaboration), LNGS-94/99 (1994)
- [106] A. Rubbia(2004), arXiv:hep-ph/0402110

- [107] A. Rubbia, J.Phys.Conf.Ser. **171**, 012020 (2009)
- [108] A. Mereaglia and A. Rubbia(2008), arXiv:0801.4035
- [109] A. Badertscher *et al.*(2008), arXiv:0804.2111
- [110] D. B. Cline, 311(2005), arXiv:astro-ph/0506546
- [111] P. Huber and T. Schwetz, Phys.Lett. **B669**, 294 (2008)
- [112] D. Autiero *et al.*, JCAP **0711**, 011 (2007), arXiv:0705.0116
- [113] D. Angus *et al.* (LAGUNA Collaboration)(2010), arXiv:1001.0077
- [114] A. Rubbia(2010), arXiv:1003.1921
- [115] A. Rubbia (LAGUNA Collaboration), Acta Phys.Polon. **B41**, 1727 (2010)
- [116] S. K. Raut, R. S. Singh, and S. Uma Sankar, Phys.Lett. **B696**, 227 (2011)
- [117] A. Joglekar, S. Prakash, S. K. Raut, and S. Uma Sankar, Mod.Phys.Lett. **A26**, 2051 (2011)
- [118] S. K. Agarwalla, T. Li, and A. Rubbia, JHEP **1205**, 154 (2012), arXiv:1109.6526
- [119] P. Coloma, T. Li, and S. Pascoli(2012), arXiv:1206.4038
- [120] P. Coloma, E. Fernandez-Martinez, and L. Labarga, JHEP **1211**, 069 (2012), arXiv:1206.0475
- [121] T. Akiri *et al.* (LBNE Collaboration)(2011), arXiv:1110.6249
- [122] LBNE Collaboration, *Conceptual Design Report* (2012)
- [123] LBNE Collaboration, *Conceptual Design Report (prior to Oct. 2012)* (2011)
- [124] LBNE Steering Committee, “LBNE Reconfiguration,” http://www.fnal.gov/directorate/lbne_reconfiguration/
- [125] D. Kelliher and J. Pasternak, “Decay Ring Status,” Talk presented during IDS-NF 10 at RAL, 6/4/2013.

- [126] S. Geer, Phys.Rev. **D57**, 6989 (1998)
- [127] A. De Rújula, M. B. Gavela, and P. Hernández, Nucl.Phys. **B547**, 21 (1999)
- [128] A. Bandyopadhyay *et al.* (ISS Physics Working Group), Rept.Prog.Phys. **72**, 106201 (2009)
- [129] S. Choubey *et al.*, “International Design Study for the Neutrino Factory, Interim Design Report,” (2011), IDS-NF-020
- [130] A. Broncano and O. Mena, Eur.Phys.J. **C29**, 197 (2003)
- [131] A. Cervera, A. Laing, J. Martin-Albo, and F. J. P. Soler, Nucl.Instrum.Meth. **A624**, 601 (2010)
- [132] A. Laing, Ph.D. thesis, University of Glasgow (2010)
- [133] I. Ambats *et al.* (MINOS Collaboration), NUMI-L-337, FERMILAB-DESIGN-1998-02(1998)
- [134] R. Bayes, “Neutrino Factory Detector Status,” Talk presented during IDS-NF 10 at RAL, 6/4/2013.
- [135] G. Acquistapace *et al.* (CMS Collaboration), CERN-LHCC-97-10(1997)
- [136] A. D. Bross, M. Ellis, S. Geer, O. Mena, and S. Pascoli, Phys.Rev. **D77**, 093012 (2008)
- [137] S. Geer, O. Mena, and S. Pascoli, Phys.Rev. **D75**, 093001 (2007)
- [138] D. Drakoulakos *et al.* (Minerva Collaboration) arXiv:hep-ex/0405002
- [139] NO ν A Collaboration, *Technical Design Report* (2007)
- [140] P. S. Baringer *et al.* (D0 Collaboration), Nucl.Instrum.Meth. **A469**, 295 (2001), arXiv:hep-ex/0007026
- [141] G. Ambrosio *et al.* (VLHC Design Study Group), SLAC-R-591, SLAC-R-0591, SLAC-591, SLAC-0591, FERMILAB-TM-2149(2001)

