A multi-signature approach to low-scale sterile neutrino phenomenology

MARK ROSS-LONERGAN

A Multi-Signature Approach to Low-Scale Sterile Neutrino Phenomenology

Mark Ross-Lonergan

A Thesis presented for the degree of
Doctor of Philosophy

Institute of Particle Physics Phenomenology
Department of Physics
University of Durham
England

February 2017

Dedicated to

Nicky Flynn,

whose scientific curiosity greatly encouraged my own

A Multi-Signature Approach to Low-Scale Sterile Neutrino Phenomenology

Mark Ross-Lonergan

Submitted for the degree of Doctor of Philosophy
February 2017

Abstract

Since the discovery of non-zero neutrino masses, through the observation of neutrino flavour oscillations, we had a plethora of successful experiments which have made increasingly precise measurements of the mixing angles and mass-differences that drive the phenomena. In this thesis we highlight the fact that there is still significant room for new physics, however, when one removes the assumption of unitarity of the 3×3 neutrino mixing matrix, an assumption inherent in the 3ν paradigm. We refit all global data to show just how much non-unitarity is currently allowed. The canonical way that such a non-unitarity is introduced to the 3×3 neutrino mixing matrix is by the addition of additional neutral fermions, singlets under the Standard Model gauge group. These "Sterile Neutrinos" have a wide range of theoretical and phenomenological implications. Alongside the sensitivity non-unitarity measurements have to sterile neutrinos, in this thesis we will study in detail two additional signatures of low-scale sterile neutrinos; the case of one or more light $\mathcal{O}(1\text{eV})$ sterile neutrinos detected by their effect on neutrino flavour oscillations, and heavier $\mathcal{O}(100\text{ MeV})$ detected via their subsequent decay to Standard Model particles. These two regimes have markedly different phenomenology, but are both measurable at terrestrial short-baseline experiments. We consistently use the Fermilab Short-Baseline Neutrino program as a concrete example which would produce world-leading bounds in both scenarios.

Declaration

The work in this thesis is based on research carried out while a PhD student at the Institute of Particle Physics Phenomenology, Department of Physics, Durham, England. 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 study of non-unitarity in Chapter 2 is based upon research undertaken by the author in collaboration with Stephen Parke, Fermilab. The current author was solely responsible for the global-fits, simulations, analysis and writing the paper. This work is contained in the following publication:

"Unitarity and the three flavour neutrino mixing matrix", Stephen Parke & Mark Ross-Lonergan, Phys.Rev. D93 (2016) no.11, 113009, arXiv:1508.05095v1 [hep-ph].

The work on the phenomenology of light oscillating sterile neutrinos at SBN, which forms the bulk of Chapter 3 is based on research undertaken by the author in collaboration with D.Cianci, A. Furmanski and G. Karagiorgi. The current author was the sole developer of the *SBNfit* framework, which was used for subsequent simulations and analysis, and lead the writing of the SBNfit paper:

"Prospects of Light Sterile Neutrino Oscillation and CP Violation Searches at the Fermilab Short Baseline Neutrino Facility", D.Cianci, A. Furmanski, G. Karagiorgi & M. Ross-Lonergan. arXiv:1702.01758 [hep-ph]. Submitted for publication in PRD.

Chapter 4 on the potential for SBN to search for heavy decaying sterile neutrinos was undertaken by the current author in collaboration with P.Ballett and S.Pascoli,

and was based on the paper:

"MeV-scale sterile neutrino decays with the Fermilab Short-Baseline Neutrino program", P. Ballett, S. Pascoli & M. Ross-Lonergan , arXiv:1610.08512 [hep-ph]. Accepted for publication in JHEP.

Acknowledgements

First and foremost I need to thank my fianceé Charlotte, whom I met on the very first day of this PhD, for the sheer amount of unconditional support over the past few years. Similarly, my parents whose unending support got me to the point I could even begin to consider starting a PhD.

I need to hugely thank my supervisor, Silvia, for her patience, guidance, insight and all the fantastic opportunities that I was opened to throughout my time at Durham. I must also give special thanks to my friends and colleagues Peter Ballett and Stephen Parke, without whom I do not think I would have learnt but a tiny fraction of what I did.

I would finally like to thank all past, present and future friends and inhabitants of my office, OC118,(especially those in the OC118 Fine Dining Consortium) for whose valued company was necessary to mentally survive the journey, in particular C.S Mann for patiently listening to my many *many* hours of late night physics ramblings.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"Duncan, have I not told you that when you think you know something, that is a most perfect barrier against learning?"

Leto II Atreides

In this introductory chapter we will explore, in brief, the turn of events that lead to the discovery of the neutrino and creation of the successful electroweak sector of the Standard Model. We will highlight the fact that non-zero neutrino masses, which have now been solidly proven by numerous experiments over the past decades, is the first concrete experimental example of new physics beyond the Standard Model. The massive nature of neutrinos was discovered through two important phenomena which we will introduce and discuss in detail, the mis-alignment of the neutrino flavour and mass states leading to the possibility of neutrino flavour oscillations. We will then give an overview of the current global experimental scenario and use this to motivate further chapters.

1.1 Historical Overview

Neutrinos have long had a history of controversial and surprising discoveries, dating back to their inception in 1930 where there were conjured up by Wolfgang Pauli to solve the serious contemporary problem surrounding the observed energy spectrum in nuclear β-decays. The emitted β-particles were observed to have a continuous spread of energies, and not a distinct line, as the conservation of energy would dictate for such a two-body decay. Pauli proposed the idea that a new neutral fermion,

the "neutrino", was also emitted in the decay [1]. This neutrino, which was extremely weakly interacting and promptly escaped detection, carried away with it sufficient energy to balance the observed β-spectrum. Although the idea of introducing a light, near-impossible to detect, invisible particle seemed drastic, it was no more desperate than the alternative; that conservation of energy did not hold for individual atomic decays.

In 1934 Enrico Fermi used the neutrino in his aptly, albeit perhaps understated, titled paper "An attempt of a theory of beta radiation" [1] which signalled the beginning of nearly a century of weak-interaction model building that continues today[2]. Fermi's theory of β-decay was so successful that the existence of the neutrino soon became a widely accepted view, although none had ever been observed. On the topic of the mass of this new particle Fermi concluded

"..that the rest mass of the neutrino is either zero, or, in any case, very small in comparison to the mass of the electron." .

Fermi's theory very successfully explained β-decay in what we now know is a correct manner; the three-body transition of a neutron to proton, electron and neutrino ($n \rightarrow p + e^- + \overline{\nu}_e$). Influenced by Fermi's theory it was suggested that to validate the model, the neutrino could be detected by searching for the reverse process, inverse β-decay, in which an anti-neutrino interacts with a proton to form a neutron and positron ($\overline{\nu}_e + p^+ \rightarrow n + e^+$) [3]. In 1956, using a nuclear reactor as an intense source of anti-neutrinos, Clyde Cowan and Fredrick Reines used inverse β decay to confirm the existence of the neutrino [4], for which Reines was awarded the Nobel Prize in physics in 1995 [3].

[1] Pauli originally named his proposed particle a *neutron*, but with the discovery of the currently-named neutron just two years later in 1932, Enrico Fermi coined the term *neutrino* for Pauli's mysterious particle, meaning "little neutral one" in Italian.

[2] For an English translation of Enrico Fermi's seminal paper see F.Wilson's article [2].

[3] Reines accepted the prize solely, as Cowan passed away in 1974.

Fermi's theory had four fermions interacting at a single local point, and was successfully used to explain muon decay and many other phenomena without a need to change the coupling constants. A term in the Lagrangian with four fermions is non-renormalizable, however, and was not a consistent theory or well behaved in many regimes, with the cross-section for such a scattering calculated in the four-fermion model scaling indefinitely with energy, $\sigma \propto E_\nu^2$, violating unitarity and breaking down at high energies. However, at low energies it was remarkably successful. If one viewed the Fermi four-fermion operator as an effective theory, in which two pairs of fermions exchanged a very heavy intermediately boson, then things could be made consistent, as we show illustratively in Figure (1.1). Such a new boson would have to be quite massive, $\mathcal{O}(100\text{ GeV})$, in order to achieve the very short range of the new force, and would have to carry ± 1 electric charge to account for the flow of charge. No such particle was known to exist, and so the search began for both theoretical models to explain such a particle and experimental hints of its existence.

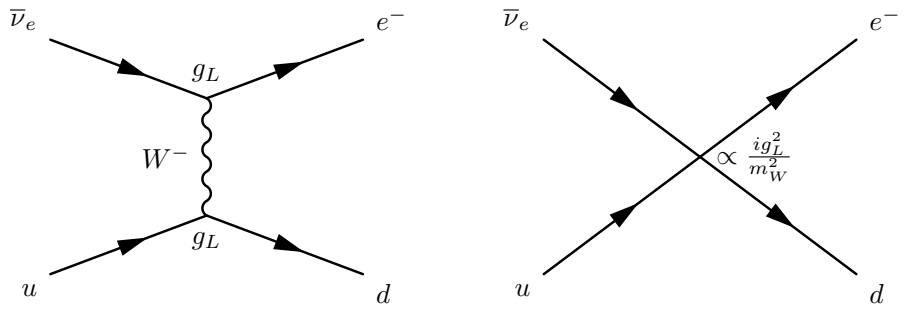

Figure 1.1: At high energies, *left*, inverse beta decay can be seen to be the exchange of a massive charged boson that couples to each of the $(\overline{\nu}_e, e^-)$ and (u, d) pairs equally with a coupling constant g_L. At low energies, *right*, the interaction is approximated by a four-fermion contact interaction whose strength is proportional to g_L^2 suppressed by the mass of the intermediate boson. The coupling of the four-fermion interaction is thus a measurement of $G_F = \sqrt{2} g_L^2/(8 m_W^2)$.

The next controversial and surprising point in the history of weak-interactions

occurred in 1956 when, motivated by recent results anomalous results in Kaon decay, T. Lee and C. Yang pointed out that there was simply no evidence whatsoever for, or against, parity conservation in any experiment involving weak-interactions thus far [5]. This was in stark contrast with Electromagnetism and the Strong force where it was seen to be conserved in all prior experiments. A parity transformation involves the flip of all spacial directions, sending a phenomena to its mirror image. It was a long held belief that the laws of nature and physics were identical under such a mirror image, and Fermi's original theory of β-decay was constructed from parity conserving vector currents. Several experiments were quickly suggested to test this and in early 1957, by studying the decay of Cobalt-60 in a magnetic field, C.S Wu showed conclusively that Parity is violated in the weak-interactions [6], breaking decades of prior belief and guaranteeing a Nobel Prize for Lee and Yang that very year.

Any Dirac fermion can be decomposed into its chiral left-handed and chiral right-handed components, $\psi = \psi_L + \psi_R$. It is useful to define the chiral projection operators $P_{L(R)} \equiv (1 \pm \gamma^5)/2$ with the property to select a given chiral component, $P_L\psi = \psi_L, P_R\psi = \psi_R$ and $P_L\psi_R = P_R\psi_L = 0$. Incorporating this parity violation required a reworking of the underlying theory. If neutrinos were indeed massless, however, it was noted that such a parity violating lagrangian could be naturally produced. This is due to the fact that, although ψ_R and ψ_L are both needed to describe a *massive* fermion, in the absence of mass a single chiral state, e.g ψ_L, is sufficient as the Dirac equation decouples into two fully independent equations, $i\partial\!\!\!/\psi_L = 0 = i\partial\!\!\!/\psi_R$. So theories in which neutrinos are described by left-handed chiral states only naturally produce what is called vector-axial vector (V-A) couplings. This name refers to how the structures, $\overline{\psi}\gamma^\mu\psi$ and $\overline{\psi}\gamma^\mu\gamma^5\psi$ produced in couplings, such as $\overline{\psi}\gamma^\mu P_L\psi$, transform under parity transformations, as a vector and an axial-vector respectively. The success of such theories led to neutrinos being treated as de-facto massless in the majority of theories of Electro-Weak interactions.

It was known by 1948, by studying the energy spectra of the emitted electrons,

that muons must decay to two neutrinos alongside the electron, not just one. Many hypothesised these neutrino might be different, with the name *neutretto* even being briefly used to distinguish them. Similarly, the neutrinos emitted alongside muons in pion decay were not known to be the same as those from the reactor experiments in the fifty's. In 1962 this was conclusively proven to be true by studying the reactions induced by a beam of, what we now know to be, ν_μ from pion decay. If they were indeed the same as ν_e the the reaction $\nu + n \rightarrow p^+ + l^-$ would only produce electrons. L. Lederman, M. Schwartz and J. Steinberger showed in such an experimental setup at Brookhaven National Labs that in a beam of ν_μ significantly more muon events were recorded that electron [7], proving the existence of the muon neutrino, for which the Nobel prize was awarded to them in 1988.

The culmination of all these experimental hints and theoretical work, along with a considerable amount more, was that the weak sector was known to be a parity violating V-A theory that conserves Lepton number and contains multiple massless neutrinos, and whose force is mediated by a massive charged boson. It was well known at this point that demanding local invariance of symmetries provided a consistent way of predicting Lagrangians, couplings and generating force carrying bosons. For example, demanding that a theory was invariant under a *Global* $U(1)_{\text{EM}}$ showed that electric charge was conserved, but demanding it was invariant under a *Local* $U(1)_{\text{EM}}$ transformation, a so called Gauge-Symmetry, generated all of the QED Lagrangian, complete with a massless photon mediator.

In this vein, in 1961 Glashow [8] put forward a theory based on an global $SU(2) \otimes U(1)$ Gauge symmetry which introduced a neutral Z gauge boson in conjunction with the, still hypothetical at the time, charged $W^\pm$ boson responsible for the Fermi theories success. Such a gauge theory could correctly account for *almost* of all the experimental observations. However, the issue that faced theorists in the mid 20th century, was that although such a theory provided gauge bosons that could explain the various interactions required, if one demanded that the theory was gauge invariant this required the gauge bosons to be massless. If massless they would cor-

respond to a long-ranged force, as the photon does, and would necessarily have been discovered already, not to mention being in stark contrast to the very massive boson needed to explain the success of the Fermi theory. Adding a mass for the bosons directly by hand provided the necessary weak short-ranged forces that agreed with experiments, but directly breaks the Gauge symmetry that introduced them.

The solution to this conundrum was developed independently by Brout, Englert [9], Higgs [10] and Kibble, Guralnik and Hagen [11] circa 1964 and involved the dynamic breaking of the $SU(2)_L \otimes U(1)_Y$ gauge symmetry down to the observable $U(1)_{EM}$ subgroup, although it required the existence of an additional field and associated scalar boson, the Higgs Boson. No fundamental scalar boson had ever been observed in nature at this point.

With this last piece of the puzzle in place, all Electro-Weak interactions were explained in a single unified theory, the Weinberg-Glashow-Salam [8, 12, 13] Model which forms the Electro-Weak sector of the Standard Model, although it took almost 10 years from the discovery of the Higgs mechanism. Of the many predictions of the theory, derived purely from the principles of the $SU(2) \otimes U(1)$ gauge symmetry, was the required existence of neutral currents via the proposed Z-boson, although there was no experimental evidence for such an interaction. The measurement, or absence, of such interactions quickly became a hot topic, whose existence would be seen as a tremendous success of the Weinberg-Glashow-Salam model. In 1974 the Gargamelle bubble chamber experiment proved conclusively, after much global scrutiny, the existence of neutral current interactions, very much in line with that predicted by the $SU(2) \otimes U(1)$ gauge symmetry [14]. This success ingrained the Weinberg-Glashow-Salam model as an accurate model of nature in the minds of physicists, and subsequently the Nobel prize in physics was awarded to Weinberg, Glashow and Salam in 1979, for their pioneering work.

It was many years later before the W-Boson and Z-Boson were directly measured themselves, using the high energy SPS accelerator at CERN in 1983 by the UA1

and UA2 experiments [15, 16], with their spokesmen Carlo Rubbia and Simon van der Meer being awarded the Nobel prize the following year for their experiments efforts. Significantly more time had to elapse, however, before the Higgs Boson was finally discovered in 2012, again using the most powerful accelerator of the time, the LHC at CERN. The 125 GeV boson was discovered by the two experiments CMS and ATLAS [17], again warranting a Nobel prize, this time to the theorists Higgs and Englert for their pioneering work that theorised a particle 48 years before its discovery.

In the intervening time the Standard Model has grown to encompass an additional generation of leptons and quarks, with the discovery of the tau-lepton [18] and its associated neutrino [19], as well as the Top [20, 21] and Bottom Quarks [22]. The Standard Model remains one of the most successful physical theories to date, making extremely precise predictions and has stood the rigours of decades of experimental tests [23].

Neutrinos, however, had one more unexpected discovery to release on the world. Despite the numerous successes of the theory of massless neutrinos, discrepancies began to emerge; Anomalies at first in the rates of observed neutrinos originating in the Sun, then in the fluxes of Atmospheric observations, followed by spectral distortions in reactor anti-neutrinos and accelerator neutrinos. All experimental results pointed to the simple and undeniable fact that neutrino are most definitely not massless. This is due to the fact that massive neutrinos exhibit the phenomena of flavour mixing and neutrino oscillations, as we will see in Section (1.3), behaviours that have been measured and quantified to high precision in the past decades. We will discuss in detail the historical measurements which first showed this in Section (1.5). Neutrino mass truly represents the first major failure of the Standard Model, the importance of which was highlighted in 2015 with the awarding of the Nobel Prize in Physics to Arthur B. McDonald and Takaaki Kajita for the discovery of Neutrino Oscillations.

1.2 The Electro-Weak Sector of the Standard Model

The Standard Model is a chiral theory, meaning it treats differently the left-handed and right-handed chiral fields of a fermion $\psi = \psi_L + \psi_R$. The left-handed chiral components of the fermions form weak isospin doublets $L_L^\alpha = (\nu_{\alpha L}, \alpha_L)^T$ for the leptons and similarly for the quarks. While the right-handed chiral fields form nine $SU(2)_L$ singlets, α_R, u_R^i and d_R^i. There are no neutral fermion right-handed fields in the Standard Model. The quantum numbers of all Standard Model field content before Electro-Weak Symmetry Breaking is given in Table (1.1).

In relation to neutrino physics we concern ourselves primarily with the Electro-Weak sector of the Standard Model, governed by the weak isospin and hypercharge symmetries of the $SU(2)_L \otimes U(1)_Y$ gauge group, and will not mention the myriad of successes the non-abelian $SU(3)_C$ gauge group has achieved in explaining the strong nuclear force[4].

It is worth taking time here to describe the process of Electro-Weak Symmetry Breaking (EWSB) that gives rise to not only the mass for the gauge bosons, but also all massive fermions in the Standard Model. Focusing solely on the Lepton sector, the Electro-Weak sector Lagrangian is the most general renormalizable Lagrangian that can be written given the field content of Table (1.1),

$$\begin{aligned}\mathcal{L}_{EW} = &\overbrace{-\frac{1}{4}A^i_{\mu\nu}A_i^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu}}^{\mathcal{L}_{\text{Gauge Kinetic}}} + \overbrace{i\overline{L_L}\not{D}L_L + i\overline{l_R}\not{D}l_R}^{\mathcal{L}_{\text{Fermion Interact+ Kinetic}}} \\ &+ \underbrace{(D_\mu H)^\dagger(D^\mu H) + \lambda H^\dagger H(v^2 - H^\dagger H)}_{\mathcal{L}_{\text{Higgs}}} - \underbrace{Y^l(\overline{L_L^l}Hl_R + \overline{l_R}H^\dagger L_L)}_{\mathcal{L}_{\text{Yukawa}}}, \qquad (1.2.1)\end{aligned}$$

where we have broadly split up the Lagrangian into the four sectors governing dif-

[4]Of course neutrinos are detected through scatterings off nuclei, so a detailed understanding of the nuclear effects involved is crucial, especially now as experiments are precise enough to need extremely accurate cross-section predictions. This is an area of intense study [24], but is a thesis worth of topics in itself and we will forgo such discussions.

ferent phenomena. The covariant derivative is given by

$$D_\mu = \partial_\mu + i\frac{g_L}{2}A^i_\mu\sigma_i + i\frac{g_Y}{2}B_\mu Y, \tag{1.2.2}$$

where σ_i are the Pauli-matrices, generators of $SU(2)$, Y is the hypercharge operator, and g_L and g_Y are the corresponding coupling constants. The field strength tensors are given by $B^{\mu\nu} = \partial^\mu B^\nu - \partial^\nu B^\mu$ for $U(1)_Y$ gauge boson and $A_i^{\mu\nu} = \partial^\mu A_i^\nu - \partial^\nu A_i^\mu - g_L\epsilon_{ijk}A_j^\mu A_k^\nu$ for the $SU(2)_L$ Bosons. Note as we have said all particles, with the exception of the Higgs, are strictly massless at this point, as any fermionic Dirac mass term $\propto\overline{\psi}_L\psi_R$ will directly violate the chiral nature of the Standard Model, with left and right handed fields transforming differently.

The Higgs Sector

The Higgs sector of the EW Lagrangian governs how the Electro-Weak gauge bosons gain mass through the Higgs mechanism. Minimising the Higgs self interaction terms leads to the Higgs obtaining a vacuum expectation value (vev) of $\langle H\rangle = 1/\sqrt{2}(0,v)^T$, by construction. Expanding out the first term in $\mathcal{L}_{\text{Higgs}}$ in this case leads to,

$$(D_\mu H)^\dagger(D^\mu H) \quad \xrightarrow[\text{EWSB}]{H\to\langle H\rangle} \quad \frac{1}{2}\left|\begin{pmatrix} g_Y B_\mu + g_L A^3_\mu & g_L A^1_\mu - i g_L W^2_\mu \\ g_L A^1_\mu + i g_L W^2_\mu & g_Y B_\mu - g_L A^3_\mu \end{pmatrix}\begin{pmatrix} 0 \\ v/\sqrt{2}\end{pmatrix}\right|^2 . \tag{1.2.3}$$

If we then define four new bosonic fields from orthogonal combinations of A^i_μ and B_μ

$$\begin{aligned} W^\pm_\mu &= \frac{1}{\sqrt{2}}\left(A^1_\mu \mp iA^2_\mu\right), \\ Z_\mu &= \frac{\sqrt{g_L^2+g_Y^2}}{g_Y g_L}\left(g_L A^3_\mu - g_Y B_\mu\right), \\ A_\mu &= \frac{\sqrt{g_L^2+g_Y^2}}{g_Y g_L}\left(g_Y A^3_\mu + g_L B_\mu\right), \end{aligned} \tag{1.2.4}$$

we see that the kinetic Higgs term, after EWSB, is precisely the mass terms for $W^\pm_\mu$ and Z_μ

$$\mathcal{L}_{\text{Higgs}} \supset (D_\mu H)^\dagger(D^\mu H) \quad \xrightarrow[\text{EWSB}]{H\to\langle H\rangle} \quad m_W^2 W^\dagger_\mu W_\mu + m_Z^2 Z_\mu Z^\mu, \tag{1.2.5}$$

with masses $m_W = g_L v/2$ and $m_Z = v\sqrt{g_Y^2 + g_L^2}/4$. The fourth field, A_μ remains massless as is the Photon field of the remaining $U(1)$ symmetry Electromagnetism[5]. This is the principle by which boson masses that break a gauge symmetry can exist in the Standard Model.

The Fermion Interaction Sector

Taking the definitions of the physical massive gauge bosons from Equation (1.2.4) we turn our attention to $\mathcal{L}_{\text{Fermion Interact}}$ expanding out the covariant derivative and dropping kinetic terms,

$$\begin{aligned}\mathcal{L}_{\text{Fermion Interact}} &= -\frac{1}{2}g_L\overline{L_L}A^i_\mu\gamma^\mu\sigma_i L_L + \frac{1}{2}g_Y B_\mu\gamma^\mu L_L + g_Y\overline{l_R}B_\mu\gamma^\mu l_R, \\ &= \underbrace{-\frac{g_L}{\sqrt{2}}\overline{\nu_{eL}}\not{W}^+ l_L + H.C}_{\mathcal{L}_{\text{CC}}} + \underbrace{e\left(\overline{l_L}\not{A}l_L + \overline{l_R}\not{A}l_R\right)}_{\mathcal{L}_{\text{NC},\gamma}} \\ &\quad \underbrace{-\frac{1}{2}e\frac{g_Y}{g_L}\left[\left(1-\left(\frac{g_L}{g_Y}\right)^2\right)\overline{l_L}\not{Z}l_L + 2\overline{l_R}\not{Z}l_R\right] - \frac{g_L g_Y}{2e}\overline{\nu_{eL}}\not{Z}\nu_{eL}}_{\mathcal{L}_{\text{NC,Z}}},\end{aligned} \tag{1.2.6}$$

where we have made the connection $e^2 \equiv g_Y^2 + g_L^2$, as this way the NC coupling of the photon field A^μ to a charged lepton reduces to the QED theory of $U(1)_{\text{EM}}$ as required. The charged current sector of the Lagrangian depends only on the $SU(2)_L$ coupling constant g_L, as the physical massive $W^\pm$ bosons are combinations of only the charged $SU(2)_L$ bosons $A^{1,2}$. The neutral current sector, however, is more complicated with the neutral $SU(2)_L$ boson A^3 mixing with the B boson of hypercharge to form the physical Z and Photon fields.

The coupling of the neutrinos to the Z-Boson in the NC Lagrangian, and coupling to the $W^\pm$-Bosons in CC sector are the only places in which neutrinos couple with the remainder of the Standard Model, and it is through these interactions that

[5]The relationship between the physical neutral gauge bosons and those of $SU(2)_L \times U(1)_Y$ is often given in terms of the Weinberg mixing angle, $Z_\mu = \sin\theta_W A^3_\mu + \cos\theta_W B_\mu$, and $A_\mu = \cos\theta_W A^3_\mu - \sin\theta_W B_\mu$ with the mixing angle given by $g_L \sin\theta_W = e = g_Y \cos\theta_W$, $\theta_W \approx 28^\circ$.

all neutrino observations must arise. This is unlike all charged fermions whose right-handed fields couple to the Higgs though Yukawa terms. Note that as the Higgs is not present in the Fermion Interaction Lagrangian, these couplings are not effected by EWSB.

The Yukawa Sector

Before an in depth discussion of neutrino flavour effects, it is worth considering the flavour and masses of the charged leptons, as the two concepts are intricately linked. Before EWSB the Standard Model does not provide a way to distinguish the three generations of charged leptons and quarks. With identical quantum numbers there is experimental signatures that can separate an electron from a muon, for example. Of course experimentally we know fermions have mass, and part of the success of the Higgs mechanism is that in breaking $SU(2)_L \otimes U(1)_Y$ masses are also generated for any fermions which couple to the Higgs through Yukawa couplings.

The general Yukawa section of the Standard Model lagrangian, $\mathcal{L}_{\text{Yukawa}}$, in Equation (1.2.1), post EWSB, takes on the form

$$\mathcal{L}_{\text{Yukawa}} \xrightarrow[\text{EWSB}]{H\to\langle H\rangle} \frac{v}{\sqrt{2}}\overline{\tilde{l_L}}\tilde{Y}^l\tilde{l_R} + H.C. \tag{1.2.7}$$

We now note the matrix of Yukawa couplings, $\tilde{Y}^l$, is not assumed to be diagonal and most generally is a complex 3×3 unitary matrix. We will use tildes to refer to the fact that in an arbitrary basis the fermion fields may not have definite masses due to these off-diagonal Yukawa elements. We are, however, free to diagonalise $\tilde{Y}^l$ through a bi-unitary transformation, $V_L^{l\dagger}\tilde{Y}^l V_R^l = Y^l$, leading to a now diagonal Y^l. Our Yukawa term in this basis then reads

$$\begin{aligned}\mathcal{L}_{\text{Yukawa}} &= \frac{v}{\sqrt{2}}\underbrace{\overline{V_L^{l\dagger}\tilde{l_L}}}_{\equiv l_L} Y^l \overbrace{V_R^{l\dagger}\tilde{l_R}}^{\equiv l_R} + H.C, \\ &= \frac{v}{\sqrt{2}}\overline{l_L}Y^l l_R + H.C.\end{aligned} \tag{1.2.8}$$

Fermions (Spin 1/2)		**Weak Isospin** I	**Hyper Charge** Y	**Electric Charge** $(Q \equiv I_3 + Y/2)$
Lepton Doublets	$L_L^\alpha = \begin{pmatrix} \nu_{\alpha L} \\ \alpha_L \end{pmatrix}$	$\frac{1}{2}$	-1	0 -1
Charged Lepton Singlets	α_R	0	-1	-1
Quark Doublets	$Q_L^i = \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix}$	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{2}{3}$ $-\frac{1}{3}$
Quark Singlets	u_R^i d_R^i	0	$\frac{4}{3}$ $-\frac{2}{3}$	$\frac{2}{3}$ $-\frac{1}{3}$
Bosons	(Spin 1)	**Weak Isospin** I	**Hyper Charge** Y	**Electric Charge** $(Q \equiv I_3 + Y/2)$
Higgs Doublet	$H = \begin{pmatrix} H^+ \\ H^0 \end{pmatrix}$	$\frac{1}{2}$	1	1 0
A Triplet	$A = \begin{pmatrix} A^1 \\ A^2 \\ A^3 \end{pmatrix}$	1	0	$+1$ -1 0
B Singlet	B^0	0	0	0

Table 1.1: The Electro-Weak sector of the Standard Model, pre EWSB. Shown are the eigenvalues of Weak Isospin (I) and Hypercharge (Y), for the Standard Model field content, as well as the Electric Charge Q of the conserved $U(1)_{\rm EM}$ remaining after EWSB. I_3 is the third component of Isospin, and takes the values $\pm I$ for each respective element of the weak isospin doublets and $+I, -I$ and 0 for the triplets, respectively. We note that unlike all charged fermions, there is no right-handed chiral fields for the neutrino. If included they would be singlets under weak isospin, and have 0 hypercharge and electric charge. The index α is over the three lepton families, electron muon and tau, which differ only in mass. The index $i = 1, 2, 3$ is over the three Quark generations, u^i over up, charm and top, and d^i over down, strange and bottom.

Which is none other than a Dirac mass term, $m\bar{l}l$, for the Dirac fields, $l \equiv l_L + l_R$, of the charged leptons of definite mass, $m_\alpha = vY^l_{\alpha\alpha}/\sqrt{2}$.

Having obtained a mass through the Higgs Mechanism, the three generations of charged leptons are no longer identical, so it is their masses we use to experimentally differentiate between them. Thus we observe an "electron flavour" lepton if we see a charged lepton with mass $m_e \approx 0.5$ MeV, a muon is a charged lepton with mass $m_\mu \approx 100$ MeV and a tau is a charged lepton with mass $m_\tau \approx 1800$ MeV. Usually these are measured via their subsequent decay products using the long-ranged electromagnetic force. However, this implies the massive fields l_L are also the fields of definite flavour, by virtue of definition.

Neutrinos, on the other hand, do not interact electromagnetically, nor have masses large enough to be measured directly using today's technology. Their only direct coupling to a charged particle is solely through CC weak interactions as described in Equation (1.2.6). As this always occurs in the presence of a charged lepton[6], we naturally choose to define the flavour of a neutrino by the *flavour of the associated charged lepton in a charged weak interaction.* As the flavour of a charged lepton is itself defined by its mass, we are in essence defining the flavour of a neutrino by the *mass* of the associated charged lepton in that $SU(2)_L$ doublet. In what follows we will explore the details and consequences of this definition.

[6]The neutral current coupling to the Z-Boson is flavour blind, so although such scatterings can identify a neutrino interacted, they cannot specify a flavour and so are not useful in determining neutrino flavour.

1.3 Neutrino Mass and Flavour

In order to define the neutrino flavour fields we turn our attention to the lagrangian term that leads to the weak charged current interaction,

$$\mathcal{L}_{\mathrm{CC}} \supset \overline{\tilde{\nu}_L}\gamma^\mu \tilde{l_L} W_\mu + H.C, \tag{1.3.9}$$

$$\supset \underbrace{\overline{\nu_L}}_{\nu_L \equiv V_L^{l\,\dagger}\tilde{\nu}_L} \gamma^\mu l_L W_\mu + H.C. \tag{1.3.10}$$

As we have shown l_L represents the physical massive eigenstates for the charged leptons, ν_L is therefore the neutrinos created alongside them in a weak interaction and by definition, the neutrino fields of distinct flavour $\nu_L \equiv (\nu_{eL}, \nu_{\mu L}, \nu_{\tau L})^T$.

In the scenario that neutrinos were massless, as they are in the Standard Model, this would be the end of the discussion with the flavour fields as defined here being massless fields, and any combination of massless fields remains a massless field, thus the flavour fields are mass-eigenstates also.

However, it is now known beyond any doubt that neutrinos do in fact have mass, albeit *very* small in comparison to all other known massive particles of the Standard Model. As we will show shortly, there is no possible mechanism to generate mass for the neutrinos in the Standard Model, and as such non-zero neutrino represents one of the few failures of the Standard Model, requiring the addition of new undiscovered physics. The details of the experiments which discovered neutrino masses through oscillation is left to Section (1.5), where we will discuss historical and contemporary experimental measurements of the neutrino sector. For now we will focus on the theory of massive neutrinos beyond the Standard Model.

As there is no right-handed chiral fields belonging to a neutral fermion in the Standard Model, it is impossible to write down a Yukawa coupling to the Higgs that post-EWSB would give rise to a Dirac mass term for the neutrinos[7] as was done

[7]With the addition of right-handed "sterile neutrinos" one can construct such a Dirac mass. This is discussed in detail in Section (2.2).

for the leptons in Equation (1.2.8). There is, however, another way to generate a mass term without the need for a right-handed chiral field, as was discovered by Majorana in 1937 [25], and we will now discuss its relevance to the neutrino sector.

Majorana Masses

It is possible to write down a solution to the Dirac equation for a massive Fermion field ψ, $(i\gamma^\mu\partial_\mu - m)\psi = 0$ where instead of ψ having two independent Weyl component spinors, $\psi = \psi_L + \psi_R$ as is the case for the massive charged fermions, but only one, ψ_L. To achieve this we would require a function of ψ_L, $\mathcal{F}(\psi_L)$, that behaved like a right-handed chiral field, in order to form a Dirac-like mass $m\overline{\psi_L}\mathcal{F}(\psi_L) + H.C$, as well as transforming like ψ_L under Lorentz transformations, so that the quantity $\overline{\mathcal{F}(\psi_L)}\psi_L$ is a Lorentz scalar. It turns out that this function is none other than the charge conjugated field

$$\mathcal{F}(\psi_L) = \psi_L^C = \mathcal{C}\overline{\psi_L}^T, \tag{1.3.11}$$

where $\mathcal{C}$ is the charge conjugation matrix. One can easily see that this behaves as a right-handed chiral field by acting on it with the left-handed projector

$$P_L(\mathcal{C}\overline{\psi_L}^T) = C\left((\underbrace{P_R\psi_L}_{=0})^\dagger\gamma^0\right)^T = 0. \tag{1.3.12}$$

and under a Lorentz transformation $\psi(x) \to \psi'(x') = \mathcal{A}\psi(x)$, it can be shown that the charge conjugated field transforms

$$\psi_L^C \xrightarrow[\text{Transformation}]{\text{Lorentz}} \mathcal{C}\left(\psi'^\dagger(x')\gamma^0\right)^T = \mathcal{C}(\mathcal{A}^{-1})^T\mathcal{C}^{-1}\psi_L^C(x) = \mathcal{A}\psi_L^C, \tag{1.3.13}$$

in the same manner. Thus, the field

$$\psi = \psi_L + \mathcal{C}\overline{\psi_L}^T, \tag{1.3.14}$$

allows us to construct a mass term using ψ_L alone,

$$\mathcal{L}_{\text{mass}}^{\text{maj}} = -\frac{1}{2}m\overline{\psi_L^C}\psi_L + H.C \quad = \frac{1}{2}m\psi_L^T C^\dagger\psi_L + H.C. \tag{1.3.15}$$

Although a simple and elegant solution, requiring no additional field content or interactions and in fact using one less Weyl spinor than a traditional Dirac mass

term, generating a Majorana mass for the neutrinos in the Standard Model fails. We can see this by noting that under a $U(1)$ transformation both ψ and ψ^T transform as $\psi^{(T)} \to e^{i\phi}\psi^{(T)}$, so a Majorana mass term would violate such a transformation. For $U(1)_{\text{EM}}$ this means ψ must be a neutral fermion in order to not violate conservation of charge. Perhaps encouragingly, as the only neutral fermions in the Standard Model, neutrinos are thus the only particles that *could potentially* have a Majorana mass term. We do note, however, that as this term violates any U(1) transformations the accidental Lepton number symmetry in the Standard Model is broken by the inclusion of such a Majorana mass term.

The fact that this term violates $SU(2)_L \otimes U(1)_Y$ is not immediately a concern as it may arise from a symmetry conserving term at higher energy scales in the same way as Fermion masses do. However, if this mass were to arise post-EWSB from a Yukawa-like coupling to the Higgs it would require a term in the Lagrangian containing two $SU(2)_L$ doublets L_L^α, which transforms as a weak isospin triplet. As the standard model does not contain an isospin triplet, it would need to be formed out of two Higgs doublets to ensure the Lagrangian remained invariant at the higher scale. This automatically leads to operators such as

$$\mathcal{L}_5 = \frac{g}{\Lambda}\overbrace{\left(L_L^T\mathcal{C}^\dagger\sigma_2\vec{\sigma}L_L\right)}^{\text{Isospin triplet}}\overbrace{\left(H^T\sigma_2\vec{\sigma}H\right)}^{\text{Isospin triplet}} + H.C, \tag{1.3.16}$$

$$\xrightarrow[\text{Breaking}]{\text{EW Symmetry}} \quad \frac{gv^2}{2\Lambda}\nu_L^T\mathcal{C}^\dagger\nu_L + H.C. \tag{1.3.17}$$

This term, however, is forbidden in the Standard Model framework as it is of energy dimension 5 and thus non-renormalizable. This further reiterates the fact that non-zero neutrino masses represent the first confirmed physics beyond the Standard Model. This is, in fact, the only unique operator that can be formed at dimension 5 using the field content of the Standard Model. It is known as the Weinberg operator [26, 27], and its sole role, post EWSB, is the generation of neutrino masses. The term Λ has dimension of energy, and if this term originated at some high scale, $\Lambda \gg$ TeV, it would be expected, by naturalness arguments, to be of the order of the scale of new physics. Thus for $\mathcal{O}(1)$ couplings and the Higgs vev as normal, it naturally would produce small neutrino masses, $m_\nu = gv^2/\Lambda$. If we explicitly add extra sterile

degrees of freedom, then additional dimension 5 operators can be formed, but we will leave this discussion until Section (4.2.2).

Despite the fact this approach fails for the Standard Model *as is*, with the inclusion of extra field content, or if one considers the Standard Model as a low-energy effective theory allowing such non-renormalizable terms in the Lagrangian, Majorana mass terms become an attractive solution. We explore this in more detail in Section (2.2) when we discuss some specific models beyond the Standard Model that can generate the Weinberg Operator.

1.3.1 Neutrino Mixing

Although the exact mechanism by which neutrino mass is generated is currently unknown, the fact that neutrinos do have mass is concretely proven. For the purposes of discussing the phenomenological effects that arise due to non-zero masses, it suffices here to treat the Standard Model as an effective theory, and include the dimension 5 Weinberg operator that generates a Majorana mass term for the neutrinos, as defined in Equation (1.3.17). In fact, one could argue that the discovery of neutrino masses shows that the Standard Model is an low-energy effective theory, and we are very justified in this assumption. In this manner we can elaborate on the effects of neutrino masses without concretely moving to a model, or adding any field content to the Standard Model.

With the inclusion of mass for the neutrinos, the definitions of charged leptons mass and flavour states remains the same, but now the neutral leptons can have a Majorana mass term in the Lagrangian,

$$\mathcal{L}_{\text{mass}} = \frac{1}{2}\tilde{\nu_L}^T \mathcal{C}^\dagger \tilde{M} \tilde{\nu_L} + H.C, \tag{1.3.18}$$

where we have extended $\tilde{M}$ as a complex symmetric Majorana mass matrix. Following the same procedure as the charged leptons we can diagonalise this to define the neutrino fields of definite mass using the unitary transformation, $V_L^{\nu T} \tilde{M} V_L^{\nu} \equiv M$. The massive chiral neutrinos are then given by $n_L = V_L^{\nu\dagger} \tilde{\nu}_L$ which form a diagonal

Majorana mass term for the neutrino

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} n_L^T \mathcal{C}^\dagger M n_L + H.C. \tag{1.3.19}$$

Thus it is $n_L \equiv (\nu_{1L}, \nu_{2L}, \nu_{3L})^T$ that represent the massive neutrinos.

Just as with the case of massless neutrinos, we now turn to the charged current lagrangian to precisely define the "flavour" fields for massive neutrinos. Rewriting the charged current lagrangian in terms of the physical massive fields for both the charged leptons and neutrinos we obtain

$$\begin{aligned}
\mathcal{L}_{\text{CC}} &\propto \overline{\tilde{\nu}_L} \gamma^\mu \tilde{l_L} W_\mu + H.C, && (1.3.20)\\
&\propto \overline{(\underbrace{V_L^{l\dagger} V_L^\nu}_{\equiv U_{\text{PMNS}}} n_L)} \gamma^\mu l_L W_\mu + H.C, && (1.3.21)\\
&\propto \overline{\nu_L} \gamma^\mu l_L W_\mu + H.C, && (1.3.22)
\end{aligned}$$

where in the final line we have the neutrino flavour states, ν_L, defined as we have chosen, to be those which are created along side charged leptons of definite mass.

$$\underbrace{\nu_L}_{\text{Flavour}} = U_{\text{PMNS}} \underbrace{n_L}_{\text{Mass}} = \begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix}. \tag{1.3.23}$$

The unitary matrix that represents the basis transformation between the neutrino mass basis and flavour basis is known as the Pontecorvo–Maki–Nakagawa–Sakata matrix or PMNS matrix [28–30]. *A priori* the Standard Model tells us nothing about the values of elements in this matrix. It could be equal to the identity matrix, meaning the flavour and mass eigenstates are equal, or it could have off-diagonal complex elements meaning flavour neutrinos would not have definite mass, but are superpositions of neutrinos with definite masses.

This is directly analogous to the quark sector, where it has been found that the quark flavour and mass eigenstates mis-align. However, it has been shown that the mixing matrix, the CKM matrix, is close to the identity matrix with only small perturbations from this, proportional to the Cabibbo mixing angle, $\theta_C \approx 13^\circ$, at leading order. As we will see in Section (1.5), the situation significantly deviates from this in the neutrino sector.

Parameterization of the PMNS Matrix

A general 3×3 unitary matrix has $3^2 = 9$ parameters and can be decomposed into three angles and six complex phases. Three of these six phases are unphysical, however, and they can be removed by rephasing the massive neutrino, $\nu_{iL} \rightarrow e^{i\phi_i}\nu_{iL}$, and charged lepton fields, $l_{\alpha L} \rightarrow e^{i\phi_\alpha} l_{\alpha L}$ appropriately. If we define $U(\theta_{ij}, \delta)$ is as a rotation of angle θ_{ij} around the $i-j$ plane, with complex phase δ, then a near ubiquitously used parameterization of U_{PMNS} is given by

$$U_{\text{PMNS}} = U(\theta_{23}, 0)U(\theta_{13}, \delta_{CP})U(\theta_{12}, 0)\mathcal{D}^M,$$

$$= \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric Sector}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{Reactor Sector}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar Sector}} \overbrace{\begin{pmatrix} e^{i\alpha_1} & 0 & 0 \\ 0 & e^{i\alpha 2} & 0 \\ 0 & 0 & 1 \end{pmatrix}}^{\text{Majorana Phases}}, \tag{1.3.24}$$

where $c_{ij} \equiv \cos(\theta_{ij})$ and $s_{ij} \equiv \sin(\theta_{ij})$ and the labels refer to the colloquial names governing oscillations at their respective scales. If neutrinos are Dirac fermions, they contain two additional Weyl-spinor degrees of freedom allowing two more phases to be removed, leaving five of the six phases unphysical, removing α_1 and α_2 in the above equation.

1.4 Oscillation Phenomenology

The fact that neutrinos are massive does not alone imply the existence of vacuum neutrino oscillations. As we will see below there needs to be a mass difference between one or more of the active neutrinos, and the mixing matrix (U_{PMNS}), between the flavour and mass states must contain off-diagonal elements. As we now know, nature has indeed chosen these two conditions, and through neutrino oscillation experiments, we have gained unprecedented knowledge of the neutrino section.

In this section we will derive the oscillation formula in detail, as it represents a key phenomena of massive neutrinos, and one that we have derived the major-

ity of knowledge in the neutrino sector from. As the more standard *"equal-energy"* derivation can be found in many textbooks and while it generates the correct answer in many cases, it contains unphysical assumptions does not account for additional phenomena, we will instead here derive the oscillation probability though a more rigorous quantum mechanical method, following the approach of C.Giunti [31].

The experimental setup we wish to model consists of the production of a neutrino in a definitive flavour eigenstate, $|\nu_\alpha\rangle = \sum_i U^\star_{\alpha i} |\nu_i\rangle$, via the decay of a particle driven by the weak interaction. The neutrino then propagates across a baseline L as a coherent superposition of its mass eigenstates, and is detected (again as a pure weak, but possibly different flavoured, state $|\nu_\beta\rangle$) at a time T through a weak scattering event. We model the propagating state as a Gaussian wave-packet. This is by no means the only shape possible, but is sufficiently simple to provide exact analytical results, yet contain enough physics to gain insight.

$$\Psi_j(k, k_j, \sigma_{k|\text{prod}}) = \frac{1}{\sqrt[4]{2\pi\sigma^2_{k|\text{prod}}}} \exp\left[-\frac{(k-k_j)^2}{4\sigma^2_{k|\text{prod}}}\right], \tag{1.4.25}$$

with mean momentum k_j and width associated with its production in momentum space of $\sigma_{k|\text{prod}}$. We additionally will assume that we are working in a highly relativistic situation and will work to leading order in $\epsilon \equiv m_i^2/E_i^2$. As this work focuses on terrestrial experiments with lowest energy's in the MeV range, and sterile mass splittings of less than 10^3 eV^2, we see that even for the most extreme cases $\epsilon \leqslant 10^{-9}$ and this is an extremely well motivated assumption. We will also assume that the wave-packets are well peaked in momentum space at production, corresponding to $E_i^2/m_i^2 \gg \sigma_{k|\text{prod}}$. These two approximations allow us to express the energy of a given mass state as $E_i(k) \approx E_i + v_i(k - k_i)$ where v_i is the corresponding group velocity, $v_i = k_i/E_i$. After expanding to leading order in ϵ these take the analytically simple

forms

$$E_i(k) \approx E_\nu + (1-\xi)\frac{m_i^2}{2E_\nu}, \tag{1.4.26}$$

$$k_i \approx E_\nu - \xi\frac{m_i^2}{2E_\nu}, \tag{1.4.27}$$

$$v_i \approx 1 - \frac{m_i^2}{2E_\nu^2}, \tag{1.4.28}$$

where ξ is a production process specific factor and E_ν is the energy of a *massless* neutrino in the same kinematic situation. Bringing this all together it allows us to write the state that is produced in the decay as

$$|\nu_\alpha(t)\rangle = \sum_j U^\star_{\alpha j} \int \mathrm{d}k\, \Psi_j(k, k_j, \sigma_{k|\mathrm{prod}})\, e^{E_j(k)t}\, |\nu_j(k)\rangle\,, \tag{1.4.29}$$

where E_j is the energy of the j$^{\mathrm{th}}$ propagating mass state as given by Equation (1.4.28). Moving to physical space, $|\nu_\alpha(x,t)\rangle = \langle x|\nu_\alpha(t)\rangle$, we have the following form

$$|\nu_\alpha(x,t)\rangle = \frac{1}{\sqrt[4]{2\pi\sigma^2_{k|\mathrm{prod}}}} \sum_j U^\star_{\alpha j} \exp\left[-iE_j t + ik_j x - \frac{(x-v_j t)^2}{4\sigma^2_{x|\mathrm{prod}}}\right], \tag{1.4.30}$$

where similarly k_j is the momentum of the j$^{\mathrm{th}}$ propagating mass state. The state that is detected, as flavour β, is constructed in a similar manner with the exception that this state does not propagate and is destroyed upon detection, and the Gaussian wavepacket will have a different width associated with the detection process $\sigma_{k|\mathrm{det}}$ which does not equal the production width.

$$|\nu_\beta(t)\rangle = \sum_j U^\star_{\beta j} \int \mathrm{d}k\, \Psi_j(k, k_j, \sigma_{k|\mathrm{det}})\, |\nu_j(k)\rangle\,, \tag{1.4.31}$$

which corresponds in physical space to

$$|\nu_\beta(x)\rangle = \frac{1}{\sqrt[4]{2\pi\sigma^2_{k|\mathrm{det}}}} \sum_j U^\star_{\beta j} \exp\left[+ik_j x - \frac{(x-v_j t)^2}{4\sigma^2_{x|\mathrm{det}}}\right]. \tag{1.4.32}$$

The amplitude in which a neutrino of definite flavour ν_α is produced at $L = 0$ and is detected as flavour ν_β at a time T later after propagating a baseline L, is then given by

$$\mathcal{A}(\nu_\alpha \to \nu_\beta : L, T) = \int \mathrm{d}x\, \langle \nu_\beta(x-L)|\, |\nu_\alpha(x,T)\rangle\,, \tag{1.4.33}$$

$$\propto \sum_i U^\star_{\alpha i} U_{\beta i} \exp\left[-iE_i T + ik_i L - \frac{(L-\nu_i T)^2}{4(\sigma^2_{x|\mathrm{prod}} + \sigma^2_{x|\mathrm{dec}})}\right]. \tag{1.4.34}$$

The probability that we are looking for is then proportional to the square of this amplitude,

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L, T) \quad \propto \quad |\mathcal{A}(\nu_\alpha \to \nu_\beta : L, T)|^2, \tag{1.4.35}$$

$$\propto \sum_{i,j} U^\star_{\alpha i} U_{\beta j} U^\star_{\beta j} U_{\alpha i} \exp\left[-i\frac{\Delta m^2_{ij}}{2E_\nu}\left\{T(1-\xi)+\xi L\right\} - \frac{(L-v_iT)^2+(L-v_jT)^2}{4\sigma_x^2}\right], \tag{1.4.36}$$

where $\sigma_x^2 \equiv \sigma^2_{x|\text{prod}} + \sigma^2_{x|\text{dec}}$. The time at which a specific neutrino arrives is not a measured quantity in any neutrino experiment, as such we integrate over all time leaving only a dependence on the baseline L, which is traditionally fixed for any given experiment. In doing this we will enforce the normalisation of the probability such that $\sum_\beta \mathcal{P}(\nu_\alpha \to \nu_\beta) = 1$. This expression notably has terms dependant on T, T^2 and independent of T,

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L) = \int_{-\infty}^{+\infty} \mathrm{d}T \mathcal{P}(\nu_\alpha \to \nu_\beta : L, T) \tag{1.4.37}$$

$$\propto \sum_{i,j} U^\star_{\alpha i} U_{\beta j} U^\star_{\beta j} U_{\alpha i} \exp\left[-i\frac{\Delta m^2_{ij}}{2E_\nu}\xi L - \frac{L^2}{2\sigma_x^2}\right] \times \tag{1.4.38}$$

$$\times \int_{-\infty}^{+\infty} \mathrm{d}T \, \exp\left[-iT\frac{\Delta m^2_{ij}(1-\xi)}{2E_\nu} + T\frac{(v_i+v_j)}{2\sigma_x^2} + T^2\frac{(v_i^2+v_j^2)}{4\sigma_x^2}\right]. \tag{1.4.39}$$

We note that the remaining time integral is of the form $\int_{-\infty}^{+\infty} e^{-ic_1x} e^{c_2x} e^{-x^2c_3}$ which has an exact solution provided $c_3 \equiv \frac{v_i^2+v_j^2}{4\sigma_x^2}$ is greater than zero,

$$\int_{-\infty}^{+\infty} \mathrm{d}x \, e^{-ic_1x} e^{c_2x} e^{-x^2c_3} = \sqrt{\frac{\pi}{c_3}} \exp\left[-\frac{(c_1+ic_2)^2}{4c_3}\right]. \tag{1.4.40}$$

This, along with the normalisation condition, then gives us the probability for detecting a neutrino of flavour β at a baseline L, in a beam of neutrinos of flavour α,

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L) = \sum_{j,m} U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m} \exp\Big[\underbrace{-2\pi i \frac{L}{L_{\text{osc}}}}_{\substack{\text{Oscillatory Term} \\ L_{\text{osc}} \equiv \frac{4\pi E_\nu}{\Delta m^2_{jm}}}} - \overbrace{\left(\frac{L}{L_{\text{coh}}}\right)^2}^{\substack{\text{Coherence Term} \\ L_{\text{coh}} \equiv \frac{4\sqrt{2}E_\nu^2 \sigma_x}{|\Delta m^2_{jm}|}}} \underbrace{-2\pi^2 (1-\xi)^2 \left(\frac{\sigma_X}{L_{\text{osc}}}\right)^2}_{\text{Localisation Term}} \Big]. \tag{1.4.41}$$

We note that if one repeats the above calculation but starting with states of anti-neutrinos

$$|\overline{\nu}_\alpha\rangle = \sum_i U_{\alpha i} |\overline{\nu}_i\rangle , \tag{1.4.42}$$

all kinematic statements and integration remains unchanged, leading to an identical result up to complex conjugation of the leading coefficients (elements of the U_{PMNS} matrix). This gives the corresponding anti-neutrino oscillation probability

$$\mathcal{P}(\overline{\nu}_\alpha \to \overline{\nu}_\beta : L) = \sum_{j,m} U_{\alpha m} U^\star_{\alpha j} U_{\beta j} U^\star_{\beta m} \exp\left[-2\pi i \frac{L}{L_{\text{osc}}} - \left(\frac{L}{L_{\text{coh}}}\right)^2 - 2\pi^2 (1-\xi)^2 \left(\frac{\sigma_X}{L_{\text{osc}}}\right)^2\right]. \tag{1.4.43}$$

Alongside the traditional oscillatory term there are two exponentially damping terms. We will now discuss the physical interpretations and effects of the three terms appearing in the exponential individually.