-
- [142] P. Huber and W. Winter, Phys.Rev. **D68**, 037301 (2003)
- [143] A. Smirnov(2006), arXiv:hep-ph/0610198
- [144] P. Huber, M. Lindner, M. Rolinec, and W. Winter, Phys.Rev. **D74**, 073003 (2006)
- [145] S. K. Agarwalla, P. Huber, J. Tang, and W. Winter, JHEP **1101**, 120 (2011)
- [146] E. Fernández Martínez, T. Li, S. Pascoli, and O. Mena, Phys.Rev. **D81**, 073010 (2010)
- [147] B. Pontecorvo, Zh.Eksp.Teor.Fiz. **33**, 549 (1957)
- [148] B. Pontecorvo, Zh.Eksp.Teor.Fiz. **34**, 247 (1958)
- [149] C. Ankenbrandt *et al.*, Phys.Rev.ST Accel.Beams **12**, 070101 (2009)
- [150] T. Li, Ph.D. thesis, Durham University (2010)
- [151] A. Dighe, S. Goswami, and S. Ray, Phys.Rev. **D86**, 073001 (2012), arXiv:1110.3289
- [152] D. Indumathi and N. Sinha, Phys.Rev. **D80**, 113012 (2009)
- [153] R. Dutta, D. Indumathi, and N. Sinha(2011), arXiv:1103.5578
- [154] K. Nakamura *et al.* (Particle Data Group), J.Phys.G **37**, 075021 (2010)
- [155] A. Donini, J. J. Gómez Cadenas, and D. Meloni, JHEP **1102**, 095 (2011)
- [156] P. Huber, M. Lindner, and W. Winter, Comput.Phys.Commun. **167**, 195 (2005)
- [157] P. Huber, J. Kopp, M. Lindner, M. Rolinec, and W. Winter, Comput.Phys.Commun. **177**, 432 (2007)
- [158] A. M. Dziewonski and D. L. Anderson, Phys.Earth Planet.Inter. **25**, 297 (1981)
- [159] F. D. Stacey, *Physics of the Earth*, 2nd ed. (Wiley, 1977)

- [160] V. Barger, P. Huber, D. Marfatia, and W. Winter, Phys.Rev. **D76**, 053005 (2007), arXiv:hep-ph/0703029
- [161] L. Bartoszek *et al.*(2004), arXiv:hep-ex/0408121
- [162] D. B. Cline, F. Raffaelli, and F. Sergiampietri, JINST **1**, T09001 (2006)
- [163] B. Baibussinov *et al.*, Astropart.Phys. **29**, 174 (2008)
- [164] M. C. Gonzalez Garcia, M. Maltoni, and J. Salvado, JHEP **1004**, 056 (2010)
- [165] T. Schwetz, M. Tórtola, and J. W. F. Valle, New J. Phys. **13**, 063004 (2011)
- [166] S. Baker and R. D. Cousins, Nucl.Instrum.Meth. **221**, 437 (1984)
- [167] S. S. Wilks, Ann.Math.Stat. **9**, pp. 60 (1938)
- [168] P. Lipari, M. Lusignoli, and F. Sartogo, Phys.Rev.Lett. **74**, 4384 (1995)
- [169] P. Lipari, Nucl.Phys.Proc.Suppl. **112**, 274 (2002)
- [170] A. Dighe, S. Goswami, and S. Ray, Phys.Rev.Lett. **105**, 261802 (2010)
- [171] P. Adamson *et al.* (MINOS Collaboration), Phys.Rev.Lett. **107**, 181802 (2011)
- [172] H. de Kerret (Double Chooz Collaboration), “The first results from the Double Chooz experiment,” Talk presented during LowNu11 at Seoul, 9/11/2011.
- [173] P. Coloma, P. Huber, J. Kopp, and W. Winter, Phys.Rev **D87**, 033004 (2012), arXiv:1209.5973
- [174] G. Altarelli and F. Feruglio, Rev.Mod.Phys. **82**, 2701 (2010), arXiv:1002.0211
- [175] H. Ishimori *et al.*, Prog.Theor.Phys.Suppl. **183**, 1 (2010), arXiv:1003.3552
- [176] S. F. King and C. Luhn, Rept.Prog.Phys. **76**, 056201 (2013), arXiv:1301.1340
- [177] P. F. Harrison, D. H. Perkins, and W. G. Scott, Phys.Lett. **B530**, 167 (2002)
- [178] A. Datta, F.-S. Ling, and P. Ramond, Nucl.Phys. **B671**, 383 (2003), hep-ph/0306002