The Coherence Term

$$L_{\text{coh}} \equiv 4\sqrt{2}\frac{E_\nu^2 \sigma_X}{|\Delta m^2_{ij}|}, \tag{1.4.44}$$

$$\approx 5.7 \times 10^{15} \left[\frac{E_\nu}{\text{GeV}}\right] \left[\frac{\sigma_X}{\text{m}}\right] \left[\frac{|\Delta m^2_{ik}|}{\text{eV}^2}\right]^{-1} \text{m}. \tag{1.4.45}$$

This term suppresses the transition probability for experiments whose baselines are greater than L_{coh}. It originates from the fact that, in a given time, wave packets associated with multiple massive neutrinos will travel different distances, with heavier mass states propagating at a slower velocity and hence not as far. If the different wave-packets representing each mass state become sufficiently separated in an experiment, $L \geqslant L_{\text{coh}}$, then they cannot all coherently interact at the detector, and thus the oscillatory effects (which arise from such coherent interference) are greatly suppressed.

The larger the mass splitting between any two neutrino, the smaller the coherence length, but this also corresponds to a proportional reduction of the oscillation length. Shown in Figure (1.2) is the region (shaded green) in which decoherence effects

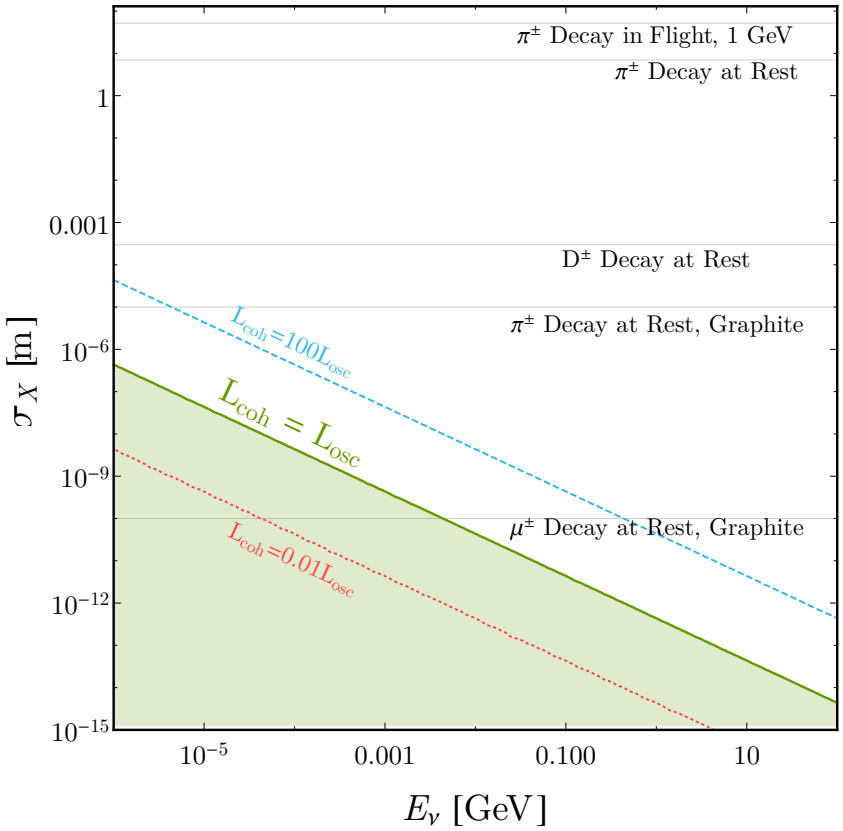

Figure 1.2: Parameter space in which the massive neutrino wave-packets are separated sufficiently in space that they do not interact coherently within the detector, suppressing the interference necessary to produce oscillations. The size of the wavepacket, σ_X, is estimated by looking at the dilated lifetime for decays in vacuum $\sigma_X \approx \tau\gamma$, and the mean free paths when decays occur in a medium. Shown are some example values for methods often used to generate a neutrino flux.

become large, relative to the oscillation length $L_{\text{osc}} = L_{\text{coh}}$. As can be easily seen, decoherence effects are not an issue for terrestrial experiments except for situations involving very small wave-packets.

The Localisation Term

The second damping term is the localisation term which suppresses oscillation if the wave-packet size begins to become comparable to the oscillation length, if $\sigma_X > L_{\text{osc}}$. If this is the case, the mass states at production are not localised to the production region and can even overlap with the detector, washing out interference between the

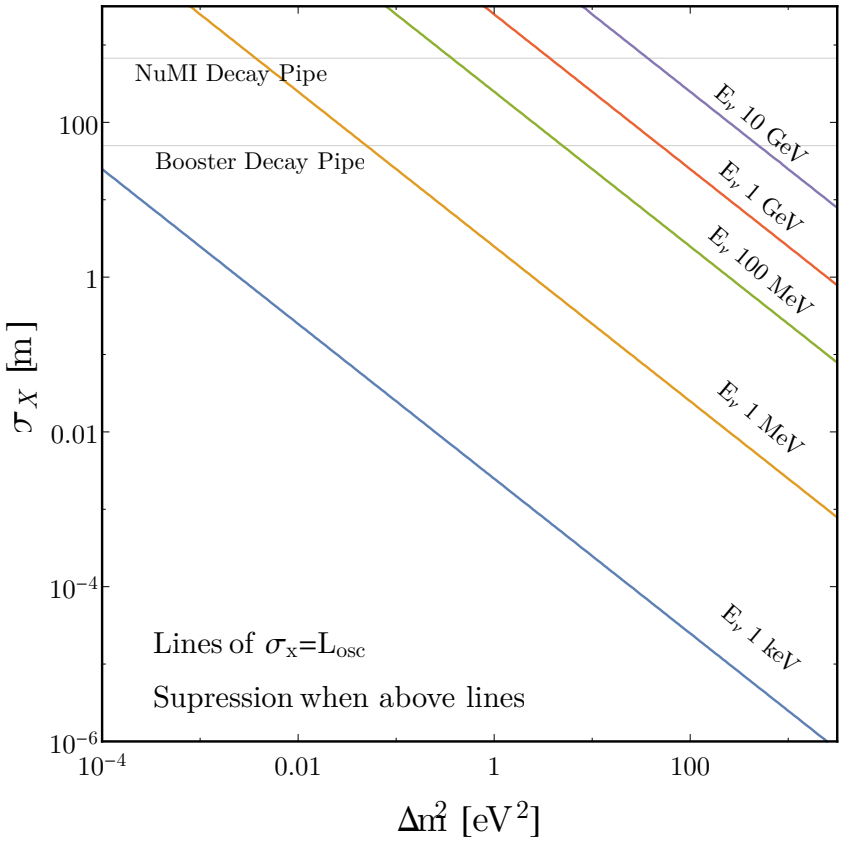

Figure 1.3: Regions in which the massive neutrino wave-packets are large enough such that the production and detection are no longer localised events (Area above each coloured line).

different mass states. It follows that the experiment can no longer even be considered as a neutrino propagation experiment.

To estimate this effect we look at what is the largest value that σ_X can reasonably take. For neutrino experiments in which the beam is produced via particle decay, even for highly relativistic parent particles, the wave-packet size is limited by the size of the decay tunnel which is usually of $\mathcal{O}$(100 meters), e.g. Booster beam tunnel at Fermilab is 50m while the NuMI beam is at the extreme edge of design space at 675m. As can be seen in Figure (1.3), for scenarios with very large mass splittings the localisation term may not be ignorable, as was pointed out in [32][8].

[8]Note in [32] what we refer to here as the localisation term is called the "decoherence" parameter.

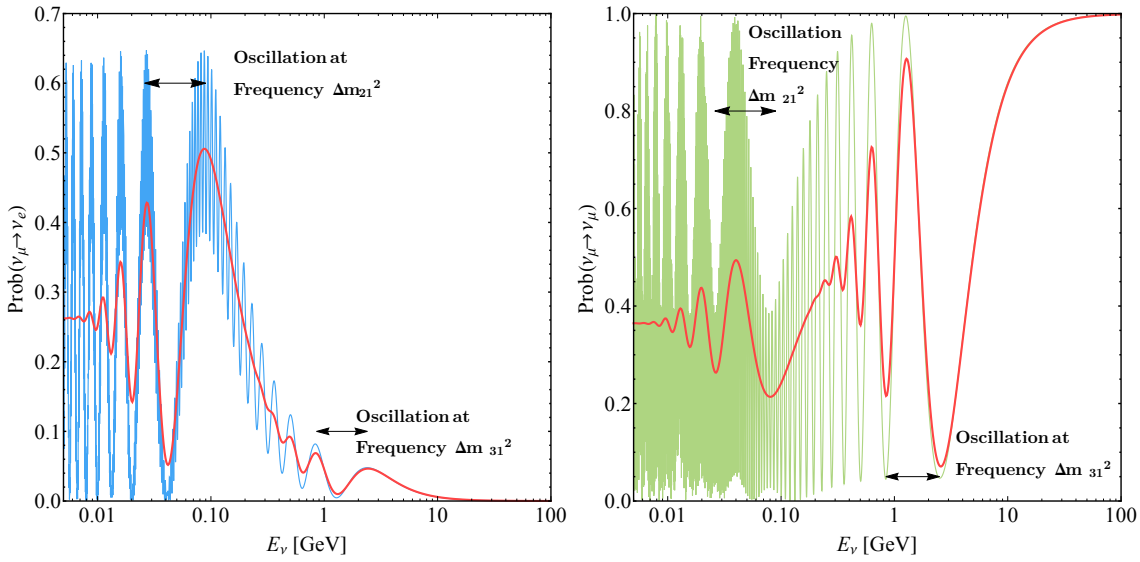

Figure 1.4: Examples of 3ν oscillation probabilities, for a hypothetical experiment with fixed 1300km baseline. On the *left* we show the $\nu_\mu \to \nu_e$ appearance oscillation probability (blue curve), and on the *right* the $\nu_\mu \to \nu_\mu$ disappearance probability (green curve). Very fast oscillations will not be resolvable in any given experiment so shown in both cases is a (red curve) where the analytic oscillation probability has been convoluted with a Gaussian to approximate finite energy resolution of 20% E_ν. This has the effect of averaging out very fast oscillations as can be seen. This plot assumes Normal Ordering and all mixing angles and mass splittings are fixed at the global best-fit values, see Table (1.2). Note that no real experiment would have sensitivity over such a wide L/E_ν, this is merely to highlight oscillation effects.

The Oscillation Term

For the majority of experiments, especially in the standard 3ν scenario, both the coherence term and localisation term effects have a negligible effect on the transition probability and can be safely ignored, with the probability reducing to the traditional form,

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L, E_\nu) = \sum_{j,m} U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m} \exp\Big[-2\pi i \frac{L}{L_{\text{osc}}} \Big]. \tag{1.4.46}$$

By expanding the exponential in terms of sines and cosines, and invoking the assumed unitarity of the U_{PMNS} matrix,

$$\sum_j |U_{\alpha j}|^2 |U_{\beta j}|^2 = \delta_{\alpha\beta} - 2\sum_{k>j} \text{Re}\left[U^\star_{\alpha k} U_{\alpha j} U^\star_{\beta j} U_{\beta k}\right], \tag{1.4.47}$$

we can then rewrite this probability into a sum of sin and $\sin^2$ frequencies

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L) = \delta_{\alpha\beta} - 4\sum_{m>j} \text{Re}\left[U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m}\right] \sin^2\left(\frac{\Delta m^2_{mj} L}{4E_\nu}\right) \tag{1.4.48}$$

$$+2\sum_{m>j} \text{Im}\left[U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m}\right] \sin\left(\frac{\Delta m^2_{mj} L}{2E_\nu}\right). \tag{1.4.49}$$

This equation governs the behaviour of all neutrino oscillations in a vacuum, or in situations where the matter effect is small enough to be ignored. Note that we have not made any assumptions on the number of neutrinos so this master equation is valid for any number of neutrino species. We plot an example of the $\nu_\mu \to \nu_e$ appearance oscillation probability and $\nu_\mu \to \nu_\mu$ disappearance probability in Figure (1.4), for the case of three neutrinos, and mixing angles according to a global fit of neutrino experiments. Neutrino experiments are not sensitive to very fast oscillations at frequencies below the achievable experimental energy resolution, we can account for this effect by convoluting the oscillation probability given in Equation (1.4.49) with a Gaussian with width appropriately chosen to represent the approximate energy resolution for a given experiment,

$$\mathcal{P}^{\text{avg}}(\nu_\alpha \to \nu_\beta : L, E_\nu, \sigma) = \int_0^\infty \text{d}E^\star \mathcal{P}(\nu_\alpha \to \nu_\beta : L, E^\star) \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{(E_\nu - E^\star)^2}{2\sigma^2}\right). \tag{1.4.50}$$

The effect of this can be clearly seen in Figure (1.4), where the fast oscillations are averaged and we are left with a constant normalisation shift in probability, effectively setting $\left\langle \sin^2 x \right\rangle \to 1/2$ for $L/E_\nu \gg \Delta m^2$, but the oscillations at smaller L/E_ν can be resolved fully.

With the understanding of how and why neutrino flavour oscillations occur, we will now turn out attention to the current global situation in the field of neutrino physics, as well as exploring the crucial results that lead to the definitive discovery of non-zero neutrino masses.

1.5 Global Situation of 3ν Paradigm

Our knowledge of the distinct PMNS neutrino mixing matrix elements comes from the plethora of successful experiments that have run since the first strong evidence for neutrino oscillations, interpreted as $\nu_\mu \rightarrow \nu_\tau$ oscillations, was discovered by Super-Kamiokande in 1998 [33]. A comprehensive review of the current status of the global situation in neutrino physics is beyond the scope of this work. It does, however, remain prudent to briefly discuss the main results of the field in order to motivate the subsequent chapters. For detailed reviews of the global oscillation fits see [34][9], [35, 36].

Experiments have long made use of naturally produced neutrinos, from nuclear reactions in the Sun, cosmic ray events in the upper atmosphere and even geo-neutrinos released in radioactive decays in the Earth's interior [37]. In the case of the two mass-splittings observed thus far, the so called Solar mass-splitting and Atmospheric mass-splitting, it was in natural sources of neutrinos that anomalies were first measured. Terrestrial man-made neutrino sources, such as neutrino super-beams [38,39] and Nuclear fission reactors, are crucial for confirming these anomalies as they often have a far better understood source with smaller systematic errors. In the three ν paradigm, there are two independent mass splittings corresponding to frequencies at which oscillations occur, as once two mass-splittings are defined the third is merely a linear combination of the others, e.g $\Delta m^2_{31} \equiv m^2_3 - m^2_1 = \Delta m^2_{32} - \Delta m^2_{21}$.

1.5.1 The Solar Parameters, $\mathbf{\Delta m^2_{21}}$ and θ_{12}

The first sign of something unusual in the neutrino sector came from the Homestake radio-chemical experiment in 1968 [40], and along with data collected over the next 25 years, showed only about a third of the expected solar neutrino flux was reaching earth, a discrepancy of over 3σ. Over the next decades many gallium based experiments, such as GALLEX [41] and SAGE [42] also saw a large deficiency in

[9] NuFIT 3.0 (2016), www.nu-fit.org.

expected ν_e rates from the sun. It was not until larger water Cherenkov detectors, such as Kamiokande and its successor Super-Kamiokande, and the heavy-water Sudbury Neutrino Observatory (SNO), began to publish results, that a clearer picture of what was happening emerged. SNO crucially was able to measure both the ν_e charged current reactions, as well as flavour independent neutral currents, of the ^{8}B ν_e solar neutrinos. The ratio of these measurement fluxes at SNO were so precise,

$$\frac{\Phi_{\rm CC}}{\Phi_{\rm NC}} = 0.340 \pm 0.023\ ({\rm stat})^{+0.029}_{-0.031}\ ({\rm sys}), \tag{1.5.51}$$

that they showed a staggering 17σ deviation from the expected value of 1 [43].

Although initial suggestions by Pontecorvo that the ν_e's were oscillating to another neutrino species as they propagated the vacuum from the Sun to the Earth [44], a full analysis of all solar data showed it preferred a solution in which the solar neutrinos underwent a strong resonant flavour transition due to the Mikheyev-Smirnov-Wolfenstein (MSW) matter effect [45, 46] as they streamed from the core of the Sun (where the vast majority of nuclear reactions take place) to the less dense outer layers of the Sun. The global best fit of all solar experiments favour a mass-splitting[10] of $\Delta m^2_{21} = 5 \times 10^{-5}$ eV^2 and a mixing angle θ_{12} of 33.8°. This oscillatory behaviour proved of course, that the neutrino mixing matrix must have off-diagonal elements.

Although this explanation provided a neat solution to the solar problem, a terrestrial experiment to measure Δm^2_{21} here on Earth in more laboratory conditions was proposed in the form of the Kamioka Liquid scintillator Anti-Neutrino Detector (KamLAND) experiment. KamLAND used MeV anti-neutrinos from Japan's 53 nuclear reactors, at a average distance of 180 km to directly measure the frequency and amplitude of the oscillation. KamLAND found values of the oscillation parameters of $\Delta m^2_{21} = 7.6 \times 10^{-5}$ eV^2 with a mixing angle $\theta_{12} = 34.2°$, in agreement with the MSW induced resonant solution of the solar experiments. In Figure (1.5) we plot the 1, 2 and 3σ ranges for both the solar experiments as well as most recent

[10] In the simplified two neutrino models originally used to analyse the data, this mass-splitting was referred to as the "solar mass-splitting", $\Delta m^2_{\rm solar}$.

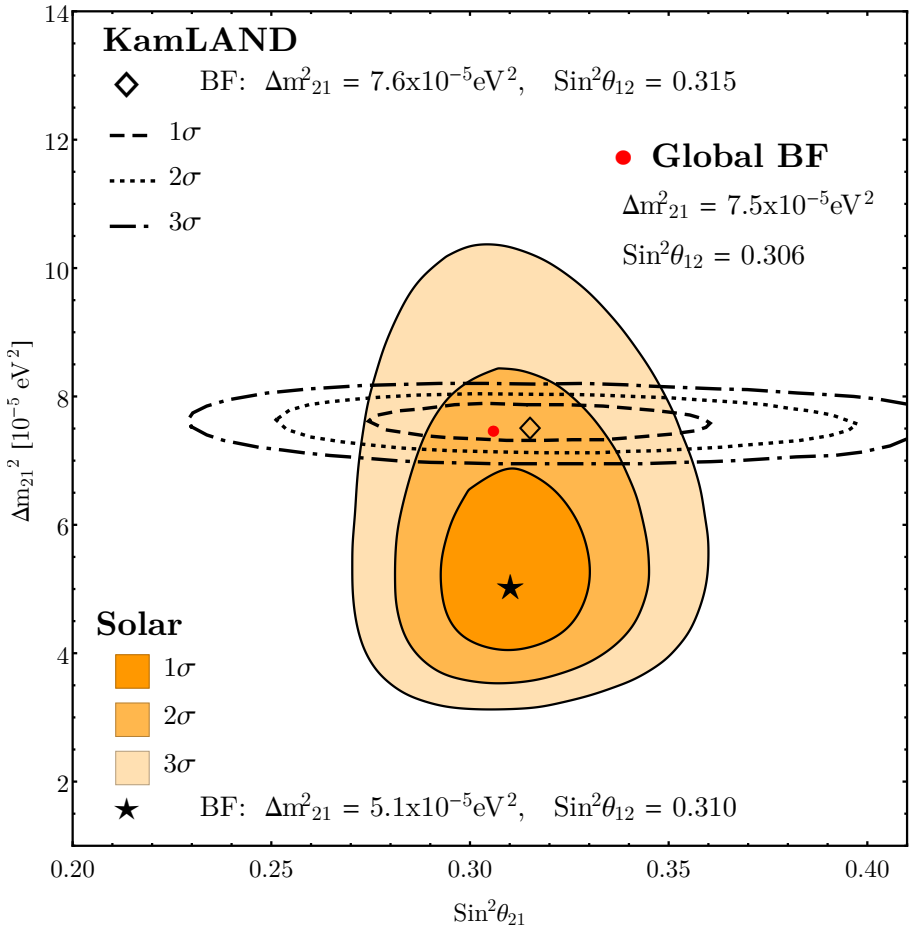

Figure 1.5: Measurements of the solar mass-splittings and θ_{12}, using both solar and terrestrial accelerator experiments. This plot assumes a fixed value of $\theta_{13} = 8.5°$. Results taken from the nu-fit global fits.

KamLAND measurements. One can see there is slight tension between the measured mass-splitting, with solar results favouring slightly lower values. However, one can also see the complementarity between the solar and long-baseline reactor results with KamLAND being stronger at measuring the frequency and the solar results bounding the amplitude to a greater extent.

1.5.2 The Atmospheric Parameters, Δm^2_{31} and θ_{23}

Cosmic rays interacting in the upper atmosphere are an important source of electron and muon neutrinos. In 1988 it was noticed that the number of ν_μ-like events expected in the Kamiokande experiment was about 60% of what was expected given

the knowledge of the atmospheric fluxes at the time [33]. This alone would not have warranted new-physics as the fluxes were not extremely known, however, the ν_e-like events observed were in complete agreement with the predictions and were highly correlated with the expected ν_μ fluxes. Disappearance of ν_μ due to neutrinos oscillations would explain this, but would require a very different frequency to that linked to the solar phenomena, one of $\mathcal{O}(100)$ times that of $\Delta m^2_{\text{solar}}$.

In 1998 Super-Kamiokande presented a $\approx 6\sigma$ result showing that ν_μ travelling upwards through the Earth disappear when compared to down-going events. This was consistent with the theory that the upward going ν_μ neutrinos were oscillating away to ν_τ neutrinos, while the downward going neutrinos do not propagate for a long enough distance to disappear. If the oscillation model was to be believed it indicated a mass-splitting of $\Delta m^2_{\text{atmos}} \approx 2.4 \times 10^{-3}$ eV^2 and maximal mixing, $\sin^2 2\theta_{\text{atmos}} \approx 1$.

Super-Kamiokande took data at L/E_ν values from 1 km/GeV all the way up to 10^4 km/GeV, and it is was in the upper half, $\geqslant 10^2$ km/GeV, in which the deficit was observed. Unlike the solar-scale oscillations, which would require baselines of $\mathcal{O}$(15,000 km) to see with standard super-beam neutrino experiments, the atmospheric oscillation region is easily probe-able using terrestrial ν_μ and $\overline{\nu}_\mu$ super-beams. The K2K experiment in Japan was the first experiment to attempt this check of atmospheric oscillations, consisting of a ≈ 1 GeV ν_μ beam fired over 250km to the Super-Kamiokande detector. KEK observed 107 events that were fully contained and reconstructable, in comparison to 151 expected in the absence of oscillation. The observed events were also crucially in spectral agreement with oscillations of maximal mixing and $\Delta m^2 = 2.8 \times 10^{-3}$eV^2 [33].

The majority of our current sensitivity to Δm^2_{32} and θ_{23} come from three accelerator ν_μ disappearance experiments; **MINOS**, an on-axis magnetised steel-scintillator detector using the Fermilab NuMI beam with a baseline of 735 km, **T2K** the spiritual successor to K2K, using a 30 GeV proton beam at J-PARC sent slightly off axis

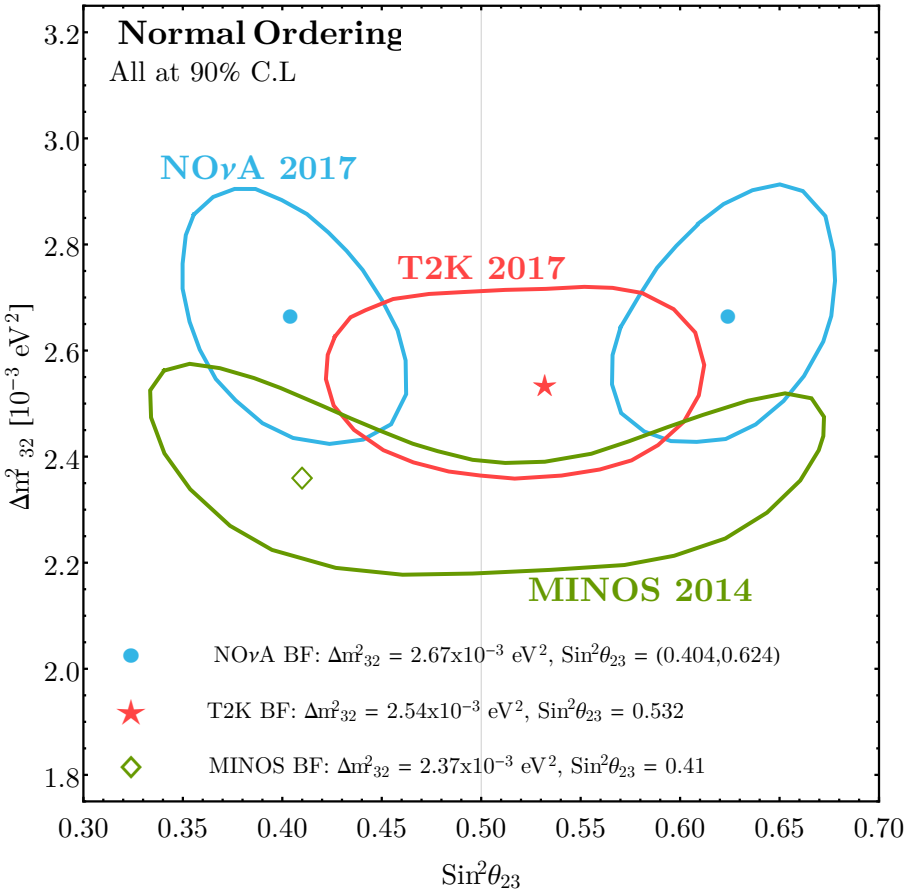

Figure 1.6: Measurements of the atmospheric mass-splitting and mixing angle, showing the slight tension in the most recent results from the three main accelerator experiments NOνA [47], T2K [48] and MINOS [49]. Although MINOS and T2K are consistent with maximal mixing, $\theta_{23} = \pi/4$, NOνA on its own disfavours it at 2.6σ significance.

(2.8 °) 285 km to the Super-Kamiokande water Cherenkov detector and **NOνA** , also using Fermilab's NuMI beam, but directed off-axis by $\approx 1^\circ$ 810km to a mineral oil scintillator detector. All three use near detectors at $\mathcal{O}(100)$m to measure the flux *before* oscillation, and can run in anti-neutrino mode to probe CP violating effects. The final run of MINOS has come to an end, but NOνA and T2K continue to collect data, with their most recent publications in January 2017 [47, 48]. The most up to date results from these three detectors can be seen in Figure (1.6), which include $\nu_\mu \to \nu_\mu$ and $\overline{\nu}_\mu \to \overline{\nu}_\mu$ disappearance as well as $\nu_\mu \to \nu_e$ appearance data. As can be seen, both MINOS and T2K are consistent with maximal mixing in the atmospheric

sector, however, NOνA favours a non-maximal value, with $\theta_{23} = \pi/4$ being excluded at the 2.6σ level.

In addition to the question of whether or not θ_{23} is maximal, there is also the question of which octant it falls into. In a simple 2 ν approximation, measurements of θ_{23} using ν_μ disappearance are completely degenerate in this regard, with the oscillation probability being proportional to $\sin^2 2\theta$. This degeneracy is weakly broken when one introduces measurements of $\nu_\mu \rightarrow \nu_e$ appearance, with the probability containing terms proportional to $\sin^2\theta_{23}$ and $\sin 2\theta_{23}$. This can be directly observed in the nearly symmetric bounds around $\sin^2\theta_{23} = 0.5$ in Figure (1.6). The determination of this octant is one of the primary remaining goals of neutrino oscillation physics.

1.5.3 The Reactor Parameter, θ_{13}

Nuclear reactors provide an intense source of anti-electron neutrinos produced by the β-decay of many neutron-rich nuclei produced in the decay chain of Uranium (235 and 238) and Plutonium (239 and 241). These anti-neutrinos have energies of $\mathcal{O}$(MeV) and stream isotropically from the reactor cores. Detection of electron anti-neutrinos takes place through the observation of a twin coincidence signal consisting of a prompt positron alongside that of a delayed neutron capture signal, both released by the electron anti-neutrino in an inverse neutrino decay reaction, $\overline{\nu}_e + p \rightarrow e^+ + n$.

KamLAND has already been mentioned as a strong measurement of the solar oscillation parameters through observation of an oscillation at a frequency Δm^2_{21}. Here we wish to consider a second generation of reactors which looked for oscillations at the much larger atmospheric mass-splitting, Δm^2_{31}. This was driven by the desire to measure the final mixing angle of the U_{PMNS} matrix, θ_{13}. θ_{13} was long known to be much smaller than the solar and atmospheric mixing angles, due to non-observation at reactor experiments, such as CHOOZ [50], and there was much theoretical motivation for a zero or near-zero values of θ_{13} through the study of

discrete symmetries [51, 52]. These $\mathcal{O}$(km) reactor measurements provide a measurement of θ_{13} that is crucially independent of other mixing angles

$$P(\overline{\nu}_e \to \overline{\nu}_e) \approx 1 - \sin^2 2\theta_{13} \sin\left(\frac{\Delta m^2_{31} L}{4E_\nu}\right). \tag{1.5.52}$$

In the case that θ_{13} was indeed very small, several experiments were designed to utilise the intense nuclear reactors at a short-baseline $\mathcal{O}$(1 km). The first measurement of a non-zero θ_{13} was announced by the Daya Bay collaboration at a significance of 5.2σ in 2012 [53], after 55 days exposure and only 5 years after the experiment was proposed in in 2007 [54]. Daya Bay achieved this rapid measurement by utilising eight identical detectors at three different locations relative to six 2.7 $\mathrm{GW}_{\mathrm{thermal}}$ fission reactors. The measured value of $\theta_{13} = 8.8^\circ$ was larger than many expected, which greatly helped the speed at which it was discovered. This was soon verified by the Double CHOOZ [55] and RENO [56] collaborations.

The most recent measurement of Daya Bay using 1230 days of exposure measure $\sin^2 2\theta_{13} = 0.0841 \pm 0.0027$ stat ± 0.0019 sys and $\Delta m^2_{32} = (2.45 \pm 0.06$ stat $\pm$ 0.06 sys) $\times 10^{-3}$ eV^2 [57], assuming Normal Ordering. This represents the most precise measurement of any neutrino parameter thus far, with over 2.5 million $\overline{\nu}_e$ events recorded.

1.5.4 The CP violating phase, δ_{CP}

The discovery of not only a non-zero θ_{13}, but a *large* non-zero θ_{13} at approximately 8.5° opened up the possibility of measuring CP violation in the neutrino sector. If one studies the mixing matrix in Equation (1.3.24) , we notice that the possible CP violating phase, δ_{CP}, always enters attached to $\sin^2\theta_{13}$. If θ_{13} was indeed zero, then all neutrino and anti-neutrino oscillation probabilities in the 3ν paradigm would be equal. One can see this explicitly by looking at an analytical approximation to the $\nu_\mu \to \nu_e$ and $\overline{\nu}_\mu \to \overline{\nu}_e$ appearance probabilities in a constant matter potential V

directly [55],

$$P(\nu_\mu \to \nu_e) \approx 4\sin^2\theta_{13}\sin^2\theta_{23}\frac{\sin^2\Delta_{31}(1-A)}{(1-A)^2} + \alpha^2\sin^2 2\theta_{12}\cos^2\theta_{23}\frac{\sin^2 A\Delta_{31}}{A^2} + \tag{1.5.53}$$

$$+ \underbrace{2\alpha\sin\theta_{13}\sin 2\theta_{12}\sin 2\theta_{23}\cos(\Delta_{31} \pm \delta_{\rm CP})}_{\text{Only term sensitive to CP violations}}\frac{\sin\Delta_{31}A}{A}\frac{\sin\Delta_{31}(1-A)}{(1-A)}, \tag{1.5.54}$$

where the $\pm$ refers to neutrino and anti-neutrino probabilities and for clarity we have defined,

$$\Delta_{31} \equiv \frac{\Delta m^2_{31}L}{4E_\nu}, \qquad A \equiv \frac{2E_\nu V}{\Delta m^2_{31}}. \tag{1.5.55}$$

The above probability an approximation to second order in the two small parameters $\alpha \equiv \Delta m^2_{21}/\Delta m^2_{31} \approx 0.032$ and $\sin\theta_{13} \approx 0.14$. As matter is made up of neutrinos and not anti-neutrinos, the matter potential, V, also changes sign when looking at anti-neutrino oscillation probabilities. This interferes with the search for non-zero $\delta_{\rm CP}$, as matter-effects can induce CP violations in the observed events, even with $\delta_{\rm CP}$ is CP conserving.

Precise measurement of $\delta_{\rm CP}$ is one of the primary physics goals of two of the largest next generation long-baseline neutrino experiments; DUNE (Deep Underground Neutrino Experiment) [58], which will fire a beam of predominantly ν_μ from Fermilab ≈1300km to the Homestake gold mine in South Dakota, and Hyper-Kamiokande [59], a gigantic 1 million metric ton water Cherenkov detector based on the highly successful Super-Kamiokande experiment. Depending on the exact value of the remaining oscillation parameters, exact detector and accelerator final design choices, DUNE should be capable of achieving $\leqslant 30^\circ(10^\circ)$ resolution of $\delta_{\rm CP}$ with a total exposure of $\approx$ 120(1000) kt-MW-years, and to exclude CP conservation in $\approx 50\%$ of $\delta_{\rm CP}$ values at $\geqslant 3\sigma$ significance [60]. DUNE is expected to achieve $\approx$ 120 kt-MW-years by ≈2035. Similarly, Hyper-Kamiokande should be able to achieve 18° resolution of $\delta_{\rm CP}$, for all possible values of $\delta_{\rm CP}$ with approximately 5 years data collection [59].

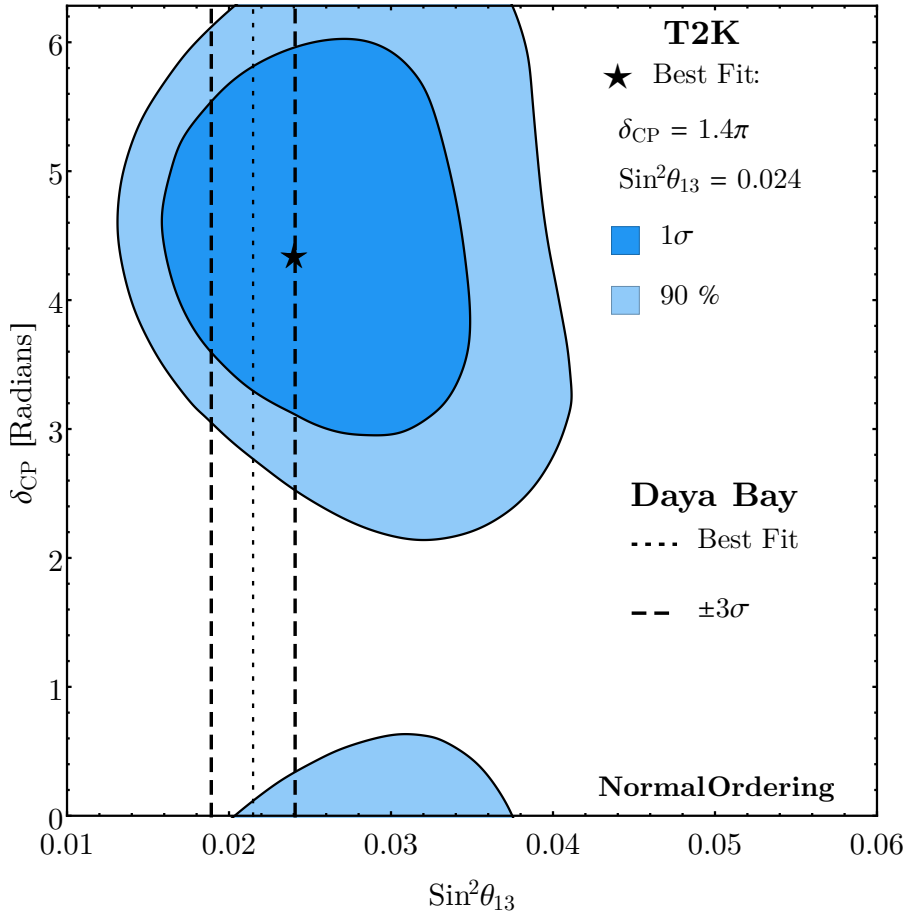

Figure 1.7: First results from T2K [48] using a combined electron-neutrino and anti-electron neutrino appearance analysis. Current preference is for values of δ_{CP} close to maximally CP violating, although significance of this statement is low. Shown also is the independent measurement of θ_{13} from reactor anti-neutrino disappearance at Daya Bay [57].

In the meantime, direct searches for δ_{CP} are undergoing. The first exciting result of a dual $\nu_\mu \to \nu_e$ and $\overline{\nu}_\mu \to \overline{\nu}_e$ accelerator search at T2K has been released, showing a weak preference to the maximally CP violating value of $\delta_{CP} = 3\pi/2$ [48]. Crucially the measured values of θ_{13} and Δm^2_{31} are consistent, and if one uses the reactor θ_{13} measurements as a prior the suggested range for δ_{CP} shrinks further, as can be seen in Figure (1.7).

We collect all the above information on the current best fit and $\pm 1\sigma$ ranges of the neutrino mixing angles and mass-splittings in Table (1.2).

Parameter	Global Best-Fit
θ_{12}	$33.56^{\circ}{}^{+0.77}_{-0.75}$
θ_{23}	$41.6^{\circ}{}^{+1.5}_{-1.2}$ $(50.0^{\circ}{}^{+1.1}_{-1.4})$
θ_{13}	$8.46^{\circ} \pm 0.15$
δ_{CP}	$261^{\circ}{}^{+51}_{-59}$
Δm^2_{21}	$7.50^{\circ}{}^{+0.19}_{-0.17} \times 10^{-5}$ eV^2
$\lvert\Delta m_{31}\rvert^2$	$2.524^{+0.039}_{-0.040} \times 10^{-3}$ eV^2

Table 1.2: Global best-fit values for neutrino mixing parameters in 3ν paradigm, as taken from the nu-fit global fit analysis [34].

1.5.5 The Mass Ordering

Although the sign of Δm_{21} is known from the MSW interpretation of solar data, it must be noted that only $|\Delta m^2_{31}|$ is known precisely. Thus there are two distinct orderings for the neutrino masses which are possible: the so-called normal ordering (NO) in which $m_1 < m_2 < m_3$, and the inverted ordering (IO) where $m_3 < m_1 < m_2$. This is shown qualitatively in Figure (1.8). Note that here we make a distinction between the mass ordering and the "mass hierarchy". The mass hierarchy is a related concept but is defined relative to the absolute neutrino mass scale and can be normal ($m_3 \gg m_2, m_1$), inverted ($m_2, m_1 \gg m_3$) or quasi-degenerate ($m_1 \approx m_2 \approx m_3$) which would occur if $m_1^2 \gg |\Delta m^2_{31}|$. [11]

1.6 Thesis Summary

There is much evidence for the validity of the 3ν paradigm, but discoveries arise from repeated strict validation of a model and any discrepancies found whilst doing so. It is crucial that we continue to probe the validity of the 3ν paradigm over the coming

[11] Note, if the mass hierarchy is inverted or normal, so too is the ordering. But one can have a degenerate hierarchy alongside either an inverted or normal ordering. This distinction is not universal in the literature.

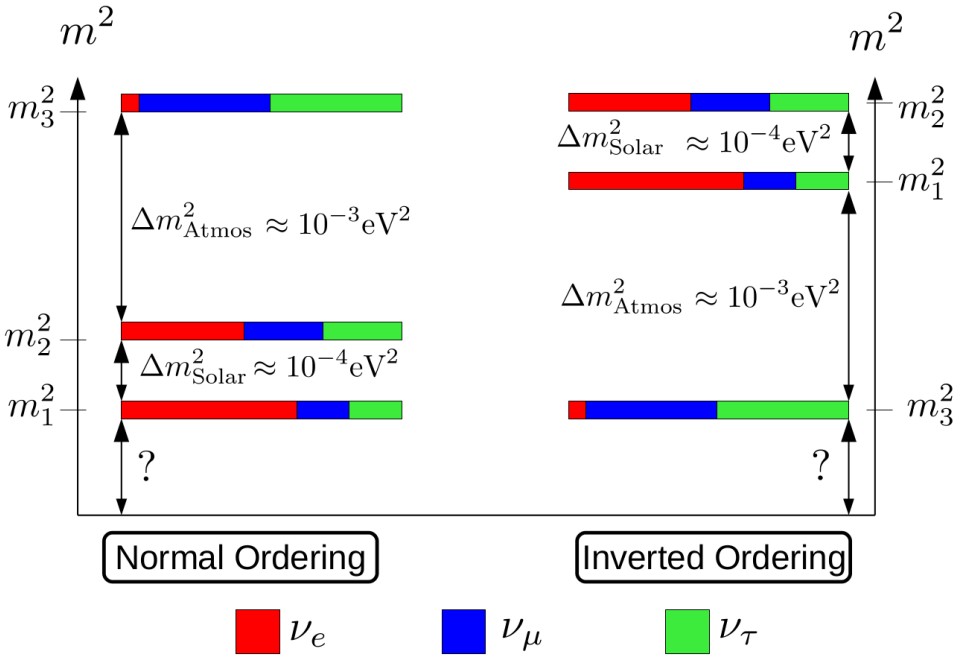

Figure 1.8: The normal (left) and inverted (right) mass orderings of the three neutrino paradigm. Neutrino oscillations are insensitive to the absolute mass scale of the neutrinos, represented by the position of the y axis relative to zero. The colours represent the approximate flavour components of each mass eigenstate.

years. In Chapter 2 we will investigate one such signature that we live in a non-3ν paradigm world, the possible non-unitarity of the 3×3 $U_{\rm PMNS}$ matrix. Unitarity is a fundamental property of any theory required to ensure we work in a theoretically consistent framework. In comparison with the quark sector, experimental tests of unitarity for the 3×3 neutrino mixing matrix are considerably weaker. We perform a reanalysis to see how global knowledge is altered when one refits oscillation results without assuming unitarity, and present 3σ ranges for allowed $U_{\rm PMNS}$ elements consistent with all observed phenomena. We calculate, for the first time, bounds on the closure of the six neutrino unitarity triangles, with the closure of the $\nu_e\nu_\mu$ triangle being constrained to be $\leqslant 0.03$, while the remaining triangles are significantly less constrained to be $\leqslant 0.1$ - 0.2. Similarly for the row and column normalization, we find their deviation from unity is constrained to be $\leqslant 0.2$ - 0.4, for four out of six such normalisations, while for the ν_μ and ν_e row normalisation the deviations are constrained to be $\leqslant 0.07$, all at the 3σ CL. We emphasise that there is signif-

icant room for new low energy physics, especially in the ν_τ sector which very few current experiments constrain directly. As the canonical method for introducing a non-unitarity is through the addition of sterile neutrinos, we end the chapter with an introduction to sterile neutrino phenomenology.

Although we argue the discrepancy between the amounts of allowed non-unitarity from experimental measurements and the requirement of strict unitarity in the 3ν paradigm is reason enough to investigate sterile neutrinos, there is also a number of experimental anomalies that hint towards the existence of such sterile states. Perhaps the strongest motivation for studying low-scale sterile neutrinos is the anomalous $\overline{\nu}_e$ appearance at the LSND and MiniBooNE experiments, consistent with neutrino flavour oscillations driven by a new mass-splitting at $\mathcal{O}(1\ \mathrm{eV}^2)$. In Chapter 3 we discuss this exciting anomaly in greater detail, exploring the phenomenology of low-scale oscillating sterile neutrinos. We proceed to investigate the ability of the Short Baseline Neutrino (SBN) experimental program being built at Fermilab to test the globally-allowed (3+N) sterile neutrino oscillation parameter space, as motivated by the LSND and MiniBooNE anomalies. We explicitly consider the globally-allowed parameter space for the (3+1), (3+2), and (3+3) sterile neutrino oscillation scenarios. We find that SBN can probe with $> 5\sigma$ sensitivity more than 85%, 95% and 55% of the parameter space allowed at 99% confidence level for the (3+1), (3+2) and (3+3) scenarios, respectively. In the case of the (3+2) and (3+3) scenarios, CP-violating phases appear in the oscillation probability terms, leading to observable differences in the appearance probabilities of neutrinos and antineutrinos. We explore SBN's sensitivity to those phases for the (3+2) scenario through the currently planned neutrino beam running, and investigate potential improvements through additional antineutrino beam running. We show that if antineutrino exposure is considered, for maximal values of ϕ_{54}, SBN could be the first experiment to directly observe $\approx 2\sigma$ hints of CP violation in the lepton sector.

When one considers "low-scale" sterile neutrinos, it is almost always light oscillating sterile signatures, such as those described in Chapter 3, that are being

discussed. In the final Chapter, we will move our discussion to an alternative signature of low-scale sterile natures involving somewhat heavier states. Nearly-sterile neutrinos with masses in the MeV range and below would be produced in the beam of the SBN program at Fermilab. In this Chapter, we study the potential for SBN to discover these particles through their subsequent decays in its detectors. We discuss the decays which will be visible at SBN in a minimal and non-minimal extension of the Standard Model, and perform simulations to compute the parameter space constraints which could be placed in the absence of a signal. We demonstrate that the SBN programme can extend existing bounds on well constrained channels, such as $N \to \nu l^+ l^-$ and $N \to l^\pm \pi^\mp$, while, thanks to the strong particle identification capabilities of liquid-Argon technology, also place bounds on often neglected channels, such as $N \to \nu\gamma$ and $N \to \nu\pi^0$. Furthermore, we consider the phenomenological impact of improved event timing information at the three detectors. As well as considering its role in background reduction, we note that if the light-detection systems in SBND and ICARUS can achieve nanosecond timing resolution, the effect of finite sterile neutrino mass could be directly observable, providing a smoking-gun signature for this class of models.

Although perhaps less theoretically motivated from a neutrino mass generation perspective, sterile neutrinos at the eV and MeV scales provide an intensely phenomenologically rich environment at short-baselines at which to probe the leptonic sector. By studying the possible signatures of one or more low-scale sterile neutrinos, be it their effect on the unitarity of the 3×3 U_{PMNS} matrix, their oscillatory effects, direct observation of their decay products or a different new anomalous signature altogether, their discovery would open up a window up to a new sector outside the current Standard Model, and revolutionise our understanding of the universe. The Fermilab SBN program is arriving at a crucial point at which it will be able confirm or deny the light sterile neutrino hypothesis for LSND and MiniBooNE anomalies, and we stress throughout that a search for decaying heavy sterile neutrinos is a complementary new physics analysis to this search for eV-scale oscillations, and would extend the Beyond the Standard Model physics program of SBN while requiring no

additional beam or detector modifications. These multi-pronged search strategies are necessary if we are to probe in depth the vast parameter space for sterile neutrinos that, we will see, our current experimental understanding of the neutrino mixing matrix allows.

Chapter 2

Unitarity of the 3×3 U_{PMNS} Matrix

"I never approve, or disapprove, of anything now.
It is an absurd attitude to take towards life."
Lord Henry Wotton

2.1 Unitarity

With the knowledge of $\sin^2 2\theta_{13}$ now at the sub 5% level, and interplay between the long baseline accelerator $\nu_\mu \rightarrow \nu_e$ appearance data [61, 62] and short baseline reactor $\overline{\nu}_e \rightarrow \overline{\nu}_e$ disappearance [53, 55, 56] data, combined with prior knowledge of θ_{23} from $\nu_\mu \rightarrow \nu_\mu$ disappearance data [63–65], suggesting tentative global hints at $\delta_{CP} \approx 3\pi/2$, there is much merit to argue that we are now in the precision measurement era of neutrino physics.

However, one must always remember that our knowledge of the matrix elements comes predominately from high statistics $\overline{\nu}_e$ disappearance and ν_μ disappearance experiments, with the concept of unitarity being invoked to disseminate this information onto the remaining elements. The concept of mixing angles is only valid as a parameterisation if the unitarity of the PMNS matrix is assumed. Figure (2.1) gives an idea of which experiments bound which U_{PMNS} *elements*, rather than the mixing *angles* as was discussed above.