-
- [179] Y. Kajiyama, M. Raidal, and A. Strumia, Phys.Rev. **D76**, 117301 (2007), arXiv:0705.4559
- [180] S. F. King and C. Luhn, JHEP **1109**, 042 (2011), arXiv:1107.5332
- [181] S. F. King, Phys.Lett. **B659**, 244 (2008), arXiv:0710.0530
- [182] S. Antusch, S. F. King, C. Luhn, and M. Spinrath, Nucl.Phys. **B856**, 328 (2012), arXiv:1108.4278
- [183] D. Hernandez and A. Y. Smirnov, Phys.Rev. **D86**, 053014 (2012), arXiv:1204.0445
- [184] C. S. Lam, Phys.Rev. **D78**, 073015 (2008), arXiv:0809.1185
- [185] C. Lam(2009), arXiv:0907.2206
- [186] S.-F. Ge, D. A. Dicus, and W. Repko, Phys.Lett. **B702**, 220 (2011), arXiv:1104.0602
- [187] S.-F. Ge, D. A. Dicus, and W. W. Repko, Phys.Rev.Lett. **108**, 041801 (2012), arXiv:1108.0964
- [188] D. Hernandez and A. Y. Smirnov, Phys.Rev. **D87**, 053005 (2013), arXiv:1212.2149
- [189] G. L. Fogli *et al.*, Phys.Rev. **D86**, 013012 (2012), arXiv:1205.5254
- [190] P. Coloma, A. Donini, E. Fernandez-Martinez, and P. Hernandez, JHEP **1206**, 073 (2012), arXiv:1203.5651
- [191] P. Huber and J. Valle, Phys.Lett. **B523**, 151 (2001), arXiv:hep-ph/0108193
- [192] A. Gago, H. Minakata, H. Nunokawa, S. Uchinami, and R. Zukanovich Funchal, JHEP **1001**, 049 (2010), arXiv:0904.3360
- [193] P. Coloma, A. Donini, J. Lopez-Pavon, and H. Minakata, JHEP **1108**, 036 (2011), arXiv:1105.5936

- [194] J. Kopp, T. Ota, and W. Winter, Phys.Rev. **D78**, 053007 (2008), arXiv:0804.2261
- [195] P. Huber, M. Lindner, and W. Winter, Comp.Phys.Comm. **167**, 195 (2005), hep-ph/0407333
- [196] P. Huber, J. Kopp, M. Lindner, M. Rolinec, and W. Winter, Comp.Phys.Comm. **177**, 432 (2007), hep-ph/0701187
- [197] F. J. P. Soler and R. Bayes(2013), private communication.
- [198] P. Ballett and S. Pascoli, Phys.Rev. **D86**, 053002 (2012), arXiv:1201.6299
- [199] S. Bertolucci *et al.*(2012), arXiv:1208.0512
- [200] F. J. P. Soler (IDS-NF Collaboration), PoS **EPS-HEP2011**, 105 (2011)
- [201] A. Longhin, <http://irfu.cea.fr/en/Phoce/Pisp/index.php?id=54>
- [202] A. Longhin, PoS **ICHEP2010**, 325 (2010)
- [203] V. Barger, D. Marfatia, and K. Whisnant, Phys.Rev. **D65**, 073023 (2002)
- [204] See talks presented at the IDS-NF 10 conference, April 2013, <https://www.ids-nf.org/wiki/RAL-2013-04-05/>