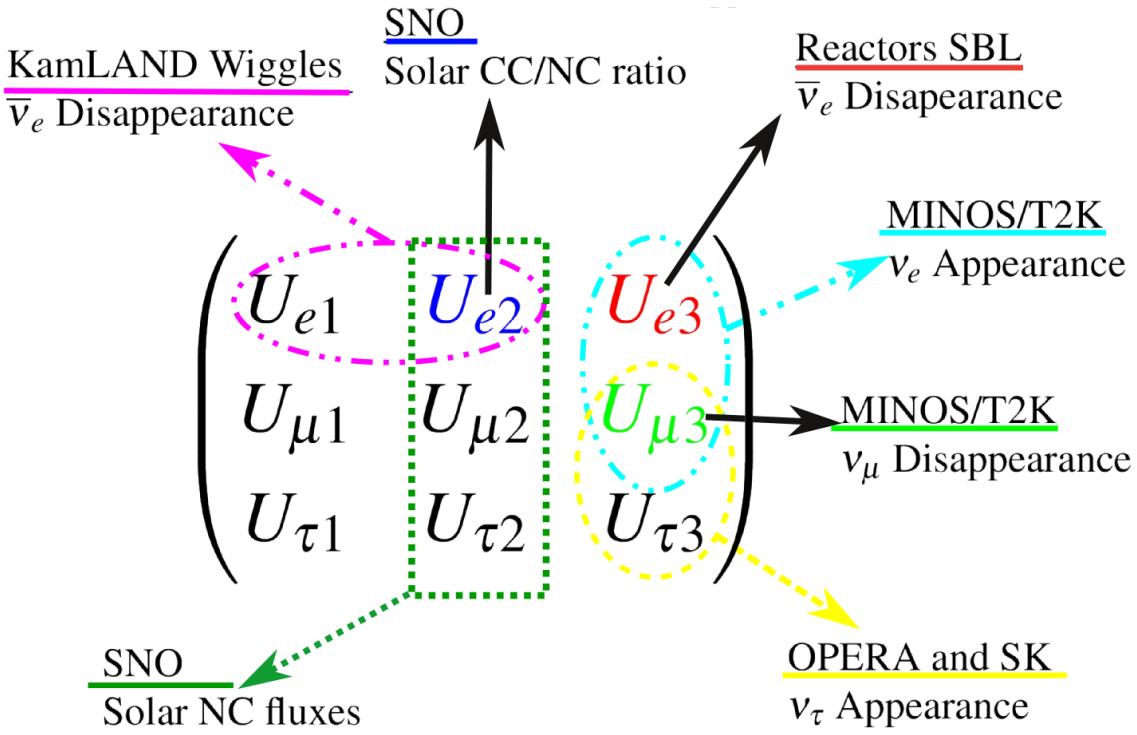

Figure 2.1: Info-graphic indicating where the strongest bounds on combinations of U_{PMNS} elements arise from. As can be seen, the majority of historic high-statistics measurements focus on the ν_e-sector (and to a lesser extent $U_{\mu 3}$), with the assumption of unitarity being invoked to disseminate this information to the remaining U_{PMNS} elements.

Unitarity of a mixing matrix is a necessary condition for a theoretically consistent description of the underlying physics, as non-unitarity directly corresponds to a violation of probability in the calculated amplitudes. In the neutrino sector unitarity can be directly verified by the precise measurement of each of the mixing elements to confirm the unitarity condition: $U^\dagger U = \mathbb{1} = UU^\dagger$. In this there are twelve conditions, six of which we will refer to as normalisations (sum of the squares of each row or column, e.g the ν_e normalisation $|U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2 = 1$) and six conditions that measure the degree with which each unitarity triangle closes (e.g the $\nu_e\nu_\mu$ triangle: $U_{e1}U^*_{\mu 1} + U_{e2}U^*_{\mu 2} + U_{e3}U^*_{\mu 3} = 0$). The concept of a unitarity triangle closure is diagrammatically explained in Figure (2.2) for the $\nu_e\nu_\mu$ triangle. See X. Qian et al. [66] for a detailed discussion of the current and future state of measurements of the ν_e normalisation. Before any measurements take place, all we know for

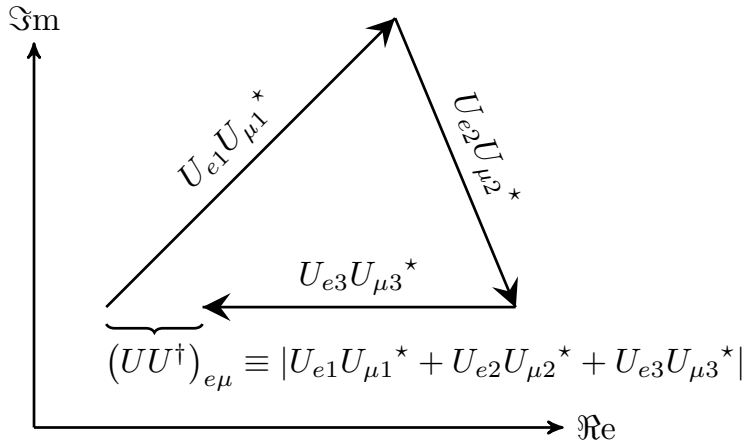

Figure 2.2: Any three complex numbers drawn in the complex plane will form a triangle if their sum is equal to zero. The *unitarity triangle closure* is the Real number defined as the absolute value of the sum, and is representative of the degree by which any pair of rows or columns violate unitarity, a value of 0 indicating the triangle closes.

certain is the ordering, $|U_{e1}| \geqslant |U_{e2}| \geqslant |U_{e3}|$, as the mass eigenstates are defined and labelled by decreasing ν_e content.

In the quark sector, the analogous situation involving the Cabibbo-Kobayashi-Maskawa (CKM) matrix has been subject to intense verification, as many experiments have access to all of the V_{CKM} elements individually. Current data shows that the assumption of unitarity for the 3×3 CKM matrix is valid in the quark sector to a high precision, with the strongest normalisation constraint being $|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 0.9999 \pm 0.0006$, and the weakest still being significant at $|V_{ub}|^2 + |V_{cb}|^2 + |V_{tb}|^2 = 1.044 \pm 0.06$ [23]. Unlike the quark sector, however, experimental tests of unitarity are considerably weaker in the 3×3 U_{PMNS} neutrino mixing matrix. It remains an initial theoretical assumption inherent in many analyses [35, 67, 68], but is the basis for the validity of the 3ν paradigm.

This non-unitarity can arise naturally in a large variety of theories. A generic feature of many Beyond the Standard Model scenarios is the inclusion of one or more new massive fermionic singlets, uncharged under the Standard Model gauge

group, $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. If these new states mix with the Standard Model neutrinos, then the true mixing matrix is enlarged from the 3×3 U_{PMNS} matrix, an example of which is to a nxn matrix,

$$U_{\text{PMNS}}^{\text{Extended}} = \begin{pmatrix} \overbrace{\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}}^{U_{\text{PMNS}}^{3\text{x}3}} & \cdots & U_{en} \\ \vdots & \vdots & \vdots \\ U_{s_n 1} & U_{s_n 2} & U_{s_n 3} & \cdots & U_{s_n n} \end{pmatrix}.$$

These so-called sterile neutrinos have been a major discussion point for both the theoretical and experimental communities for decades. *A priori* these new states can sit at practically any mass as there is no known symmetry to dictate a scale. Any given subset of the total extended mixing matrix, such as the 3×3 subset involved in neutrino oscillations, will thus not necessarily be unitary as the known neutrino mass eigenstates may contain some admixture of the new sterile states. This is the canonical model of how new physics, introduced at any scale, breaks observed unitarity in the neutrino sector. We will discuss sterile neutrinos in detail in Section (2.2), and although they are the chief method invoked in the literature to generate a Non-Unitarity, we believe focusing on the unitarity of the U_{PMNS} matrix alone is a worthwhile exercise.

If this physics enters solely at a high scale, as in the Minimal Unitarity Violation (MUV) scheme [69], then one can utilise weak decays, rare lepton decays (e.g. $\mu \to e\gamma$) and EW precision measurements to bound the amount of non-unitarity to the level of 0.5%, with 90% C.L ranges of [70];

$$\left|UU^\dagger\right|_{\text{MUV}} = \begin{pmatrix} 0.9979 - 0.9998 & < 10^{-5} & < 0.0021 \\ < 10^{-5} & 0.9996 - 1.0 & < 0.0008 \\ < 0.0021 & < 0.0008 & 0.9947 - 1.0 \end{pmatrix}. \tag{2.1.2}$$

Here we consider the alternative case in which the new physics that provides this non-unitarity enters at a relatively low scale, as several current experimental hints suggest, with anomalous results from LSND [71], MiniBooNE [72], theggallium anomaly [41, 73] and the Reactor anomaly [74]. In this regime neutrino oscillations are the most important experimental probe we have access to. The most convincing means of verification of unitarity in the neutrino sector would be analogous to the quark sector, via direct and independent measurement of all the U_{PMNS} elements, to overconstrain the parameter space and confirm that the 12 unitarity constraints hold to within experimental precision. However, we do not currently have access to enough experiments in the ν_μ and ν_τ sectors to bound all of the elements to a sufficient degree to verify all 12 conditions. Thus we must look for alternative ways to constrain the U_{PMNS} elements.

One can perform indirect searches of unitarity by searching for mixing elements outside those of the 3ν mixing regime. These class of searches do not measure the 3×3 mixing elements per se, but rather by looking for additional states one can constrain the violations they would induce in the 3×3 subset. One proceeds by noting all null results at frequencies distinct to those of the 3ν paradigm. We do not wish to perform a global fit for new physics, as this has been well covered in the literature [75, 76], instead we focus on what unresolved physics can do to our current precision, hence we do not include any positive signals such as LSND or the MiniBooNE anomaly.

Such a sterile driven approach requires additional assumptions on the exact origin of the non-unitarity, thus losing some model-independence. However, as an extended U_{PMNS} matrix encompasses many beyond the Standard Model scenarios, it is natural to include this in our analysis. To proceed one must then consider what scale the new physics enters at. However, as we do not focus on the origin of such non-unitarity we choose to marginalise over the new scale(s) assuming the possibility that they enter in at an oscillating scale, with at least $|\Delta m^2| \geqslant 10^{-2}$ eV^2. Below this scale, states degenerate with Standard Model neutrinos leads to requirements

of a much more detailed analysis.

A non-unitary mixing matrix can be parameterised as a 3×3 matrix hosting 9 complex non-unitary elements, 5 phases of which can be removed by rephasing the lepton fields, leaving 13 parameters: 9 real positive numbers and 4 phases. There are many ways to parametrise this matrix, e.g [77], however, for clarity we choose to keep it directly in terms of its matrix elements. The oscillation probability for a neutrino (anti-neutrino) of initial flavour α and energy E_ν to transition to a neutrino (anti-neutrino) of flavour β after a distance L with such a non-unitary mixing matrix is given by

$$\begin{aligned} P\left(\overset{(-)}{\nu}_\alpha \rightarrow \overset{(-)}{\nu}_\beta\right) = & \left|\sum_{i=1} U^*_{\beta i} U_{\alpha i}\right|^2 \\ & - 4\sum_{i<j} \mathrm{Re}(U_{\beta i} U^*_{\beta j} U_{\alpha i} U^*_{\alpha j}) \sin^2\left(\Delta m^2_{ji} \frac{L}{4E_\nu}\right) \\ & \overset{(-)}{+}\ 2\sum_{i<j} \mathrm{Im}(U_{\beta i} U^*_{\beta j} U_{\alpha i} U^*_{\alpha j}) \sin\left(\Delta m^2_{ji} \frac{L}{2E_\nu}\right), \end{aligned} \tag{2.1.3}$$

where now, without assuming unitarity, the leading term is not a function of $\Delta m^2 L/E_\nu$ and is also not necessarily equal to 1 or 0 in neutrino disappearance and appearance experiments respectively. This term has been called the "zero-distance" or "instantaneous oscillation probability" in the literature.

Although violations of unitarity such as these modify the oscillation amplitudes and total normalisation of the probability, they do not have any effect on the oscillation frequency, which remains a function of the mass differences and L/E_ν only (ignoring higher order non-unitary matter effects). Thus, for simplicity of analysis, the global best fit values for the mass squared differences are assumed ($\Delta m^2_{21} = 7.6 \times 10^{-5}\mathrm{eV}^2, |\Delta m^2_{31}| = 2.4 \times 10^{-3}\mathrm{eV}^2$) [23].

For each observed oscillation one can then directly compare the measured amplitude with the non-unitary expression for the oscillation probability. It is this amplitude-matching that we use to undertake a global-fit and provides us the ranges for U_{PMNS} that would successfully reproduce the measured oscillation amplitudes and

normalisations. A table containing some example non-unitary amplitudes, as well as the corresponding unitary values, for a variety of appearance and disappearance neutrino experiments, is shown in Table (2.1). Also given is the row normalisation, or unitarity triangle-closure normalisation that the corresponding experiment can probe if the total flux uncertainty is sufficiently known. As a concrete example, take the $\overline{\nu}_e \to \overline{\nu}_e$ oscillations as observed by short-baseline reactor experiments, such as Daya Bay and Reno

$$P_{\overline{\nu}_e \to \overline{\nu}_e} = \left(|U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2\right)^2 \left(1 - \underbrace{\frac{4(|U_{e1}|^2 + |U_{e2}|^2)|U_{e3}|^2}{\left(|U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2\right)^2}}_{\text{Visible Amplitudes}} \sin^2\left(\frac{L\Delta m^2_{31}}{4E_\nu}\right) \right),$$

$$\text{if Unitary} = \quad 1 - \overbrace{\sin^2 2\theta_{13}}^{\text{Compare Directly}} \sin^2\left(\frac{L\Delta m^2_{31}}{4E_\nu}\right). \tag{2.1.4}$$

Provided the normalisation is within experimental error, it is impossible to distinguish a measured amplitude that originated from a single degree of freedom, θ_{13}, or a degenerate combination of mixing elements.

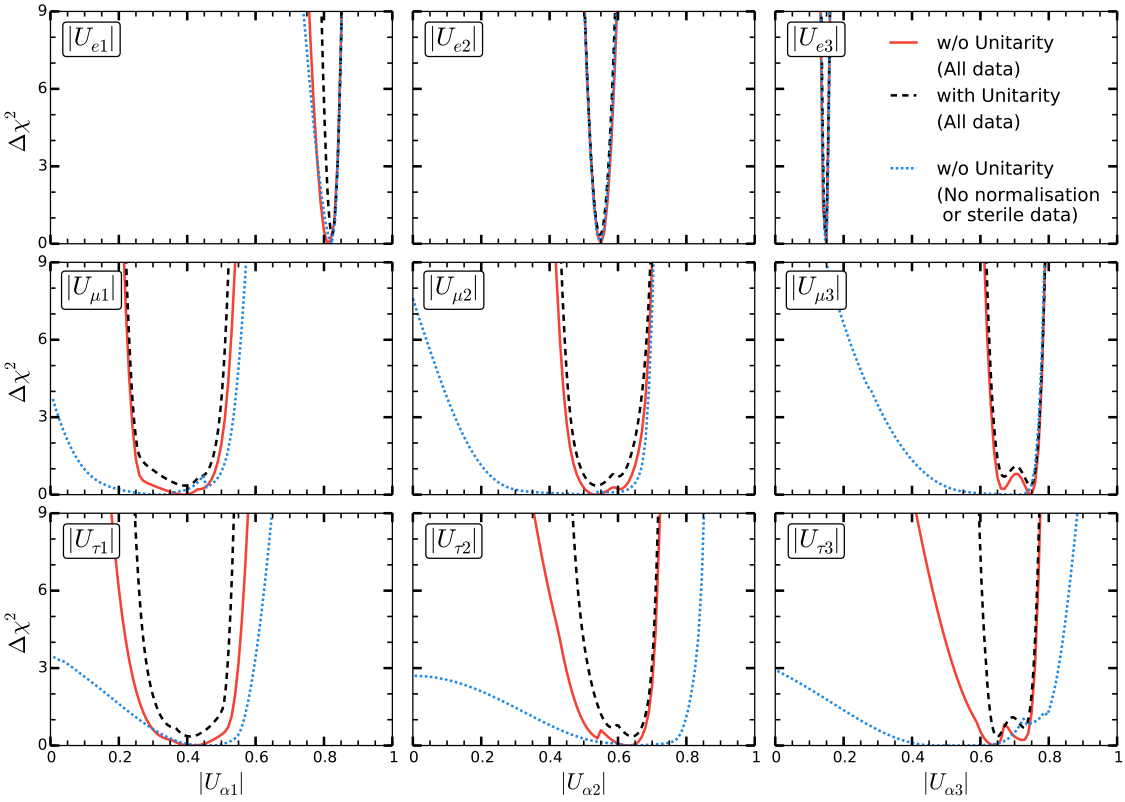

Figure 2.3: Marginalised 1-D $\Delta\chi^2$ for each of the magnitudes of the 3×3 neutrino mixing matrix elements, without (red solid) and with (black dashed) the assumption of unitarity. In order to highlight the importance of normalisation and sterile search data on these non-unitarity studies, also shown is the results of the fit when no normalisation data is used (blue dotted). Note in this scenario while the ν_e row worsens slightly, the ν_μ and ν_τ sectors lose almost all sensitivity. The x-axis is the magnitude of each individual matrix element, and the y-axis is the associated $\Delta\chi^2$ after marginalisation over all parameters other than the one in question.

Experiment	**Measured quantity with unitarity**	**Without unitarity**	**Normalisation**
Reactor SBL ($\overline{\nu}_e \to \overline{\nu}_e$)	$4\lvert U_{e3}\rvert^2 \left(1 - \lvert U_{e3}\rvert^2\right) = \sin^2 2\theta_{13}$	$4\lvert U_{e3}\rvert^2 \left(\lvert U_{e1}\rvert^2 + \lvert U_{e2}\rvert^2\right)$	$\left(\lvert U_{e1}\rvert^2 + \lvert U_{e2}\rvert^2 + \lvert U_{e3}\rvert^2\right)^2$
Reactor LBL ($\overline{\nu}_e \to \overline{\nu}_e$)	$4\lvert U_{e1}\rvert^2\lvert U_{e2}\rvert^2 = \sin^2 2\theta_{12} \cos^4\theta_{13}$	$4\lvert U_{e1}\rvert^2\lvert U_{e2}\rvert^2$	$\left(\lvert U_{e1}\rvert^2 + \lvert U_{e2}\rvert^2 + \lvert U_{e3}\rvert^2\right)^2$
SNO (ϕ_{CC}/ϕ_{NC} Ratio)	$\lvert U_{e2}\rvert^2 = \cos^2\theta_{13} \sin^2\theta_{12}$	$\lvert U_{e2}\rvert^2$	$\lvert U_{e2}\rvert^2 + \lvert U_{\mu 2}\rvert^2 + \lvert U_{\tau 2}\rvert^2$
SK/T2K/MINOS ($\nu_\mu \to \nu_\mu$)	$4\lvert U_{\mu 3}\rvert^2 \left(1 - \lvert U_{\mu 3}\rvert^2\right) =$ $4\cos^2\theta_{13} \sin^2\theta_{23} \left(1 - \cos^2\theta_{13} \sin^2\theta_{23}\right)$	$4\lvert U_{\mu 3}\rvert^2 \left(U_{\mu 1}\rvert^2 + \lvert U_{\mu 2}\rvert^2\right)$	$\left(\lvert U_{\mu 1}\rvert^2 + \lvert U_{\mu 2}\rvert^2 + \lvert U_{\mu 3}\rvert^2\right)^2$
T2K/MINOS ($\nu_\mu \to \nu_e$)	$4\lvert U_{e3}\rvert^2\lvert U_{\mu 3}\rvert^2 = \sin^2 2\theta_{13} \sin^2\theta_{23}$	$-4\,\mathrm{Re}\{U_{e3}U_{\mu 3}^* \left(U_{e1}U_{\mu 1}^* + U_{e2}U_{\mu 2}^*\right)\}$	$\lvert U_{e1}U_{\mu 1}^* + U_{e2}U_{\mu 2}^* + U_{e3}U_{\mu 3}^*\rvert^2$
SK/OPERA ($\nu_\mu \to \nu_\tau$)	$4\lvert U_{\mu 3}\rvert^2\lvert U_{\tau 3}\rvert^2 = \sin^2 2\theta_{23} \cos^4\theta_{13}$	$-4\,\mathrm{Re}\{U_{\tau 3}U_{\mu 3}^* \left(U_{\tau 1}U_{\mu 1}^* + U_{\tau 2}U_{\mu 2}^*\right)\}$	$\lvert U_{\tau 1}U_{\mu 1}^* + U_{\tau 2}U_{\mu 2}^* + U_{\tau 3}U_{\mu 3}^*\rvert^2$

Table 2.1: Example experiments and the leading order functions of U_{PMNS} matrix elements they measure, in both the unitary and non-unitary case, as well as in the standard three neutrino parametrization in the case of unitarity. The third column shows the normalisation that can be bound if the experimental measurements of the fluxes and backgrounds are known to a high enough degree.

We focus on the physically motivated subclass of unitarity violations such that $|U_{\alpha 1}|^2+|U_{\alpha 2}|^2+|U_{\alpha 3}|^2 \leqslant 1$, for $\alpha = e, \mu, \tau$, and $|U_{ei}|^2+|U_{\mu i}|^2+|U_{\tau i}|^2 \leqslant 1$ for $i = 1, 2, 3$. One must also use the knowledge of the unitarity of the true extended mixing matrix to invoke Cauchy-Schwartz inequalities and place six geometric constraints on the mixing elements [69]. Without the unitarity of the extended mixing matrix the strongest statement about the two rows of 3×3 subset would be

$$\left|\sum_{i=1}^{3} U_{\alpha i} {U_{\beta i}}^*\right|^2 \leqslant \sum_{i=1}^{3} |U_{\alpha i}|^2 \sum_{i=1}^{3} |U_{\beta i}|^2, \qquad (2.1.5)$$
$$\text{for} \quad \alpha, \beta = (e, \mu, \tau), \quad \alpha \neq \beta.$$

As one expects the L.H.S. to be small if unitarity violations are small, and the R.H.S to be $\mathcal{O}(1)$ this statement represents a very weak bound, and does not influence the available parameter space for the mixing elements. However, by forming the same inequality using only the extended mixing elements (note the sum is now from 4 to N elements),

$$\left|\sum_{i=4}^{N} U_{\alpha i} {U_{\beta i}}^*\right|^2 \leqslant \sum_{i=4}^{N} |U_{\alpha i}|^2 \sum_{i=4}^{N} |U_{\beta i}|^2, \qquad (2.1.6)$$
$$\text{for} \quad \alpha, \beta = (e, \mu, \tau), \quad \alpha \neq \beta,$$

and by the virtue of the assumed unitarity of the extended mixing matrix, $\sum_{i=4}^{N} U_{\alpha i} {U_{\beta i}}^* = -\sum_{i=1}^{3} U_{\alpha i} {U_{\beta i}}^*$ and $\sum_{i=4}^{N} |U_{\alpha i}|^2 = 1 - \sum_{i=1}^{3} |U_{\alpha i}|^2$, we can rewrite as two much stronger bounds on the rows and columns of the 3×3 subset,

$$\left|\sum_{i=1}^{3} U_{\alpha i} {U_{\beta i}}^*\right|^2 \leqslant \left(1 - \sum_{i=1}^{3} |U_{\alpha i}|^2\right)\left(1 - \sum_{i=1}^{3} |U_{\beta i}|^2\right),$$
$$\text{for} \quad \alpha, \beta = (e, \mu, \tau), \quad \alpha \neq \beta,$$
$$\left|\sum_{\alpha=e}^{\tau} U_{\alpha i} {U_{\alpha j}}^*\right|^2 \leqslant \left(1 - \sum_{\alpha=e}^{\tau} |U_{\alpha i}|^2\right)\left(1 - \sum_{\alpha=e}^{\tau} |U_{\alpha j}|^2\right),$$
$$\text{for} \quad i, j = (1, 2, 3), \quad i \neq j. \qquad (2.1.7)$$

These Cauchy-Schwartz constraints enable precision measurements in a single sector to be passed subsequently to all elements of the mixing matrix[1].

[1]These Cauchy-Schwartz inequalities are analogous to the commonly used statement that one

2.1.1 What we *really* know about the U_{PMNS} matrix

To perform the analysis, for each experiment considered[2] we take the observed amplitude of the $\nu_\alpha \to \nu_\beta$ (or $\overline{\nu}_\alpha \to \overline{\nu}_\beta$) oscillation alongside its published uncertainty and construct a chi-squared for the associated non-unitary amplitudes, along with any necessary normalisation systematics as pull factors. For short-baseline (SBL) sterile searches, if an experiment publishes the resultant χ^2 surface of their analyses in a 3+N format then this is used as a prior to bound any non-unitarity. Otherwise an appropriate prior is estimated by performing a 3+N fit to published data.

We minimize the constructed χ^2 over all parameters, satisfying the Cauchy-Schwartz constraints, using a Markov chain Monte Carlo minimizer. The results of the analyses are shown in Figure (2.3), without unitarity (red solid line) and with the assumption of unitarity (black dashed line). The non-unitary analysis was performed under the strict assumption that any non-unitarity comes solely from an extended U_{PMNS} and that no new interactions, such as an additional $U(1)'$ which can lead to strongly modified matter effects, are active at oscillation energies. Total event rate normalisation in a given experiment often has significantly large theoretical and experimental uncertainty, but is crucial for the measurement of the zero-distance unitarity effect, corresponding to the case in which a sterile neutrino has averaged out before being detected in the detector. The results of a fit in which only the spectrally observed amplitudes associated with the known two mass differences are used, with no normalisation or associated sterile data, is also shown in Figure (2.3) as the blue curve.

---

can bound $\nu_\mu \to \nu_e$ appearance by the associated ν_μ and ν_e neutrino disappearance limits, in 3+N sterile neutrino scenarios. Indeed, in the case of unitarity the inequality is saturated as an equality, reducing the number of degrees of freedom.

[2]The experimental data considered in this analysis is: Bugey [78], CCFR [79–82], CDHS [83], CHORUS [84], CHOOZ [50], Daya Bay [85, 86], Double Chooz [55], ICARUS [87], KARMEN [88], KamLAND [89, 90], MINOS [62, 64, 91, 92], NOMAD [93, 94] , NOνA [95], NuTeV [96], OPERA [97, 98], RENO [56], SNO [43], SciBooNE [99], Super-Kamiokande [63, 100–102], T2K [65, 103].

The effects of a sterile neutrino on unitarity measurements can be reduced to one of three ways. If an experiment is taking place at an L/E_ν, such that oscillations due to light steriles have not yet taken place (the limit $\Delta m_{41} \to 0$) then there will be no change in oscillation probability from that of the 3ν equations, and will only be detectable by its impact on the unitarity of the 3×3 subset of U_{PMNS}. In the context of a ν_e disappearance example, this would require precise enough measurements to discern that $|U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2 \neq 1$. The light almost-sterile mass states are still produced in the coherent superposition, but have no effect on the probability.

If, however, the sterile is sufficiently heavy, such that the oscillations have averaged out by the time of detection, $\left\langle \sin^2\left(\Delta m^2_{41}/4E_\nu\right)\right\rangle = 1/2$, this would lead to an additional contribution to the normalization of the oscillation probability, and correspondingly different amplitude. If we let the violation of unitarity be defined as $1 - \rho^2 = |U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2$ for brevity, then we can rewrite Equation (2.1.4) as

$$\begin{aligned} P^{\text{Non-Unitary}}_{\overline{\nu}_e \to \overline{\nu}_e} &= \left(1-\rho^2\right)^2 \left(1 - \frac{4(1-|U_{e3}|^2-\rho^2)|U_{e3}|^2}{\left(1-\rho^2\right)^2} \sin^2\left(\frac{L\Delta m^2_{31}}{4E_\nu}\right)\right), \\ &= \underbrace{1 - 2\rho^2\left(1-\rho^2/2\right)}_{\text{in the SBL limit } \Delta m_{31} \to 0}. \end{aligned} \tag{2.1.8}$$

Where as a sterile neutrino that has averaged out will produce a ν_e disappearance probability of

$$\begin{aligned} P^{\text{Averaged Sterile}}_{\overline{\nu}_e \to \overline{\nu}_e} &= \left(1-|U_{e4}|^2\right)^2 \left(1 - \frac{4(1-|U_{e3}|^2-|U_{e4}|^2)|U_{e3}|^2}{\left(1-|U_{e4}|^2\right)^2} \sin^2\left(\frac{L\Delta m^2_{31}}{4E_\nu}\right)\right) + |U_{e4}|^4, \\ &= \underbrace{1 - 2|U_{e4}|^2(1-|U_{e4}|^2)}_{\text{in the SBL limit } \Delta m_{31} \to 0}. \end{aligned} \tag{2.1.9}$$

So we see, that in the SBL limit $\rho^2 \approx |U_{e4}|^2 - |U_{e4}|^4/2 + \mathcal{O}(|U_{e4}|^6)$ and non-unitarity is approximately measuring $|U_{e4}|^2$ up to corrects of order $|U_{e4}|^4$. This holds for any number of sterile neutrinos.

If the sterile neutrino is even heavier still, too heavy to be kinematically produced in meson decay, then the superposition formed will not contain the sterile

mass eigenstate and the sum over massive states in Equation (1.4.49) would not contain the almost-sterile state.

As the use of non-observation of additional light oscillating frequencies is more model specific, and does not measure the mixing elements directly, the red curve of Figure (2.3) where these bounds are included should be considered a best-case scenario in which the new states introduced have no additional interactions whatsoever. In the case of null appearance short baseline experiments, any interactions that increase the decay of the light sterile state, through decays to a dark sector for example, can drastically reduce sensitivity. In comparison, by measuring the mixing elements via only the direct observation of oscillating amplitudes, the blue-dotted curve, one can see that although a much more model independent and generic result, the precision is dramatically worse, with only the ν_e sector having any degree of certainty of the values of individual mixing elements.

Upon minimization the best-fit points agree in the unitary and both non-unitary fits. From this point on we focus predominately on the results including null short-baseline light sterile experiments, as without their inclusion there is little or no sensitivity. To compare how the precision varies we consider the frequentist 3σ ranges of the one-dimensional $\Delta\chi^2$ projections without unitarity assumed (with unitarity), where we marginalise over all parameters except the one in question, we obtain

$$|U_{\text{PMNS}}|_{3\sigma}^{\substack{\textbf{w/o Unitarity}\\(\text{with Unitarity})}} = \begin{pmatrix} 0.76 \to 0.85 & 0.50 \to 0.60 & 0.13 \to 0.16 \\ (0.79 \to 0.85) & (0.50 \to 0.59) & (0.14 \to 0.16) \\ 0.21 \to 0.54 & 0.42 \to 0.70 & 0.61 \to 0.79 \\ (0.22 \to 0.52) & (0.43 \to 0.70) & (0.62 \to 0.79) \\ 0.18 \to 0.58 & 0.38 \to 0.72 & 0.40 \to 0.78 \\ (0.24 \to 0.54) & (0.47 \to 0.72) & (0.60 \to 0.77) \end{pmatrix}. \quad (2.1.10)$$

The ranges for the individual elements, assuming unitarity (bracketed numbers in above expression), are in good agreement with published results in contemporary global fits such as ν-fit [67].

As a further metric, if we define the shift in range of allowed values as the ratio of the difference in 3σ ranges without and with unitarity, to that derived with unitarity, the increase in parameter space for $|U_{ei}|, i = 2, 3$ and $|U_{\mu i}|, i = 1, 2, 3$ are all $\leqslant 10\%$ (4%, 8%, 8%, 7% and 4% respectively), with $|U_{e1}|$ taking the majority of the discrepancy in the ν_e sector, with an increase of allowed range of 68%, primarily due to the weaker bounds from KamLAND compared to the SBL reactors. The entire ν_τ sector, however, may contain substantial discrepancies from unitarity with shifts in allowed regions of 37%, 46% and 104% respectively.

2.1.2 Bounds on unitarity violation

We must stress that even if the 3σ ranges of the U_{PMNS} elements agree closely with the unitarity case, as is the case with the ν_e sector, this does not equate to the neutrino mixing matrix being unitary. In the unitary case the correlations are much stronger, and choosing an exact value for any one of the mixing elements drastically reduces the uncertainty on the remaining elements, a fact which is hidden in the 1-D projections in Figure (2.3) . One can address this issue by looking at the row and column unitarity triangle closures and the row and column normalisations to better understand the level at which we know unitarity is violated or not.

For the case of the six neutrino unitarity triangles, we present, for the first time, the allowed ranges for their closures in Figure (2.4). For the three row unitarity triangles the bounds originate from a combination of the corresponding Cauchy-Schwartz inequalities along with appearance data in the respective channel. The column unitarity triangles, being bound primarily by the geometric constraints and not direct measurement, are less known. Only one unitarity triangle does not contain a ν_τ element, the $\nu_e\nu_\mu$ unitarity triangle, and hence it is the only unitarity triangle in which it is constrained to be closed by $\leqslant 0.03$ at the 3σ C.L, compared to $\leqslant 0.1$ - 0.2 at the 3σ C.L for the remaining unitarity triangles. This hierarchical situation will not improve unless precise measurements can be made in the ν_τ sector.

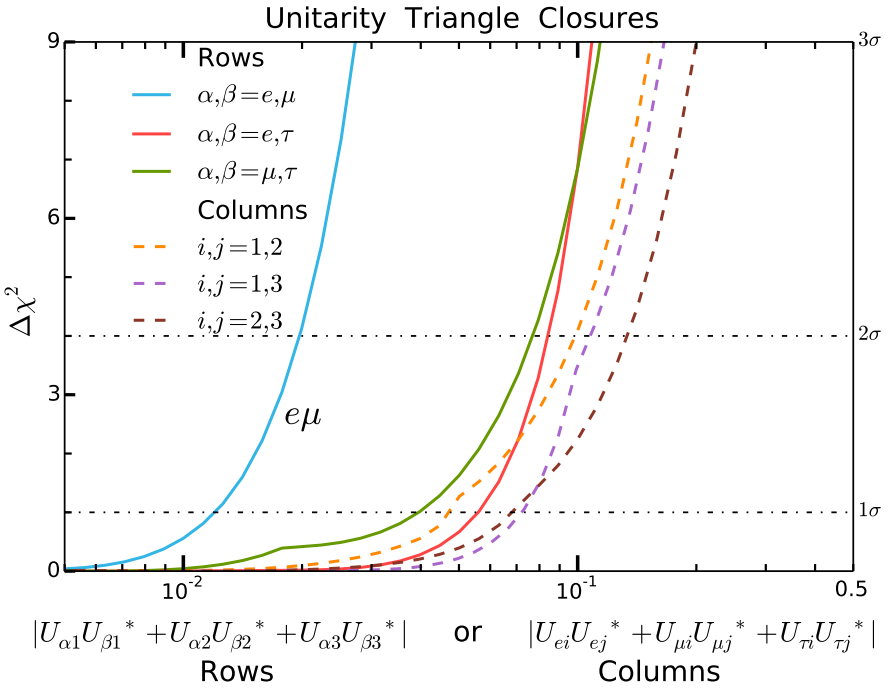

Figure 2.4: 1-D $\Delta\chi^2$ for the absolute value of the closure of the three row (solid) and three column (dashed) unitarity triangles when considering new physics that enters above $|\Delta m^2| \geqslant 10^{-2}$ eV^2. There is one unique unitarity triangle, the $\nu_e\nu_\mu$ row unitarity triangle, in that it does not contain any ν_τ elements, and hence is constrained to be unitary at a level half an order of magnitude better than the others. By comparison to Figure 2.5, one can clearly see the Cauchy-Schwartz constraints are satisfied.

We also plot the resultant ranges for the normalisations in Figure (2.5). We see that the ν_e and ν_μ normalisation deviations from unity are relatively well constrained ($\leqslant 0.06$ and 0.07 at 3σ C.L respectively), primarily by reactor fluxes and a combination of precision measurements of the rate and spectra of upward going muon-like events observed at Super-Kamiokande [100]. We note the ν_μ normalisation deviation from unity is constrained slightly ($\approx 1\%$) better than the ν_e normalisation. This is due to the large theoretical error, 5%, on total flux from reactors assumed [104]. The remaining normalisation deviations from unity are all constrained to be $\lesssim 0.2$ - 0.4 at 3σ C.L.

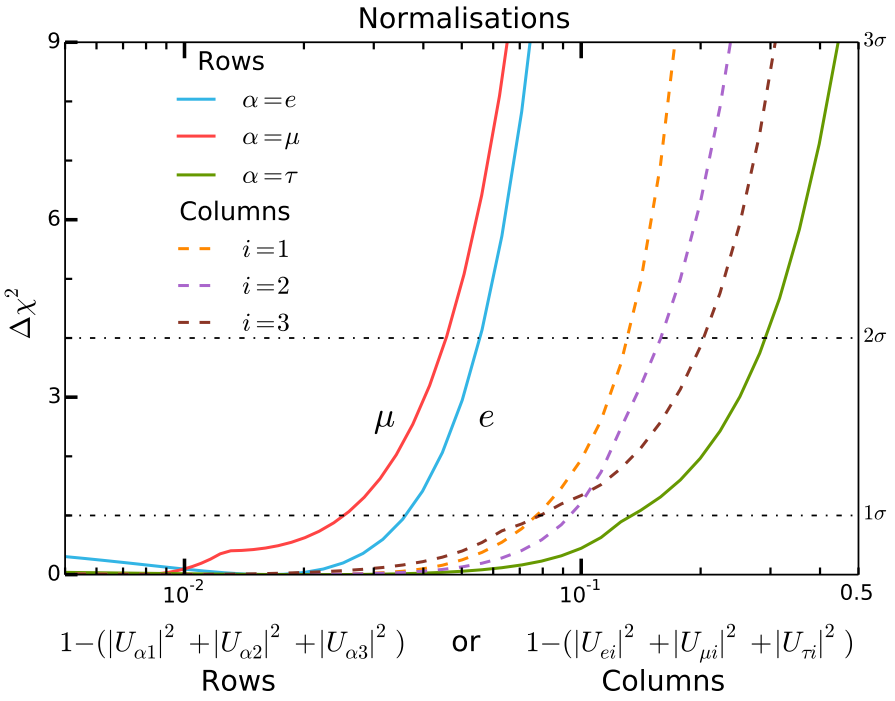

Figure 2.5: 1-D $\Delta\chi^2$ for deviation of both U_{PMNS} row (solid) and column (dashed) normalisations, when considering new physics that enters above $|\Delta m^2| \geqslant 10^{-2}\text{eV}^2$.

If one wishes to proceed with measurements of unitarity, without the assumption of an extended U_{PMNS} matrix and its subsequent Cauchy-Schwartz constraints, then prospects for improvement are essentially limited to measuring the ν_e normalisation. Improvement of all ν_e elements is possible, especially if the new generation reactor experiments, JUNO [105] and RENO50 [106], proceed as planned, see reference [66]. In order to qualitatively develop new probes of the U_{PMNS} mixing elements, in the ν_μ and ν_τ sector, one would have to develop high statistics $\nu_e \to \nu_\tau$ and $\nu_\mu \to \nu_\tau$ appearance experiments in which one could ascertain an oscillation amplitude, rather than the tail of the oscillation probability, as is the case with OPERA. Even more useful, although perhaps more wishful, would be a muon disappearance experiment on the solar mass scale $\approx 15,000$ km/GeV. However, this is well beyond what is currently technologically feasible except for the scenario of a full scale neutrino factory [107]. Possibilities of directly probing the τ sector are even further disjointed from the present, with a ν_τ disappearance experiment defying anything more advanced than speculation at the moment.

Improvements due to indirect sterile neutrino searches are more promising. The Fermilab Short Baseline Neutrino [108] program consisting of the SBND, MicroBooNE and ICARUS experiments on the Booster beam, will be capable of probing a wide range of parameter space for 3+N models, increasing both the appearance and disappearance bounds. Subsequently, the long baseline program DUNE [109] will also be able to significantly extend the constrained region of $\nu_\mu \rightarrow \nu_e$ appearance to lower mass differences, leading to increased constraints on the $\nu_e\nu_\mu$ unitarity triangle in this regime. An understanding of the neutrino flux and cross sectional uncertainties are crucial for unitarity measurements. However, no one experiment can probe all scales and complementarity is vital to definitively make a statement about unitarity from new low-energy physics. Perhaps crucially for ν_τ measurements, Hyper-Kamiokande [59] will be quite sensitive to atmospherically averaged steriles, $\geqslant 0.1$ eV^2, and will significantly improve the current bounds on $|U_{\tau 1}|^2+|U_{\tau 2}|^2+|U_{\tau 3}|^2$ in this regime, to approximately $1-|U_{\tau 1}|^2+|U_{\tau 2}|^2+|U_{\tau 3}|^2 \leqslant 0.07$ at the 99% C.L [110], which would bring all sectors inline with each other.

We strongly emphasise the fact that current experimental bounds on unitarity within the 3ν paradigm allows for considerable violation, and without the unitarity assumption, the precision on the individual U_{PMNS} elements can vary significantly (up to 104% in the case of $|U_{\tau 3}|$). However, we find no evidence for non-unitarity. The prospects of directly measuring all the 12 unitarity constraints with high precision are poor, currently we can only constrain the amount of non-unitarity to be $\lesssim$ 0.2 - 0.4, for four out of six of the row and columns normalisations, with the ν_μ and ν_e normalisation deviations from unity constrained to be $\leqslant 0.07$, all at the 3σ C.L, see Figure 2.5. Similarly, five out of six of the unitarity triangles are only constrained to be $\lesssim$ 0.1 - 0.2, with opening of the remaining $\nu_e\nu_\mu$ unitarity triangle being constrained to be $\leqslant 0.03$, again at the 3σ C.L, see Figure 2.4. One must be careful when assessing the current experimental regime with the addition of new physics we are currently insensitive to, as without the assumption of unitarity there is much room for new effects, especially in the ν_τ sector where currently significant

information comes from the unitarity assumption and not direct measurements.

2.2 Sterile Neutrinos

Although verifying the unitarity of the 3×3 U_{UPMNS} matrix should be an aim of the neutrino community, regardless of the possible sources of the non-unitarity, in order to confirm we truly are in a 3ν paradigm, we turn our attention now to specific models Beyond the Standard Model that can extend the neutrino sector mixing matrix in such a way to break the unitarity of the 3×3 U_{UPMNS}. We use the bounds derived from the current global experimental situation as concrete motivation to study such BSM physics. For the remainder of this thesis we will focus on these BSM physics in a less model-independent manner, discussing additional motivation for their inclusion and the potential for next generation facilities to probe such scenarios.

As mentioned in the previous Section, the canonical method for introducing non-unitarity to the 3×3 subset of the U_{UPMNS} matrix is the introduction of additional fermionic degrees of freedom which mix with the standard model neutrinos. LEP showed that there are no additional active neutrinos below half the mass of the Z-Boson [111], so if light degrees of freedom do exist they must be singlets of the Standard Model gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. To investigate the effects of such a state, we introduce now a single right-handed chiral state, n_R , which *a priori* can mix with the neutral neutrinos of the Standard Model though Yukawa terms such as ,

$$\mathcal{L}_{\text{yuk}} \supset -y_\alpha \overline{n}_R \tilde{H}^\dagger L_{\alpha L} + H.C, \tag{2.2.11}$$

$$\xrightarrow[\text{Breaking}]{\text{EW Symmetry}} = -\frac{y_\alpha v}{\sqrt{2}} \overline{n}_R \nu_{\alpha L} + H.C. \tag{2.2.12}$$

where $L_{\alpha L}$ is a single left-handed lepton doublet of flavour α such that $L_\alpha = (\nu_{\alpha L}, \alpha_L)^T$ containing the left-handed chiral states $\nu_{\alpha L}$, and in the second line the Higgs has obtained vev post-Electroweak (EW) symmetry breaking in the manner described in Section (1.2), $\langle H \rangle = \frac{1}{\sqrt{2}} (0, v)^T$. This is of course the same Yukawa

structure in the charged fermion sector that leads to the creation of a Dirac mass term for the charged leptons. Recall that such a term could not be used to generate masses in the neutrino sector, as in the Standard Model there is no right-handed neutrinos, and we were forced to investigate Majorana mass terms for the neutrinos instead. The sterile right-handed state n_R, however, assumes this role and Equation (2.2.12) can be rewritten as $\mathcal{L} = -m\overline{\nu}\nu$, a pure Dirac mass term, with $m = y_\alpha v/\sqrt{2}$ for the Dirac neutrino field $\nu = \nu_{\alpha L} + n_R$.

These right-handed states, n_R, are colloquially called "Sterile Neutrinos". We use this term when referring to *any* Standard Model $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ gauge singlets, even if no mixing is present. In this extreme case they truly are completely sterile, having no interactions with the Standard Model particles, other than gravitational . Other terms which have been used in the literature are "Heavy Neutral Leptons" (HNLs) or "inert neutrinos".

As the inclusion of these states allows for the introduction a mass term for the neutrinos, via the same Higgs mechanism that generates masses for all charged leptons and quarks in the Standard Model. This looks like a promising manner to solve the problem of neutrino masses. However, the smallness of the observed neutrino masses raises a naturalness issue with this, as it would require Yukawa couplings *far* smaller in magnitude than those which generate the masses for the leptons and quarks. If the same Higgs field and boson is responsible for both neutrino and charged fermion masses, then the ratio of masses is equal to the ratio of Yukawa couplings. Looking at this ratio for the electron and top quark, the two most extreme masses in the Standard Model, we have (assuming ≈ 0.1 eV neutrino masses),

$$\frac{m_\nu}{m_e} = \frac{y_\nu}{y_e} \approx 10^{-7}, \tag{2.2.13}$$

$$\frac{m_\nu}{m_{\text{top}}} = \frac{y_\nu}{y_{\text{top}}} \approx 5 \times 10^{-13}. \tag{2.2.14}$$

Although the Standard Model already requires ≈ 6 orders of magnitude in Yukawa's to successfully explain the hierarchy of masses between electron and top quark masses, we would require a void of an additional 7 orders of magnitude before we

can explain neutrino masses. Due to the vast differences in scale, many feel this is evidence that neutrino masses are generated by an entirely different process to that of the Standard Model Higgs mechanism. Possibilities as to what this mechanism could be arise when one studies further the possible additions to the Standard Model Lagrangian, with the inclusion of one or more sterile neutrinos.

In addition to the Dirac mass term generated by introducing n_R, the sterile neutrino itself can form additionally a Majorana mass term, as described in Equation (1.3.15), expanding the mass Lagrangian (post EWSB) to

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} m_R n_R^T C^\dagger n_R - m_D \overline{n}_R \nu_L + H.C, \tag{2.2.15}$$

where $m_D \equiv y_\alpha v/\sqrt{2}$. When combined with a possible Majorana mass term for the left handed neutrinos, $1/2 m_L \nu_L^T C^\dagger \nu_L$, as referenced in Chapter 1, we can rewrite the mass Lagrangian as

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} \mathcal{N}_L^T C^\dagger M \mathcal{N}_L + H.C, \tag{2.2.16}$$

where $\mathcal{N}_L$ is a now a single column matrix of *left-handed chiral fields* only

$$\mathcal{N}_L \equiv \begin{pmatrix} \nu_L \\ C\overline{n}_R^T \end{pmatrix}, \tag{2.2.17}$$

and M is a generic symmetric mass matrix given by

$$M = \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix}. \tag{2.2.18}$$

Studying and diagonalising this matrix allows us to calculate the physical masses for the neutrinos for a wide variety of scenarios.

2.2.1 The See-Saw Mechanism

It is instructive here to focus, as a simple but canonical example, on a standard model-like scenario containing one active neutrino and one additional sterile state. As mentioned in Chapter 1, the active left-handed chiral neutrino that is involved

in Standard Model weak interactions is forbidden under the Standard Model gauge group to have a Majorana mass term, as the term $\nu_L^T C^\dagger \nu_L$ has total hypercharge $Y = -2$, thus m_L is required to be vanishing to preserve the symmetry. The sterile state n_R, however, being a singlet under the Standard Model, can have a non-zero Majorana mass m_R, and in fact unless an additional symmetry is imposed there is no reason not to write down such a term. At a cursory glance, a direct Dirac mass term $\approx m_d \overline{n}_R \nu_{\alpha L}$ is also forbidden by Standard Model as it also violates $SU(2)_L \otimes U(1)_Y$, however, it is trivial to arrive at such a term post-electroweak symmetry breaking using the Standard Model Higgs mechanism as described in Eq.(2.2.12). One would expect this mass term to be of order the symmetry breaking scale, by arguments of naturalness, and so in the Standard Model should be $\approx \mathcal{O}(100\text{GeV})$. As the sterile is a singlet under the Standard Model there is no scale, *a priori*, which could dictate where the Majorana mass, m_R, should lie. If we expect it to be generated around the symmetry breaking scale, one possibility is that it is generated at the scale that the Standard Model breaks down, the scale of grand unification $10^{13} - 10^{16}$ GeV [112, 113].

This well studied limit, $m_L = 0$ and $m_R \gg m_D$, which is known as the famous "See-Saw" mechanism [114–117], once can show that Equation (2.2.18) reduces to,

$$M = \begin{pmatrix} 0 & m_D \\ m_D & m_R \end{pmatrix}, \tag{2.2.19}$$

which, after diagionalization, leads to the following physical massive eigenstates,

$$m_{\text{light}} \approx m_D \left(\frac{m_D}{m_R} \right), \tag{2.2.20}$$

$$m_{\text{heavy}} \approx m_R. \tag{2.2.21}$$

Thus the heavier the right handed Majorana mass term, the lighter the neutrino mass being suppressed by the small factor m_D/m_R, giving rise to to the aptly named "see-saw" moniker. For $m_D \approx 1$ GeV, a Majorana mass term for the sterile neutrino of $m_R \geqslant 10^9$ GeV is required to generate sub-eV neutrinos masses as required by experimental bounds. For a Dirac mass equal to that of the top-quark, also thought to be generated by the same Standard Model Yukawa terms, m_R is required to be

$\geqslant 10^{13}$ GeV.

This simple example provides a succinct and compelling method to generate the smallness of observed neutrino masses. The addition of multiple generations of active neutrinos and sterile states is necessary to provide multiple masses consistent with experimental observations. It is easily generalised, however, to three active and three sterile neutrinos by promoting the mass terms to matrices leading to the three light physical states of $m_{\text{light}} \approx M_D^T M_R^{-1} M_D$.

If the right-handed sterile neutrino mass, m_R, is indeed very large, then one can integrate the massive degree of freedom out of our theory, in the same vein that one removes the massive W-boson leading to Fermi's four-fermion interaction. This leaves an effective term in our Lagrangian, formed from the $\overline{n_r} L_L$ Yukawa terms that directly couples two left-handed doublets L_L and two Higges H,

$$\mathcal{L}_{\text{eff}} \approx \frac{y_\alpha^2}{m_R} \left(L_L^T \mathcal{C}^\dagger \sigma_2 \vec{\sigma} L_L \right) \left(H^T \sigma_2 \vec{\sigma} H \right) + H.C, \tag{2.2.22}$$

$$\xrightarrow[\text{Breaking}]{\text{EW Symmetry}} \quad \frac{m_D^2}{m_R} \nu_L^T \mathcal{C}^\dagger \nu_L + H.C. \tag{2.2.23}$$

This is none other than the Weinberg dimension 5 operator, as discussed in Equation (1.3.17), which generates a Majorana mass term for the Standard Model neutrinos, where we have now identified the heavy Majorana mass, m_R, with the high scale Λ, and the coupling constant g with the neutrino Yukawas y_α^2. The addition of a right-handed sterile neutrino with a large Majorana mass is thus one way to generate the Weinberg operator. It is said to be a *UV completion* of the effective theory, where now all terms are indeed renormalizable, and is referred to as the Type I completion.

There are two other ways in which one can complete the non-renormalizable Weinberg operator, in addition to the Type I completion. One can introduce an additional scalar triplet (Type II) or a fermionic triplet (Type III). We will not go into detail about the Type II or Type III completions of the Weinberg operator,

nor the masses they can generate, however, we show all three diagrammatically and briefly discuss them in Figure (2.6).

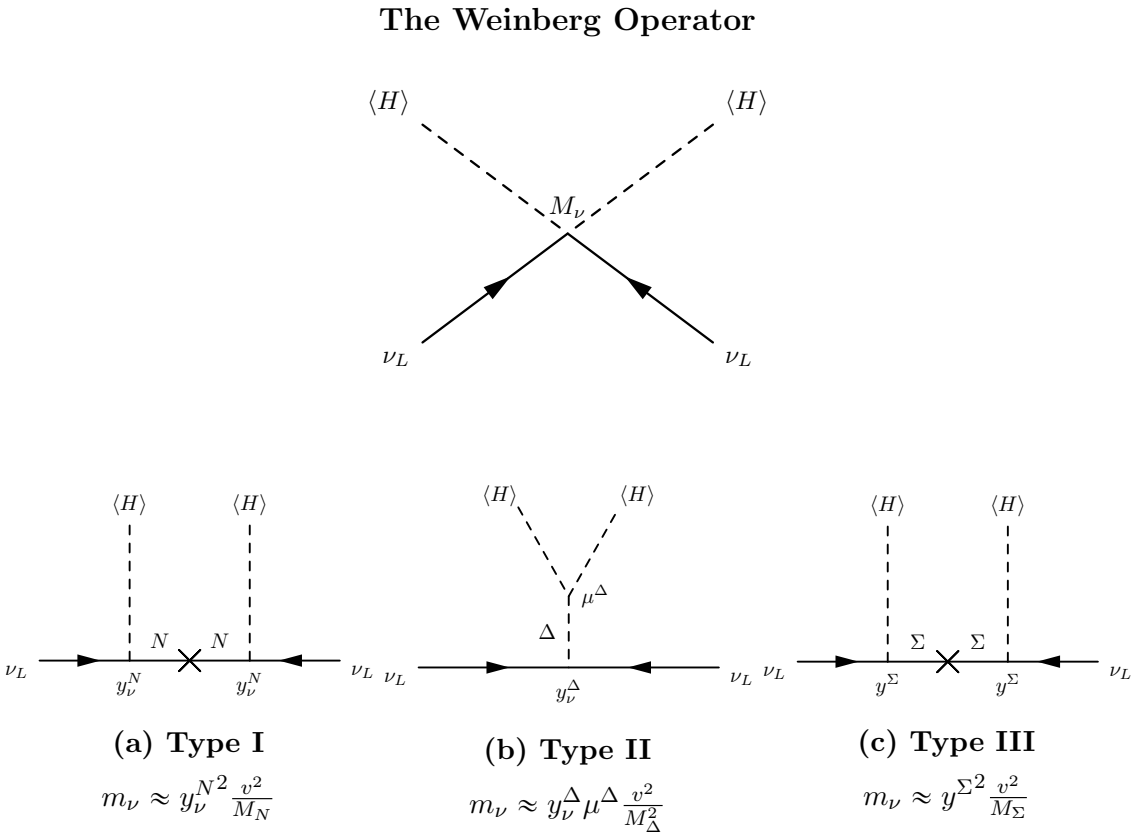

Figure 2.6: Three UV completions of the four-point dimension 5 Weinberg operator (schematically given on top), that can all lead to the formation of a Majorana neutrino mass. In the Type I completion, we introduce a right-handed sterile neutrino, N, with a Yukawa coupling to the Higgs. In the Type II a scalar triplet, Δ, has a direct coupling to both the Higgs and the left-handed neutrinos. In the Type III completion we introduce a fermionic triplet, Σ, that couples to both the Standard Model neutrinos and Higgs. Shown also is the approximate Majorana mass that the active neutrinos receive due to the dimension 5 operator, post EWSB.

2.2.2 The Many Scales of Sterile Neutrinos

The Type I See-Saw mechanism is but one example of a theoretically motivated scale for which sterile neutrinos might exist, however, the associated mixing matrix

required to diagonalise this mass matrix contains mixing angles of

$$\theta \approx \frac{m_D}{m_R}, \tag{2.2.24}$$

which for the values discussed above $\theta \ll 10^{-10}$, almost completely decoupling the heavy states from the light active states, meaning that they are not observable at the energies of terrestrial experiments. As they lack a known symmetry to dictate their mass, sterile neutrinos can potentially have masses varying over tens or orders of magnitude. Such a vast range of potential energy scales leads to an equally vast range of qualitatively distinct phenomena that they are involved with, with many having much more observable impact, and larger mixing, than the aforementioned high-scale "see-saw" sterile neutrinos.

In Table (2.2) we collate a selection of example masses and the corresponding phenomenology sterile neutrinos at a scale that it might be involved with. This is by no means an exhaustive list, but helps highlight the vast amount of experimental and theoretical potential that sterile neutrinos encompass. Note that such sterile neutrinos need not have any connection with the generation of light neutrino masses, indeed their own masses may be due by an entirely separate phenomena, although one would have to be careful as to explain why the Majorana and Dirac mass terms as described above do not arise, perhaps by imposing a global symmetry that forbids them, such as the case of a $U(1)$ B-L symmetry which forbids the Majorana mass term for the sterile neutrinos, $m_R n_R^T C^\dagger n_R$.

For the remainder of this thesis, we will focus entirely on two of the most phenomenologically rich regimes of sterile neutrino behaviour in the context of current and near-future Short-Baseline neutrino experiments. In Chapter 3, we will explore in detail the potential of Fermilab's Short Baseline Neutrino Program (SBN) to probe light oscillating sterile neutrinos at 0.01 eV^2 to 100 eV^2 mass-splittings in the (3+1), (3+2) and (3+3) scenarios, with a special focus on the CP violating phases introduced in the (3+2) sterile neutrino scenario. Although not involved with the generation of neutrino masses, sterile neutrinos at this scale are motivated by several low-energy anomalies, as was discussed briefly already.

And in Chapter 4, we explore the idea of using the same SBN program to probe heavy $\mathcal{O}(100\text{ MeV})$ non-oscillating sterile neutrinos, which through their subsequent decay to visible particles can be observed. In this chapter we will delve into more detail about the possible interactions that a sterile neutrino can have, rather than focus on the generation of light neutrino masses as we have in this chapter. We stress these two separate regimes can be searched for simultaneously, without the need for any detector or beam modifications allowing the physics reach of such a short-baseline program to cover a wide range of possible mass scales beyond what it was designed for.

Mass Scale	Associated Phenomenology
$\ll eV$	Superlight sterile neutrinos $\Delta m^2_{41} \leqslant \Delta m^2_{\text{sol}}$, almost degenerate with lightest active neutrino. Explains absence of upturn at low energies of the solar neutrino energy spectrum [118, 119], and could possible be measured at future facilities [120]. If degenerate with Δm^2_{31} can interfere with measured value of θ_{13} [121].
$\mathcal{O}$(eV)	Active oscillations at facilities of L/E_ν of $\approx \mathcal{O}(100\text{m})/\text{GeV}$. Experimentally motivated, as possible to explain LSND [71], MiniBooNE [72], the gallium anomaly [41, 73] and the Reactor anomaly [74]. See Chapter 3 for more details.
$\leqslant 100$ keV	Produces kinks in the β-decay spectrum of a wide array of elements such as ^{3}H , ^{187}Re, ^{63}Ni, ^{35}S, ^{20}F [122–126].
$\mathcal{O}$(keV)	keV sterile neutrinos are a potential warm dark matter candidate [127].
MeV $-$ GeV	Produced in the decay of $\pi^\pm$, K$^\pm$ and D$^\pm$ mesons. Results in monochromatic lines in energy spectrum of associated lepton [128] or by their subsequent decay to visible particles. See Chapter 4 for more details.
GeV $-$ TeV	Can be produced in LHC collisions, can be detected through their decay to dileptons/jets [129, 130] or missing energy and displaced vertices if long-lived [131–133]. Theoretically TeV scale steriles might be sufficient to produce low-scale Leptogenesis [134].
$10^9 - 10^{16}$ GeV	Generate light neutrino masses naturally via see-saw mechanisms [135–137] . Produce leptogenesis from asymmetric decays in early universe [138].

Table 2.2: A summary list of scales at which sterile neutrino masses could be realised and some of the corresponding phenomenology associated with those scales, along with corresponding references (and references therein) for further reading.

Chapter 3

Light Sterile Neutrinos, $\mathcal{O}(\mathrm{eV}^2)$

"A honey bee never volunteers an attack, or acts on the offensive, when it is gorged or filled with honey."

L.L Langstroth; Father of Modern beekeeping,
Great-Grand Father of Clyde Cowan,
Discoverer of the Neutrino

3.1 Motivation: LSND and MiniBooNE

During the past two decades, concurrently with the experimental exploration of the three-neutrino oscillation paradigm, several additional oscillation-like anomalous experimental signatures have surfaced, which may require new physics to interpret. One possible such new physics interpretation is that of additional, light sterile neutrinos [139] as was briefly alluded to in the previous Chapter. In this section we will investigate these light states further, and go into detail about their phenomenological effects at modern short-baseline experiments. The mass range we will focus on is of order $0.1 - 10$ eV, leading to small-amplitude neutrino oscillations at relatively small $L/E \sim 1$ m/MeV. As we have seen, in the context of muon neutrino and electron neutrino mixing, the constraints imposed by unitarity of the overall neutrino mixing matrix, together with existing experimental bounds on the elements of the neutrino mixing matrix (see, e.g. [140, 141]), lead us to assume at most $\mathcal{O}(1\%)$ level mixing. If one assumes a neutrino beam of $\approx$ GeV energy, then the corresponding baselines at which mass-splittings of order $\Delta m^2_{41} \approx 0.1 \rightarrow 10$ eV^2 are spectrally active is $\leqslant 1$ km, hence this signature is often referred to as *short-baseline oscillations*.

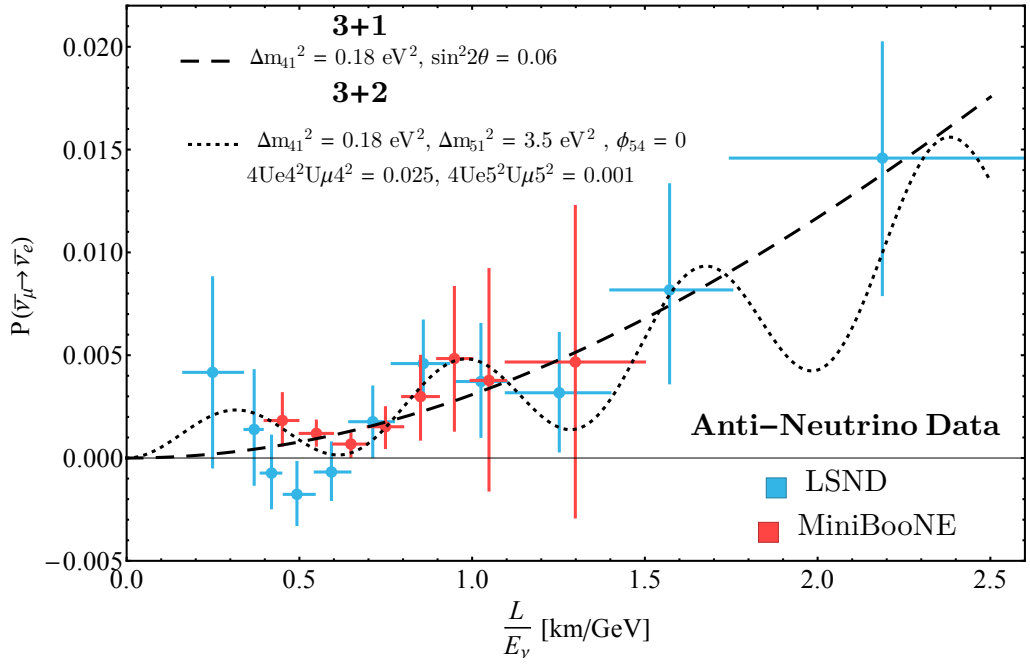

Figure 3.1: The antineutrino appearance oscillation probabilities as measured by both LSND (blue) and MiniBooNE (red) alongside two example oscillation probabilities under the 3+1 (dashed) and 3+2 (dotted) light oscillating sterile neutrino hypothesis. Agreement is good between both antineutrino data sets.

The first hint of a third frequency at which oscillations were active in the neutrino sector came from the Liquid Scintillator Neutrino Detector (LSND), which ran at the Los Alamos National Laboratory from 1993-1998. The LSND beam primarily consisted of anti-muon neutrinos produced in the decay-at-rest of μ^+. The anti-muons are produced from a very high energy proton beam, 798 GeV, producing a intense source of π^+. The majority of π^- are captured without the release of neutrinos, leading to a very small fraction of $\overline{\nu}_e$ in a well understood $\overline{\nu}_\mu$ beam. The detector consisted of 167 tons of liquid scintillator, inside a cylindrical tank approximately 8.3m long with 2.35m radius [142]. This was situated $\approx$ 30m from the primary proton target [1]. The signal of $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ oscillations was through the reaction $\overline{\nu}_e p \rightarrow e^+ n$, in which the emitted positron is observed as well as an additional 2.2

[1]Note that the exact target composition and location was changed over the running time of the LSND experiment, although 30m baseline is correct for the majority of protons-on-target (POT).

MeV photon from the re-absorption of excited neutron via $np \rightarrow \gamma d$.

The first publications from LSND in 1996 showed the existence of 22 anomalous events over an expected background of 4.6 ± 0.6 events, consistent with $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ oscillations [142]. After further data taking and background re-evaluation, LSND published a final excess of $87.9 \pm 22.4 \pm 6.0$ events in 2001 [71]. At such a short-baseline the effect of the known three neutrino mixing angles and mass-splittings are negligible and are ignored, this is the so called *short-baseline approximation*, and the analysis was performed under a two neutrino approximation in which there is only one mass-splitting, Δm^2_{41}, and one mixing angle, $\theta_{\mu e}$. The values of Δm^2_{41} and $\theta_{\mu e}$ that the LSND anomaly favoured can be seen in Figure (3.11) below, with best fit values of

$$\sin^2 2\theta_{\mu e}^{\text{LSND}} = 0.003, \qquad \Delta m^2_{41}{}^{\text{LSND}} = 1.2 \text{ eV}^2. \tag{3.1.1}$$

The significance of this anomaly is $\approx 3.8\sigma$, although later independent re-analyses argue a lower estimate of $2.3\sigma \rightarrow 2.9\sigma$ [143,144] due to larger backgrounds. Irrespective of the exact value, the LSND anomaly represented a very strong signal of new physics in the neutrino sector and the neutrino community adapted accordingly to investigate this new phenomena. One of the primary tools in this endeavour was to the the MiniBooNE detector.

The MiniBooNE detector was designed and built with the goal of definitely verifying, or bringing into question , the $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ oscillation interpretation of the LSND anomaly. The MiniBooNE experiment was built at a different baseline and neutrino energy than the LSND experiment, but at the same L/E_ν ratio, as to be sensitive to the same oscillation frequency and thus perform an independent search. Unlike the LSND neutrino beam which was primarily from decay-at-rest μ^+, the $\mathcal{O}(1$ GeV) neutrino beam that impinged on MiniBooNE was produced from the ≈ 8 GeV Booster proton beam at Fermi National Accelerator Laboratory [145], that struck a Beryllium target to produce pions and kaons which were subsequently focused and allowed to decay in a 50m decay pipe, to form the predominately ν_μ beam. The MiniBooNE detector consisted of a 12.2m diameter spherical tank of mineral oil

(CH_2) located approximately 500m from the proton target. A total of 1280 8 inch PMT's covered the outside of the detector, and utilised the Cherenkov light emitted by charged particles to detect electrons and muons that were produced though neutrino charged current scattering.

The MiniBooNE experiment could look for ν_e appearance, ν_μ disappearance, and though the reversal of the magnetic horn current, run in anti-neutrino running mode to perform $\overline{\nu}_e$ appearance and $\overline{\nu}_\mu$ disappearance also. The first $\nu_\mu \rightarrow \nu_e$ results using 6.46e20 POT in neutrino mode observed an excess of $128.8 \pm 20.4\text{stat} \pm 38.3\text{sys}$ electron-like events [146], spectrally consistent with ν_e charge current scattering. Unexpectedly, however, this excess did *not* strongly favour the LSND anomaly, with the excess occurring within reconstructed neutrino energy of $200 \rightarrow 475$ MeV, whereas oscillations matching the LSND parameter space would be expected to appear at higher energies. This excess at a lower energy than expected is often called the MiniBooNE "low-energy excess" to distinguish its possible origins to that of the LSND anomaly.

MiniBooNE continued to collect data in anti-neutrino running mode, resulting in a combined $\nu_\mu \rightarrow \nu_e$ and $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ oscillation analysis [147] using an additional 11.27e20 POT of anti-neutrino running. An excess of 78.4 ± 28.5 events was observed in antineutrino running mode, alongside an updated 162 ± 47.8 excess events in neutrino mode. The combined excess of 240.3 ± 62.9 represented a 3.8σ significant signal consistent with sterile neutrino oscillations. Spectrally the excess in anti-neutrino running mode was at a higher energy, allowing for greater agreement with the LSND anomaly, although due to the much lower energy anomaly in neutrino mode, MiniBooNE data favours a best fit corresponding to a much lower mass-splitting with maximal mixing,

$$\sin^2 2\theta_{\mu e}^{\text{MiniBooNE}} = 1.0, \qquad \Delta m^2_{41}{}^{\text{MiniBooNE}} = 0.037 \text{ eV}^2. \tag{3.1.2}$$

As the LSND search was performed using antineutrinos, it is worth pointing out MiniBooNEs antineutrino data is in good agreement with LSND on its own. We show the exact L/E_ν dependence of the LSND and MiniBooNE $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ in Fig-

ure (3.1). As we will discuss in detail later, the addition of more than 1 sterile neutrino allows for CP violating effects so the tension between neutrino and antineutrino running mode results can be alleviated by the including of additional light sterile neutrinos. As LSND and MiniBooNE antineutrino data show very similar oscillation probabilities, this is strong motivation for considering more than one light sterile neutrino, as one can maintain the good agreement with antineutrino data sets and alleviate the tension with neutrino mode using CP violating phases.

Although *consistent* with the LSND anomaly, the MiniBooNE ν_e and $\overline{\nu}_e$ appearance data certainly does not confirm with certainly the $\nu_\mu \rightarrow \nu_e$ sterile neutrino oscillation interpretation. Crucially, MiniBooNE being a Cherenkov light detector, cannot distinguish electrons from photons, meaning the low-energy electron-like excess might be a previously undiscovered nuclear effect or background.

It is worth briefly mentioning there exists additional motivation for a $\mathcal{O}(\text{eV}^2)$ sterile neutrino. A third observation consistent with short-baseline oscillations has been provided in the ν_e disappearance channel from calibration measurements employing intense radioactive sources of high ν_e flux in radiochemical experiments, during the mid 1980's [42, 148], the so-called "gallium anomaly". A similar fourth hint had been provided by past reactor-based short-baseline oscillation searches; specifically, recent reactor data re-analyses using updated reactor flux predictions showed evidence of a deficit in the reactor electron antineutrino event rates measured collectively by several experiments at L/E ranging between 2-20 m/MeV. This has been referred to as the "reactor anomaly" [74]. However, recent realisations that large and unaccounted-for systematic uncertainties are associated with reactor neutrino flux predictions (see, e.g. [104, 149, 150]) dictate that the reactor anomaly cannot yet be interpreted decisively as a light sterile neutrino oscillation signature; such interpretations should await either improved reactor antineutrino flux modelling or dedicated searches for light sterile neutrino oscillations at reactor short baselines which are sensitive to distortions in reconstructed event spectra that are L/E dependent. Such searches are now under way with a number of ex-

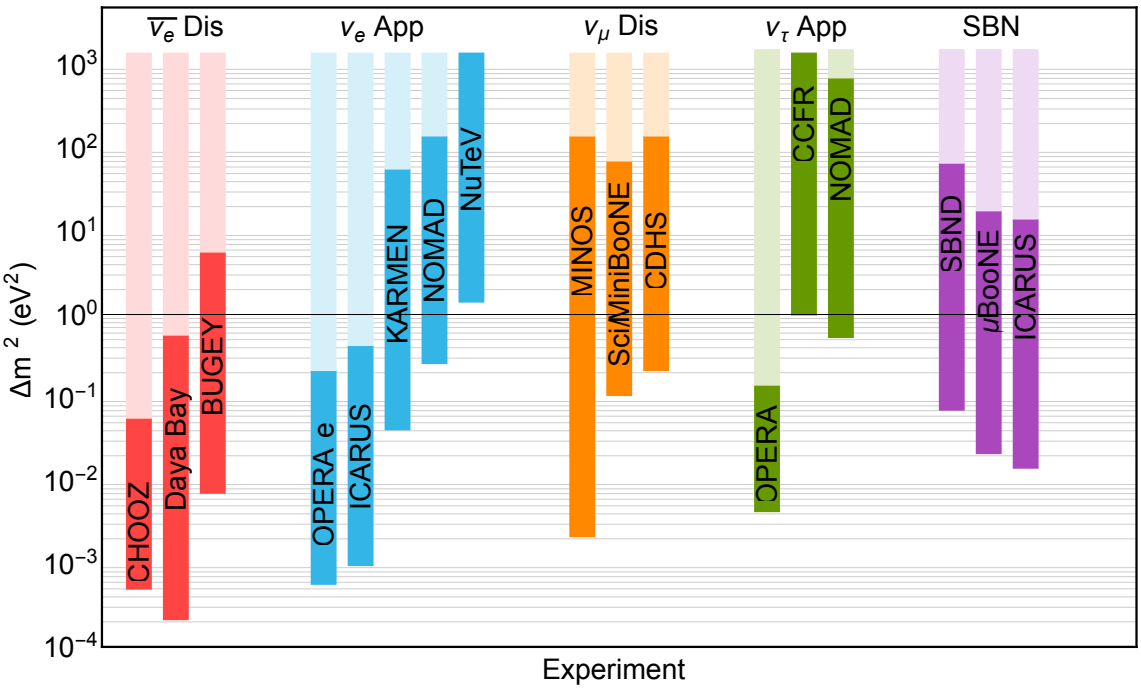

Figure 3.2: The variety of mass-splittings that recent (and future in the case of SBND and ICARUS) experiments are sensitive to light neutrino oscillations. Each experiment has two scales of interest; a point below which oscillations have not yet happened, indicated by the lower end of each coloured band. And a point at which the oscillations are averaged out and thus are only sensitive to normalisation shifts, indicated by the light coloured bands stretching upwards. These points are not clear and distinct and are a function of energy resolution and other systematic uncertainties, and are representative values for qualitative information.

periments [151–157]. In this work we predominantly focus on SBL accelerator effects.

Interpreting the above $\nu_\mu \to \nu_e$ appearance and ν_e disappearance observations as sterile neutrino oscillations would imply large ν_μ disappearance observable at short baselines, as the mixing angle $\sin^2 2\theta_{\mu e}$ being non-zero demands a non-zero $U_{\mu 4}$ in the (3+1) sterile neutrino interpretation. Thus, a $\nu_\mu \to \nu_\mu$ disappearance search can be used as a probe if combined with prior ν_e or $\overline{\nu}_e$ measurements. Such disappearance signatures have not yet been observed; on the contrary, multiple experiments have imposed stringent bounds on sterile neutrino mixing parameters involved in the ν_μ disappearance channel including MiniBooNE itself where the disappearance

searches show no evidence of oscillations [99, 158].

This tension between appearance and disappearance results brings the viability of sterile neutrino models into question [159]. The most recent ν_μ disappearance data sets include IceCube [160] and MINOS+ [161]. The most up to date global fits and results, incorporating IceCube constraints, are presented in Ref. [162]. Despite the strong disappearance constraints, the MiniBooNE, LSND, and arguably the calibration source experimental results still stand as significant anomalous observations that require independent direct tests. Tension in global fits alone, even strong tension, will not be sufficient to rule out the sterile neutrino interpretation.

3.2 The Fermilab Short-Baseline Neutrino Program

To definitively address these collective anomalies, the Short Baseline Neutrino (SBN) experimental program [108] was successfully proposed and is now under construction in the Booster Neutrino Beamline (BNB) at Fermilab. The BNB provides a high intensity, sign-selected, primarily (>99%) muon neutrino (and muon antineutrino) flux [146]. Three liquid argon time projection chamber (LArTPC) detectors, comprising the already operating MicroBooNE detector, the SBND detector which is under construction, and the ICARUS detector which is under refurbishment, sit within the ν_e and $\bar{\nu}_\mu$ flux content at three distinct baselines. This allows SBN to perform an electron neutrino appearance search and a muon neutrino disappearance search with highly competitive sensitivity coverage [108]. Note, however, that the discovery potential of SBN has only been considered for the simplest sterile neutrino scenario in the proposal, the (3+1) model.

ICARUS is the first large-scale LArTPC neutrino detector ever constructed, and has previously operated at Gran Sasso National Laboratory in Italy, starting in 2010. It is presently being refurbished and prepared for transit to Fermilab in Spring of

2018. It has an active mass of 476 tons of liquid argon and will be placed 600 meters from the neutrino production, forming the far detector of the SBN program. MicroBooNE is the mid detector, and it has already began operations in the BNB, in October 2015, and already presented Michel electron results at Neutrino 2016 in London. The MicroBooNE active mass is 89 tons, and it is located at 470 meters from neutrino production, at roughly the same baseline as the MiniBooNE experiment. SBND will act as a near detector of the SBN program, located at 110 meters from neutrino production and with an active mass of 112 tons. It is currently under construction and is scheduled to begin taking data with ICARUS and MicroBooNE in late 2018 [108].

In Figure (3.2) we show the regions of light sterile neutrino mass scales that the three detectors of the SBN program are sensitive to, alongside several other contemporary and historical neutrino oscillation experiments. The dark shaded regions represent the regions that spectral distortions of the neutrino flux are to be expected at the L/E_ν of the experiment due to light sterile neutrino induced oscillations. Below these regions oscillations from small mass-splittings have not occurred yet. The light shaded regions above show the mass-splittings that produce too fast an oscillation to be probed spectrally in any given experiment, but would average out to produce a possibly detectable normalisation shift. As can be seen, all three detectors that make up the SBN program sit directly around the $0.1 \rightarrow 10$ eV^2 regions motivated by the LSND anomaly.

LArTPCs represent a significant advancement forward in neutrino detection technology for the study of both the MiniBooNE low-energy anomaly and the LSND anomaly. LArTPCs have extremely good energy and angular resolution [108] and crucially allows for the distinguishing of electromagnetic showers originating from an electron/positron from those from photons that pair produces. This is achieved utilising the calometric capabilities of LArTPCs to studying the rate of energy deposition of the initial few cm's of a candidate track, the dE/dx. Photons that convert to e^+e^- pairs will deposit roughly twice as much energy compared to a single elec-

tron track [163]. In addition to this, if a vertex is identified by hadronic activity, LArTPC's have sufficient position resolution that it is possible to see the non-ionizing conversion distance that would accompany photons but not electron/positrons. ArgoNeut has shown that LArTPC's can potentially be sensitive to protons as low as 20 MeV [164] allowing for very good vertex tagging.

The strength of the SBN program comes from the utilisation of each of these three detectors in concert. SBND in particular will be recording incredibly high statistics of interactions of the (mostly unoscillated) neutrino flux, and thus will be capable of constraining flux and cross section systematic uncertainties for the farther detectors. Since all three detectors share the same detector technology, their detector uncertainties can also be correlated. This will grant unprecedented sensitivity to short-baseline neutrino oscillations, allowing for the verification or ruling out of a large area of parameter space for (3+N) sterile neutrino oscillations.

In this chapter, we perform an independent phenomenological study where we expand beyond the (3+1) scenario and evaluate SBN's sensitivity to sterile neutrino oscillation models with two and three additional sterile neutrinos, referred to as (3+2) and (3+3), respectively. Furthermore, for the (3+1) scenario, we re-evaluate SBN's sensitivity to electron neutrino appearance without the explicit assumption of negligible disappearance of intrinsic ν_e backgrounds, unlike what has been followed by the SBN collaboration [108]. Because of the large (3+N) parameter space dimensionality for $N = 2, 3$, in order to provide definitive statements on SBN's sensitivity reach with respect to (3+N) models, we exploit existing experimental constraints to sterile neutrino oscillation scenarios, provided in the form of global fits to a representative sample of short-baseline oscillation data sets (both signal and null results).

The Chapter is organised as follows: In Section 3.3, we will discuss the phenomenology of light sterile neutrino oscillation within the short-baseline approximation invoked in this work. In Section 3.4 we give the prescription used to fit global sterile neutrino oscillation data to reduce the parameter space over which the

SBN sensitivity is subsequently quantified. We also summarise the results of fits performed under each oscillation hypothesis in Sections 3.4.1-3.4.3. In Section 3.5, we introduce the SBN experimental facility in detail. In Section 3.5.1, we describe the analysis method followed to estimate SBN's sensitivity to (3+N) sterile neutrino oscillations; more specifically, in Section 3.5.2 we describe the method used to predict the SBN measureable event spectra given any set of (3+N) oscillation parameters, and in Section 4.3 we describe the SBN fitting framework and χ^2 calculation. We present sensitivity results for (3+1), (3+2) and (3+3) in Section 3.6, and we further explore SBN's sensitivity to CP-violating phases measurable in the (3+2) and (3+3) scenarios in Section 3.7. Finally, a summary and conclusions are provided in Section 3.8.

3.3 Light Sterile Neutrino Phenomenology at Short-Baselines

When we derived the oscillation probability in Equation (1.4.49) we noted that no assumptions were made about the number of neutrino species. Thus this equation is the same probability that governs neutrino oscillations, whether it be due to active to sterile neutrinos. Let us recap it here for convenience,

$$\mathcal{P}(\nu_\alpha \to \nu_\beta : L) = \delta_{\alpha\beta} - 4 \sum_{m>j} \mathrm{Re}\left[U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m}\right] \sin^2\left(\frac{\Delta m^2_{mj} L}{4E_\nu}\right) \tag{3.3.3}$$

$$\pm 2 \sum_{m>j} \mathrm{Im}\left[U^\star_{\alpha m} U_{\alpha j} U^\star_{\beta j} U_{\beta m}\right] \sin\left(\frac{\Delta m^2_{mj} L}{2E_\nu}\right), \tag{3.3.4}$$

where the $+(-)$ in the last term describes neutrino (antineutrino) oscillation.

As the existence of a light sterile neutrino is motivated directly by experimental anomalies, in this Chapter we will assume that the first sterile neutrino mass state will be on the order of 1 eV^2, which follows from past and recent global fits [76, 162]. The two lowest mass-squared splittings, Δm^2_{21} and Δm^2_{32}, are both well-established and of order 10^{-5} eV^2 and 10^{-3} eV^2. As both are sufficiently small, we apply the

short baseline approximation, wherein the three lowest mass states are set to be degenerate at $m_1 \sim m_2 \sim m_3 \sim 0$ eV. This also assumes a hierarchy where the ν_1, ν_2 and ν_3 mass states are the lightest. As we will be solely in the short-baseline approximation, working with sub kilometre baselines, we ignore the effect of matter on the oscillation probability. We will consider three scenarios, the (3+1), (3+2) and (3+3) sterile neutrino scenarios corresponding to the addition of one,two and three additional sterile states on top of the three active neutrinos of the Standard Model.

Although Equation (3.3.4) is all one needs to calculate the oscillation probabilities in all scenarios, we expand and study each one in turn here for reference. For a (3+1) sterile neutrino model, the oscillation probabilities for appearance and disappearance are given by

$$P(\nu_\alpha \to \nu_\beta) = 4|U_{\alpha 4}|^2|U_{\beta 4}|^2 \sin^2 x_{41}, \tag{3.3.5}$$

and

$$P(\nu_\alpha \to \nu_\alpha) = 1 - 4|U_{\alpha 4}|^2(1 - |U_{\alpha 4}|^2) \sin^2 x_{41}, \tag{3.3.6}$$

respectively, where $x_{ij} \equiv 1.27\Delta m^2_{ij} L/E$. Thanks to the short baseline approximation and the unitarity of the PMNS matrix, this case reduces to the well known approximation of two neutrino oscillations, where appearance and disappearance mixing angles are often defined as $\sin^2 2\theta_{\alpha\beta} \equiv 4|U_{\alpha 4}|^2|U_{\beta 4}|^2$ and $\sin^2 2\theta_{\alpha\alpha} \equiv 4|U_{\alpha 4}|^2(1-|U_{\alpha 4}|^2)$ respectively. The (3+1) sterile neutrino scenario is CP conserving, so $\overline{\nu}_\alpha$ appearance and disappearance probabilities are identical to those given above for neutrino oscillations. [2]

[2]Note that the (3+1) sterile neutrino scenario is only CP conserving when working in the short-baseline approximation. If you are probing oscillations at long baseline, such as DUNE at $\approx$ 1300km, an additional CP violating phase ϕ_{42} is present on top of the standard 3ν phase, δ_{CP}. See Gandhi et al [165] for more details.

For a (3+2) sterile neutrino model, the oscillation probability is given by

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) = &- 4|U_{\alpha5}||U_{\beta5}||U_{\alpha4}||U_{\beta4}| \cos\phi_{54} \sin^2 x_{54} \\ &+ 4(|U_{\alpha4}||U_{\beta4}| + |U_{\alpha5}||U_{\beta5}| \cos\phi_{54})|U_{\alpha4}||U_{\beta4}| \sin^2 x_{41} \\ &+ 4(|U_{\alpha4}||U_{\beta4}| \cos\phi_{54} + |U_{\alpha5}||U_{\beta5}|)|U_{\alpha5}||U_{\beta5}| \sin^2 x_{51} \\ &+ 2|U_{\beta5}||U_{\alpha5}||U_{\beta4}||U_{\alpha4}| \sin\phi_{54} \sin 2x_{54} \\ &+ 2|U_{\alpha5}||U_{\beta5}||U_{\alpha4}||U_{\beta4}| \sin\phi_{54} \sin 2x_{41} \\ &- 2|U_{\alpha4}||U_{\beta4}||U_{\alpha5}||U_{\beta5}| \sin\phi_{54} \sin 2x_{51}, \end{aligned}$$

in the case of appearance ($\beta \neq \alpha$). Note that this is the sum of not only two individual oscillations at frequencies driven by Δm^2_{41} and Δm^2_{51}, but also an interference term between the two. This interference term contains the CP violating phase, ϕ_{54}. Indeed, for all (3+N) sterile neutrino models with N $\geqslant$ 2, one must consider the complex phases of the associated mixing matrix. These CP violating phases ϕ_{ij} in the oscillation probability are defined as $\phi_{ij} = \arg\{U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}\}$ for neutrino oscillations, and $\phi_{ij} = \arg\{U^*_{\beta i} U_{\alpha i} U_{\beta j} U^*_{\alpha j}\}$ for antineutrino oscillations. This is equivalent to substituting ϕ_{ij} with $-\phi_{ij}$ when considering antineutrino appearance probabilities.

The (3+2) sterile neutrino scenario is of particular interest to us as it is the simplest CP violating theory, when working in the short base-line approximation. It has one phase, ϕ_{54}, which represents a possible source of observable CP violation in the lepton sector. One of the main parameters of interest in studying this is the CP asymmetry defined as

$$\mathcal{A}^{\rm CP}_{54} \equiv \frac{P(\nu_\mu \to \nu_e) - P(\overline{\nu}_\mu \to \overline{\nu}_e)}{P(\nu_\mu \to \nu_e) + P(\overline{\nu}_\mu \to \overline{\nu}_e)}. \tag{3.3.7}$$

We show in Figure (3.3) typical values for the CP asymmetry parameter for L/E_ν appropriate for the three SBN detectors. The colour indicates weather more events would be seen in neutrino mode (red) or anti-neutrino running mode (blue). We highlight the approximate L/E_ν that each detector of SBN covers in the right-hand panel. As can be seen, SBND alone covers a much smaller L/E_ν than MicroBooNE and ICARUS. SBND's role is predominantly that of a very accurate measurement

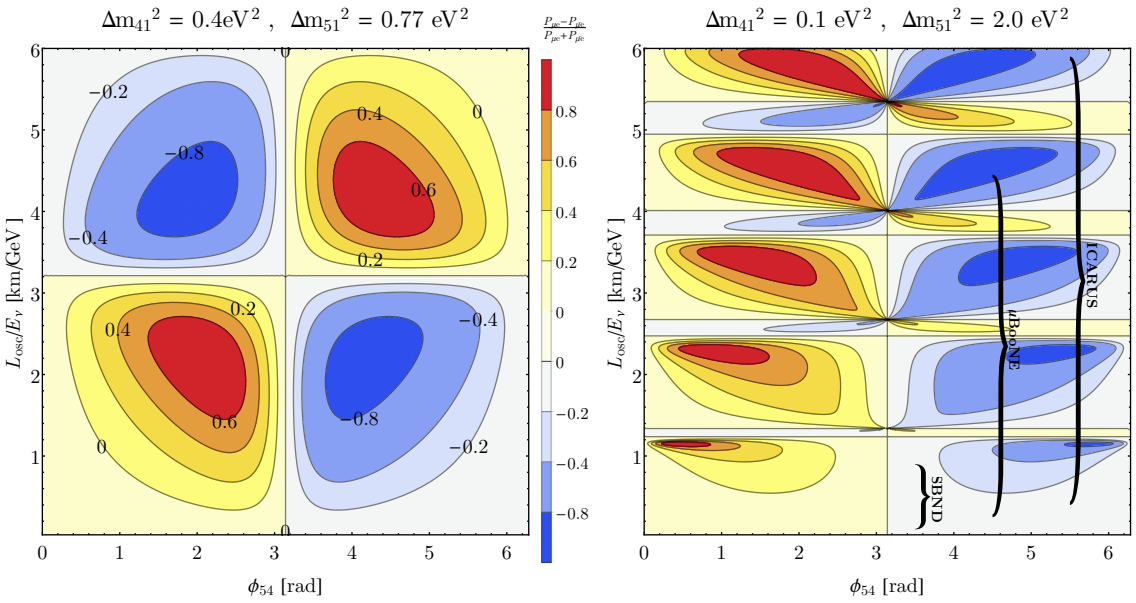

Figure 3.3: Contours of CP Asymmetry as defined in Equation (3.3.7). The same colour scale is in both figures. Highlighted in the right hand figure is the approximate L/E_ν that each of the three SBN detectors cover.

of the fluxes at a smaller baseline. At the increased baselines, and thus larger L/E_ν coverage for a given flux, ICARUS and MicroBooNE will see larger variances in the probabilities. We do stress, however, that statistics at any given L/E_ν, especially larger values, might be very low with poor resolution. Each experiment will generally be most sensitive to one or two oscillation peaks at the centre of their applicable L/E_ν's.

In the case of (3+2) sterile neutrino models, the disappearance probability takes the form

$$\begin{aligned} P(\nu_\alpha \to \nu_\alpha) =& 1 - 4|U_{\alpha 4}|^2|U_{\alpha 5}|^2 \sin^2 x_{54} \\ & - 4(1 - |U_{\alpha 4}|^2 - |U_{\alpha 5}|^2)|U_{\alpha 4}|^2 \sin^2 x_{41} \\ & - 4(1 - |U_{\alpha 4}|^2 - |U_{\alpha 5}|^2)|U_{\alpha 5}|^2 \sin^2 x_{51}. \end{aligned} \tag{3.3.8}$$

Lastly, the (3+3) sterile neutrino model oscillation probability is given by

$$\begin{aligned}
P(\nu_\alpha \to \nu_\beta) = &- 4|U_{\alpha5}||U_{\beta5}||U_{\alpha4}||U_{\beta4}|\cos\phi_{54}\sin^2 x_{54} \\
&- 4|U_{\alpha6}||U_{\beta6}||U_{\alpha4}||U_{\beta4}|\cos\phi_{64}\sin^2 x_{64} \\
&- 4|U_{\alpha5}||U_{\beta5}||U_{\alpha6}||U_{\beta6}|\cos\phi_{65}\sin^2 x_{65} \\
&+ 4(|U_{\alpha4}||U_{\beta4}| + |U_{\alpha5}||U_{\beta5}|\cos\phi_{54} + |U_{\alpha6}||U_{\beta6}|\cos\phi_{64})|U_{\alpha4}||U_{\beta4}|\sin^2 x_{41} \\
&+ 4(|U_{\alpha4}||U_{\beta4}|\cos\phi_{54} + |U_{\alpha5}||U_{\beta5}| + |U_{\alpha6}||U_{\beta6}|\cos\phi_{65})|U_{\alpha5}||U_{\beta5}|\sin^2 x_{51} \\
&+ 4(|U_{\alpha4}||U_{\beta4}|\cos\phi_{64} + |U_{\alpha5}||U_{\beta5}|\cos\phi_{65} + |U_{\alpha6}||U_{\beta6}|)|U_{\alpha6}||U_{\beta6}|\sin^2 x_{61} \\
&+ 2|U_{\beta5}||U_{\alpha5}||U_{\beta4}||U_{\alpha4}|\sin\phi_{54}\sin 2x_{54} \\
&+ 2|U_{\beta6}||U_{\alpha6}||U_{\beta4}||U_{\alpha4}|\sin\phi_{64}\sin 2x_{64} \\
&+ 2|U_{\beta6}||U_{\alpha6}||U_{\beta5}||U_{\alpha5}|\sin\phi_{65}\sin 2x_{65} \\
&+ 2(|U_{\alpha5}||U_{\beta5}|\sin\phi_{54} + |U_{\alpha6}||U_{\beta6}|\sin\phi_{64})|U_{\alpha4}||U_{\beta4}|\sin 2x_{41} \\
&+ 2(-|U_{\alpha4}||U_{\beta4}|\sin\phi_{54} + |U_{\alpha6}||U_{\beta6}|\sin\phi_{65})|U_{\alpha5}||U_{\beta5}|\sin 2x_{51} \\
&+ 2(-|U_{\alpha4}||U_{\beta4}|\sin\phi_{64} - |U_{\alpha4}||U_{\beta5}|\sin\phi_{65})|U_{\alpha6}||U_{\beta6}|\sin 2x_{61},
\end{aligned} \tag{3.3.9}$$

in the case of appearance, where now we have three CP violating phases, ϕ_{54}, ϕ_{64} and ϕ_{54}. The corresponding disappearance probability is given by

$$\begin{aligned}
P(\nu_\alpha \to \nu_\alpha) = &1 - 4|U_{\alpha4}|^2|U_{\alpha5}|^2\sin^2 x_{54} - 4|U_{\alpha4}|^2|U_{\alpha6}|^2\sin^2 x_{64} - 4|U_{\alpha5}|^2|U_{\alpha6}|^2\sin^2 x_{65} \\
&- 4(1 - |U_{\alpha4}|^2 - |U_{\alpha5}|^2 - |U_{\alpha6}|^2)(|U_{\alpha4}|^2\sin^2 x_{41} + |U_{\alpha5}|^2\sin^2 x_{51} + |U_{\alpha6}|^2\sin^2 x_{61}).
\end{aligned} \tag{3.3.10}$$

We note here that for all $(3 + N)$ sterile scenarios, in both appearance and disappearance, the neutrino oscillation probabilities take the form of a sum of defined oscillations with amplitudes given by the mixing matrix elements, at frequencies of either $\sin^2 x_{ij}$ or $\sin 2x_{ij}$. We will use this fact in Section (4.3) to aid in in the simplification of computation and calculation of SBN's expected spectra. This is true for all neutrino oscillations, as can be seen easily in Equation (1.4.49), but is also true when one includes a matter potential, as the probability can be re-parameterized in the same function form but with shifted mass-splittings and effective mixing angles [166, 167].

3.4 Globally Allowed (3+N) Parameter Space

Although we have qualitatively discussed the parameter space that drives the LSND and MiniBooNE anomalies, it is instructive here to study in more detail the global situation of light sterile neutrinos. We will use this to quantitatively motivate what exact parameter space future experiments, such as SBN, should be sensitive to in order to have the strongest capability to confirm or rule out the LSND and MiniBooNE anomaly. To such an end, for each (3+N) scenario under consideration, a fit is performed over all existing short-baseline neutrino experiment data, to extract the globally-allowed 90% and 99% C.L region. The experimental data sets included in the global fit are summarised in Table 3.1. We omit the recent MINOS+ [161] and IceCube [160] constraints from the global fits, although we note that in the future those constraints should be included for more quantitatively accurate results. We expect that the qualitative conclusions drawn in this work stand regardless of inclusion of those more recent constraints

.

Given the sheer scale of the parameter space, particularly for (3+3) oscillations, which features twelve fully independent mixing parameters, a grid scan of any reasonable resolution would be computationally impossible. Instead, we elect to perform a more efficient scanning of the mixing parameters for each oscillation scenario using the a Markov chain minimization routine, following the method employed in Reference [76]. We define range of intererest in which each oscillation parameter is allowed to vary as follows:

$$0 \leqslant U_{\alpha i} \leqslant 0.5, \qquad 0.01 \leqslant \Delta m^2_{i1} \leqslant 100 \text{ eV}^2, \qquad 0 \leqslant \phi_{ij} < 2\pi. \tag{3.4.11}$$

where $\alpha = e, \mu$ and $i, j = 4, ..., 3 + N$. Initial values for the N additional neutrino mass states, mixing matrix elements and CP violating phase(s) are generated randomly from within their appropriate ranges. Then, each parameter θ is generated for each successive step using the recursive Markov chain condition $\theta_{\text{new}} = \theta_{\text{old}} + (R - 0.5)(\theta_{\text{max}} - \theta_{\text{min}})s$ where R is a random number in (0,1) and s is a configurable maximum step size. Additionally, we apply further unitarity constraints on $U_{\alpha i}$, by rejecting parameter points where any of the following definitions

Dataset	Oscillation Channel
Appearance	
KARMEN [88]	$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$
LSND [71]	$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$
MiniBooNE - BNB [72, 146, 168, 169]	$\overset{(-)}{\nu}_\mu \rightarrow \overset{(-)}{\nu}_e$
MiniBooNE - NuMI [170]	$\nu_\mu \rightarrow \nu_e$
NOMAD [93]	$\nu_\mu \rightarrow \nu_e$
Disappearance	
KARMEN, LSND (xsec) [171]	$\nu_e \rightarrow \nu_e$
Gallium (GALLEX and SAGE) [42, 148]	$\nu_e \rightarrow \nu_e$
Bugey [74, 172]	$\bar{\nu}_e \rightarrow \bar{\nu}_e$
MiniBooNE - BNB [99, 173]	$\overset{(-)}{\nu}_\mu \rightarrow \overset{(-)}{\nu}_\mu$
MINOS-CC [174, 175]	$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$
CCFR84 [79]	$\nu_\mu \rightarrow \nu_\mu$
CDHS [83]	$\nu_\mu \rightarrow \nu_\mu$
Atmospheric Constraints [176–180]	$\nu_\mu \rightarrow \nu_\mu$

Table 3.1: The short-baseline oscillation data sets included in global fits to (3+N) sterile neutrino oscillation scenarios, and used to provide allowed regions over which SBN's sensitivity is quantified.

are invalid $\sum_{i=4,\ldots,3+N} |U_{\alpha i}|^2 \leqslant 0.3$ for $\alpha = e, \mu$, and $\sum_{\alpha=e,\mu} |U_{\alpha i}|^2 \leqslant 0.3$ for $i = 4, 5, 6$.

For each given step, a χ^2 is calculated for the entire set of oscillation parameters and is compared against the χ^2 from the previous step, χ^2_{old}, to determine the probability P_T of accepting this new point into the Markov chain. This probability is given by $P_T = \min(1, e^{-(\chi^2 - \chi^2_{\rm old})/T})$, where T is the "temperature" of the Markov chain, which can be static or vary itself. The case of slowly dropping temperature, or simulated annealing, is a very powerful technique to help ensure the Markov chain reaches the global minimum χ^2 while evading local minima. Multiple starting chains, with randomly chosen intial R, s and T further ensure that the true global minima is located.

For each experimental data set included in the global fit, a Monte Carlo prediction is calculated using the oscillation probability derived for a given set of sterile neutrino oscillation parameters, and compared against observed data from the experiment. The resulting χ^2 for each experimental data set is summed to form a global χ^2 for each sterile neutrino model, assuming that there are no correlations among data sets.

We use the resulting global χ^2 to determine the parameter space allowed at a certain confidence level, using a $\Delta\chi^2$ cut relative to the global χ^2 minimum, $\chi^2_{\rm min}$, that the Markov Chain found. Once a globally-allowed region for a certain scenario is obtained, the region gets discretized over a grid of 100^n spacepoints, where n is the number of oscillation parameters in the given scenario. The spacepoints are evenly distributed over the ranges defined above, and in a linear scale in mixing elements $U_{\alpha i}$ and a logarithmic scale in Δm^2_{i1}. For illustrating two-dimensional projected allowed regions, we profile over the oscillation parameter space and thus a $\Delta\chi^2$ cut of 4.61 (90% C.L) and 9.21 (99% C.L) using 2 degrees of freedom (dof) is applied. However, to extract the n-dimensional phase-space over which we later quantify the SBN sensitivity, the $\Delta\chi^2$ cuts applied more appropriately correspond to n dof, where $n = 3, 7$ and 12 dof for (3+1), (3+2) and (3+3), respectively.

In the following sections we will provide a summary of the global fit results we obtain using this method. These regions will then subsequently be used as input to further detailed sensitivity studies.

3.4.1 (3+1) Globally Allowed Regions

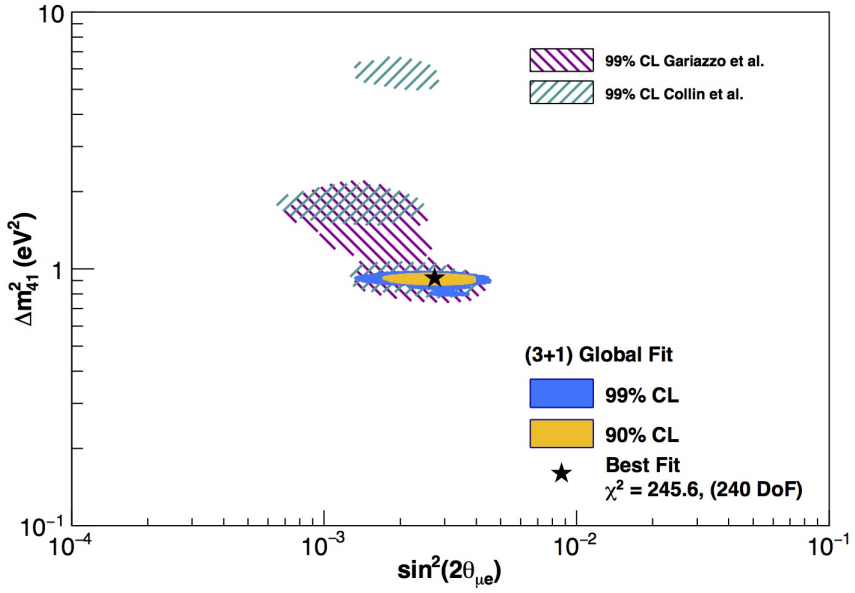

Figure 3.4: The 90% and 99% C.L regions allowed by a simultaneous fit to all data sets listed in Tab. 3.1 and following the prescription in Section 3.4, under a (3+1) sterile neutrino oscillation hypothesis. Overlaid are results from other recent global fit analyses, which do includethe new constraints from the IceCube experiment [160]. There are three free oscillation parameters in this fit, but here we profile over them to provide 2D projections in regions of Δm^2_{41} and $\sin^2 2\theta_{\mu e} = 4|U_{e4}|^2|U_{\mu 4}|^2$ that are allowed at the chosen confidence levels, assuming 2 dof.

In this section, we summarise the results of the global fit to all data sets listed in Tab. 3.1 under the (3+1) oscillation hypothesis. The best fit parameters obtained in this fit, and corresponding $\chi^2_{\rm min}/dof$, are provided in Tab. 3.2. A two-dimensional allowed region profiled into Δm^2_{41}-$\sin^2 2\theta_{\mu e}$ is illustrated in Figure (3.4), where $\sin^2 2\theta_{\mu e} = 4|U_{e4}|^2|U_{\mu 4}|^2$. The region at around 1 eV^2 is largely driven by the

LSND and MiniBooNE anomalies, the two strongest signals of new physics. Note, however, that the recent IceCube constraints appears to shift this allowed region slightly, to higher Δm^2_{41} and slightly low mixing amplitudes. The χ^2 difference between the $\Delta m^2_{41} \sim 1$ eV^2 and $\Delta m^2_{41} \sim 2$ eV^2 regions has been reported to be very small, which is why we have proceeded to carry out sensitivity studies without the IceCube constraints included.

3.4.2 (3+2) Globally Allowed Regions

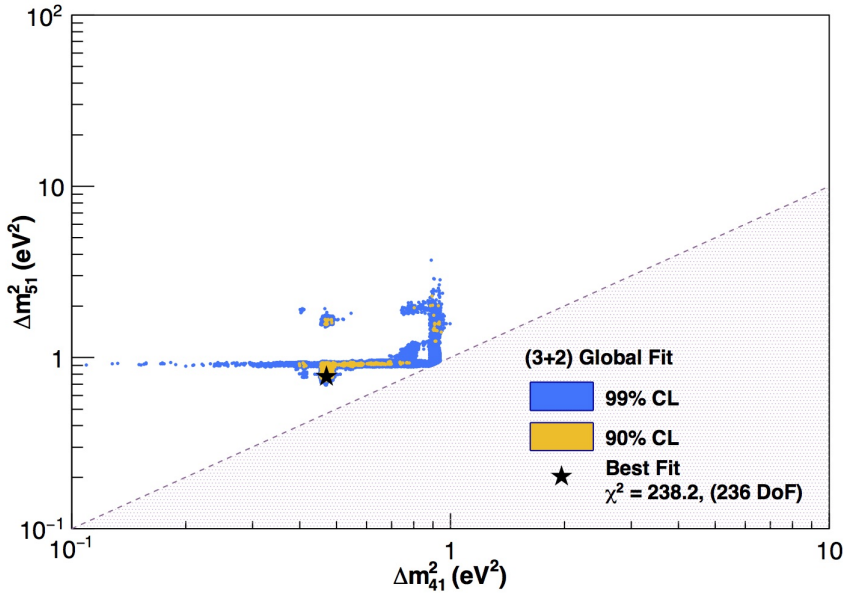

Figure 3.5: The 90% and 99% C.L regions allowed by a simultaneous fit to all data sets listed in Tab. 3.1, and following the prescription in Section 3.4, under a (3+2) sterile neutrino oscillation hypothesis. There are seven free oscillation parameters in this fit, but here we profile over them to provide 2D projections in regions of Δm^2_{41} and Δm^2_{51} that are allowed at the chosen confidence levels, assuming 2 dof.

In this subsection, we summarise the results of the global fit to all data sets listed in Tab. 3.1 under the (3+2) oscillation hypothesis. The best fit parameters obtained in this fit, and corresponding $\chi^2_{\rm min}/dof$, are provided in Tab. 3.2. A two-dimensional allowed region profiled into Δm^2_{51}-Δm^2_{41} is illustrated in Figure (3.5).

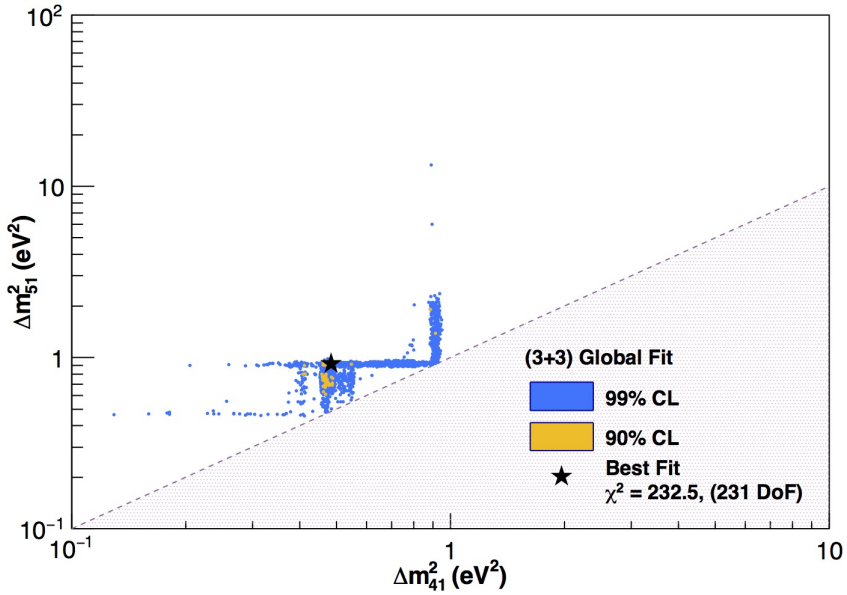

Figure 3.6: The 90% and 99% C.L regions allowed by a simultaneous fit to all data sets listed in Tab. 3.1, and following the prescription in Section 3.4, under a (3+3) sterile neutrino oscillation hypothesis. There are twelve free oscillation parameters in this fit, but here we profile over them to provide 2D projections in regions of Δm^2_{41} and Δm^2_{51} that are allowed at the chosen confidence levels, assuming 2 dof.

As mentioned above, in adding a second light sterile neutrino one also adds the CP-violating phase, ϕ_{54}. This additional phase can be influential at short baselines and can relieve some of the tension between neutrino and antineutrino data sets, providing a better overall fit to global data. This improvement has been demonstrated to be the case in particular when considering appearance-only data sets (see, e.g. [75, 76, 159]). As can be seen in Figure (3.5), one mass-splitting at the $\approx$ 1eV^2 level is required to explain the LSND and MiniBooNE anomalies in the same was as the (3+1) scenario, however, the second mass splitting can vary $\approx 0.1 \rightarrow 4$ eV^2.

3.4.3 (3+3) Globally Allowed Regions

Finally we summarise the results of the global fit to all data sets listed in Tab. 3.1 under the (3+3) oscillation hypothesis. The best fit parameters obtained in this fit, and corresponding χ^2_{min}/dof, are provided in Tab. 3.2. Two-dimensional allowed

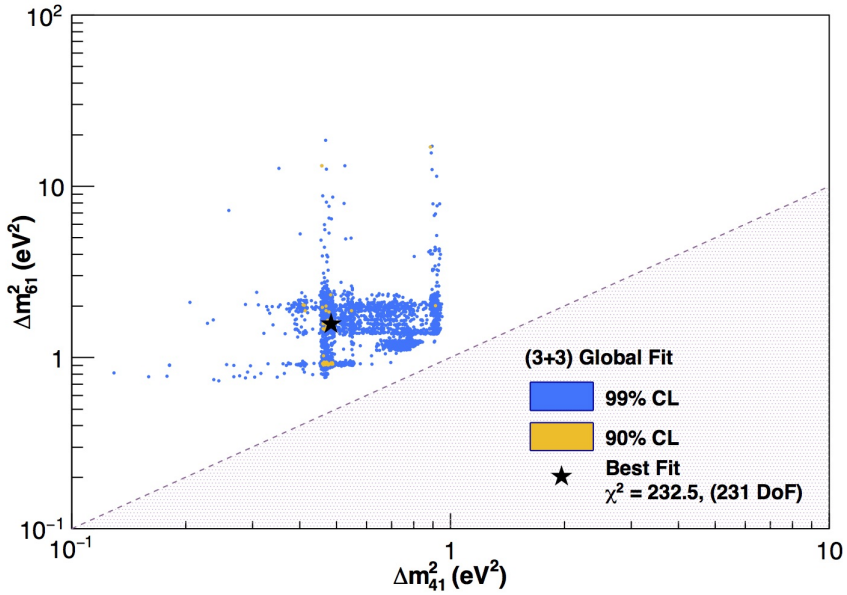

Figure 3.7: The 90% and 99% C.L regions allowed by a simultaneous fit to all data sets listed in Tab. 3.1, and following the prescription in Section 3.4, under a (3+3) sterile neutrino oscillation hypothesis. There are twelve free oscillation parameters in this fit, but here we profile over them to provide 2D projections in regions of Δm^2_{41} and Δm^2_{61} that are allowed at the chosen confidence levels, assuming 2 dof.

regions profiled into Δm^2_{51}-Δm^2_{41} and Δm^2_{61}-Δm^2_{51} are illustrated in Figure (3.6) and Figure (3.7), respectively.

With the addition of yet another light sterile degree of freedom comes with five additional parameters, including an additional independent mass splitting, two additional mixing elements, and two additional CP-violating phases. This further increases the hypervolume of parameter space allowed under the global data sets, although the preference for one of the best fit Δm^2_{i1} being close to $\mathcal{O}(1\text{eV}^2)$ evident in the (3+1) and (3+2) hypotheses seems to persist. Furthermore, as in the (3+2) case, the additional CP-violating phase(s) has been shown to lead to a further reduction in tension between neutrino and antineutrino data sets as well as an overall lessening of the disagreement between appearance-only and disappearance-only fits (see, e.g., Refs. [75, 76, 159]).

(3+1)	Δm^2_{41}	$U_{\mu 4}$	U_{e4}	χ^2/dof
Best Fit	0.92	0.17	0.15	245.6/240

(3+2)	Δm^2_{41}	$U_{\mu 4}$	U_{e4}	Δm^2_{51}	$U_{\mu 5}$	U_{e5}	ϕ_{54}	χ^2/dof
Best Fit	0.46	0.15	0.13	0.77	0.13	0.14	5.56	238.2/236

(3+3)	Δm^2_{41}	$U_{\mu 4}$	U_{e4}	Δm^2_{51}	$U_{\mu 5}$	U_{e5}	Δm^2_{61}	$U_{\mu 6}$	U_{e6}
Best Fit	0.68	0.18	0.12	0.90	0.13	0.14	1.55	0.03	0.12
	ϕ_{54}	ϕ_{64}	ϕ_{65}	χ^2/dof					
	5.60	4.31	3.93	232.5/231					

Table 3.2: Global best-fit parameters obtained under the (3+1) (top), (3+2) (middle) and (3+3) (bottom) oscillation hypothesis. Mass-squared splittings are presented in eV^2 and CP violating factors are given in in radians. The null hypothesis has a χ^2/dof of 299.5/243.

3.5 SBN Sensitivity to (3+N) Oscillations

The global fits as calculated above are extremely powerful tools for guiding the parameter space that experimentalists focus on, but they do not replace, nor should they be thought of, a single dedicated experiment. Global fits can involve a tremendous amount of simplifications, approximations and assumptions which may vary significantly and unexpectedly from those of the original experiments which they attempt to combine. No matter how much global combined data disagrees with the LSND anomaly, until a single experiment can definitely rule out or confirm the light sterile oscillation hypothesis it will still be considered an open anomaly by many. None the less, they do provide a useful reference for the parameter space of interest. For the remainder of this Chapter we will focus on determining how sensitive SBN is to this parameter space associated with the LSND and MiniBooNE light sterile

oscillations, and attempt to answer the question "Will SBN *definitively* be able to cover the light oscillating sterile anomaly?".

3.5.1 Sensitivity Analysis Method

In order to evaluate SBN's sensitivity to $(3 + N)$ sterile neutrino oscillations, we consider the oscillation-induced fluctuations that are measurable in the exclusive ν_e (and $\bar{\nu}_e$) charged-current (CC) and ν_μ (and $\bar{\nu}_\mu$) CC event spectra of each of the SBN detectors. The event spectra are provided in terms of reconstructed neutrino energy, and are estimated as described in Section 3.5.2.

The ν_e CC spectrum is sensitive to potential ν_μ to ν_e appearance in the ν_μ-dominated BNB beam. For this sample, because background contributions are comparable to signal contributions for most of the globally-allowed (3+N) oscillation parameter space, we additionally consider the effects of (1) disappearance of the ν_e (and $\bar{\nu}_e$) intrinsic background in the beam; and (2) ν_μ (and $\bar{\nu}_\mu$) disappearance of the mis-identified background from ν_μ (and $\bar{\nu}_\mu$) CC interactions. We assume that the mis-identified background from neutral-current (NC) interactions will be measured and constrained independently and *in situ* for each of the SBN detectors, and therefore we ignore any oscillation variations on that particular background in these fits.

The ν_μ CC spectrum, on the other hand, is sensitive to exclusively ν_μ disappearance. In this case, we ignore not only oscillation variations on any backgrounds, but also background contributions from NC $\pi^\pm$ production events altogether. Based on Ref. [108], this background contribution has negligible effect on the SBN sensitivity.

Combining ν_e and ν_μ CC measurements thus allows one to simultaneously constrain both appearance and disappearance probabilities for ν_e and ν_μ oscillations. The fit method is described in detail Section 3.5

3.5.2 Predicting SBN Event Spectra

The SBN ν_e and ν_μ CC event spectra were fully simulated on an event-by-event basis. The raw rates of each flavor of neutrino impinging on the three SBN detectors were evaluated using the flux predictions in Reference [181]. Events are generated in GENIE 2.8.6 [182] (default settings used) separately for each neutrino type (ν_e, ν_μ, $\bar{\nu}_e$, $\bar{\nu}_\mu$) and for the beam polarity in both neutrino and antineutrino mode.

Ten million events were generated for each flavor, detector, and beam polarity. This corresponds to 8×10^{20} POT for the SBND neutrino mode ν_μ flux, and significantly more for all other samples. Weights were applied to all events to normalise them to the rates predicted by GENIE for the expected exposure and each detector mass.

Subsequent to event generation, events were processed further to emulate the reconstructed and selected ν_e CC and ν_μ CC spectra presented in Ref. [108]. More specifically, to estimate detector effects without the need for a full detector simulation, neutrino interaction product energies were smeared according to a Gaussian around their true value, to emulate expected detector energy resolution discussed in Ref. [108]: electrons and photons received a $15\%/\sqrt{E}$ smearing, whereas muons and pions received a $6\%/\sqrt{E}$ smearing; all protons with true kinetic energy below 21 MeV were assumed to be non-reconstructable, while those above this threshold as well as other charged hadrons had their kinetic energies smeared by 5%. All smeared hadronic energies were added to form the hadronic activity, and the reconstructed neutrino energy was then defined as the total sum of visible (smeared) lepton or photon energy and hadronic activity, as well as the rest masses of all leptons and non-proton charged hadrons. A lower threshold of 100 MeV was placed on electron and photon energies in order for them to be defined as reconstructable. This is quite a conservative estimate, as the true threshold in liquid argon should be substantially lower than this, but was chosen to be in line with the SBN proposal.

The fiducial volume cut efficiency for each detector was then emulated by ran-

domising the neutrino interaction vertex position within the predefined active volume, and applying geometric cuts, with the position and direction of all simulated muons and e/γ showers accounted for to accurately estimate backgrounds and efficiency's. This is of utmost importance to the ν_e appearance signal as $\pi^0 \rightarrow \gamma\gamma$ decays, in which only one photon is reconstructed successfully, are a major background. The following contributions are included explicitly in the ν_e CC sample:

- Intrinsic ν_e CC events are the largest contribution to the ν_e CC sample. All appearance signal (from potential $\nu_\mu \rightarrow \nu_e$ oscillations) and intrinsic beam ν_e CC events producing an electron with reconstructed neutrino energy $E_{\text{reco}} \geqslant$ 200 MeV are included with an overall 80% identification efficiency. This value of 80% was taken from the SBN proposal, where it was estimated from hand-scanning MC simulated events.

- NC single photon events, from either NC Δ production followed by radiative decay, or π^0 production followed by decay into two photons where only one photon is reconstructable, are also considered as potential background contribution in the ν_e CC sample. In particular, events in which the photon is reconstructed too close (within 3 cm) to a vertex identified by significant hadronic activity (defined as $E_{\text{visible hadronic}} \geqslant 50$ MeV), or in which no hadronic activity is visible, are included as backgrounds if the reconstructed event energy satisfies the 200 MeV threshold. Those selected events receive an additional scaling assuming a 94% photon rejection efficiency.

- ν_μ CC events in which the muon is mis-identified as a pion and simultaneously an additional photon (e.g from π^0 decay) mimics the electron from a ν_e CC event are also included as a background contribution to the ν_e CC sample. To quantify this background, all ν_μ CC events with a track length $\geqslant$ 1 m are assumed to be identifiable as muons and are rejected. Those with a track length below 1 m are accepted as potential mis-identified events, if any photons in the event are accepted under the same conditions as in the NC single photon events, above.

- Interactions outside of the TPC producing photons that propagate inside the

active volume are a potential source of backgrounds as well. These "Dirt" backgrounds are included with rates taken directly from Ref. [108]. We assume that independent measurements of these backgrounds at each detector location will render this contribution insensitive to any oscillation effects.

- Cosmogenic backgrounds are expected to be well constrained by topological, calometric and timing cuts, with the background contribution scaling linearly with beam exposure (POT). The numbers we use are taken directly from Ref. [108] and correspond to 146, 88 and 164 cosmogenic background events for SBND, MicroBooNE and ICARUS, respectively. Although significantly smaller than the intrinsic ν_e CC backgrounds, as they tend to accumulate at low energy, they are included in our analysis following the approach in Ref. [108].

Similarly, for the ν_μ CC sample, intrinsic beam ν_μ CC events are assumed to be selected with an 80% reconstruction and identification efficiency. Potential background contributions would result from NC $\pi^\pm$ interactions where the $\pi^\pm$ can be mis-identified as a muon. This is mitigated by requiring that all contained muon-like tracks have a track length larger than 50 cm, and that all escaping tracks that have a track length of less than 1 m are rejected. This is the same methodology as what was followed in Ref. [108].

We show the results of the ν_e CC and ν_μ CC simulations for all SBN detectors in Figure (3.8), along with an estimated appearance-only signal prediction for a benchmark (3+1) sterile neutrino oscillation model with $\Delta m^2_{41} = 0.39\text{eV}^2$ and mixing for ν_e appearance of $\sin^2 2\theta_{e\mu} = 0.013$ in the upper plot, and ν_μ disappearance for a sterile neutrino with $\Delta m^2_{41} = 1.1\text{eV}^2$ and $\sin^2 2\theta_{\mu\mu} = 0.1$ in the lower panels.

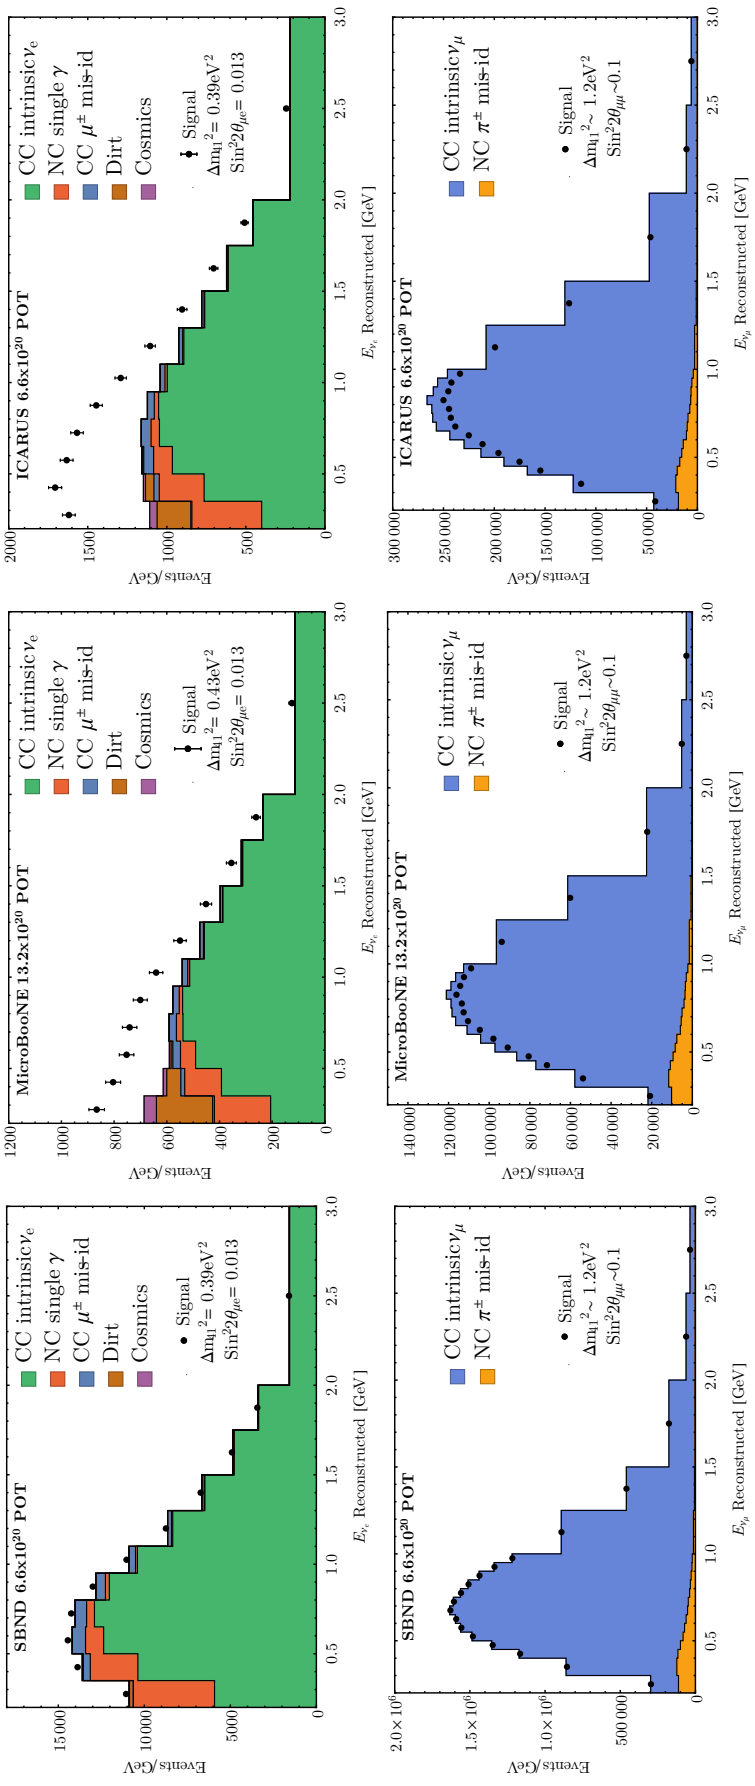

Figure 3.8: ***Top:*** The ν_e CC inclusive sample used in SBN sensitivity studies, shown for SBND, MicroBooNE and ICARUS detectors respectively. Expected intrinsic and mis-identified backgrounds to ν_e appearance/disappearance are shown in stacked, coloured histograms. Shown also is the expected signal for a benchmark sterile neutrino oscillation model with $\Delta m^2_{41} = 0.39\text{eV}^2$ and $\sin^2 2\theta_{e\mu} = 0.013$, for comparison. ***Bottom:*** The ν_μ CC inclusive sample used in SBN sensitivity studies, shown for SBND, MicroBooNE and ICARUS detectors respectively. We only consider true ν_μ events and $\pi^\pm$ mis-identified events as a possible background. Shown in black point is a benchmark sterile neutrino oscillation model with $\Delta m^2_{41} = 1.2\text{eV}^2$ and $\sin^2 2\theta_{\mu\mu} = 0.1$.

In quantifying sensitivity, we consider primarily two fitting methods:

- ν_e appearance-only fits, where $N_k^{\text{osc}}(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})$ is evaluated assuming only $\nu_\mu \to \nu_e$ oscillations, and no ν_e or ν_μ disappearance; this is the method followed by past MiniBooNE oscillation searches [72] as well as in Ref. [108]; and

- combined ν_e dis/appearance and ν_μ disappearance fits, where $N_k^{\text{osc}}(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})$ is evaluated assuming $\nu_\mu \to \nu_e$ oscillations, ν_e disappearance, as well as ν_μ disappearance. We note that this is the first time that SBN sensitivities are evaluated without the implicit assumption of no significance ν_e or ν_μ disappearance; we note, as demonstrated in the results section, this implicit assumption can have a significant effect on the SBN sensitivity.

Finally, when considering the ν_e and ν_μ CC inclusive samples for antineutrino running mode, we follow the same methodology as above. It should be noted that none of the SBN detectors are assumed to be capable of differentiating between a neutrino and an antineutrino interaction on an event by event basis.

3.5.3 SBN χ^2 Calculation

To facilitate a multi-baseline, multi-channel, and multi-mode (neutrino and antineutrino running) oscillation search with the SBN detectors, we use a custom fitting framework to generate, and subsequently fit the reconstructed ν_e CC and ν_μ CC inclusive spectra expected at each detector with and without oscillations, and for each running mode, simultaneously. This framework is refered to as "SBNfit". This simultaneous, side-by-side fit of multiple event samples by way of a full covariance matrix which contains statistical and systematic uncertainties as well as systematic correlations among the different samples, baselines, and running modes, is an approach that has been followed by the MiniBooNE collaboration for several analyses, e.g. [72, 99, 146, 168, 169, 173], as well as by the SBN collaboration (although not in a multi-channel fitting manner as we utilise here). We have chosen this approach specifically so that we may exploit powerful correlations shared within and among the spectra that are measurable by each of the three detectors.

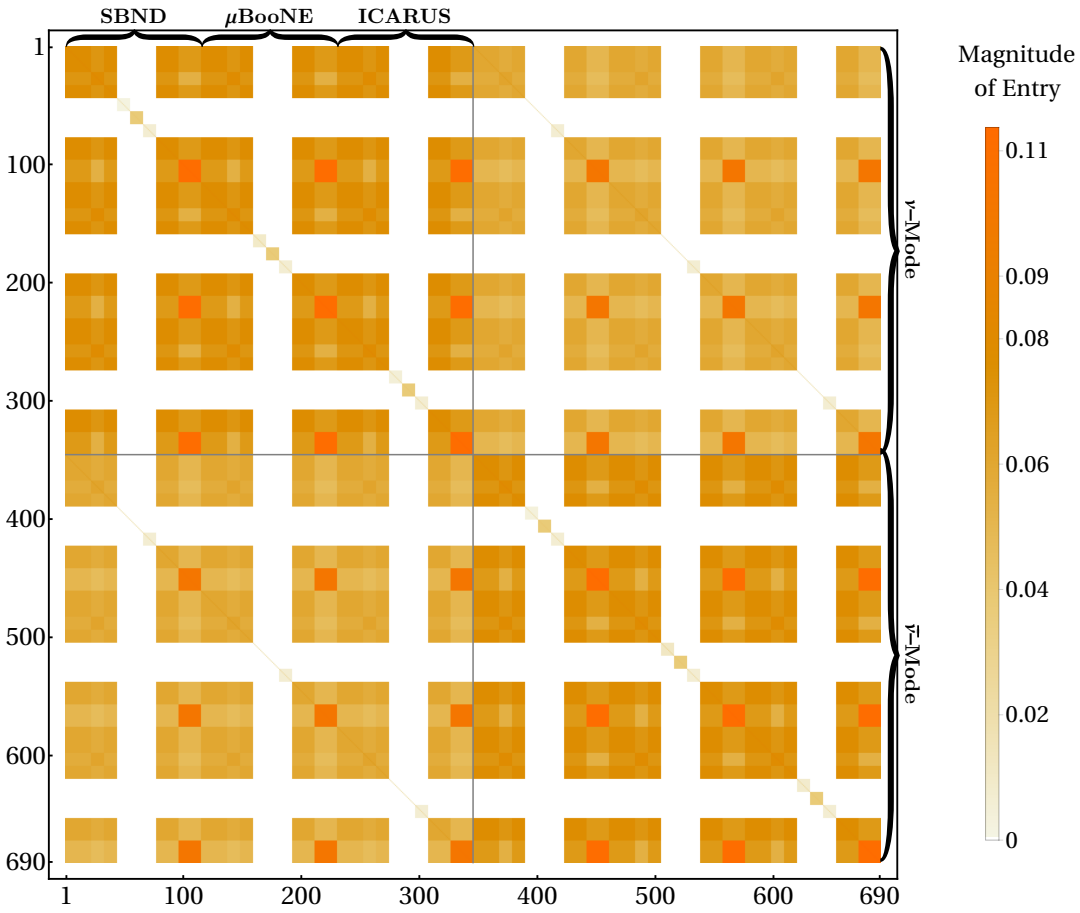

Figure 3.9: 690×690 fractional correlation matrix for combined ν_e-like and ν_μ-like signals at all three SBN detectors, in a dual neutrino and antineutrino running mode configuration. As this is a direct visualisation of the correlation matrix, the x and y axis just refer to each matrix element. For details of construction see text.

The SBN fit quality is quantified over an n-dimensional oscillation parameter space volume $(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})$ by way of a χ^2. The χ^2 is calculated over concatenated ν_e CC inclusive and ν_μ CC inclusive spectra for all three detectors, as

$$\chi^2(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij}) = \sum_{k=1}^{M}\sum_{l=1}^{M}\left[N_k^{\text{null}} - N_k^{\text{osc}}(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})\right] E_{kl}^{-1}\left[N_l^{\text{null}} - N_l^{\text{osc}}(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})\right] , \tag{3.5.12}$$

where N_k^{null} is the number of events expected under the no oscillation hypothesis (defined as $U_{\alpha i} = 0 \quad \forall \quad \alpha, i, j$) in the k^{th} bin of reconstructed neutrino energy; $N_k^{\text{osc}}(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})$ is the number of events predicted to be observed in reconstructed neutrino energy bin k under an oscillation hypothesis described by the

set of parameter values $(\Delta m^2_{i1}, U_{\alpha i}, \phi_{ij})$; and E_{kl} is a full $M \times M$ covariance matrix containing the total systematic and statistical uncertainty, including systematic correlations between any two bins k and l. The ν_e CC and ν_μ CC samples for each detector location are binned in 11 and 19 bins of reconstructed neutrino energy, respectively, as shown in Figure (3.8). This binning is the same as used in all previous MiniBooNE analysis, as well as SBN proposal, although due to improvements in resolution the effect of modifying binning should be investigated going forward. Thus, for all three detector locations, the concatenated spectra N_k^{null} and N_k^{osc} consist of a total of $M = 90$ bins for neutrino mode only, and $M = 180$ bins for neutrino and antineutrino combined fits.

The covariance matrix, which is a 90×90 matrix for neutrino only fits, and a 180×180 matrix for combined neutrino and antineutrino fits, is calculated as the sum of individual sources of systematic and statistical uncertainties,

$$E = E^{\text{stat}} + E^{\text{flux}} + E^{\text{cross section}} + E^{\text{cosmic}} + E^{\text{dirt}} + E^{\text{detector}} . \quad (3.5.13)$$

Table (3.3) summarises the assumed variations on specific contributions to the inclusive ν_e and ν_μ CC samples due to different sources of systematic uncertainty; those variations were used to calculate the fractional systematics covariance matrix. The assumed numbers are based on Ref. [108]. Flux systematic uncertainties are estimated by assuming an overall 20% normalization uncertainty fully correlated among the intrinsic ν_e (background and signal) and ν_μ events, with the exception of exclusive samples that are assumed to be constrained *in situ*; namely, dirt, cosmogenic, and NC backgrounds in the ν_e CC sample. A 60% $\nu_e - \nu_\mu$ flux correlation coefficient is assumed among ν_e and ν_μ events. Cross section systematic uncertainties are estimated by assuming an overall 20% normalization uncertainty fully correlated among CC-only events, and a corresponding 30% normalization uncertainty among NC-only events. Again, dirt, cosmogenic, and NC backgrounds in the ν_e CC sample are exempted from this uncertainty. A 50% CC-NC cross section correlation coefficient is assumed among CC and NC events. Furthermore, neutrino and antineutrino run CC cross section uncertainties are assumed to be 100% correlated, and likewise for NC cross-section uncertainties. Detector systematics are assumed to be fully

Source of Uncertainty	Assumed variation
ν_e flux	15.3% on ν_e events
ν_μ flux	15.1% on ν_μ events
CC cross section	20% on CC events
NC cross section	30% on NC events
detector effects	2.5% on all events

Table 3.3: Assumed variations on exclusive event samples due to different systematic uncertainties, used to evaluate the total systematics covariance matrix. See text for more details.

uncorrelated among different detectors, and contribute to the overall uncertainty at the level of 2.5%. These are taken to be fully correlated for neutrino and antineutrino run samples in any given detector.

The dirt event rate uncertainty is assumed to be constrained through *in situ* dirt-enhanced sample measurements at each detector and in each running mode. A 15% normalization uncertainty is assumed for dirt events, taken to be uncorrelated between the different detectors and the neutrino and antineutrino run samples. Similarly, the cosmogenic background uncertainty is assumed to be constrained through *in situ* off-beam high-statistics rate measurements at each detector. A 1% normalization uncertainty is assumed for cosmic backgrounds, assumed to be uncorrelated between different detectors, but fully correlated between neutrino and antineutrino run samples within any given detector. Finally, NC backgrounds are also assumed to be constrained through an situ NC π^0 event rate measurements in each detector, thus the estimated statistical uncertainty of the *in situ* measurement is taken as the systematic uncertainty on these backgrounds. This corresponds to a 0.24%, 1.3%, and 5% normalization uncertainty for the SBND, MicroBooNE, and ICARUS NC background rates, respectively. This systematic uncertainty is assumed to be uncorrelated for neutrino and antineutrino run samples. We show the resultant fractional correlation matrix in Figure (3.9). When operating in neutrino running mode (or

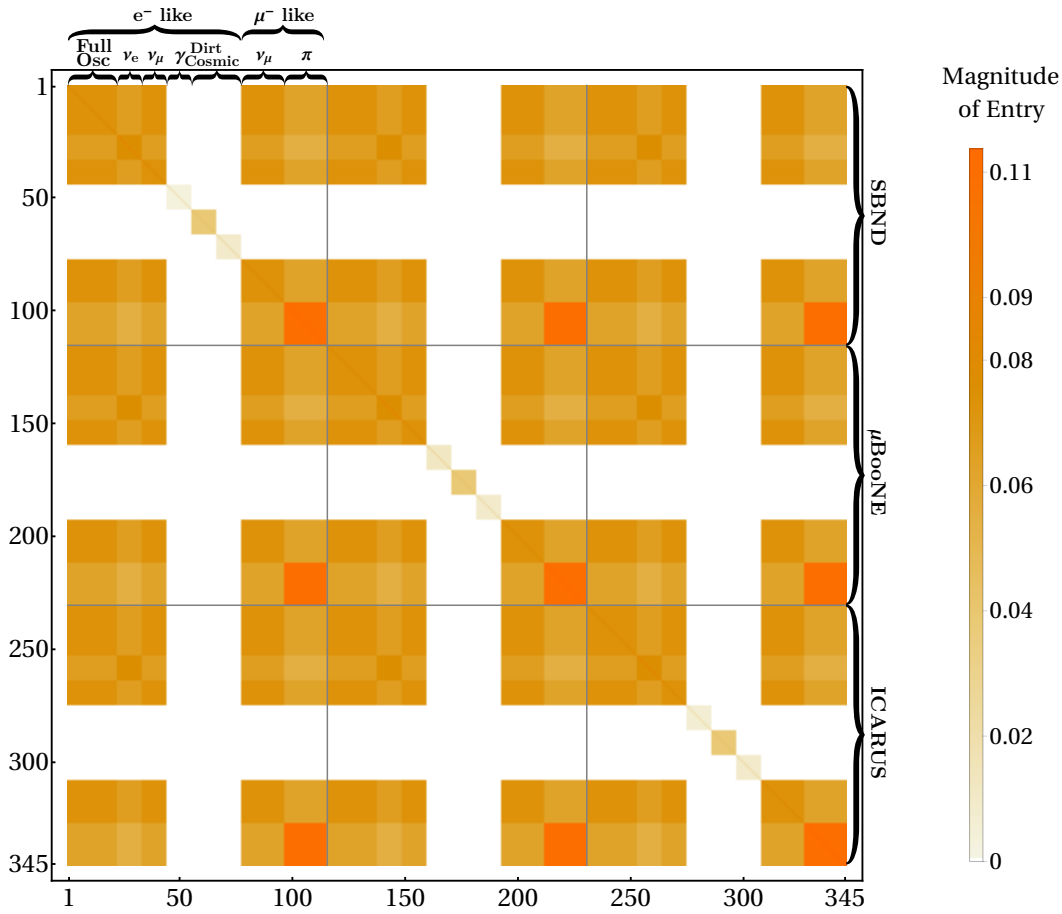

Figure 3.10: 345 × 345 fractional correlation matrix for combined ν_e-like and ν_μ-like signals at all three SBN detectors in neutrino mode only. This is a zoomed in detail of Figure (3.9). Full Osc refers to both the fully oscillated $\nu_\mu \to \nu_e$ and $\overline{\nu}_\mu \to \overline{\nu}_e$ fluxes.

anti-neutrino running mode) only the appropriate 345×345 subset is used. This subset is shown in detail in Figure (3.10) where we highlight each of the electron like spectra and each of the muon-like spectra.

As we have previously noted in Equation (3.3.9) both the neutrino appearance and disappearance oscillation probabilities, for any number of sterile neutrinos, are the sum of a finite number of frequencies of either $\sin^2 x$ or $\sin x$ functional form. When calculating the χ^2, rather than compute the event spectra at SBN for every set of oscillation parameters by running over each event in the 40 million event input

flux files, a very lenghtly and computatenly intense proceedure, we elect to pre-compute six sets of 100 spectra corresponding to the expected events at each SBND, MicroBooNE and ICARUS for both a $\sin x$ and $\sin^2 x$ frequency oscillation with unit amplitude. These 100 spectra correspond to 100 mass-squared differences ranging in a log-scale from 0.01 eV^2 to 100 eV^2. The calculation of a given χ^2 then involves the loading of the expected spectra from the pre-computed frequencies corresponding to input Δm^2. This is then weighted by input U_{PMNS} elements accordingly to form each amplitude. The final spectra at each SBN detector site is the sum total of each individual spectra calculated this way.

3.6 SBN Sensitivity to Sterile Neutrino Oscillations: Results

3.6.1 (3+1) Scenario at SBN

Throughout this analysis we will use the globally allowed regions of sterile neutrino parameter space, as described in Section 3.4, to investigate which parameter space areas SBN should be strongest at probing. For reference, we first explore SBN's sensitivity reach in neutrino running mode under three separate oscillation cases:

- $\nu_\mu \rightarrow \nu_e$ appearance-only (assuming no ν_μ or ν_e disappearance whatsoever). We note that this case requires an unphysical assumption in a (3+1) oscillation hypothesis, as $\nu_\mu \rightarrow \nu_e$ appearance implies both ν_μ and ν_e disappearance. However, in the past this case has been applied to MiniBooNE searches to a reasonably valid approximation, and has furthermore been applied to SBN sensitivity studies in Ref. [108]. We therefore consider it only as an instructive example, and to further argue that it is not a reasonably valid approximation to use for SBN.

- ν_μ disappearance-only (assuming no ν_e dis/appearance). We consider this case only as an instructive scenario, as the interpretation of short-baseline positive

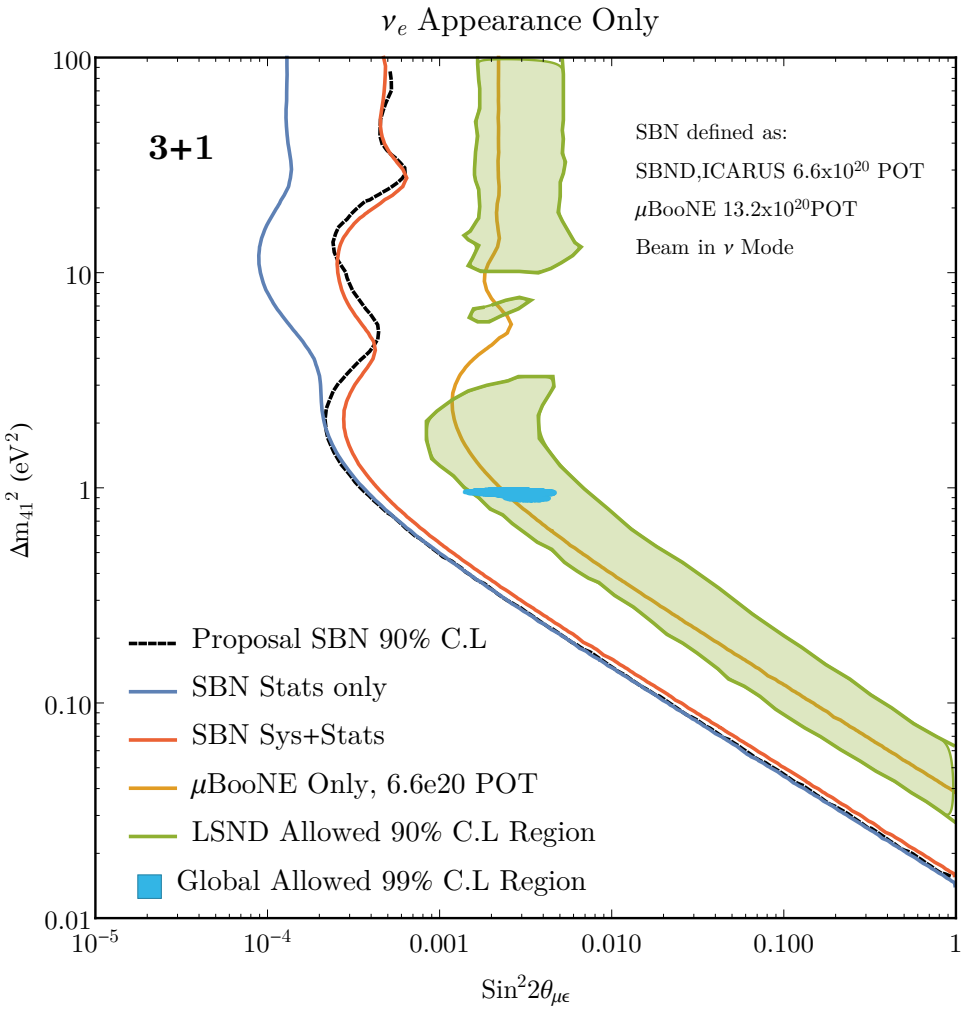

Figure 3.11: The estimated 90% C.L exclusion contours for the entire SBN program for ν_e appearance only (yellow solid line) with full detector, flux and cross-section systematics included as well as statistic only (blue). The same contour as estimated in the SBN proposal is shown in (black dashed) line. This vastly covers the current 99% (3+1) allowed regions (crimson) and LSND 90% allowed region (green). Shown also is the μBooNE only contour (orange) which can probe a large fraction of the global allowed region.

signals also require ν_e dis/appearance. However, ν_μ disappearance only is physically allowed if $U_{ei} = 0$, unlike appearance only.

- ν_e disappearance-only (assuming no ν_μ disappearance or ν_e appearance). We also consider this case only as an instructive scenario, as the interpretation of short-baseline positive signals require both ν_e and ν_μ disappearance (and ν_e

appearance).

Figure (3.11) shows the SBN appearance-only sensitivity reach in Δm^2_{41} vs. $\sin^2 2\theta_{e\mu}$ under a (3+1) hypothesis obtained using the χ^2 definition described in Section 4.3 and applying a "raster scan" over this reduced two-dimensional parameter space. The appearance-only sensitivity is provided here merely for comparison to the sensitivity presented in the SBN proposal [108], which uses the same assumption of no ν_e background disappearance, as a means of validating our methodology. The results of this work when incorporating full detector, cross-section and flux systematics (yellow curve) is consistent with the results as published in the SBN proposal (black curve). The statistics-only sensitivity curve obtained in this work is shown in blue. Comparing the blue and red curves demonstrates the effect of systematic uncertainties on the sensitivity, which is to diminish sensitivity to higher-Δm^2_{41} oscillations. This is due to the fact that the dominant systematic is the flux and cross-section normalization uncertainty. The comparison also demonstrates the power of exploiting correlations that exist among multiple baselines and multiple interaction channels. Accounting for these correlations leads to an effective cancellation of systematic uncertainties in particular in the low-Δm^2_{41} region. Shown also is the MicroBooNE-only (μBooNE) result after its first run corresponding 6.6e20 POT. Overlaid over all these curves is the LSND 90% C.L allowed region (shaded green area) as well as the (3+1) 99% C.L globally allowed region from Figure (3.4). The raster scan sensitivities are obtained using a $\Delta\chi^2$ cut for 1 dof, while the globally allowed region corresponds to a global scan using a $\Delta\chi^2$ cut for 2 dof.

The SBN ν_μ disappearance-only search gives the sensitivity curve shown in Figure (3.12) (red curve). As the sensitivity presented in the SBN proposal (black curve) does not include detector systematics, it outperforms the one obtained in this work. This is expected, as detector systematics across the three detectors are taken to be fully uncorrelated in our fits. As a cross check, we compare to the statistics-only sensitivity obtained in this work (blue curve), which is found to lie mostly to the left of both other curves, as expected. Shown also is the prediction for MicroBooNE (μBooNE) after its first 6.6e20 POT exposure.

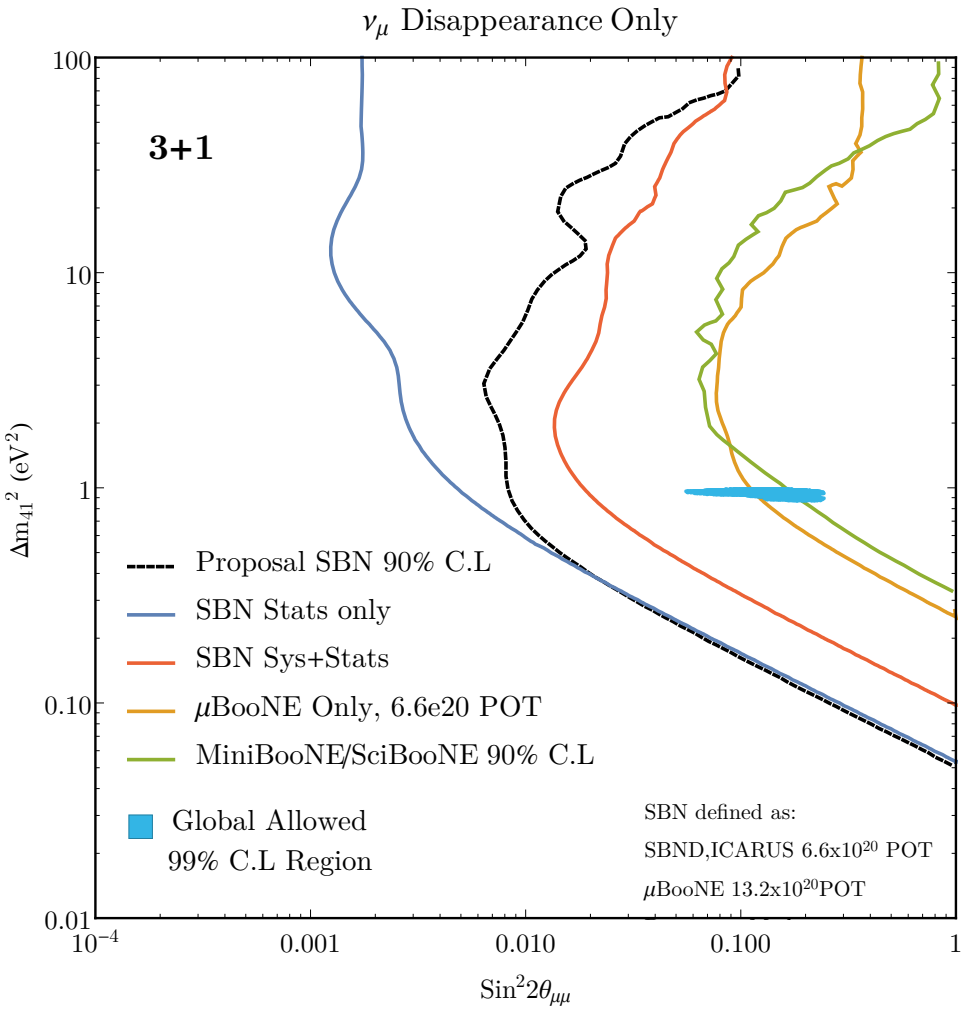

Figure 3.12: The estimated 90% C.L contours for the combined SBN using ν_μ disappearance only. The globally allowed region in Δm^2_{41} and $\sin^2 2\theta_{\mu\mu}$ is completely covered. Shown also is the prediction for MicroBooNE after 6.6e20 POT.

Due to the proximity of the SBND experiment to the BNB target, the flux of intrinsic ν_e at the detector is extremely large. Specifically, SBND expects to record over 35,000 CC events in 6.6e20 POT. This allows for an additional oscillation channel to be probed, that of ν_e disappearance. The SBN ν_e disappearance-only sensitivity reach is shown in Figure (3.13) (red curve). We note that this is the first time that SBN's sensitivity to ν_e disappearance is being explored. Although this search is less sensitive to the 1eV^2 region, due to the fact that the ν_e flux has a higher mean energy, at a higher Δm^2_{41} values it is comparable in $\sin^2 2\theta_{ee}$ reach to

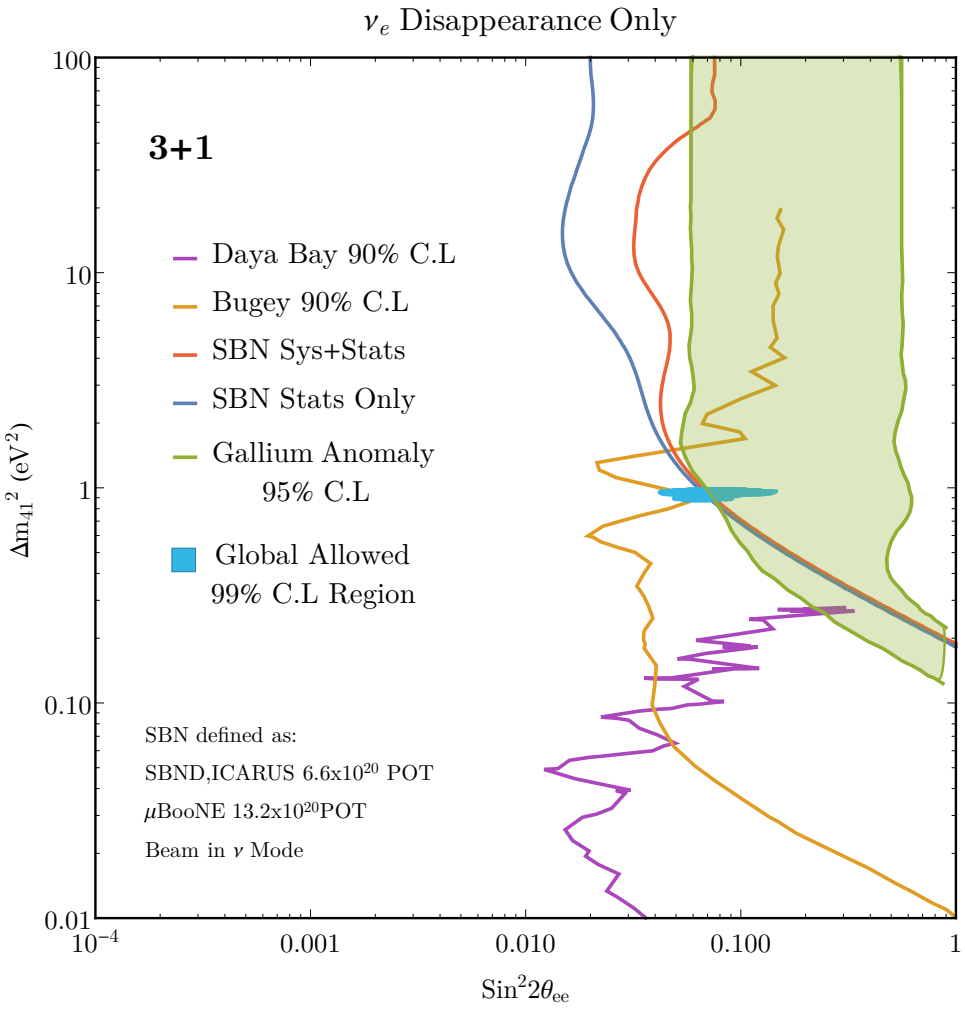

Figure 3.13: Due to the very large intrinsic ν_e component of the beam at SBND, one can also perform a ν_e disappearance only analysis directly probing $\sin^2 2\theta_{ee}$ at high $\Delta m^2 \geqslant 0.2\text{eV}^2$. This is traditionally probed using reactor antineutrinos at a much lower MeV scale energy, and so would provide yet another way of probing the low-energy sterile neutrino anomalies. This is a direct probe of $\sin^2 2\theta_{ee}$ using a neutrino beam rather than the lower energy (MeV) reactor antineutrinos.

that of reactor short-baseline $\bar{\nu}_e$ disappearance bounds. It is also a direct probe of $\sin^2 2\theta_{ee}$ using a high-energy neutrino beam in complementarity with the MeV-scale antineutrino reactor flux.

We note that, although instructive, none of the above three cases are appropriate for an SBN oscillation search if one believes the sterile neutrino contains mixing to

both the electron and muon sectors. Instead, a proper search for oscillations at SBN should consider the simultaneous effects of both ν_e disappearance and ν_μ disappearance and, consequently, ν_e appearance. We therefore adopt this case, referred to as ν_e dis/appearance and ν_μ disappearance, as the proper SBN sensitivity search method, and we present results corresponding to this case throughout the following sections.

As the primary physics goal of the SBN programme is to *definitively* probe the light sterile neutrino sector that could be responsible for the low-energy anomalies, we define here a new metric used to quantify how well SBN can achieve this goal under each of the (3+1), (3+2) and (3+3) scenarios. This metric is referred to as *Global X% C.L Coverage*, and it refers to the fraction of hypervolume of the $X\%$ C.L globally-allowed oscillation parameter space that can be ruled out by SBN with a certain confidence level, if SBN observed no oscillations. To estimate global coverage, we first discretize the sterile neutrino parameter space in 100 points in each independent mass-squared difference and mixing element. The mass-squared difference is discretized over the range of 0.01 eV^2 to 100 eV^2 (in grid points that are equidistant in logarithmic scale), while the mixing elements $|U_{\alpha 4}|$ are discretized in 100 linearly spaced grid points ranging from 0 to 0.5, and the CP violating phases in 100 points ranging from 0 to 2π. This allows to calculate a hypervolume represented by the number of space points or the "*size*" of parameter space that is preferentially allowed by global data at a given confidence interval (i.e. 99%). We can then express SBN's sensitivity reach as the fractional number of space points or fraction of this hypervolume that SBN can exclude at any given confidence level.

A concrete example of this methodology is shown in Figure (3.14), where we show the percent of the 99% C.L allowed region that SBN can exclude at a given $\Delta\chi^2$ in a ν_e appearance only (dotted line), a ν_μ disappearance only (dashed line), as well as a ν_e dis/appearance and ν_μ disappearance (solid line) fit, assuming 6.6E20 POT collected concurrently with all three SBN detectors, after the first MicroBooNE-only run of 6.6E20 POT. The results for the (3+1) scenario are shown in the top panel.

Shown also are the results for the (3+2) and (3+3) scenarios, in the middle and bottom panels, which will be discussed in their respective sections below.

From the top panel, it is evident that the best performance is possible in the case of a ν_e dis/appearance and ν_μ disappearance search (solid line). In that case, SBN can cover close to 100% of the 99% C.L globally allowed (3+1) parameter space at 3σ, and similarly 85% of the parameter space at 5σ. In contrast, an appearance-only search can only cover 85% of the parameter space at 3σ, and only 50% of the parameter at 5σ. We note that in drawing these comparisons we use $\Delta\chi^2$ cuts corresponding to three *dof* for all three cases (ν_e appearance, ν_μ disappearance, and ν_e dis/appearance and ν_μ disappearance.

Nevertheless, although a ν_e dis/appearance and ν_μ disappearance search provides a more powerful sensitivity to the (3+1) parameter space, one would like to see a strong exclusion in both the exclusive ν_e appearance search and the exclusive ν_μ disappearance and ν_e disappearance searches individually in order to conclusively rule out any light sterile neutrino oscillation hypothesis. The POT at which such a statement can be made is explored in Figure (3.15), which shows the SBN 3σ and 5σ coverage (in yellow and red, respectively) as a function of POT delivered to the SBN program. As we assume that MicroBooNE has already ran for 6.6e20 POT by the time that the three-detector SBN program commences, the x axis corresponds explicitly to the POT delivered for the three-detector operations, and the plot by construction demonstrates the MicroBooNE-only (6.6E20 POT) coverage at $x = 0$. We note that even a MicroBooNE-only combined ν_e dis/appearance and ν_μ disappearance search would yield a 3σ coverage of 25% of the (3+1) globally-allowed parameter space. In general, the total coverage is driven primarily by the ν_μ disappearance channel, as evident by the dotted line(s) lying close to the solid line(s).

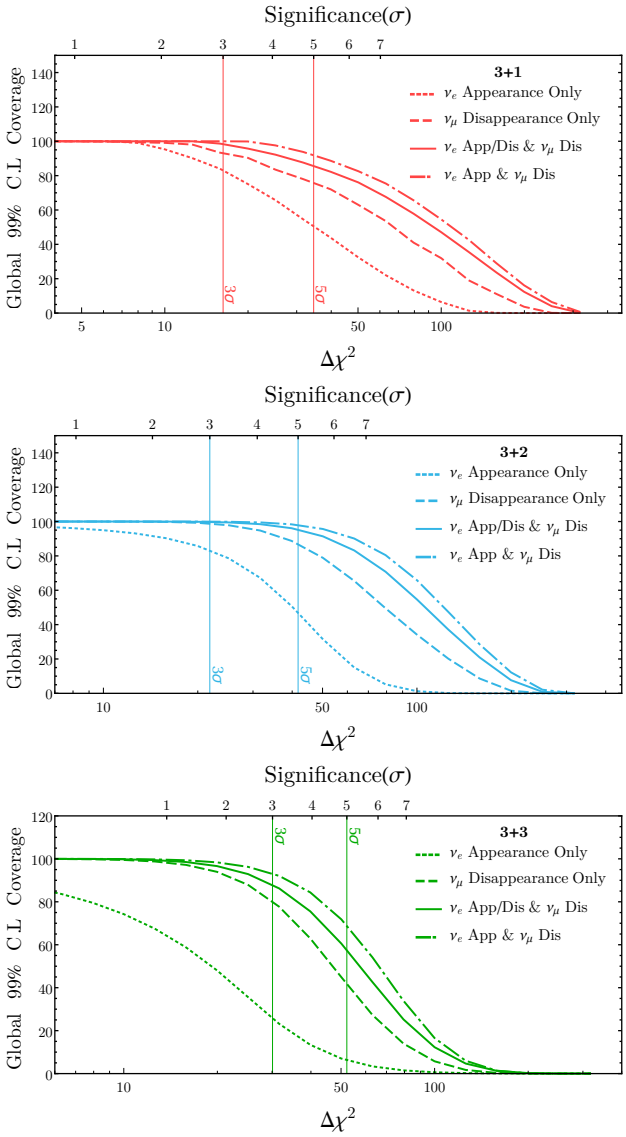

Figure 3.14: SBN coverage, showing the fraction of 99% C.L allowed global fit region that SBN can exclude at any given $\Delta\chi^2$, for the (3+1) (red, top) (3+2) (blue, middle) and (3+3) (green, bottom) sterile neutrino oscillation scenarios. The dotted curves correspond to ν_e appearance only searches, the dashed curves correspond to ν_μ disappearance only searches, and the solid curves correspond to a combined ν_e dis/appearance and ν_μ disappearance search, which provides the highest sensitivity overall. The percentage covered is shown as a function of $\Delta\chi^2$ on the bottom x-axis and as a function of significance on the top x-axis, assuming 3, 7 and 12 dof for (3+1), (3+2), and (3+3) fits, respectively.

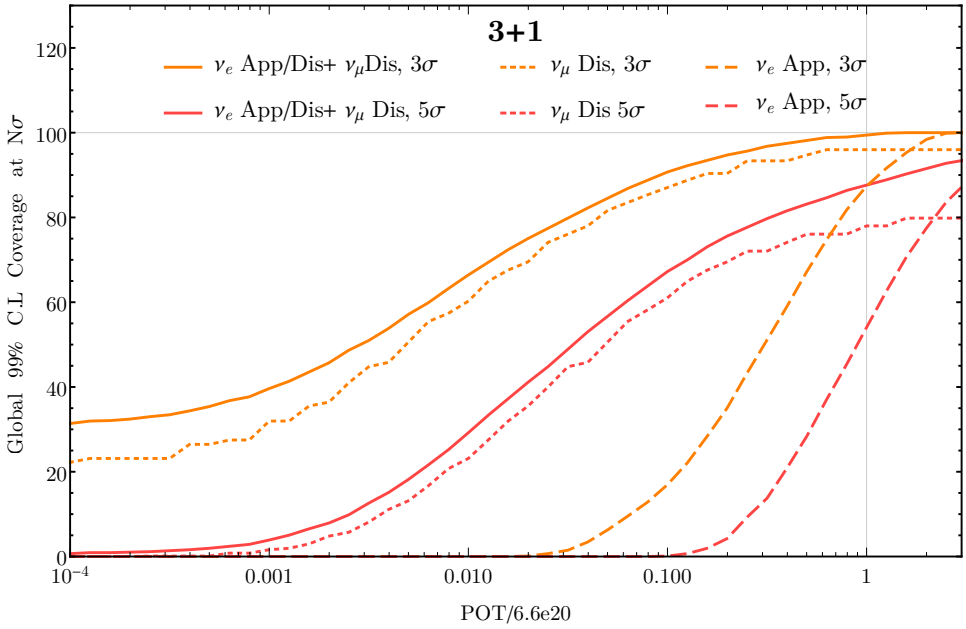

Figure 3.15: The percentage of 99% C.L globally allowed (3+1) parameter space that SBN can exclude at the 3 (orange) and 5 (red) σ C.L for ν_e appearance only (dotted), ν_μ disappearance only (dashed) and a combined appearance and disappearance fit (solid), as a function of POT. We assume that MicroBooNE has already obtained 6.6e20, hence the plateau at low POT.

3.6.2 (3+2) Scenario at SBN

To achieve its goal of definitely addressing these oscillations, SBN will need to have extensive coverage of the (3+2) (and similarly (3+3)) sterile neutrino oscillation parameters as well. In the case of the (3+2) scenario, the additional parameters introduced when one adds another light sterile neutrino happen to enlarge the size of the parameter space that is preferred by the global fits. Nevertheless, as can be seen in the middle panes of Figure (3.14), the percentage of globally allowed (3+2) parameter space (at 99% C.L) that SBN can cover at any given confidence level is generally comparable to that of the (3+1) scenario. SBN is able to cover 100% (95%) of parameter space the 3(5)σ level in a combined ν_e dis/appearance and ν_μ disappearance under the (3+2) scenario. In contrast, using ν_e appearance-only fits, SBN is limited to a maximum of 82(46)% possible coverage at 3(5)σ, assuming a nominal exposure of 6.6E20 POT. The SBN 3σ and 5σ coverage of the (3+2)

parameter space as a function of POT can be shown in Figure (3.16). We note that in drawing these comparisons we use $\Delta\chi^2$ cuts corresponding to seven *dof* for all three cases.

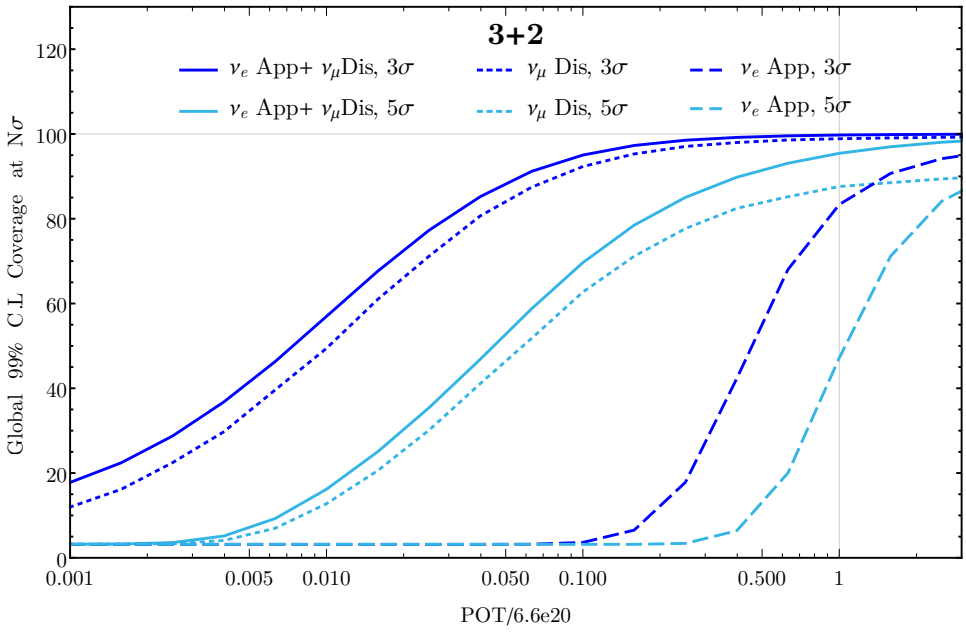

Figure 3.16: The same as Figure (3.15) but for the 3+2 light sterile neutrino scenario. The percentage of 99% C.L globally allowed 3+2 parameter space that SBN can exclude at the 3 (dark blue) and 5 (light blue) σ C.L for ν_e appearance only (dotted), ν_μ disappearance only (dashed) and a combined appearance and disappearance fit (solid), as a function of POT.

3.6.3 (3+3) Scenario at SBN

The (3+3) scenario represents the most difficult scenario for the SBN program to definitively rule out, containing a total of three independent CP violating phases and twelve independent mass and mixing parameters. As can be seen in Figure (3.14), bottom panel, at its full planned exposure of 6.6E20 POT, the SBN program can cover only 90(57)% of the globally allowed 99% C.L region at greater than 3(5)σ, and only with a combined $\nu-e$ dis/appearance and ν_μ disappearance search. In a ν_e appearance-only search, SBN only covers 25(5)% of the globally allowed parameter space at 3(5)σ. The SBN coverage of (3+3) allowed regions as a function of delivered POT is shown in Figure (3.17).

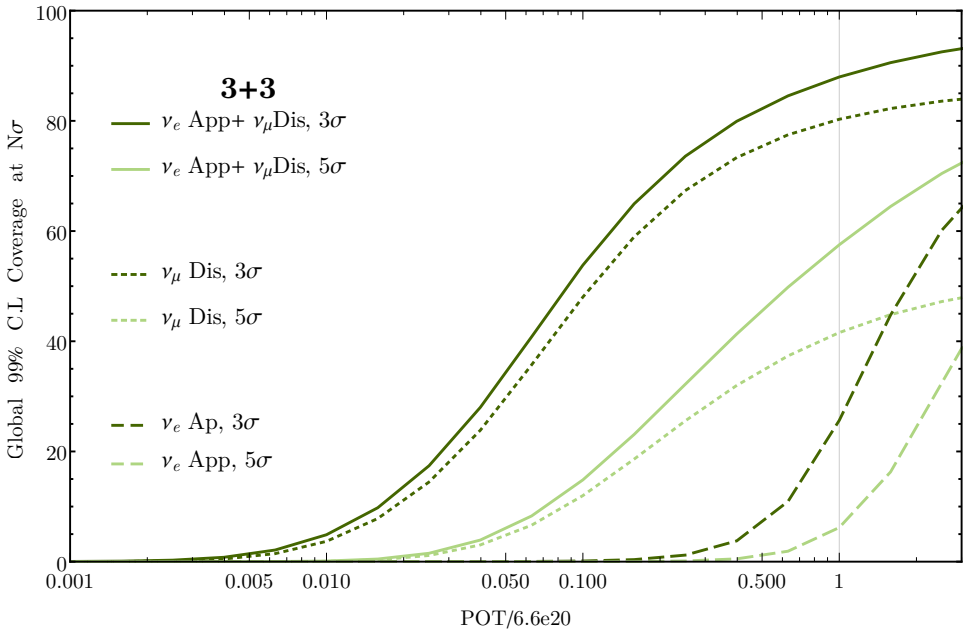

Figure 3.17: The same as Figure (3.15) but for the 3+3 light sterile neutrino scenario. The percentage of 99% C.L globally allowed 3+3 parameter space that SBN can exclude at the 3 (dark green) and 5 (light green) σ C.L for ν_e appearance only (dotted), ν_μ disappearance only (dashed) and a combined appearance and disappearance fit (solid), as a function of POT.

3.6.4 ν_e disappearance effects at SBN

As this is the first time that SBN's sensitivity to ν_e disappearance has been demonstrated, we find it interesting to consider explicitly the effect of ignoring ν_e disappearance effects in the measured ν_e CC spectra, when performing combined ν_e appearance and ν_μ disappearance fits. As such we additionally show, in Figure (3.14), the SBN coverage under the (3+1), (3+2), and (3+3) scenarios in a combined ν_e appearance and ν_μ disappearance only search (dot-dashed line). By comparing this to the scenario in which the ν_e background is allowed to oscillate away, it is evident that performing an SBN search for sterile neutrino oscillations without the explicit assumption of negligible disappearance of intrinsic ν_e backgrounds has a significant effect on SBN's sensitivity, and warrants careful consideration of systematic correlations among exclusive samples measurable at SBN.

3.7 CP violating phases at SBN

The addition of CP violating phases in the (3+2) and (3+3) sterile neutrino scenarios introduces the potential of new probable phenomena at SBN. Although there is currently no planned antineutrino running for SBN, when one begins to consider the possibility of the existence of sterile neutrino CP violating phases it is natural to ask the question if SBN's sensitivity coverage improves with the inclusion of a combination of neutrino and antineutrino running. In particular does SBN's ability to rule out the low-energy anomaly if one fails to observe a signal increase with the addition of antineutrino running? Similarly if a potential signal is indeed observed, would antineutrino running allow for more precise measurements of these new neutrino mass splittings and mixing and any CP violating phases associated with the N additional states.

3.7.1 Antineutrino exposure in the absence of a signal

To initially investigate the impact of antineutrino running at SBN, we expand the fit as described in Section 4.3 to include observable ν_e CC and ν_μ CC spectra at the three SBN detectors in antineutrino running mode, as well as neutrino mode. The same background definitions are considered as in neutrino mode, and the backgrounds are re-evaluated assuming no right- or wrong-sign discrimination. Cosmics and dirt background contributions are taken to be identical to the neutrino running mode samples (scaled only according to POT). Statistical and systematic uncertainties on the antineutrino ν_e CC and ν_μ CC spectra are also considered, and the covariance matrix is expanded to include both those as well as correlations between neutrino and antineutrino spectra.

First, coverage is evaluated for a variety of beam POT exposures assumed for each running mode. Figure (3.18) shows the exposure in POT for both neutrino and antineutrino running mode that the SBN program requires in order to probe the 99% C.L globally allowed regions at 3σ (solid curves) and 5σ (dashed curves) for

the (3+1) scenario, at a percentage coverage as indicated explicitly on each curve. We focus on the strongest excluding case as shown in Sec (3.6.1), corresponding to combined ν_e dis/appearance and ν_μ disappearance fit, as such a combined search at SBN has significantly better coverage than appearance-only search for the same POT. We highlight that it is far more efficient to cover a given fraction of parameter space with neutrino-only rather than antineutrino-only or any combination of neutrino plus antineutrino running. This is evident from these figures as no point on any curve deviates from the origin by a distance smaller than the curve's x-coordinate for $y = 0$. This is not unexpected as neutrino and antineutrino oscillation probabilities under the two-neutrino oscillation approximation we've employed are identical by construction in the (3+1) scenario. Therefore, antineutrino running offers no additional information, and is generally less efficient due to the lower flux and cross-section, and, hence, event statistics.

Figures (3.19) and (3.20) show the same information for the (3+2) and (3+3) scenarios, respectively. Interestingly, just as in the (3+1) case, we observe again that it is far more efficient to cover any given fraction of parameter space with neutrino-only rather than antineutrino-only or any combination of neutrino plus antineutrino running. At first this may seem counter-intuitive, as it may be expected that antineutrino running would provide visibly more coverage due to enhanced sensitivity to CP-violating phases in these scenarios. However, the increased statistics per POT that are available in neutrino mode running are far more efficient in constraining all other mixing parameters and masses allowed in each oscillation hypothesis. Since these plots quantify overall coverage of the n-dimensional phase-space in each scenario, it is quite reasonable (and arguably expected) that antineutrino running proves less effective in terms of this metric.

In the absence of a possible signal, additional POT in antineutrino mode does *not* help to rule out the null hypothesis faster. Improvements due to antineutrino running tends to suffer from the significant wrong sign neutrino contribution inherent in the beam. Further studies into methods to mitigate this by differentiating

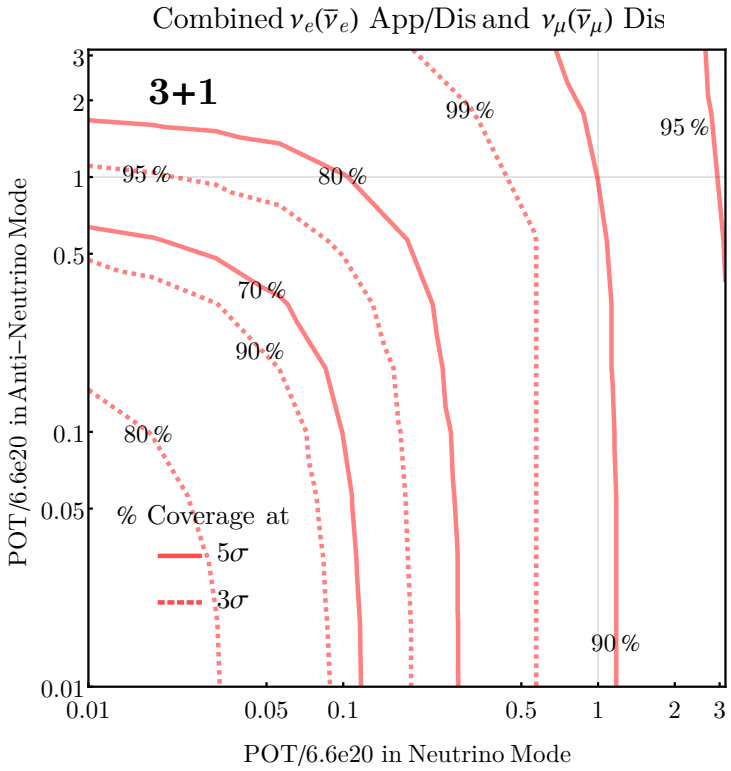

Figure 3.18: The amount of POT required in neutrino and antineutrino running modes for SBN to cover a given percentage of the 99% C.L globally allowed regions at 3σ (dashed curves) and 5σ (solid curves) in the (3+1) light sterile neutrino scenario. This corresponds to a combined ν_e dis/appearance and ν_μ disappearance search. Note that, as MicroBooNE will have already collected 6.6e20 POT in neutrino mode before SBN begins its run, the x-axis refers to additional POT beyond this 6.6E20 POT collected for MicroBooNE-only neutrino mode running.

between neutrino and anti-neutrino events, such as μ^- absorption rates on Argon or exploratory studies into the difference in Q^2 distributions from ν and $\overline{\nu}$ driven scatterings, would be especially useful at this crucial time in LArTPC development.

If, on the other hand, SBN *does* observe a sterile neutrino-like signal the focus would quickly turn to the subsequent measurement of the new parameters. Here the impact of SBN antineutrino running is more complicated, giving access to a distinctly different oscillation probability than purely neutrino running mode would allow. In the (3+2) and (3+3) scenarios the inclusion of CP violating phases could

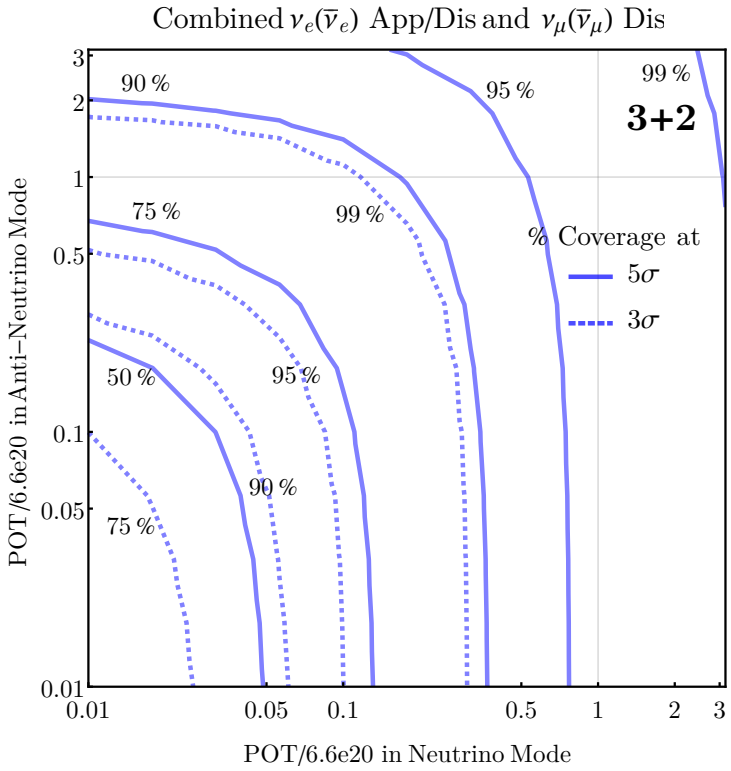

Figure 3.19: The same as Figure (3.18) but for the (3+2) light sterile neutrino scenario.

convolute the measurement of these parameters, as with a finite neutrino energy resolution, they become degenerate with many of the neutrino mixing parameters. In what follows we focus solely on the (3+2) scenario with a single CP violating phase ϕ_{54}, for simplicity, but note that these metrics can be applied to the (3+3) scenario with minimal expansion.

3.7.2 Sensitivity to ϕ_{54}

The sensitivity of SBN to the CP violating phases is studied under the hypothesis that SBN observes a signal consistent with multiple light sterile neutrinos. To analyse this sensitivity we inject potential signals, for a given set of oscillation parameters into the fit. These injected parameters are labelled as "true" parameters, and the spectra produced when one assumes these parameters take the place of the background only spectra in the χ^2 and covariance matrix as described in Section

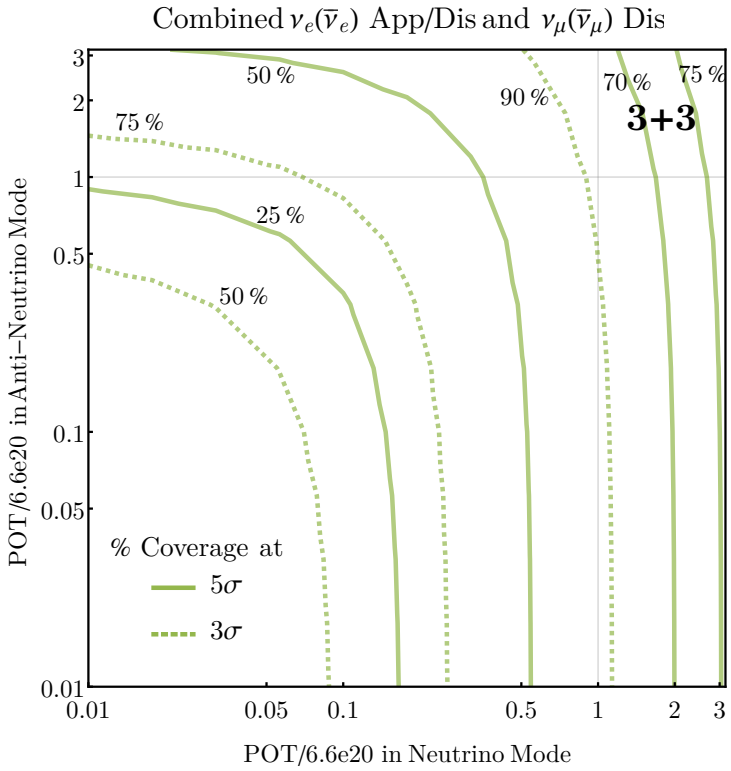

Figure 3.20: The same as Figure (3.18) but for the (3+3) light sterile neutrino scenario.

(4.3). This χ^2 will then allow us to gauge the potential sensitivity of SBN to sterile-active mixing parameters given a hypothetical signal. Unlike the case of ruling out the null hypothesis, in which ν_μ-disappearance is a driving force, sensitivity to ϕ_{54} is due solely to the ν_e-appearance channel in which it uniquely appears. Due to this as well as the large number of degree's of freedoms in (3+2) and (3+3) sterile neutrino scenarios we make here the simplifying assumption $U^2_{e4}U^2_{\mu 4} = U^2_{e5}U^2_{\mu 5}$ and analyse under the assumption of ν_e appearance only, so as to better understand and convey the behaviour in 2D of the main parameter of interest, ϕ_{54}. Although allowing all parameters to vary uniquely does indeed change the quantitative results, the qualitative phenomenology remains consistent.

In Figure (3.21) we show a sample scenario in which we inject a true ϕ_{54} of $3\pi/2$, for values of mass splittings closest to the global best fit that we simulate on our grid, $\Delta m^2_{41} = 0.48$ eV^2 and $\Delta m^2_{51} = 0.83$ eV^2. We then vary the strength of

the active neutrino-sterile neutrino mixings, $U_{e4}^2 U_{\mu 4}^2$ and show the range of possible reconstructed ϕ_{54}, at a given confidence level, all the while profiling over remaining mixing elements.

For a sterile neutrino of $\Delta m_{41}^2 \approx 1$ eV^2, in order to explain the LSND anomaly, requires mixings of order $U_{e4}^2 U_{\mu 4}^2 \approx 10^{-4} \rightarrow 2 \times 10^{-3}$. We note that ϕ_{54} resolution in this region varies from no-sensitivity to $\pm$ 40° at the 1σ level. Under the standard exposure of 6.6e20 POT in neutrino mode alone (red solid line) one can see there is no sensitivity for even the largest values of mixing parameters consistent with the (3+2) global data, $U_{e4}^2 U_{\mu 4}^2 \approx 2 \times 10^{-3}$. As such we concentrate on whether of not it is advantageous to run further in neutrino mode (red dashed line) or a combination of neutrino and antineutrino running mode (purple shaded regions). As can be seen, for unrealistically large mixings, SBN can strongly pick out the true ϕ_{54}, but as the mixings drop the resolution on ϕ_{54} reduces until one reaches $U_{e4}^2 U_{\mu 4}^2 \approx 4 \times 10^{-4}$, by which all values of ϕ_{54} are indistinguishable. We also show the 2σ contour for the case in which we run entirely in neutrino-mode (red dashed lines) and note that for the majority of the parameter space, is worse than a combined neutrino and antineutrino exposure.

The exact sensitivity of ϕ_{54} depends not only on the magnitude of mixings, on the assumed mass splittings also. In Figure (3.22) we repeat the same analysis for $\phi_{54} = \pi/2$, $\Delta m_{41}^2 = 0.16$ eV^2 and $\Delta m_{51}^2 = 1.0$ eV^2. This point corresponds to the point with largest mixings allowed in our (3+2) global fit at the 99% C.L. The green shaded region assumes 6.6e20 POT in both neutrino and anti-neutrino running and shows sensitivity to ϕ_{54} for values of $U_{e4}^2 U_{\mu 4}^2$ as low as 10^{-4}. Again we see that running in 50:50 neutrino and antineutrino running mode, over pure neutrino mode (red lines), allows one to exclude the same regions at a higher significance level, going from $\leqslant 2\sigma$ to $\geqslant 3\sigma$ for a wide regions of parameter space.

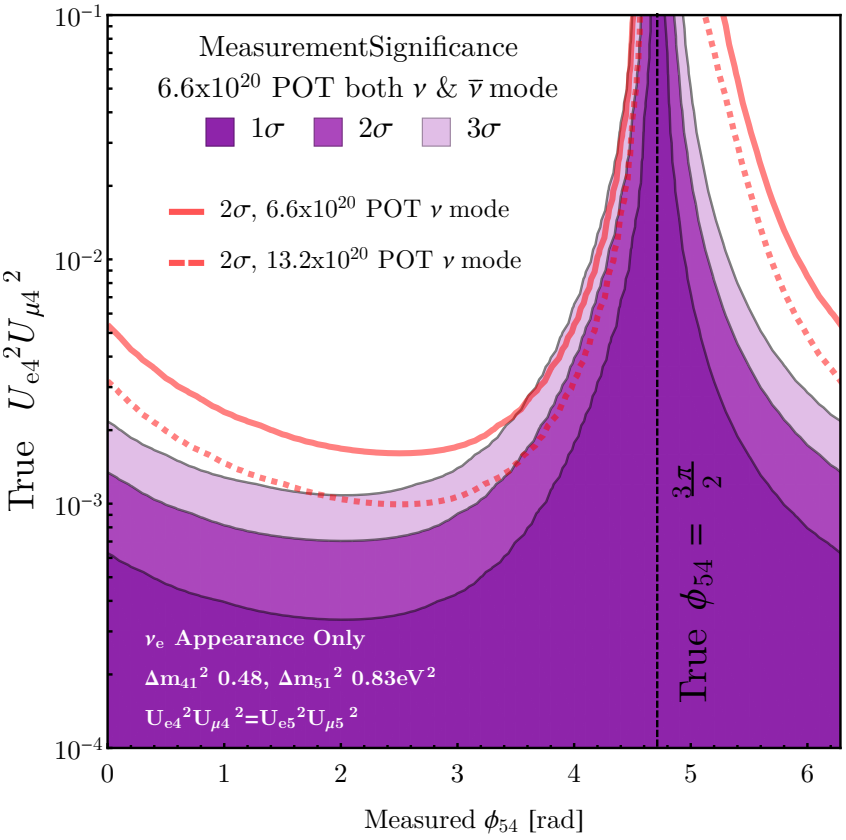

Figure 3.21: Sensitivity of SBN to a (3+2) scenario sterile neutrino signal, as a function of true mixing $U_{e4}^2 U_{\mu 4}^2$ for $\phi_{54}^{\text{true}} = 3\pi/2$. We show the regions of reconstructed ϕ_{54} that is in agreement at 1,2 and 3σ significance in purple shaded regions for a combined 6.6e20 POT neutrino running mode and 6.6e20 POT antineutrino running mode. In dashed red we also show the equivalent 2σ contour for 13.2e20 POT neutrino running mode only. The mass splittings correspond to the global (3+2) best fit point. As the true mixings are fixed in each test case, the contours are drawn at $\Delta\chi^2$'s of 1, 4 and 9, corresponding to the 1 remaining *dof*, ϕ_{54}, after profiling over all other parameters.

3.7.3 Prospects for CP violation searches

A related measurement to that of determining the value of ϕ_{54} given an observed signal, is the significance with which SBN could potentially rule out CP conserving values of ϕ_{54}, 0 and π. Definitive CP violation in the sterile neutrino sector would be a crucially important discovery and of potential worry to future experiments looking to measure the standard 3ν phase δ_{CP} [165]. To estimate this, for a given injected

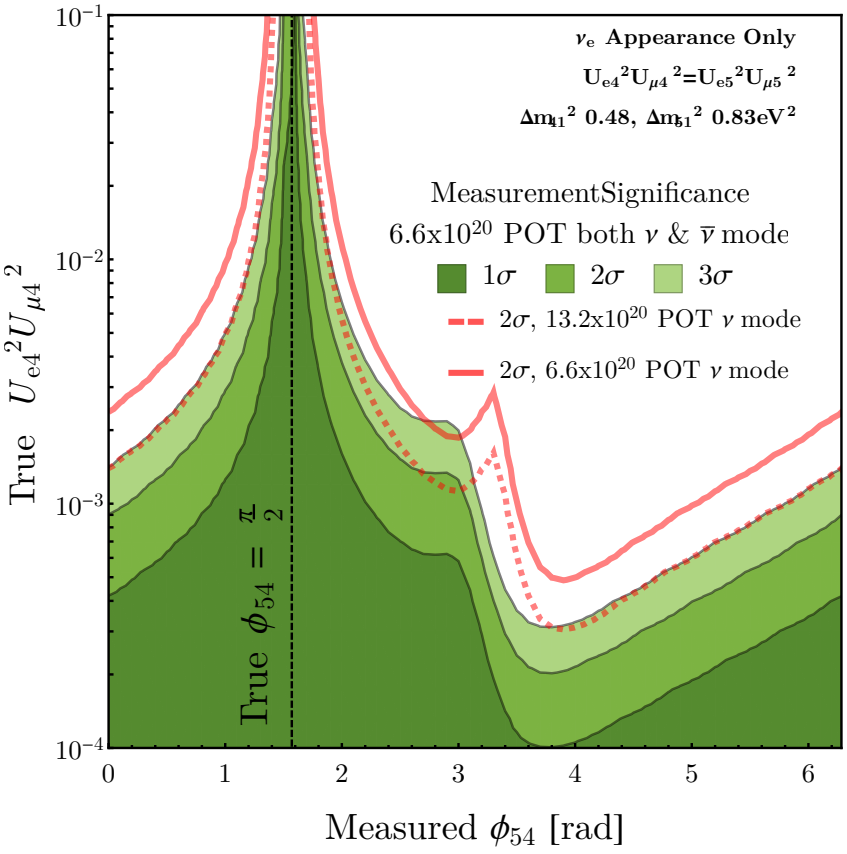

Figure 3.22: Same as Figure (3.21) but for injected $\phi_{54}^{\text{true}} = \pi/2$ and mass splittings corresponding to the largest mixing allowed by current global (3+2) best fit point. See text for more details.

signal with $\phi_{54} = \phi_{\text{true}}$ and fixed values of Δm^2's and mixings $U_{\alpha i}$'s, we form the metric

$$\chi^2_{\text{CP}}(\phi_{\text{true}}) \equiv \text{Min}[\chi^2(\phi_{54} = 0|\phi_{\text{true}}), \chi^2(\phi_{54} = \pi|\phi_{\text{true}})].$$

In each component χ^2 all active-sterile neutrino mixing elements are then varied in order to find the set which minimises the χ^2 under consideration, to account for possible degeneracies in the observed spectra. To get as realistic a measurement as possible we relax the simplifying constraint that $U_{e4}^2 U_{\mu 4}^2 = U_{e5}^2 U_{\mu 5}^2$ and allow all parameters to vary, fitting to a combined ν_e appearance and ν_μ and ν_e disappearance.

We show in Figure (3.23) is results of such an analysis for the same two possible injected signals, the global (3+2) best fit point (red lines) as well as the aforementioned maximum allowed mixing point (blue lines). For smaller values of mixings,

corresponding to the best fit point, little or no spectral shifts can be measured due to varying ϕ_{54}, and as such even for maximally violating CP angles, ϕ_{54} can always be mis-reconstructed to one of the CP conserving value, with shifts in U's to compensate for the rate. The standard 6.6e20 POT in neutrino mode is shown by the solid line and shows no sensitivity to CP violation, similarly if we assume an additional 6.6e20 POT in neutrino mode, the situation does not change (dotted line) significantly. Although the inclusion of 6.6e20 POT in antineutrino mode (dashed line) does double the potential sensitivity, this remains a sub 1σ effect and thus it is clear that within reasonable exposure SBN is completely insensitive to CP violation if nature does choose sterile neutrinos at this mass splitting.

As the strength of mixing increases, individual variations in the energy spectrum due to ϕ_{54} driven oscillations becomes harder for degeneracies in mixing to explain and the significance at which certain CP violating phases are in disagreement with $\phi_{54} = 0$ or π increases. This is evident when we look at the CP violation curves assuming the "maximum allowed mixing" sterile neutrino parameters. If we again assume a standard exposure of 6.6e20 POT in neutrino mode (solid blue line) it is evident that SBN has no sensitivity to CP violation, with significance's of less than 1σ even with maximum CP violation. Doubling the POT in neutrino mode (dotted blue line) gives an effectively negligible increase, but it is here that the benefits of including antineutrino running over purely additional neutrino mode POT is most evident. An additional 6.6e20 POT in antineutrino mode allows for 2σ significance at maximal mixing, with $> 1\sigma$ significance over 68% of ϕ_{54} parameter space. Although certainly not enough to claim discovery, SBN could provide the first hints of CP violation in the sterile neutrino sector.

It is worth clarifying that although if nature is kind enough to choose a maximally CP violating phases, $\phi_{54} = \pi/2$ or $3\pi/2$, SBN could indeed potentially observe CP violation at the $\approx 2\sigma$ significance level, this requires large mixings already somewhat in tension with global data $U^2_{\mu 5} \approx 0.0038$, and is not true for all sterile neutrino mass splittings. For non-maximally violating CP phases the significance at which SBN

can make statements diminishes rapidly, and for the majority of the parameter space motivated by the low-energy anomalies, the potential for SBN to measure a CP violating phase to the accuracy necessary to rule out CP conservation is very low and insignificant. Conversely, for values of active-sterile neutrino mixings and Δm^2 splittings outside that of those considered here, namely ones which help less to explain the low-energy anomalies but could be interesting models none the less, the sensitivity to CP violation could be significantly greater than those presented here.

3.8 Summary of SBN capability

In this chapter we have considered SBN's sensitivity to extended light sterile neutrino oscillation scenarios. We find that, in the case of a (3+1) oscillation scenario, SBN is capable of definitively exploring (i.e. with 5 σ coverage) 85% of the 99% C.L parameter space region which is allowed by global short-baseline oscillation data (for 3 dof). This is possible after a three year neutrino mode run with all three SBN detectors, running concurrently to collect data corresponding to 6.6×10^{20} POT, and with a combined ν_e dis/appearance and ν_μ disappearance search. Furthermore, by performing such a combined search, MicroBooNE alone, during its first three years of running prior to the SBN program commencing, will be able to test 25% of the globally allowed (3+1) oscillation parameter space at 3σ.

In the case of a (3+2) scenario, in its three year neutrino run, SBN can definitively explore 95% of the 99% C.L allowed parameter space (7 dof). In this scenario, a single CP-violating phase, ϕ_{54}, enters in the $\nu_\mu \rightarrow \nu_e$ appearance probability and leads to differences in neutrino and antineutrino appearance probabilities. Dedicated BNB antineutrino mode running for three years (6.6×10^{20} POT), beyond the currently planned neutrino mode running, does not significantly expand SBN's 5 σ sensitivity coverage. Overall, by performing a multi-baseline and multi-channel oscillation search with sign-selected neutrino and antineutrino beams, the SBN experiment will be able to, within six years of operation, overconstrain a significant

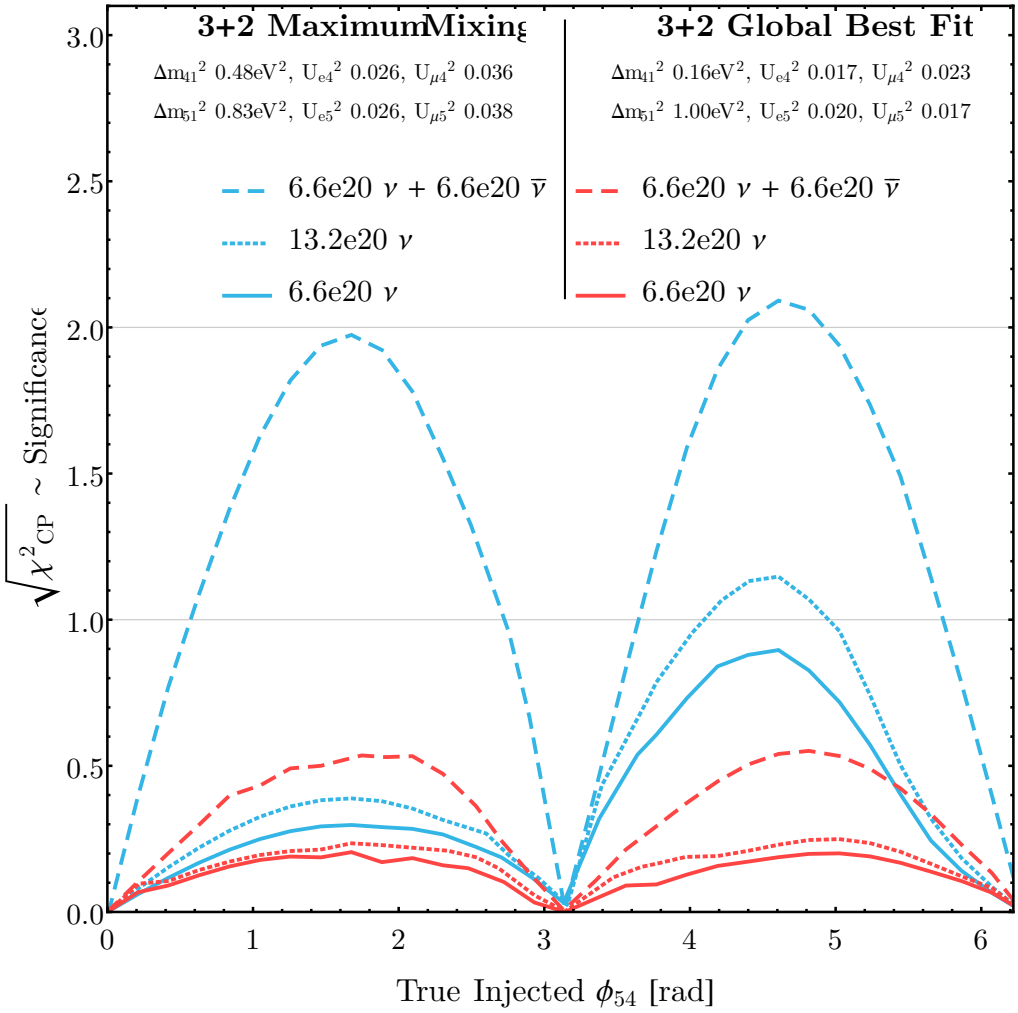

Figure 3.23: Significance at which SBN can observe CP violation in the (3+2) sterile neutrino scenario, as a function of true ϕ_{54}, for two injected signals corresponding to the global (3+2) best fit point (red lines) as well as the parameter point with largest total mixings (blue lines), for a variety of POT in neutrino and antineutrino running modes. Unlike previous plots we make no assumption on mixing and fit to ν_e appearance and both ν_e and ν_μ disappearance simultaneously, profiling over $U_{e4}, U_{\mu4}, U_{e5}$ and $U_{\mu5}$. As all remaining parameters are profiled over, and only 1 *dof* remains, the $\sqrt{\chi^2}$ will approximate the significance of the measurement.

fraction of parameter space which is currently allowed by global fits to sterile neutrino oscillation.

Although the addition of antineutrino running mode POT does not aid in the

exclusion of the null hypothesis in the absence of an oscillation signal, in the case where a potential signal consistent with multiple light sterile-neutrinos occurs, dedicated antineutrino running at SBN proves to be of substantial value in increasing the significance of an observation of CP violation. For the (3+2) sterile neutrino scenario, an additional 6.6e20 POT in antineutrino running mode could allow SBN to provide the first 2σ hints of CP violation in the lepton sector, provided nature chooses maximal CP violating phases $\phi_{54} = \pi/2$ or $3\pi/4$, and mixing elements consistent with global data at the 99% C.L, $\Delta m^2_{41} = 0.16$ eV^2, $\Delta m^2_{51} = 1.0$ eV^2, $U^2_{e4} = U^2_{e5} = 0.026, U^2_{\mu 4} = 0.036$ and $U^2_{\mu 5} = 0.0038$. For SBN to be able to observe CP violation at a greater significance than this would require active-sterile mixing already in significant tension with global data.

Finally, we must point out a caveat in these studies, in that the data sets used to constrain the (3+N) oscillation parameter suffer from large apparent incompatibility in the parameter space they seem to prefer. We argue that this incompatibility leads to overestimated in size global allowed regions, and it is arguably a more conservative approach to consider these larger allowed regions provided by these fits in exploring SBN's discovery reach.

Chapter 4

Heavy Sterile Neutrinos, $\mathcal{O}$(MeV)

"Don't worry, dear Pamela,
I'll do my scientific best *to command your fleet."*
Diana

4.1 Introduction

As we have discussed at length in Chapter 1, the neutrino sector of the Standard Model is known to be incomplete. The observation of oscillatory behaviour between neutrino flavour states [33] suggests that neutrinos possess a mass matrix with off-diagonal terms in the flavour basis. There are many models that have been invoked in the literature to explain this observation as well as the lightness of neutrino masses, ranging from the ever popular See-Saw mechanisms [135–137] to radiative mass generation [183, 184] or even more involved constructions such as neutrino masses originating from extra-dimensions [185]. It will ultimately be the role of phenomenology to find ways to distinguish between potential candidate models, and explore what can be deduced about the completion of the neutrino sector from the analysis of contemporary experiments. In Chapter 2 we introduced a common, although not necessary, feature in many Beyond the Standard Model models which successfully account for neutrino masses. The presence of sterile neutrinos. We have already discussed their role in the generation of neutrino masses, and the introduction of a non-unitarity to the 3×3 U_{PMNS} matrix, but in this Chapter we will delve further into the possible interactions of these sterile neutrinos, such as their production and decay to Standard Model particles through mixing-suppressed

gauge interactions.

In the preceding chapter we have discussed at length one of the best known examples, the short-baseline oscillation signature associated with a sterile neutrino mass around the eV-scale. We have shown that the Fermilab SBN program will be able to extend the current bounds on light oscillating sterile neutrinos, thoroughly exploring the eV-scale sterile neutrino mass region, whilst also pursuing many other physics goals [108]. In this Chapter, we assess SBN's potential to contribute to the search for much heavier sterile neutrinos, in a manner complementary to the oscillatory analysis. The new fermions in our study are assumed to have masses around the MeV scale. These particles are light enough to be produced in neutrino beams via meson decay, but have masses sufficiently large to prevent oscillatory effects with the active neutrinos through loss of coherence (see *e.g.* Ref. [186]), instead propagating long distances along the beamline. Due to the presence of mixing they are unstable, and their subsequent decay products can be observed in neutrino detectors. We stress that the search for MeV-scale sterile neutrinos is entirely compatible with the primary goals of SBN, and requires modification of neither the beam nor detector designs.

The same properties of LArTPC detectors that allow for such precision in the (3+N) sterile neutrino searches apply here too. The reconstruction [187, 188], energy resolution [189] and excellent calorimetric particle identification capabilities of LAr [190] technology means the SBN program provides an ideal scenario to study this "decay-in-flight" of sterile neutrinos. This technology allows for a high degree of background suppression on well studied decay modes while also allowing the study of channels which have been poorly bounded by similar experiments due to large backgrounds and challenging signals. For example, the differentiation between an electron- or photon- induced EM shower can be achieved by studying their rate of energy loss in the first 3 cm of their ionising track [163]. Furthermore, as we discuss in Section (4.3.4), if a sufficiently good timing resolution of scintillation light is achieved, the timing structure of markedly sub-luminal sterile neutrinos can be

	PS-191	SBND	MicroBooNE	ICARUS
POT	0.86×10^{19}	6.6×10^{20}	13.2×10^{20}	6.6×10^{20}
Volume	216m^3	80m^3	62m^3	340m^3
Baseline^{-2}	$(128\text{m})^{-2}$	$(110\text{m})^{-2}$	$(470\text{m})^{-2}$	$(600\text{m})^{-2}$
Ratio/PS-191	-	38.5	3.3	5.5
S/$\sqrt{B}$ Ratio	-	16.3	1.8	1.1

Table 4.1: A comparison of the relative exposure at each SBN detector compared to PS-191, the experiment which resulted in the current best bounds on heavy sterile neutrinos. One would expect all three SBN detectors to see increased numbers of events than PS-191 did, with SBND seeing the largest enhancement of a factor of 38.5. The final row takes into account the scaling in masses leading to increased backgrounds, although the achievable reconstruction of LAr should reduce these significantly. Despite being a smaller detector, the POT and baseline of SBND results in significantly larger rate of expected events.

utilised as both a rejection mechanism for beam related backgrounds as well as a further aid for model discrimination and mass measurement.

We restrict our analysis to sterile neutrino masses below the kaon mass. Kaons and pions are produced in large numbers at BNB, and their subsequent decays will generate a flux of sterile neutrinos. In this mass range, the strongest bounds on sterile neutrinos which mix with electron and muon neutrinos come from PS-191 [191, 192], a beam dump experiment which ran at CERN in 1984. PS-191 was constructed from a helium filled flash chamber decay region, followed by interleaved iron plates and EM calorimeters. It was located 128 m downstream of a beryllium target and 2.3° (40 mrads) off-axis, obtained 0.86×10^{19} POT over the course of its run-time, and had a total detector volume of $6 \times 3 \times 12 = 216$ m^3. We can estimate the sensitivity of the three SBN detectors and how they will compare to PS-191 by estimating the experiments' *exposure*, defined here as POT $\times$ Vol $\times$ R^{-2}. We compare the three detectors to PS-191 in Table (4.1), which indicates that all detectors

of the SBN complex expect a larger exposure, with SBND seeing the greatest enhancement by a factor of around 40. In addition to the larger exposure, there is also an enhancement of the expected decay events at SBN due to its lower beam energy. The sterile neutrinos at SBN are produced by the 8 GeV BNB beam and have a softer spectrum than those produced by the 19.2 GeV CERN Proton Synchrotron beam used at PS-191. As we discuss in more detail in Section (4.2), the probability that the sterile neutrino decays inversely scales with momentum, $1/|P_N|$, and we would therefore expect any BNB detector to see more events than PS-191 even if considering equivalent neutrino exposures.

However, exposure alone does not dictate the resulting sensitivity. PS-191 was purposefully built to search for such decays of heavy fermions. To minimise the background induced by active neutrino scattering, the total mass of the detector (and therefore number of target nuclei) was chosen to be small (approximately 20 ton). Conversely, the SBN detectors were designed to search for neutrino interactions and thus have significantly larger masses (112, 66.6 and 476 tons respectively). SBN will not only see a greater number of decay events than PS-191 but also a greater background for a given exposure. Therefore, the degree of background reduction will be crucial in determining its ultimate performance. We return to this issue in Section (4.3.2).

This Chapter is structured as follows. In Section (4.2) we present an overview of sterile neutrino decay in minimal and non-minimal models relevant for beam dump experiments. We then present the details of our simulation in Section (4.3) and show illustrative event spectra for some channels of interest. In Section (4.4), we present and discuss the exclusion contours that SBN could place on the model in the absence of a signal. We then study how the event timing information could be used to test the hypothesis of sterile neutrino decay-in-flight and to help constrain the particle masses if a positive signal were detected. We make some concluding remarks in Section (4.5).

4.2 Sterile neutrino production and decay

The most general renormalizable lagrangian extending the Standard Model to include a new gauge-singlet fermion N is given by

$$\mathcal{L}_N = \mathcal{L}_{\text{SM}} + \overline{N} i\partial\!\!\!/ N + \left(\frac{m_N}{2}\overline{N^c}N + y_\alpha \overline{L}_\alpha H N + H.C\right), \tag{4.2.1}$$

where N represents a massive Majorana sterile neutrino with mass m_N, as described in Chapter 1. The extension to multiple new fermions involves promoting y and m_N to matrices with indices for the new states, but will offer no real phenomenological difference in the following analysis.[1] Much work has been done understanding the phenomenology of such novel neutral states, which varies significantly over their large parameter spaces. Lagrangians similar to this have been used in the literature for a wide range of purposes. If the new particle has a mass around 10^{12}-10^{15} GeV it could provide a natural way to suppress the size of active neutrino masses through the Type I or III see-saw mechanisms [135–137]. A lighter neutral fermion, with a mass around the keV scale, remains a promising dark matter candidate [127]. A synthesis of these ideas is found in the so-called νMSM which simultaneously can explain dark matter, neutrino masses and successful baryogenesis [193]. If we consider sterile neutrinos at even lower energy scales, with masses at the eV scale or below, these particles can alter the neutrino oscillation probability, leaving observable signatures at oscillation experiments. Indeed, such particles have been proposed to alleviate short-baseline oscillation anomalies; although, no minimal solution seems to provide a compelling universal improvement to the current data [75, 76].

A key feature of models of sterile neutrinos are the weaker-than-weak interactions which arise from mass mixing. In the minimal lagrangian in Equation (4.2.1), the only direct couplings to new sterile flavour eigenstates are neutrino–Higgs interactions. However, these couplings generate off-diagonal neutrino bilinears below the electroweak symmetry breaking scale, leading to mixing-mediated interactions with Standard Model gauge bosons for the mostly neutral mass eigenstate. This allows

[1] This minimal single N extension does not allow for the observed masses of the neutrinos, as the mass matrix is rank 1. We assume that an appropriate extension has been introduced to satisfy neutrino oscillation data while introducing no new dynamics at the lower energy scales of interest.

them to be produced in and decay via Standard Model gauge interactions, albeit suppressed by the mixing angle.

The possibility remains that extra particles exist beyond the minimal lagrangian and these mediate other interactions, either directly with Standard Model fields or, as before, via mixing. Throughout our work, we assume that the production of N, described in Section (4.2.1), is generated by the interactions in Equation (4.2.1). However, we will return to the idea of a non-minimal lagrangian in Section (4.2.2) when discussing the decay modes of N.

4.2.1 Production at BNB

For sterile neutrinos which are light enough to be produced from a meson beam, there is a qualitative divide in the phenomenology somewhere between keV and eV masses[2]. If the sterile neutrinos are massive enough for their mass-splittings with the light neutrinos to be larger than the wavepacket energy-uncertainty associated with the production mechanism, they no longer oscillate, as we showed in Equation (1.4.41) [186]. Neutral particles produced in the beam will propagate towards the detector and may be observed by their subsequent decay into Standard Model particles. Experiments seeking to measure such decays are generally known as beam dump experiments, where proton collisions with a target produce particles to be observed down-wind of the source [191, 192, 194–198]. It has been pointed out that the difference between a beam dump and a conventional neutrino beam is more a matter of philosophy than design, and we can expect many experiments to have some sensitivity to novel heavy states [109, 199, 200].

For the BNB, we can estimate the mass at which the oscillatory behaviour is suppressed as follows: the decay pipe for BNB is around 50 m in length, which is considerably shorter than the decay lengths of the mesons in the beam, and we assume that this length defines the wavepacket width at production. The relevant

[2]The precise mass range depends on details of the process under consideration.

parameter is the localisation parameter $\xi = 2\pi\frac{\lambda_{\rm d}}{\lambda_\nu}$, where $\lambda_{\rm d} = 50$ m and λ_ν is the standard neutrino oscillation length $\lambda_\nu = \Delta m^2/4E_\nu$. For $\xi \gg 1$ the wave packet is insufficiently broad to accommodate a coherent superposition of the heavy and light neutrino states. We estimate that this occurs for the BNB at $\Delta m^2 \gtrsim 100$ eV^2.

In a conventional neutrino beam, most neutrinos are derived from meson decay, and we assume in this work that the sterile neutrinos are produced from the decays of pions and kaons, and subsequent muons , restricting our sterile neutrino mass to $m_N \leqslant m_K$. Larger sterile neutrino masses could be probed by working at higher energies in the initial proton beam, where the neutral fermions could come from decays of charmed mesons such as $D^\pm$. This strategy will be used by the upcoming SHiP experiment [197,198] but will not be considered further in the present work as D mesons are produced in extremely small numbers due to the relatively low energy of the BNB beam [145]. As such we restrict ourselves to the naturally defined mass range of interest for SBN, eV $\ll m_N \lesssim 494$ MeV. We focus on $m_N \gtrsim$ MeV scale states where the prospects for detection are greatest due to enhanced decay rates.

Although novel dynamics may lead to enhanced production rates of sterile neutrinos by alternative unconventional means, we neglect this possibility and assume that the sterile component of the BNB flux arises solely from meson (or secondary $\mu^\pm$) decays. This process requires only mass-mixing from the N-ν Yukawa term in Equation (4.2.1). It follows that the amplitudes for these decays are related to those of the standard leptonic decays of mesons via an insertion of the mixing matrix element $U_{\alpha 4}$, and to leading order in the mass of the sterile neutrino over the meson mass, the N-fluxes will be a rescaling of the fluxes for the active neutrinos. However, in order to account for flavour-specific effects, it is necessary to go beyond this approximation and consider the kinematic differences of heavy sterile neutrino production. The flux of sterile neutrinos produced from the decay of a given meson M is approximated by

$$\phi_N(E_N) \approx \phi_{\nu_\alpha}(E_{\nu_\alpha})|U_{\alpha 4}|^2 \frac{\rho\left(\delta_M^a, \delta_M^i\right)}{\delta_M^a\left(1-\delta_M^a\right)^2}, \tag{4.2.2}$$

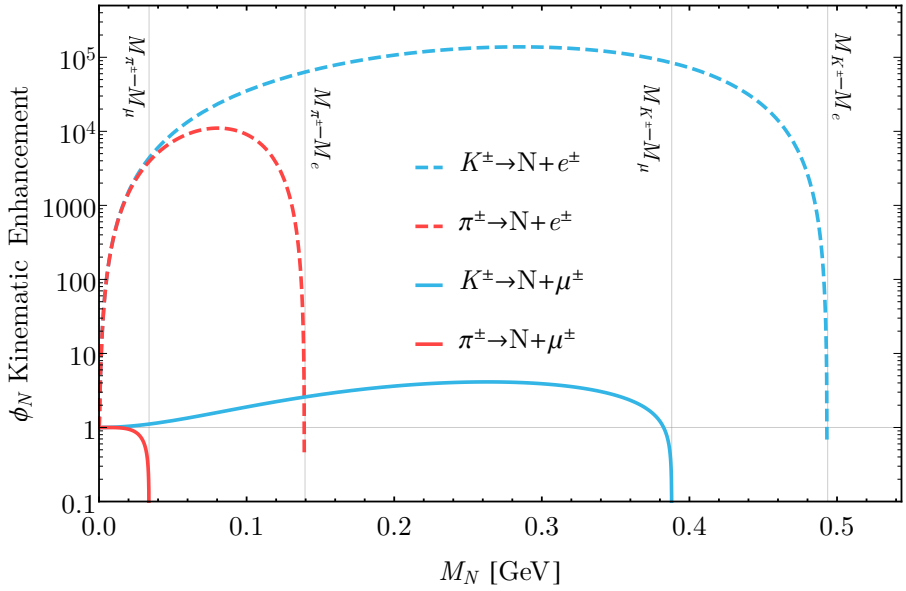

Figure 4.1: Total kinematic enhancement of flux from meson decay when decaying to massive sterile neutrinos. Red shows decays originating from parent pions, blue from parent kaons. Solid lines represent $U_{\mu 4}$ dominant mixing producing a muon alongside the sterile neutrino, where as dashed U_{e4} dominant mixing associated with electron production. The huge enhancement due to helicity un-suppression of $\pi^\pm \to e^\pm \nu$ can clearly be seen for both parent pion and kaon.

where $\rho(a,b) = \mathcal{F}_M(a,b)\lambda^{\frac{1}{2}}(1,a,b)$ is a kinematic factor consisting of a term proportional to the two body phase space factor, $\lambda(x,y,z) = x^2+y^2+z^2-2(xy+yz+xz)$ and a term proportional to the matrix element, $\mathcal{F}_M(a,b) = a+b-(a-b)^2$, with $\delta_M^{a(i)} = m^2_{l_a(\nu_i)}/M^2$ [201]. We plot this enhancement for sterile neutrinos of mass $\leqslant M_{K^\pm}$, from parent pion and kaon decays in in Figure (4.1).

The kinematic factor leads to two effects. First, it provides a threshold effect of suppressing the production when the phase space decreases near a kinematic boundary. Secondly, it allows for the helicity un-suppression of channels which in a conventional beam are highly suppressed. For example, the decay $\pi^\pm \to e^\pm \nu_e$ which is significantly suppressed compared to the muonic channel, sees no such suppression when the neutrino is replaced with N. This kinematic effect for the pion and kaon can be very substantial. For $\pi \to e\nu$ this factor can be as large as 10^5, which more than compensates for the significantly smaller intrinsic flux of ν_e intrin-

sic to the BNB, which is around 0.52% of the total flux [145]. The approximation in Equation (4.2.2) starts to fail as the mass of the sterile neutrino increases, and we begin to see components of the active flux having energies less than the sterile mass which are truncated by the kinematic factor. In order to keep the normalisation of total neutrino events constant, before $U_{\alpha 4}$ and kinematic scaling, any events which are below the sterile neutrino mass threshold are removed and the remaining flux is scaled accordingly.

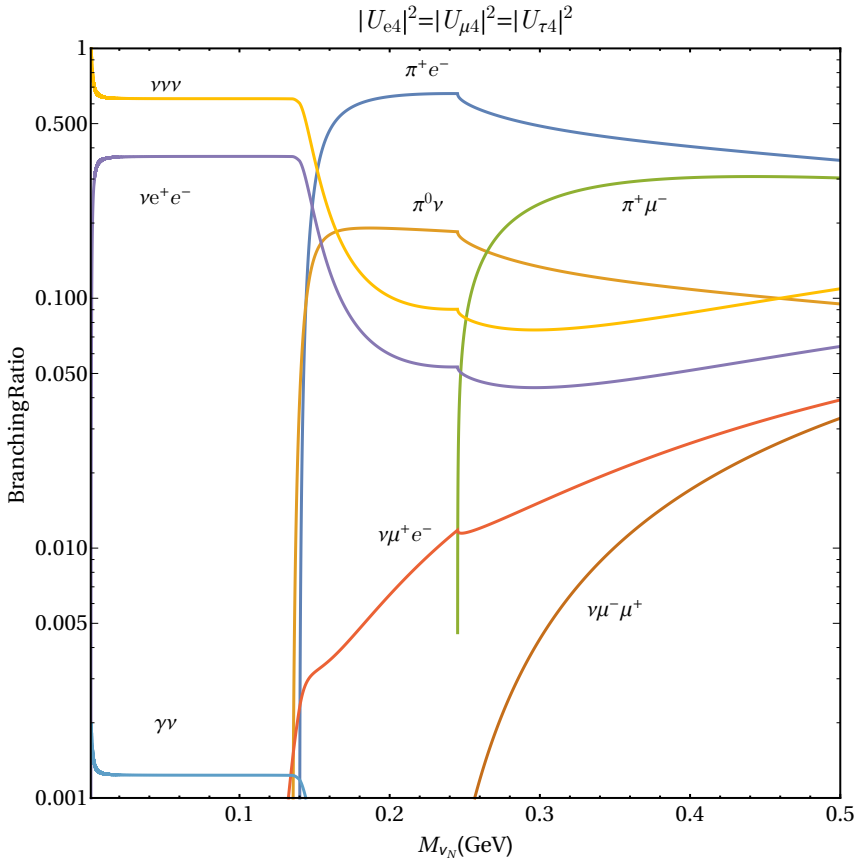

Figure 4.2: The branching ratios for heavy sterile neutrino decays in the minimal 3 sterile neutrino Standard Model extension, with masses between 1 and 500 MeV. We assume a flavour independent mixing pattern, with $U_{e4} = U_{\mu 4} = U_{\tau 4}$.

4.2.2 Decay at SBN

The fermions N will generally be unstable, albeit possibly long-lived, allowing for decays-in-flight into Standard Model particles. In what follows, we try to keep an open mind about the interactions of the sterile neutrino and consider all kinematically possible tree-level decays to visible Standard Model particles for sterile fermions produced by pion and kaon decays, 10 MeV $\lesssim m_N \lesssim m_K$. The precise decay rates and branching ratios for these channels are model dependent. In this section, we discuss the decay rates for a minimal extension of the Standard Model, as well as the implications of a non-minimal model.

Minimal model

We define the minimal sterile neutrino model by the Lagrangian in Equation (4.2.1). This encompasses the best known model of sterile neutrino phenomenology — the UV-complete Type I see-saw (and its low-scale variants) — but also provides an effective description of more complicated extensions of the Standard Model in which the additional field content does not directly affect the neutrino sector at low energies. Decays of sterile neutrinos in such a model proceed via Standard Model interactions suppressed by the mixing angle and have been studied in Ref. [201–203]. We have plotted the branching ratios for sterile neutrinos in our mass range in Figures (4.2) and (4.3), and we will now briefly summarise the decay rates most important for the present study.

The decays of the minimal model depend only on the mass of the N and the size of neutrino mixing to various flavours, parameterized by the elements of an extended PMNS matrix, *e.g.* for one additional particle $U_{\alpha 4}$ for $\alpha \in \{e, \mu, \tau\}$. The branching ratios for these decays are shown in Figures (4.2) and (4.3), as a function of mass for two sets of assumptions about the PMNS matrix. In Figure (4.2), we show the branching ratios if all mixing elements are of a similar size, whereas in Figure (4.3) we assume that only $U_{\mu 4}$ or U_{e4} are non-zero. This leads to certain semi-leptonic decays being forbidden, significantly changing the phenomenology of the model for some masses.

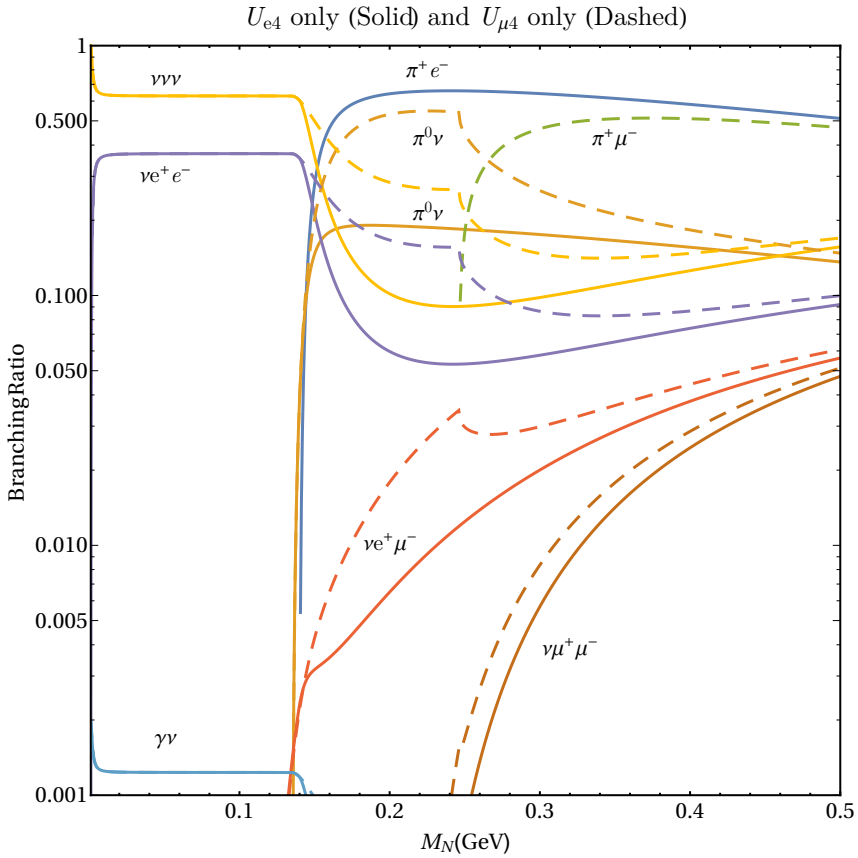

Figure 4.3: The same as Figure (4.2) but where we assume there is a hierarchical flavour scenario in which either U_{e4} (solid lines) or $U_{\mu 4}$ (dashed lines) is the dominant mixing-matrix element.

For sterile neutrino masses less than the pion mass, the dominant decay is into three light neutrinos. This channel is for all practical purposes unobservable and we will not consider it further. The dominant decay into *visible* particles will be into an electron-positron pair with a branching fraction of around 38%. This is true regardless of the flavour structure of the mixing matrix; although, this decay channel is not flavour-blind. If the sterile neutrino mixes primarily through U_{e4}, the decay proceeds via both neutral and charged currents, but for $U_{e4} = 0$, this channel occurs via neutral current only. We illustrate the Feynman diagrams of the studied decay rates in Figure (4.4).The decay rates we use follow the notation of Ref. [202]. We repeat them here in the interests of clarity. The decay rate for this channel is

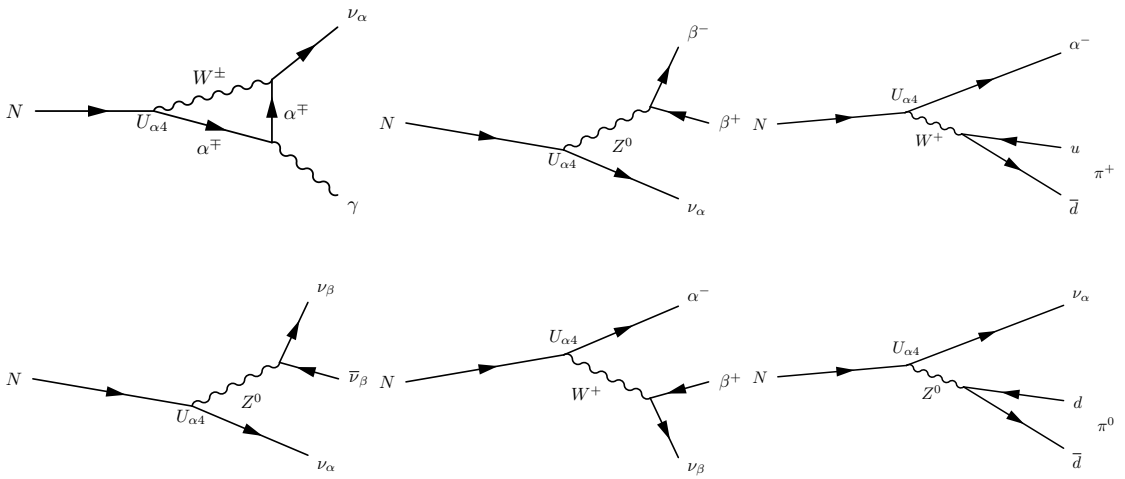

Figure 4.4: *Left Column*: On top we have the loop level photonic decay, with the emission of the photon off the lepton or charged W. On bottom we have the 3ν decay which, although a large branching fraction at low masses, is unobservable. *Middle Column*: The two charged leptonic decays, on top we have the neutral current componant, and on the bottom the charged current, depending on the exact flavour structure one of these may be forbidden. *Right Column*: On top we have the semi-leptonic decay involving a charged pion, while on bottom we have the neutral current driven π^0 production.

given by

$$\Gamma\left(N \to \nu_\alpha e^+ e^-\right) = \frac{G_{\rm F}^2 m_N^5}{96\pi^3} |U_{\alpha 4}|^2 \left[\left(g_L g_R + \delta_{\alpha e} g_R\right) I_1\left(0, \frac{m_e}{m_N}, \frac{m_e}{m_N}\right) + \left(g_L^2 + g_R^2 + \delta_{\alpha e}(1 + 2g_L)\right) I_2\left(0, \frac{m_e}{m_N}, \frac{m_e}{m_N}\right) \right],$$

where $g_L = -1/2 + \sin^2\theta_{\rm W}$, $g_R = \sin^2\theta_{\rm W}$. The two functions, $I_1(x,y,z)$ and $I_2(x,y,z)$ are integrals over phase space such that $I_1(0,0,0) = 1$ and $I_2(0,0,0) = 0$, and

$$I_1(x,y,z) = 12 \int_{(x+y)^2}^{(1-z)^2} \frac{ds}{s}(s - x^2 - y^2)(1 + z^2 - s)\sqrt{\lambda(s, x^2, y^2)}\sqrt{\lambda(1, s, z^2)},$$

$$I_2(x,y,z) = 24yz \int_{(y+z)^2}^{(1-x)^2} \frac{ds}{s}\left(1 + x^2 - s\right)\sqrt{\lambda\left(s, y^2, z^2\right)}\sqrt{\lambda\left(s, y^2, z^2\right)},$$

$$\lambda(a,b,c) = a^2 + b^2 + c^2 - 2ab - 2bc - 2ca.$$

Although the electron-positron channel dominates the visible decays at $m_N \leqslant m_\pi^0$, we also consider the radiative decay $N \to \nu_i \gamma$ which would generate an obser-

vationally challenging single photon signal [204]. In the minimal model the decay occurs via a charged-lepton/W loop, as shown in Figure (4.4), and has a rate given by

$$\Gamma(N \to \nu_i \gamma) = \frac{G_{\rm F}^2 m_N^5 |U_{\alpha 4}|^2}{192\pi^3} \left(\frac{27\alpha}{32\pi}\right).$$

This decay channel is significantly suppressed by the light mass of the sterile neutrino, the mixing-matrix elements and the loop factor. It can be estimated at around $\Gamma(N \to \nu_i \gamma)/(\text{GeV}) \approx 10^{-20}(m_N/\text{GeV})^5$. We see in Figure (4.2) that this leads to a branching ratio of around 10^{-3}.

Additional leptonic decays open up for sterile neutrino masses which satisfy $m_N \geqslant m_\mu + m_e$. Although with a smaller branching ratio, decays involving muons are clean signatures at LAr detectors. In the case of $N \to \nu_\alpha \mu^+ \mu^-$ the decay occurs by both neutral and charged currents and follows from the $N \to \nu_\alpha e^+ e^-$ decay given above with the replacement $m_e \to m_\mu$. The mixed-flavour decays, *e.g.* $N \to \nu_\alpha \mu^\pm e^\pm$, occur by charged current only and are given by

$$\Gamma(N \to \nu_\alpha \beta^- \alpha^+) = \frac{G_{\rm F}^2 m_N^5 |U_{\beta 4}|^2}{192\pi^3} I_1\left(\frac{m_\beta}{m_N}, \frac{m_\alpha}{m_N}, \frac{m_\alpha}{m_N}\right),$$

with $\{\alpha, \beta\} = \{e, \mu\}$. The next thresholds lie just above the pion mass, where two further decays become possible: $N \to \nu\pi^0$ and $N \to e^\pm \pi^\mp$. These processes quickly become the dominant decays at this mass range. The decay rate for the first process is given by

$$\Gamma\left(N \to \nu_i \pi^0\right) = \sum_\alpha \frac{G_{\rm F}^2 f_\pi^2 m_N^3 \left|U_{\alpha 4}\right|^2}{64\pi} \left[1 - \left(\frac{m_\pi}{m_N}\right)^2\right].$$

The decay into a charged pion and a lepton is an important channel, and one of the most constrained in direct decay experiments due to its clean two-body signal. Its decay rate has a similar algebraic form to the rate of $N \to \nu\pi^0$ with the addition of a CKM matrix element arising from the W-vertex,

$$\Gamma\left(N \to l^\pm \pi^\mp\right) = |U_{l4}|^2 \, \frac{G_{\rm F}^2 f_\pi^2 |V_{ud}|^2 m_N^3}{16\pi} I\left(\frac{m_l^2}{m_N^2}, \frac{m_\pi^2}{m_N^2}\right), \tag{4.2.3}$$

with

$$I(x, y) = \left[(1 + x + y)(1 + x) - 4x\right] \lambda^{\frac{1}{2}}(1, x, y).$$

For $N \to e^{\pm}\pi^{\mp}$ the kinematic function $I(x, y)$ produces only weak suppression ($I(x, y) \geqslant 0.5$) for sterile neutrino masses above $m_N \gtrsim 150$ MeV, whilst for $N \to \mu^{\pm}\pi^{\mp}$ the equivalent threshold is $m_N \gtrsim 270$ MeV.

If it is allowed by the flavour structure, the $N \to e^{\pm}\pi^{\mp}$ channel dominates the branching ratios for sterile neutrino masses which satisfy $m_{\pi^{\pm}} \lesssim m_N$. However, as it is mediated by a W boson, in the absence of U_{e4} mixing, this decay would be forbidden and the decay into a neutral π^0 and a light neutrino would be dominant. Once the mass of the sterile fermion is above $m_N \gtrsim 235$ MeV, the $\mu^{\pm}\pi^{\mp}$ charged-lepton pion decay opens up. This is another strongly constrained channel, and its decay rate is already given in Equation (4.2.3) with $m_l = m_\mu$. Although this decay rate can also be arbitrarily suppressed by reducing the size of $U_{\mu 4}$, due to the constraint that all sterile neutrinos must be produced through $U_{\mu 4}$ or U_{e4} mixing, in no case will both of the $l^{\pm}\pi^{\mp}$ channels be absent. As can be seen in Figure (4.3), we can expect one of them to dominate for higher masses.

Non-minimal models

In the previous section we have discussed the decay rates for the minimal model of Equation (4.2.1). Although such low-scale see-saw models provide a viable and phenomenologically interesting region of parameter space for both masses and mixing, they lack a theoretically appealing mechanism to explain the sub-electroweak sterile neutrino mass scale and the sizes of active neutrino masses. Alternative models exist which feature light neutral particles, but these rely on additional fields or symmetries to help explain these scales. Indeed it has been stressed before [205] that the discovery of a light sterile neutrino would necessitate not just the addition of new neutral fermions to the Standard Model but would be a sign of the existence of some non-trivial new physics with which to stabilise the mass scale.

If heavy novel fields are present in the full model, we can view Equation (4.2.1) as the renormalizable terms of an effective lagrangian, just as we did when considering neutrino masses in Chapter 1. The effective field theory of a Standard

Model extension involving new sterile fermions has been considered at dimension 5 [205, 206], dimension 6 [205] and dimension 7 [207]. We extend the field content of the Standard Model by a set of sterile fermions N_i. The lagrangian can then be decomposed as a formal power series of terms of increasing dimension d, suppressed by a new physics energy scale Λ,

$$\mathcal{L} = \mathcal{L}_N + \sum_{d=5}^{\infty} \frac{1}{\Lambda^{d-4}} \mathcal{L}_d,$$

where $\mathcal{L}_N$ is given by Equation (4.2.1) as the sum of the Standard Model and renormalizable terms including N_i. In Ref. [206] the phenomenology of the effective operators at dimension 5 are considered in detail. Along with the Weinberg operator, which could be the source of a light neutrino Majorana mass term [208], the authors find two effective operators: an operator coupling the sterile neutrino to the Higgs doublet and a tensorial coupling between the sterile neutrino and the hypercharge field strength

$$\mathcal{L}_5 \supset \frac{a_{ij}}{\Lambda} \overline{N_i^c} N_j (H^\dagger H) + \frac{b_{ij}}{\Lambda} \overline{N_i^c} \sigma^{\mu\nu} N_j B_{\mu\nu}.$$

At energies below the electroweak scale, and after diagonalisation into mass eigenstates for the neutrinos, these operators generate novel couplings, for example a vertex allowing $N \to h\nu$ ($N_1 \to hN_2$), $N \to \nu Z$ ($N_1 \to ZN_2$) and $N \to \nu\gamma$ ($N_1 \to N_2\gamma$) at a rate governed by the scale of new physics suppressing these operators. Of particular interest is the electroweak tensorial operator, which induces a rich range of phenomena [206]. In the mass range of interest in the present work, bounds on such an operator are fairly weak: strong constraints from anomalous cooling mechanisms in astrophysical settings apply only for lower sterile neutrino masses, whilst collider bounds only become competitive for higher masses. This could also be related to the enhanced $N \to \nu\gamma$ decay rate introduced in Ref. [209, 210] to explain the short-baseline anomalous excesses. See also Ref. [211] for a discussion of decay rates in the effective sterile neutrino extension up to dimension 6.

If light degrees of freedom are present in addition to (or instead of) heavy ones, the predictions could be very different from those derived from the minimal model or the low-energy effective theory. For example, models with sterile neutrinos that also

feature novel interactions can have significantly different decay rates and branching fractions, strengthening some bounds and invalidating others [212]. As an example, a model with a leptophilic Z' [213] could enhance the magnitudes of some leptonic decay rates, such as $N \to \nu e^+ \mu^-$, while leaving unchanged semileptonic processes like $N \to e^\pm \pi^\mp$. Often, the bounds on masses and mixing angles in these models need to be reconsidered.

For the reasons discussed so far, it is desirable to place bounds on all possible decays of a neutral fermion allowing for non-standard decay rates to visible particles. The main consequence of this is that there is *a priori* no known relationship between the magnitude of the different decay rates — a single channel may be enhanced beyond its value from Section (4.2.2) — and bounds inferred from the non-observation of a given channel may not hold in a non-minimal model when applied to another channel. We therefore do not restrict our study to those decays which lead to the most stringent bounds on the parameters of the minimal model, instead studying all kinematically viable decays independently.

4.2.3 Existing bounds on $U_{\alpha 4}$

The minimal lagrangian in Equation (4.2.1) has been the basis of many prior experimental searches for heavy sterile fermions, leading to a variety of bounds on the magnitude of the active-sterile mixing relevant for sterile neutrino masses around the MeV-scale. In this section we discuss the relevance of three key bounds on our model: peak searches, beam dumps and non-terrestrial considerations.

An established way to find strong model independent bounds on heavy sterile neutrinos is through the study of two-body decays of mesons, particularly pions and kaons [214, 215]. Due to the two-body kinematics, the magnitude of the neutrino mass manifests itself as a monochromatic line in the charged lepton energy spectrum at $E_l = \left(m^2_{\pi(K)} + m^2_l - m^2_N\right)/m^2_{\pi(K)}$. These peak searches provide strong bounds on the sterile-active mixing, while remaining agnostic as to the ultimate fate of the

sterile neutrino, which may be extremely long lived[3]. Meson decay peak searches have taken place for $\pi \to \nu e(\mu)$ and $K \to \nu e(\mu)$ and strongly bound active-sterile mixing angles at low masses. The strength of these bounds is not a function of sterile neutrino decay-rate, and as such, peak searches tend to perform worse at higher masses in comparison to bounds from experiments which derive their signal from large sterile neutrino decay rates.

The tightest bounds on MeV scale sterile neutrinos come from beam dump experiments. Beam dump experiments study the particles emitted during proton collisions with a target. Although BSM particles may be produced directly [216, 217], sterile neutrinos would predominantly arise as secondary decay products of mesons produced in the initial collision. The set-up required for such an experiment is quite minimal — a proton beam, a target and a down-wind detector — and for this reason searches of this type have taken place at many accelerator complexes, taking advantage of preexisting proton beams in their design. Seeking to produce and observe the subsequent decay of the sterile neutrinos, the sensitivity of beam dump experiments is driven by both flux intensity and the decay rate of the heavy sterile neutrino, which typically scales as ($\Gamma \propto m_N^3$) $\Gamma \propto m_N^5$ for (semi-) leptonic decays. As such they typically set tighter bounds as the sterile neutrino mass increases. As discussed in the introduction, PS-191, which ran in parallel with the BEBC bubble chamber, provides the strongest limits on active-sterile mixing for masses below the kaon mass. Above this mass, a higher energy proton beam is needed to further the same strategy. This was achieved by moving from the CERN PS to the SPS proton beam in both the CHARM [195] and NA3 [218] experiments. Beam dumps are incredibly sensitive to active-sterile mixing and limits $|U_{e4}|^2 \leqslant 10^{-8}$–$10^{-9}$ were set for $m_N \geqslant 200$ MeV.

Results from beam dump experiments are most often presented, as we did above,

[3]If, on the other hand, the sterile neutrino is extremely short-lived, these bounds may in fact be weakened. If the particle decays on the scale of the experiment, it may produce a multi-lepton final state and escape observation by the single-lepton analysis cuts.

as upper limits on active-sterile mixing in the context of the minimal model. However, beam dump experiments actually set two bounds: there is also a lower bound on the mixing-matrix elements, where the decay rate is so large that the sterile neutrino beam attenuates en route to the detector. In the minimal model, this lower bound is often at very large (or unphysical $\geqslant 1$) values of $|U_{\alpha 4}|^2$, presenting consistency issues with unitarity data, and is justifiably ignored. If one considers enhanced decay rates in a non-minimal model, however, care must be taken with existing bounds as an enhanced decay rate would modify both bounds. This can reduce the applicability of certain bounds to non-minimal models. It is instructive to discuss how to scale existing bounds on the minimal model, or indeed the bounds we will present in Section (4.4), to an extended model which has an enhancement in the decay rate for one or more channels. By comparing the flux-folded probabilities to decay inside a detector for a beam dump experiment of baseline L and detector length λ, we can map the published lower bound, given by tilded mixing parameters $|\widetilde{U}_{\alpha 4}|^2$, to both the new upper and lower bound on the mixing matrix element in a non-minimal model, $|U_{\alpha 4}|$.

For a generic non-minimal model in which the total decay rate is scaled by a factor A with respect to the minimal model, and the decay rate into the channel of interest is scaled by a factor B, the constraint takes the form of Lambert's equation (at leading order in λ/L), and the bounds on the non-minimal mixing-matrix element are given by the two real branches of the Lambert-W function,

$$\frac{|\widetilde{U}_{\alpha 4}|^2}{B\kappa}\mathcal{W}_{-1}\left(\exp^{\kappa}\frac{B}{\sqrt{A}}\kappa\right) \quad \leqslant \quad |U_{\alpha 4}|^2 \quad \leqslant \quad \frac{|\widetilde{U}_{\alpha 4}|^2}{B\kappa}\mathcal{W}_{0}\left(\exp^{\kappa}\frac{B}{\sqrt{A}}\kappa\right),$$

where $\kappa \equiv -\Gamma_{\mathrm{T}} L/(2\gamma\beta)$ with Γ_{T} the total decay rate calculated with $|\widetilde{U}_{\alpha 4}|^2$. The upper bound is primarily dependent on the decay rate into the channel of interest, governed by the parameter B, whilst the lower bound is predominantly sensitive to the total decay rate and the parameter A. Physically, the upper bound is seen to depend on how many decays are produced and is sensitive to the (possibly enhanced) decay rate into that channel, while the lower bound arises when the beam attenuates due to decay before the detector, the rate of which is governed by the

total decay rate.

Although the exact behaviour of the bounds for a non-minimal BSM extension are model-dependent due to correlations between A and B, in many situations the upper bound can be significantly simplified. We consider two distinct scenarios depending on whether the enhancement affects the decay rate of the channel being observed, or another decay channel. We write the total decay rate as $\Gamma_{\mathrm{T}} = \Gamma_{\mathrm{o}} + \Gamma_{\mathrm{c}}$, where Γ_{c} denotes the channel whose decay products are being measured and Γ_{o} the sum of all other decay rates. In our first scenario, the only enhancement is to the channel of interest, and the total decay rate can be written as $\Gamma_{\mathrm{T}} = \Gamma_{\mathrm{o}} + B\Gamma_{\mathrm{c}}$. In this case, the upper bound from Equation (4.2.3) can be simplified by expanding in the published bound[4], $|\widetilde{U}_{\alpha 4}|^2$. In this approximation, the new bound is seen to be a simple scaling of the old bound

$$|U_{\alpha 4}|^2 \leqslant \frac{|\widetilde{U}_{\alpha 4}|^2}{\sqrt{B}}.$$

This follows our naïve expectations: a larger decay rate produces more events and so bounds are proportionally stronger.

The lower bound on $|U_{\alpha 4}|^2$ has no corresponding simple form, but numerically can be seen to follow a similar scaling relationship: as the enhancement goes up, the bound moves to lower values of the mixing-matrix element. In this case, apart from a replacement of the minimal $|\widetilde{U}_{\alpha 4}|^2$ by an effective mixing-matrix element $|U_{\alpha 4}|^2/\sqrt{B}$, the bounds are to a good approximation unchanged. The situation is qualitatively different in our second scenario, however. In this case, we consider an enhancement to Γ_{o}, so that $\Gamma_{\mathrm{T}} = A\Gamma_{\mathrm{o}} + \Gamma_{\mathrm{c}}$. We find that the upper bound is unchanged to leading order, $|U_{\alpha 4}|^2 \leqslant |\widetilde{U}_{\alpha 4}|^2$. However, the lower bound moves to smaller values as the enhancement increases. For large enhancements, this can significantly reduce the region of parameter space in which an experiment can bound the model. We will return to these simplified models of decay rate enhancement in Section (4.4), and

[4]These are typically of the order 10^{-4}–10^{-8} and so such an expansion is an extremely good approximation.

give diagrammatic examples.

We note in passing that the limit of large λ/L can also be relevant for Equation (4.2.3). This corresponds to experiments where production and detection occur inside the detector, which can be seen as zero baseline beam-dumps. We find that these experiments produce only an upper bound on the mixing angle, as the number of incoming sterile neutrinos can no longer be attenuated through decays occurring before the detector.

Although peak searches and beam dumps set some of the most stringent upper limits on mixing, non-terrestrial measurements may also place bounds on such long lived sterile neutrinos due to their effect on the evolution of the early universe. Heavy sterile neutrinos can have a strong impact on the success of Big-Bang Nucleosynthesis (BBN) by both speeding up the expansion of the universe with their additional energy, and thus effecting an earlier freeze out of the neutron-proton ratio, as well as potentially modifying the spectrum of active neutrinos via their subsequent decays. If, however, the sterile has sufficiently short lifetime then their effect on BBN is mitigated as the bulk of thermally produced sterile neutrinos have decayed long before $T_{\rm BBN} \approx 10$ MeV [219]. The strength of these bounds have been estimated conservatively for a single sterile neutrino, $m_N < m_{\pi^0}$, as [220, 221]

$$\begin{aligned} \tau_{\rm N} &< 1.287 \left(\frac{m_N}{\rm MeV}\right)^{-1.828} + 0.04179 \text{ s} \qquad \text{for } U_{\mu 4} \text{ or } U_{\tau 4} \text{ mixing,} \\ \tau_{\rm N} &< 1699 \left(\frac{m_N}{\rm MeV}\right)^{-2.652} + 0.0544 \text{ s} \qquad \text{for } U_{e4} \text{ mixing,} \end{aligned}$$

at the 90% C.L Although the scenario for $m_N > m_{\pi^0}$ has not been studied in as much detail, an often quoted bound is that $\tau_N > 0.1$ s is excluded under current BBN measurements [220]. In the minimal model, this upper bound on the sterile neutrino lifetime is directly mapped to a minimum bound on the active-sterile mixing elements $U_{\alpha 4}$. However, even a modest increase in the total sterile neutrino decay rate, for example by additional interactions in the sterile neutrino sector leading to decays that are not mixing suppressed, pushes the total sterile neutrino lifetime below 0.1 s and avoids these bounds, while still leaving channel specific signatures

observable at SBN as the upper bounds are independent on total decay rate. Similarly in a non-standard model of the early universe, these bounds may not apply. Therefore, although setting important complementary bounds on models of sterile neutrino decay, model dependent factors make it possible for discrepancy between peak search, beam dump and cosmological constraints. As such a wide program of experimental work is desirable, with as varied a methodology as possible, to best identify new physics.

4.3 Simulation of SBN

SBN consists of three LArTPC detectors (SBND, MicroBooNE and ICARUS) located in the Booster neutrino beam. The Booster neutrino beam is a well understood beam, having been recently studied by the MiniBooNE experiment. For the purposes of this analysis each detector is assumed to be identical apart from their geometric dimensions. We simulate the event numbers and distributions at each detector site using a custom Monte Carlo program which allows efficiency's to be taken into account arising from experimental details such as energy and timing resolution in a fully correlated way between observables, and provides us with event level variables for use in a cut-based analysis. We compute the total number of accepted events in channel "c" via the following summation,

$$N_{\mathrm{c}} = \sum_i \left.\frac{\mathrm{d}\phi}{\mathrm{d}E}\right|_{E_i} P_{\mathrm{D}}\left(E_i\right) W_{\mathrm{c}}\left(E_i\right),$$

where $P_{\mathrm{D}}(E)$ is the probability for a sterile neutrino of energy E_i to travel the baseline distance and then decay inside the detector labelled D. The simplest approximation is to ignore all geometric effects, so that every particle travels exactly along the direction of the beam line, which gives the following probability

$$P_{\mathrm{D}}\left(E\right) = e^{-\frac{\Gamma_{\mathrm{T}} L}{\gamma\beta}} \left(1 - e^{-\frac{\Gamma_{\mathrm{T}} \lambda}{\gamma\beta}}\right) \frac{\Gamma_{\mathrm{c}}}{\Gamma_{\mathrm{T}}},$$

where Γ_{T} (Γ_{c}) denotes the rest-frame total decay width (decay width into channel c), and L (λ) the distance to (width of) the detector. The combination $\gamma\beta$ is the usual special relativistic function of velocities of the parent particle and provides

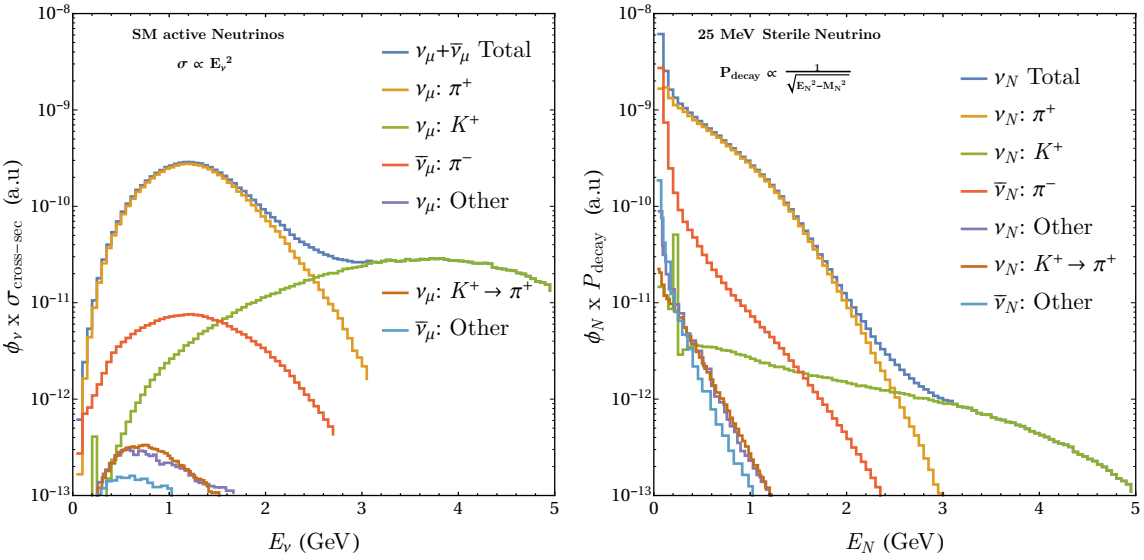

Figure 4.5: ***Left:*** The composition of fluxes of ν_μ and $\overline{\nu}_\mu$ at MicroBooNE with horn in positive polarity (neutrino mode), scaled by an illustrative scattering cross-section, $\sigma \propto E_\nu^2$, in order to stress that the low-energy flux is not very important for traditional scattering experiments. "Other" refers to contributions primarily from meson decay chains initiated by meson-nucleus interactions. ***Right:*** Sterile neutrino fluxes weighted by the probability to decay inside MicroBooNE, for a sample 25 MeV sterile neutrino with equal $|U_{e4}|^2 = |U_{\mu 4}|^2$. Requiring that the sterile neutrino decays inside the detector has the effect of vastly increasing the importance of lower energy bins. The difference in importance of high and low energy fluxes in scattering and decay experiments respectively is dramatic, leading to very different kinematics.

the sole dependence on energy and sterile neutrino mass m_N of the expression

$$\frac{1}{\gamma\beta} \equiv \frac{m_N}{\sqrt{E^2 - m_N^2}}.$$

As we are exploring a large parameter space, often this expression takes a simplified form depending on the size of $\Gamma_{\mathrm{T}}\lambda/\gamma\beta$:

$$\Gamma_{\mathrm{T}}\lambda \ll 1 \qquad P_{\mathrm{D}} = e^{-\frac{\Gamma_{\mathrm{T}} L}{\gamma\beta}} \frac{\Gamma_{\mathrm{c}}\lambda}{\gamma\beta} + \mathcal{O}\left(\Gamma_{\mathrm{T}}^2\lambda^2\right), \tag{4.3.4}$$

$$\Gamma_{\mathrm{T}}\lambda \gg 1 \qquad P_{\mathrm{D}} = e^{-\frac{\Gamma_{\mathrm{T}} L}{\gamma\beta}} \frac{\Gamma_{\mathrm{c}}}{\Gamma_{\mathrm{T}}} + \mathcal{O}\left(\frac{1}{\Gamma_{\mathrm{T}}\lambda}\right), \tag{4.3.5}$$

where the rate for slowly decaying particles can be seen to grow with detector size until a width of $\lambda \sim \gamma\beta\Gamma_{\mathrm{T}}^{-1}$. For detectors longer than this scale, the event rate

becomes independent of detector size, as most sterile neutrinos decay within a few decay lengths.

The spectral flux of sterile neutrinos impinging on a SBN detector, $\mathrm{d}\phi/\mathrm{d}E$, is estimated as described in Section (4.2.1). Of crucial importance to this is accurate knowledge of active neutrino fluxes at all three SBN detectors. These are calculated from published MiniBooNE fluxes [145], after scaling by appropriate $1/r^2$ baseline dependence, *e.g.* $(470/540)^2 \approx 1.3$ at MicroBooNE. This is similarly scaled by $1/r^2$ for ICARUS at 600 m, however, an additional energy dependent flux modifier is applied for SBND at 110 m to account for the softer energy spectrum due to the proximity of the detector to the production target [108]. We consider sources of neutrinos that are relevant including wrong sign neutrinos, smaller sub-dominant $K^+ \to \pi^+ \to \nu_\alpha$ sources as well as other contributions, predominantly from meson decay chains initiated by meson-nucleus interactions.

In Figure (4.5) we highlight a crucial point about the neutrino and sterile neutrino fluxes. We note that the flux of impinging sterile neutrinos is not the exact parameter of interest, rather what we are interested in is the flux of sterile neutrinos that subsequently decay. In the right panel, we show the effective spectrum of *decaying* particles at MicroBooNE. As the decay probability for any given sterile neutrino scales as $1/|P_N|$, we see an enhancement of the lowest energy parts of the spectrum. This is in stark contrast to standard neutrino interaction cross sections, which tend to scale as approximately E_ν^2 at low energies. In the left hand panel of Figure (4.5) we show the active neutrino spectrum at MicroBooNE where we have weighted the flux by E_ν^2 in order to stress the difference in regions of importance between scattering and decaying neutrinos. This, low-energy bias for decays naturally exaggerates further the kinematic differences between our decay-in-flight signal and the dominant background events, the majority of which are neutrino induced scattering events.

Finally, the function $W_\mathrm{c}(E)$ is a weighting factor which accounts for all effects which reduce the number of events in the sample: for example, analysis cuts or

detector performance effects. To compute these factors, we run a Monte Carlo simulation of the decays for a large number of sample events with a given energy. Each sterile neutrino event is associated with a decay of type c. We then apply experimental analysis cuts to the decays based on our assumptions about the detector's capabilities and backgrounds, to produce a spectrum representing the final event sample when considering events in the bucket timing window (See Section (4.3.2) for details of the background analysis). The percentage of accepted events defines the weight factor for that energy. In this manner the full spectral shape of the signal is included in the total rate analysis.

4.3.1 PS-191 Bound Reproduction

As a consistency check of our methodology as described above we reproduce here the bounds on $|U_{e4}|$ and $|U_{\mu 4}|$ for sterile masses below m_π as published by PS-191. The detector geometry is assumed to be 6m × 3m × 12m and was located 128m downstream of the Beryllium target using 19.2 GeV protons from the PS proton beam. Fluxes of all neutrinos produced from pion sources at PS-191 were obtained from [222]. No accurate kaon sources could be obtained and as such only low mass bounds are reproduced here. It must be noted that PS-191 ignored all neutral current contributions to $N \to \nu_\alpha e^+ e^-$ and assumed the sterile neutrinos were Dirac particles; the effect of this is that the bounds published are not directly comparable to the minimal model discussed above, and must be scaled appropriately. The bounds reproduced are in good agreement with published data.

4.3.2 Background reduction

In order to estimate the impact of potential backgrounds we have performed a Monte-Carlo analysis using the neutrino event generator GENIE [182]. This provides generator level information about the kinematics of the beam-driven backgrounds, with rates normalised by expected NC and CC inclusive values as published in the SBN proposal. Energy and angular smearing is then implemented to allow for approximate estimates of the effects of detector performance to the level necessary for

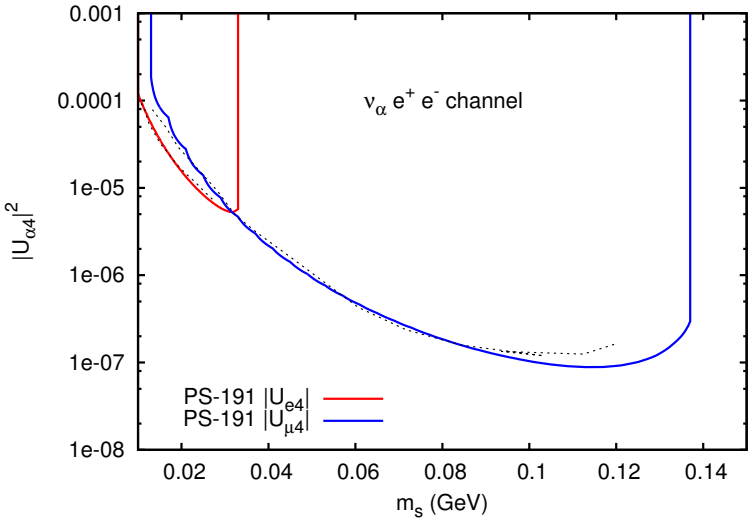

Figure 4.6: Estimated 90% confidence limits on the sterile neutrino mixing parameters $|U_{e4}|^2$ and $|U_{\mu 4}|^2$ for a heavy sterile neutrino decaying to $\nu_\alpha e^+ e^-$ at PS-191. The dotted black lines are the 90% C.L results as published by PS-191, and the blue and red curves are the results of our simulation for 0.86×10^{19} POT, using the above mentioned methodology and code.

this analysis, without the need for a full GEANT detector simulation. Energies are smeared according to a Gaussian distribution around their true MC energies, with a relative variance $\sigma_E/E = \xi/\sqrt{(E)}$, where ξ is a detector dependent resolution. For this study we take the energy resolution for EM showers, muons and protons to be 15%, 6% and a conservative 15% respectively, alongside an angular resolution in LAr of 1.73° [108].

Of utmost importance in all studied channels is the identification of a scattering vertex, which cleanly indicates that the process is not a decay-in-flight event. Any hadronic activity localised at the beginning of the lepton track is a smoking gun for a deep-inelastic or quasi-elastic beam-related scattering event. Therefore we reject any event containing one or more reconstructed protons or additional hadrons. For counting this proton multiplicity we assume a detection threshold of 21 MeV on proton kinetic energy in liquid Argon [164], after smearing. Background events with energies below this threshold and events that do not contain any protons (such as

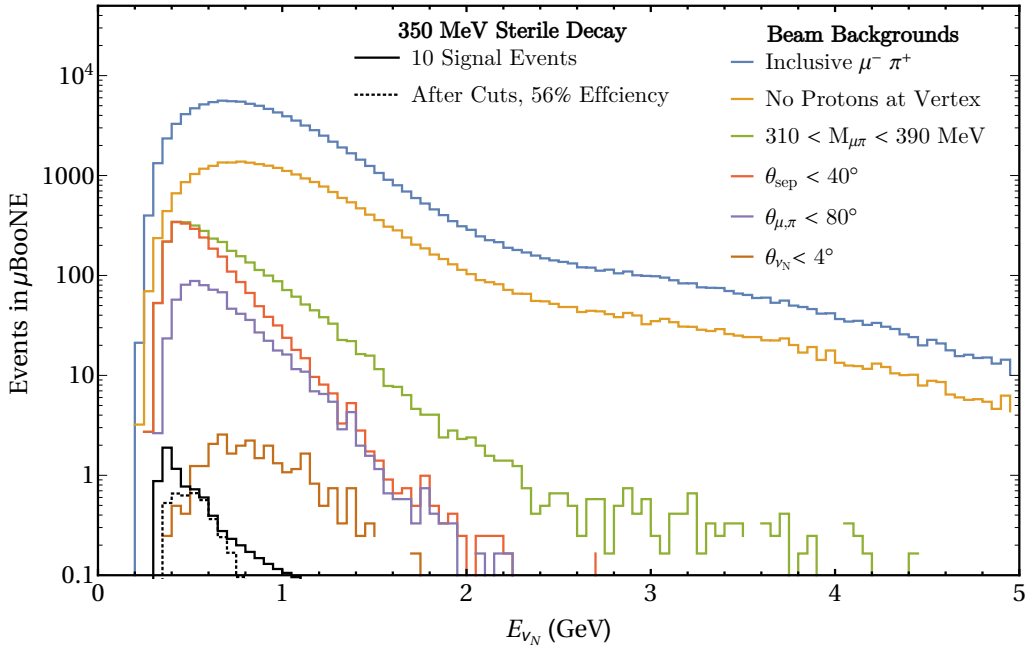

Figure 4.7: Reconstructed sterile neutrino energy spectra for CCν_μ backgrounds in comparison to a 350 MeV decaying sterile neutrino at MicroBooNE, normalised to 10 signal events. Total expected background of 98,013 events is reduced to $\approx$ 27 by successive kinematic cuts (as listed in legend) which utilise the stark differences between decay-in-flight and scattering kinematics. Further cuts on energy would allow for even greater reduction.

events originating from coherent pion production) remain a viable background and further rejection must come from the kinematics of the final state particles only. The kinematics of such daughter particles originating from decay-in-flight and backgrounds from scattering events, however, have strikingly different behaviour leading to strong suppression capabilities.

In all channels a cut on vertex activity is applied, in addition other channel specific cuts are used to estimate sensitivity. We will now provide a brief description of the backgrounds and additional cuts considered for each channel.

$N \to \mu^{\pm}\pi^{\mp}$

The decay $N \to \mu^{\pm}\pi^{\mp}$ represents the channel with largest expected beam related backgrounds in all SBN detectors, the dominant component of which arises from genuine charged current $\pi\mu$ production. These events can be produced incoherently, often with large hadronic activity and so will greatly be reduced by the cut on a scattering vertex, or from coherent scattering, where the neutrino scatters from the whole nucleus. Coherent cross-sections for these processes have been studied in MiniBooNE [223], MINERνA [224] and lately T2K and cross-sections appear to agree with Monte Carlo calculations based on the Rein-Sehgal model [225] and generally do not have an additional hadronic component to cut on. Kinematics of the daughter particle alone but be used for background rejection. There has been a noted deficit of forward going muons [225] in these coherent cross-sections, which is in stark comparison to the relatively forward behaviour of sterile neutrino decays.

Furthermore, this channel, and indeed $e^{\pm}\pi^{\mp}$, has a powerful discriminator in the reconstructed invariant mass of the charged particle pair, *e.g.* $M^2_{l^{\pm}\pi^{\mp}} = m_l^2 + m^2_{\pi^{\pm}} + 2(E_l E_\pi - |P_l||P_\pi|\cos\theta_{\text{sep}})$ for $N \to \pi^{\pm} l^{\mp}$, which sum to that of the the parent sterile neutrino (within detector resolution), whereas the background forms a broad spectrum across the energies of the incoming neutrinos. On top of this strong invariant mass discriminator, these two body decays allow for reconstruction of the parent sterile neutrino angle with respect to the beamline which is very close to on-axis, as opposed to the more isotropic backgrounds. We find that approximately 95% of the reconstructed sterile neutrino angles from these decays are inside a 4° cone centred on the beamline.

We show the effect of our cuts for this channel in Figure (4.7), which ultimately leads to a reduction in the inclusive $\mu\pi$ event rate at SBND (MicroBooNEand ICARUS) from 1,530,900 (98,013 and 164,716) to 323 (27 and 46) while maintaining a signal efficiency of 56%. This level of background suppression crucially relies on the angular and energy resolution of LAr detectors, but requires no modification to the current design.

$N \to e^{\pm}\pi^{\mp}$

The expected numbers of $e\pi$ events in the SBN detectors is significantly smaller than that of the $\mu\pi$ channel, as the fraction of intrinsic ν_e in the BNB beam is of $\mathcal{O}(1\%)$ level in comparison to ν_μ. However, additional backgrounds to the $e\pi$ channel originate from the dominant ν_μ component of the beam. CC ν_μ events which contain an additional photon $(\mu + \gamma)$ have the potential to be be mis-identified as an (πe) event, provided the muon has a sufficiently short track length, < 0.5 m, in order to mimic a π^-. Additionally the photon must be mis-identified as an electron, with an efficiency of 94% using dE/dx measurement, and must convert to an e^+e^- pair close enough to the interaction vertex as so there is no visible gap, $\leqslant 3$ cm. As energy resolution for EM showers is lower than muons, the invariant mass cut is less powerful requiring all events have an invariant mass below 500 MeV. A cut on the opening angle between lepton on meson, $\theta_{l\pi} < 40^\circ$ as well as individual emission angles, $\theta_{l,\pi} < 80^\circ$ further reduces the potential background. The $e^-\pi^+$ channel is one of the cleanest channels under consideration in this analysis, with 9,223 events in SBND reducing to 22 expected events post cuts, and with MicroBooNE and ICARUS expecting a reduction of 784 (1,317) events to 2 and 3 respectively, with a signal efficiency of 71%.

$N \to \nu_\alpha e^+e^-$ and $N \to \gamma\nu_\alpha$

A sufficiently boosted, and thus overlapping, e^+e^- pair is topologically indistinguishable from a converted photon in a LAr detector. Additional, non-topological measures such as the rate of energy loss, dE/dx, is also identical to a pair-converted photon. Thus we split this channel into two sub categories, when the e^+e^- is overlapping and photon-like, defined to be all events whose angular separation is $\leqslant 3^\circ$ [226] and all remaining separable two track events. The opening angle between the e^+e^- in a photon pair production scales roughly as $\approx m_e/E_\gamma$, with 3° corresponding to 100 MeV and used as a lower bound on energy. These backgrounds are also applicable to the $N \to \gamma\nu_\alpha$ channel.

The predominant source of backgrounds is the decay of a neutral pion in which a single photon is not resolved or escapes the fiducial volume. This background, however, is relatively isotropic in distribution in stark contrast to the very forward signal arising from the decay in flight of N. We place cuts on visible photon energy, $E_\gamma \geqslant 300$ MeV, and angle of the observed photon to the beamline, $\theta_\gamma \leqslant 5^\circ$. This reduces the number of expected events from 42,580 (3,620 and 6,082) to 176 (46, and 110) events in SBND (MicroBooNE and ICARUS), while retaining a signal efficiency of 93%.

For the opposite scenario both daughter electrons have a well defined and large separation and thus can cleanly be identified as two distinct single electron showers. There are few significant processes that produce high energy, distinguishable e^+e^- pairs. Instead the majority of the backgrounds are due to misidentifying two photons. We apply the same photon cuts as defined above. To further reject backgrounds in this channel, we apply a cut on the angle of separation between the distinct e^+e^- tracks of $\theta_{\text{sep}} \leqslant 40^\circ$ and total energy, $E_{e^+} + E_{e^-} \geqslant 100$ MeV. This reduces the number of expected background events from 173 to 5 for SBND.

$N \to \pi^0 \nu_\alpha$

Single neutral pions are produced in great numbers at the three SBN detectors, so the lack of any nuclear recoil is crucial in eliminating the incoherent neutral pion production background. Only events in which two photons convert inside the fiducial volume and reconstruct the pion invariant mass are accepted. We further require the reconstructed pion is within 10° of the beamline and has an energy above 500 MeV. SBND expects 127,211 π^0 events, of which ≈ 602 survive all cuts with a signal efficiency of 32% for a sample 350 MeV sterile neutrino. MicroBooNE (ICARUS) sees a similar reduction, from 10,813 (18,172) events to 51 (86) post cuts.

$N \to \nu\mu^{\pm}\mu^{\mp}$

The primary background for this channel is genuine ν_μ CC events in which a $\pi^{\pm}$ is also produced and is misidentified as a secondary muon. All pions with reconstructed tracks longer than 50 cm are considered a potential muon. After this track length cut, there is 2,044,380, 177,972 and 292,034 events in SBND, MicroBooNE and ICARUS respectively. As we cannot directly reconstruct the parent sterile neutrino mass or angle, we again rely on the kinematical difference between scattering events and decays. After these cuts, significant backgrounds remain, and we use a multivariate analysis, an adaptive boosted decision tree (BDT), in order to further reduce them. We use five parameters in this analysis, the energy and angle with respect to the beamline of each muon, as well as the angular separation between both muons. We take a minimum muon energy of 200 MeV. Cutting on the BDT response variable allows for background efficiency of 0.13%, with a corresponding signal efficiency of 44%. This allows for a $S/\sqrt{S+B} \approx 8$ with approximately 1000 sterile neutrino events. Similar performance is achievable at MicroBooNE and ICARUS, with 117,972 and 292,034 events being reduced to 534 and 876 events respectively.

$N \to \nu e^{\pm}\mu^{\mp}$

We consider here two potential sources of backgrounds: the first derives from true ν_μ CC events in which a single photon, either from nuclear processes or from the decay of a π^0 in which only photon converts inside the fiducial volume, subsequently mimics the electron. We apply the same cuts on the photon as in previous channels. Secondly we consider intrinsic ν_e CC events in which a final state $\pi^{\pm}$ is misidentified as a muon due to a long ($\geqslant$ 50 cm) track in the TPC. In conjunction with the requirement of no visible scattering vertex we expect 7,103, 618 and 1,014 events in SBND, MicroBooNE and ICARUS, respectively. To reduce this further we employ the same multivariate analysis as described for the $N \to \nu\mu\mu$ channel above, assuming a representative 250 MeV sterile neutrino decaying. A cut on the BDT allows for a background efficiency of 0.5% , signal efficiency of 36% with a resultant $S/\sqrt{S+B}$ of 7.9. For MicroBooNE and ICARUS the backgrounds, 618 and 1,014 respectively, can be brought down to sub 10 events.

4.3.3 Non-Beam related backgrounds

Cosmogenic events are a potential source of background for any analysis at SBN as all three detectors have minimal overburden. In the case of cosmic muons, ICARUS expects to see approximately 2.5×10^6 cosmic events in the 211 second beam spill, which are reduced to approximately 5 events expected after utilising the spill structure, scintillation light patterns and cuts on $\frac{dE}{dx}$ [108]. Alongside this impressive cosmic rejection, our signal events are focused heavily along the beamline with distinct kinematics, hence we do not expect cosmics to be a major source of background to any channel. *In situ* beam-off cosmic studies will also allow potential backgrounds to be extremely well understood by the time of an analysis such as this, and for these reasons, we do not include cosmogenic backgrounds in our analysis.

4.3.4 Role of event timing

On top of the impressive background rejection capabilities of LAr from kinematic cuts, there is the potential for an even greater background suppression by considering the time of arrival of observed events. Although the drift time of electrons in LAr can be as large as 100 μs, the ionisation and excitation of Argon from the passage of a charged particle also produces 128 nm scintillation light of which there is a nano-second scale contribution from the decay of the excited state Ar_2^* [227]. LAr is transparent to this light, and if the light detection system (LDS) employed by the SBN detectors has a nanosecond resolution, this can allow for precise timing to be attached to each TPC triggered event.

Light neutrinos propagate and reach the furthest detector of the SBN complex after approximately 2μs. In the conventional physics program of the SBN, the timing of these events play an important role in the analysis of backgrounds, tight timing windows are placed around the 19.2 μs beam spill to limit constant rate backgrounds such as cosmogenic events [108]. The LDS of both SBND and ICARUS, however, are potentially able to achieve significantly better timing resolution than this, around 1–2 ns depending on the exact technology used, which potentially allows for the use

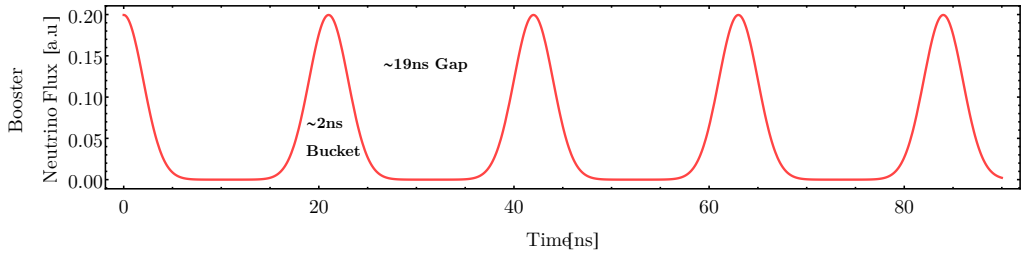

Figure 4.8: An illustrative picture of the Booster Neutrino Beam bucket cycle. Most neutrinos are expected to be in a roughly $\approx$ 2 ns bucket. If SBN detectors can observe events outside of this bucket, they are less likely to be beam-induced backgrounds.

of both bucket and spill structure in the background analysis. The BNB consists of 81 Radio-Frequency buckets of approximately 2 ns length, separated by 19 ns, to form the 19.2 μs spill with a frequency of 3Hz [108]. We illustrate this in Figure (4.8) If this nano-second resolution is indeed achieved, it allows for events in individual buckets to be identified. Such a nano-second resolution was achieved previously by the PMTs used in MiniBooNE [108], with potential for improvement in the next generation SBN detectors.

As particles with finite rest mass, heavy sterile neutrinos will propagate at sub-luminal speeds which can produce observable timing delays. This effect begins to become relevant when the sterile neutrinos have MeV-scale masses and above. As the flux of decaying sterile neutrinos is inversely proportional to its momentum after convolving with their decay probability, many of these low energy sterile neutrinos are travelling at sufficiently slower speeds than their light counterparts to be distinguishable. Shown in Figure (4.10) is the fraction of events that are expected to fall outside the bucket window in both SBND and ICARUS. For the purposes of this plot we define the beam-correlated window to be a 6 ns period, 2 ns each side of the 2 ns beam bucket. The exact width of the beam bucket window can be modified and optimised if studying channels with low expected backgrounds. In this section, we consider only the timing of events relative to the bucket window[5], a structure which

[5]Absolute arrival times could in principle be used, but this would require good synchronisation

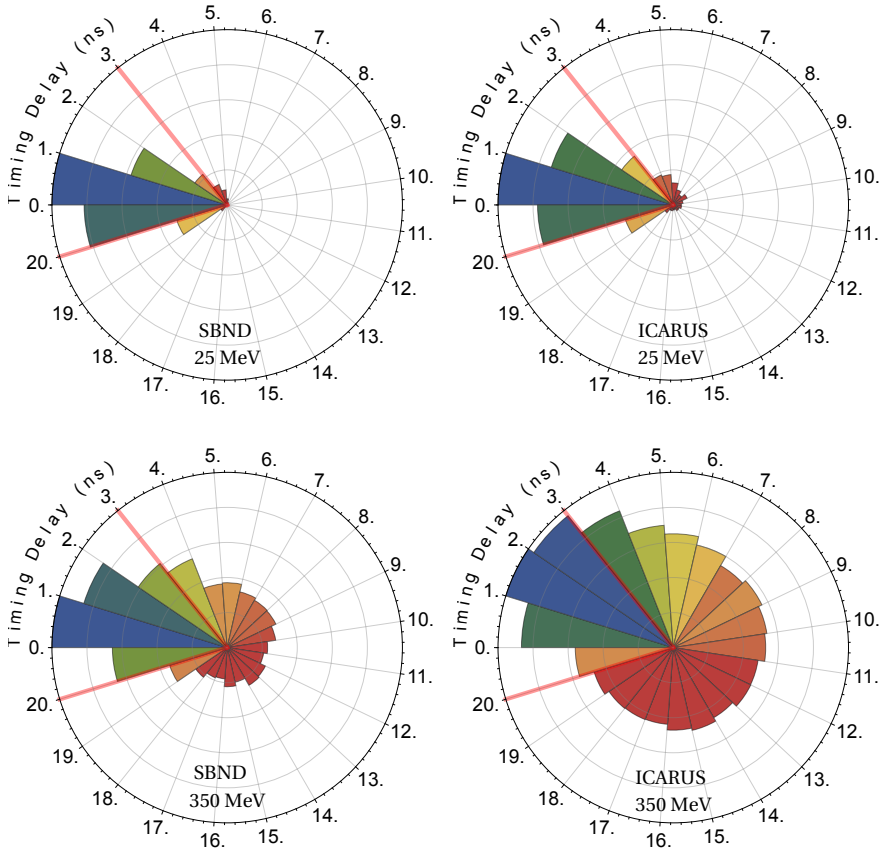

Figure 4.9: The timing delay of sterile neutrino decays in nano-seconds for both a 25 MeV (top) and 350 MeV (bottom) sterile neutrino at the SBND and and ICARUS detectors (110 and 600 m respectively). A 4 ns beam bucket window is shown highlighted in red from 0 to 4 ns, followed by an additional 17 ns gap. The timings are calculated as a difference to the time of flight of a active neutrino, assuming the decay occurred in a uniform sample across the 50 m BNB decay pipe. A timing resolution of 1 ns is assumed to smear the observed events.

repeats every 21 ns. Delayed events can be observed in any subsequent window, producing a 21-fold degeneracy in their reconstructed arrival time. This lends a cyclical nature to the timing information, with a distinctive structure at the different detector sites for larger masses. Some illustrative timing distributions are shown in Figure (4.9) for two sterile neutrinos at the extreme ends of the masses studied,

between geographically separated clocks. Alternatively, the relative timing between signal and beam-related backgrounds could be used. However, we do not consider these options further.

a 25 MeV sterile neutrino and a 350 MeV sterile neutrino.

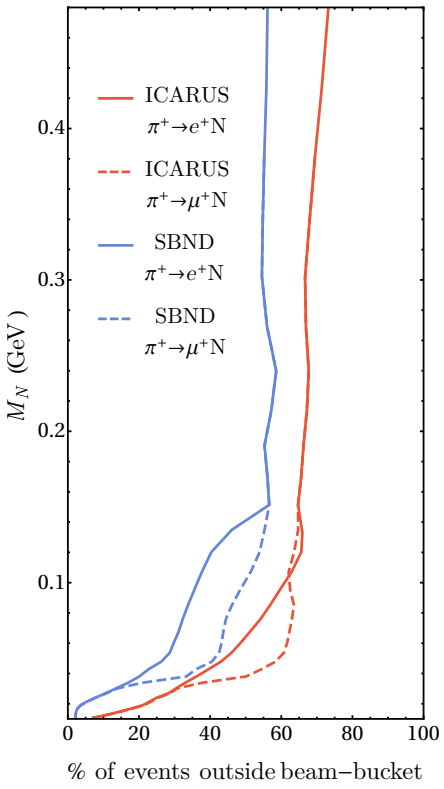

Figure 4.10: The percentage of sterile neutrino decay events that fall into the inter-bucket region as a function of sterile neutrino mass for SBND and ICARUS, assuming a flux derived from U_{e4} ($U_{\mu 4}$) mixing in solid (dashed) lines. Both SBND and ICARUS see a sizeable fraction of total events outside the beam bucket windows when the sterile neutrino mass exceeds $\approx$ 10 MeV.

We find a significant proportion of sterile neutrino events distributed throughout the inter-bucket region. Events which fall into the beam-bucket timing window have to be analysed on top of all known beam-related backgrounds, but events in the inter-bucket window have significantly reduced beam-correlated backgrounds. For larger masses, we have shown that the majority of events fall into these regions, and this may allow for a low background search strategy for decaying sterile neutrinos. Instead of beam-correlated backgrounds, the constant rate backgrounds will limit the sensitivity for this analysis. Understanding these backgrounds in detail is beyond the scope of this work; however, we expect the strongly forward kinematics,

combined with *in situ* beam-off measurements will allow for a very low backgrounds to be obtained.

In the following sections, timing information will inform our work in three ways. First we will compute SBN's sensitivity to decaying sterile neutrinos assuming the full backgrounds, reduced only by the cut-based analysis discussed previously. This is a proven sensitivity, applicable for all sterile neutrino masses and detectors and is independent of the attainable timing resolution. Secondly, we compute a backgroundless sensitivity. This can be seen as either the result of improved analysis techniques, or as the inclusion of timing information at SBND and ICARUS for the largest masses. Finally, in Section (4.4.2), we will study the use of the timing information itself to constrain the underlying model of decaying sterile neutrinos.

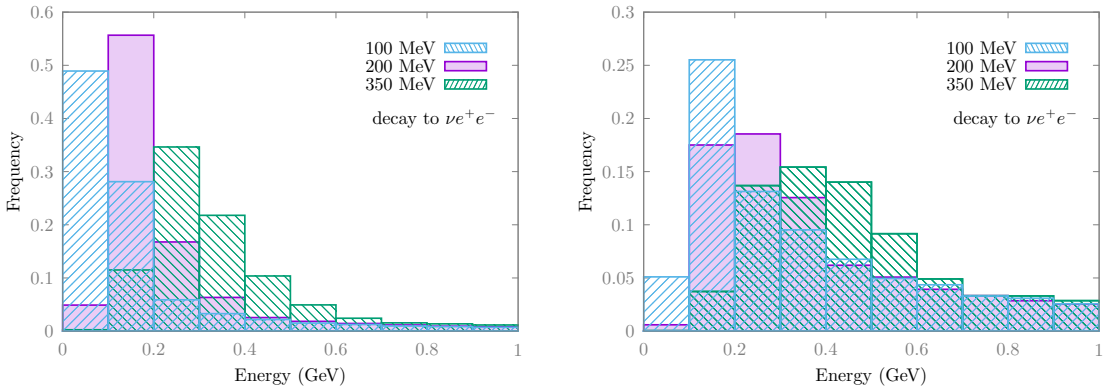

Figure 4.11: Characteristic spectra for the total energy of observed e^+e^- pairs seen at MicroBooNE produced in the $N \to \nu e^+e^-$ decay mode, for three representative masses. In the left panel, the spectra have no analysis cuts or detector reconstruction effects applied, while on the right these are included, reducing the number of lowest-energy events.

4.3.5 Event spectra

The differential distributions from heavy sterile neutrino decay tend to produce distinctive low-energy distributions of events with an appreciably forward direction. The tendency towards low energies is predominantly due to the higher decay rates of

low-energy particles, which leads to factors of $1/E_\nu$ in the event rate formula Equation (4.3.4). The forward trajectory is inherited from the kinematics of a boosted object decaying in flight. However, this effect is slightly mitigated by the preference for lower energy decays, meaning that lower energy sterile neutrinos are more likely to decay, which are the least boosted objects.

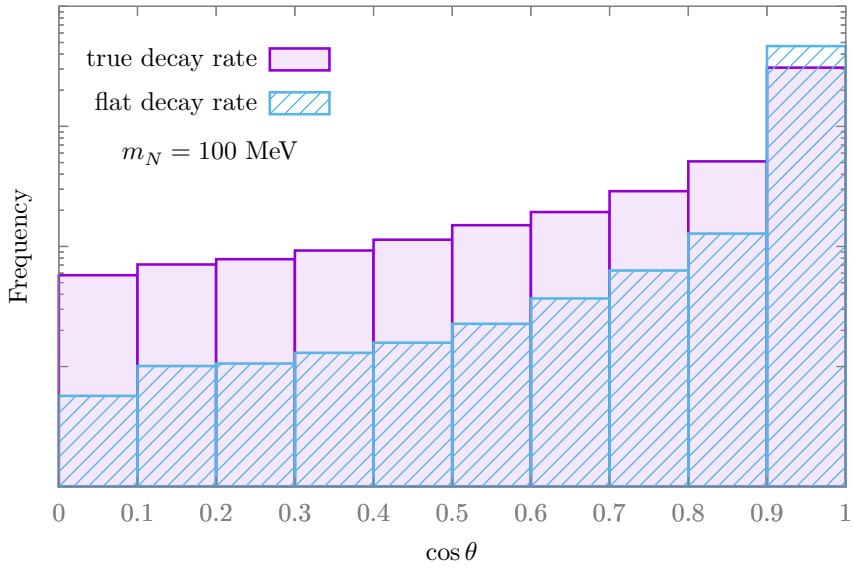

Figure 4.12: The expected angular distributions for the e^+e^- pair from a sterile neutrino of mass $m_N = 100$ MeV. The red histogram shows the true expected distribution, while the blue histogram shows the distribution if we do not take into account the preferential decay rate for lower energy particles, instead using an energy independent decay rate.

We show an example of a distribution for electron-positron production in the left panel of Figure (4.11). For the lowest masses that we consider, almost all events have energies below 0.5 GeV, in this case illustrated by the blue histogram. The distribution tends towards larger energies as the mass of the sterile neutrino increases, but for sterile neutrino masses less than the kaon mass, never produces significant numbers of events above 1 GeV. As we can see in Figure (4.11), the number of events in the lowest energy bin is strongly indicative of the mass of the parent particle, and therefore the lowest energy events will play a strong role in model discrimination. However, in the cut-based analysis which we outlined in Section (4.3.2) the lowest en-

ergy event distribution is significantly reduced due to poor efficiency's at low-energy in our cuts, as can be seen in the right panel of Figure (4.11). In Figure (4.13) we show the same distributions for electron-pion production, noting similar spectral features of the electron-positron channel. Indeed this behaviour qualitatively exists in all channels studied. In the inset of Figure (4.13), we highlight the differences that an accurate timing resolution can give, with the in-bucket and out of bucket spectra showing very significant differences. Through optimisation of this part of the analysis, we expect the sensitivity to these models can be improved; however, this is beyond the scope of the present work.

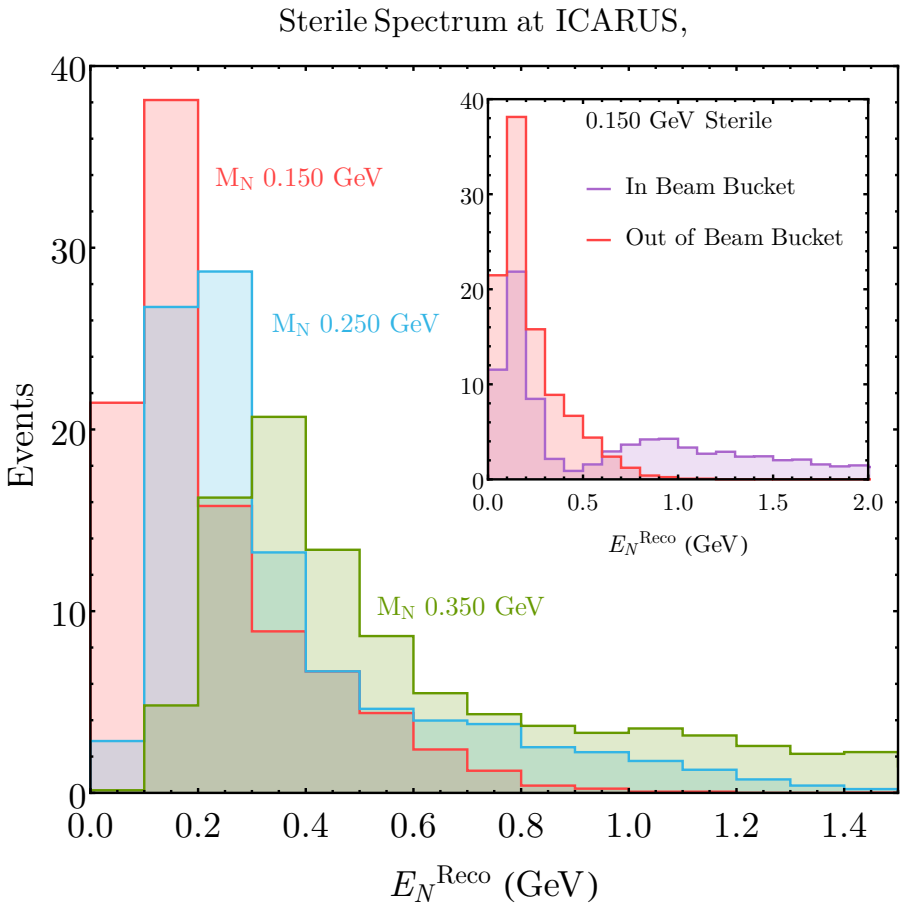

Figure 4.13: Characteristic spectra for the reconstructed neutrino energy for $N \to e^{\pm}\pi^{\mp}$ and a sterile neutrino masses of 150, 250 and 350 MeV. The insert shows the stark differences in spectrum when one considers events falling within the beam bucket window and without.

The angular spectrum is expected to be very informative in these models, and the events are predict to align with the beam direction. The red histogram in the lower panel of Figure (4.12) illustrates an expected distribution for the four momentum of a e^+e^- pair in the decay $N \to \nu e^+e^-$. We compare it to the expected distribution found for events without the low-energy biasing effect of decay-in-flight, with an unphysical energy independent decay rate (denoted 'flat', shown in blue). Not only does the decay-in-flight probability lead to a lower energy events, but it also makes the angular distribution less forward.

4.4 Results

In this section, we present the results of our simulation for two analyses. In the first, we compute exclusion contours which could be expected to be set by SBN if no signal is seen. We compute these for all decay modes presented in Figure (4.2). Our second analysis considers the phenomenological potential of energy and timing spectral information at the SBN experiment if a potential signal is observed.

Due to its proximity to the BNB target, SBND provides the majority of the statistics, and hence the sensitivity, to sterile neutrino decays. The addition of MicroBooNE and ICARUS increases the event rate by approximately 6%. However, the power of the three detector SBN setup arises not from the increased statistics, but rather from the additional phenomenology of a multi-baseline experiment. We show below that ICARUS, being the furthest detector, can play an important role in study of any observed signals through precision timing measurements. Similarly, although MicroBooNE contributes a small fraction to the raw number of sterile neutrino events expected, the MicroBooNE experiment is significantly more advanced than its two SBN counterparts. At the time of writing, MicroBooNE has already collected close to 50% of its planned POT (around 3.5×10^{20} POT) and has already presented its first results on ν_μ CC inclusive and ν_μ CCπ^0 interactions [228]. As such, MicroBooNE is in a unique position in that it has the potential to observe any excess in advance of SBND or ICARUS and thus to inform a possible search using

the full SBN programme. Non-observation of any excess at MicroBooNE would not negatively effect the subsequent search for new physics at SBND or ICARUS significantly, however, as a large fraction of the testable parameter space is only accessible through the enhanced exposures of the full SBN programme.

4.4.1 Limits on sterile neutrino mixing

We have computed the bounds that SBN could place on sterile neutrino mixing-matrix elements for all kinematically accessible visible decays. Figure (4.14) presents the results of our analysis assuming a combined 6.6×10^{20} POT at SBND and ICARUS, and 13.2×10^{20} POT at MicroBooNE. We plot the predicted upper limits on sterile neutrino mixing for the leptonic decay channels $N \to \nu_\alpha e^+ e^-$, $N \to \nu_\alpha \mu^+ \mu^-$ and $N \to \nu_\alpha \mu^\pm e^\mp$ as well as the semi-leptonic and photonic channels $N \to l^\mp \pi^\pm$, $N \to \nu_\alpha \pi^0$ and $N \to \nu_\alpha \gamma$. The plot on the right (left) assumes that the mixing-matrix element with the electron (muon) flavour is dominant. The y-axis refers to a single mixing element, $|U_{\alpha 4}|^2$, but each bound is equally applicable to the combination $|U_{\alpha 4}||U_{\tau 4}|$, as although production must proceed through electron-neutrino or muon-neutrino mixing, the decay can take place through $U_{\tau 4}$ driven processes. The lower solid coloured lines are the background-less 90% C.L upper limit contours defined as 2.44 events following the procedure of Ref. [231]. This represents the best expected sensitivity at the SBN program, assuming perfect signal efficiency. We also present the results of the cut based background analysis discussed in Section (4.3.2) (upper solid coloured lines). Depending on the optimisation of the analysis, including the possibility of using improved timing information, we expect the ultimate sensitivity to be within the solid-shaded region, lying between the proven cut-based sensitivity and the backgroundless one.

The increased event rates at SBN compared to those of PS-191 allows for an improvement of the bounds on $|U_{e4}|^2$ and $|U_{\mu 4}|^2$ in all channels studied over wide regions of parameter space. The strongest bounds come from the semi-leptonic $N \to l^\pm \pi^\mp$ searches, where mixing-matrix elements greater than $|U_{e4}|^2 \leqslant 10^{-9}$ can be excluded at the 90% C.L for $m_N \approx 0.350$ GeV. The bounds have the potential

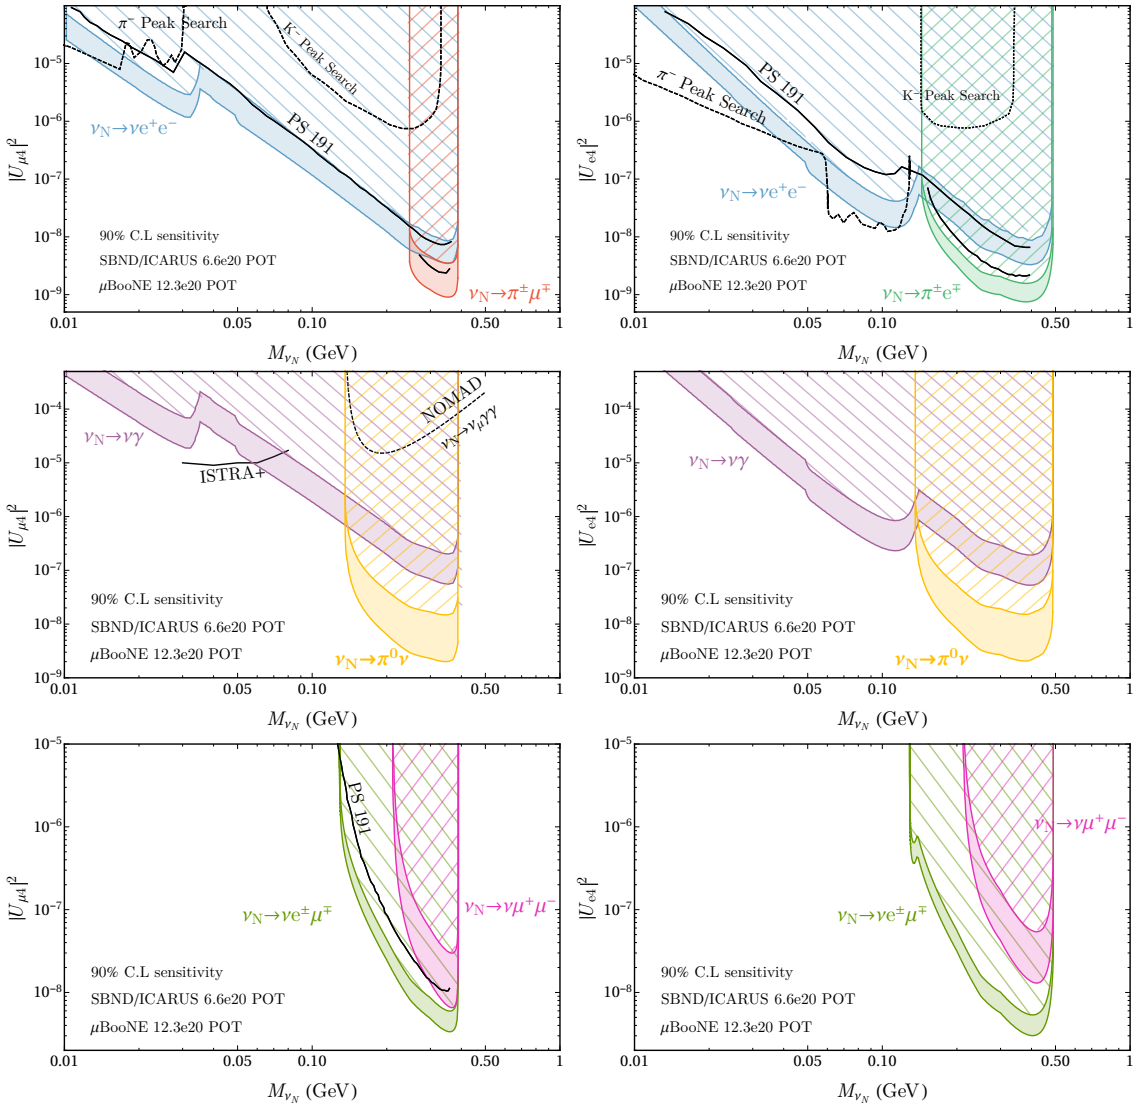

Figure 4.14: The predicted 90% C.L upper limit contours for the combined SBN detectors. Shown also in black solid lines is the prior best bounds from PS-191, scaled to show bounds on the minimal extension as discussed here, as well as bounds from lepton peak searches in pion and kaon decay [214,215] (dashed black lines), although kaon peak searches are not competitive in the minimal model. Note that the peak searches are only valid when bounding pure mixing combinations, *e.g.* $|U_{\mu 4}|^2$ and not $|U_{\mu 4}||U_{\tau 4}|$. The photonic channels have little or no direct bounds, with ISTRA+ bounding the radiative decay [229] and reinterpreted $N \to \nu\gamma\gamma$ bounds at NOMAD on $N \to \nu\pi^0$ [230]. In all panels, the mixing matrix elements not shown on the y-axis are zero.

to improve upon the π^- peak search bounds for $m_N \leqslant 0.033$ GeV and $m_N \leqslant 0.138$ GeV for muon and electron mixing respectively, if the backgrounds can be further suppressed, possibly through the use of timing information.

Additionally, we show that the previously poorly bounded photonic-like channels $N \to \nu_\alpha \pi^0$ and $N \to \nu_\alpha \gamma$ can be probed across the entire parameter space, providing new constraints on exotic sterile neutrino signatures. The potential beam-related backgrounds are large for these photonic channels, the effect of which is a much wider separation between our cut-based limits and the optimal ones. These photonic channels allow SBN to probe the electromagnetic nature of the sterile neutrino, placing bounds on any models containing enhanced couplings to photons. For sterile neutrinos whose mass lies $m_\pi^0 \leqslant m_N \leqslant m_\pi^\pm + m_\mu$ and mix primarily with muons, the $N \to \nu_\mu \pi^0$ channel can extend the limits beyond that of the traditional e^+e^- searches to probe new parameter space, even in the purely minimal model. For sizeable U_{e4}, the π^0 bounds are less powerful than that of the semi-leptonic $N \to e^\pm \pi^\mp$ when one assumes the minimal model.

Although we have plotted the limits on mixing angles in Figure (4.14) in terms of the parameters of the minimal model, they are model independent in the sense that an enhanced decay rate in that channel would only alter the interpretation of the y-axis. If the enhancements to the decay rates are modest, to reinterpret any bounds on Figure (4.14) in the context of a non-minimal extension in which the channel of interest is enhanced by $(1 + \alpha)$ then the quantity bounded on each vertical axis is given approximately by $|U_{\alpha 4}|^2/\sqrt{1+\alpha}$ as discussed in Section (4.2.3).

However, for larger enhancements, the lower-bound on the mixing-matrix element must also be considered. In Figure (4.14), this bound lies at large values of $|U|^2$, and is not shown in the plots, but it is also affected by an enhanced decay rate and can become relevant of reasonable enhancements. This can be seen in Figure (4.15), where we show the region of parameter space that SBN could exclude when studying the decay mode $N \to \nu e^+ e^-$ as we increase its decay rate by

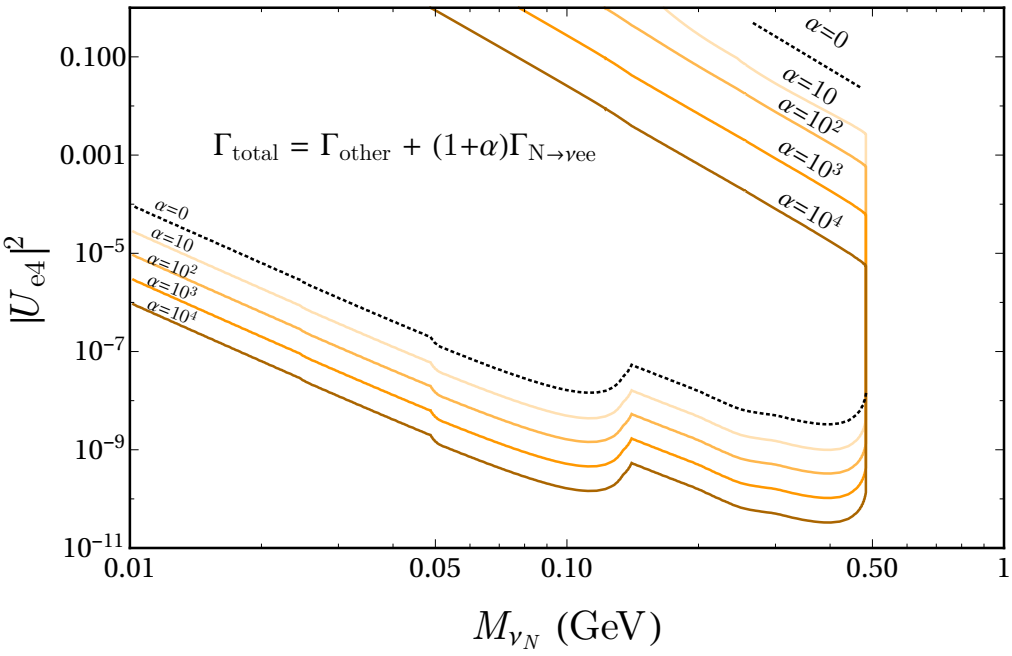

Figure 4.15: The 90% C.L contours for the decay $N \to \nu e^+ e^-$ assuming dominant U_{e4} mixing in non-minimal scenarios, at leading order in λ/L. We enhance the rate of the e^+e^- channel itself by factors shown in the labels while keeping all other decay rates constant. The excluded region remains roughly constant but shifts downwards.

factors of $10, 10^2, 10^3$ and 10^4. As was shown analytically in Section (4.2.3), the upper bound scales as $1/\sqrt{1+\alpha}$ as the number of events in the detector increases. However, the enhancement eventually leads to significant beam attenuation before the detector. This alters the lower bound, which begins to move to smaller values of the mixing-matrix element, opening up a region of parameter space in the top-right of the plot.

In Figure (4.16), we show an alternative non-minimal model in which the decay rate $\Gamma(N \to \nu e^+ e^-)$ is held constant, but the total decay rate is enhanced. This could be due to the enhancement of a decay to visible or invisible final states. In this scenario, the upper bound remains unchanged as the rate is enhanced (to leading order in λ/L), but the enhanced total decay rate leads to beam attenuation and fewer sterile neutrinos reach the detector. Eventually, the lower bound is reduced significantly, and the experiment loses sensitivity over much of the parameter space

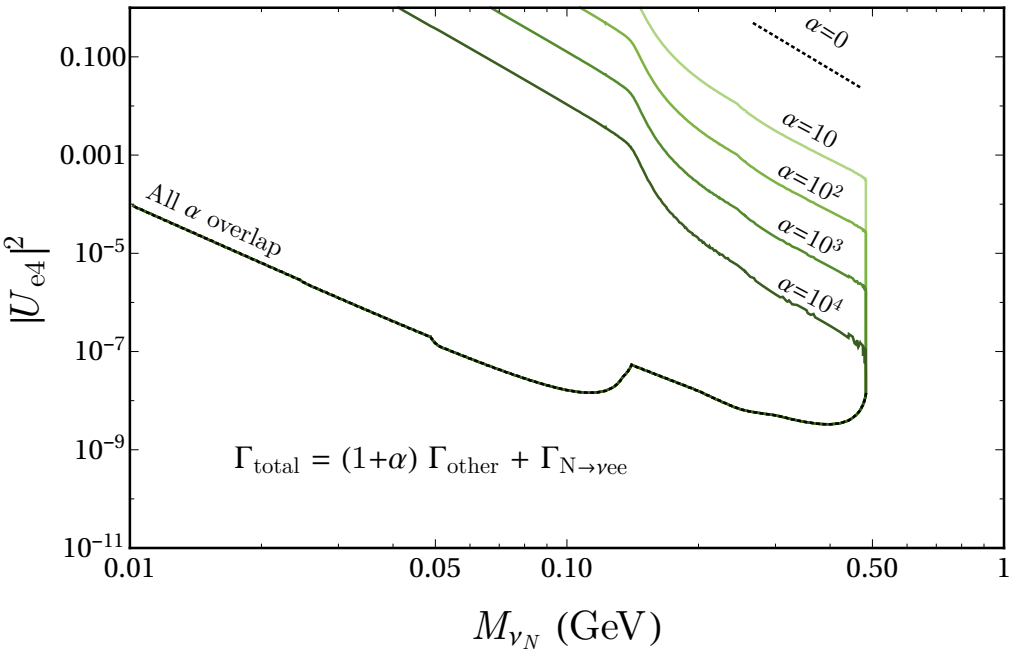

Figure 4.16: The same as Figure (4.15), except where we show the effects of keeping the $N \to \nu e^+ e^-$ decay rate constant, but enhancing the total decay rate. The sensitive region shrinks quickly as α increases, allowing non-minimal models to escape detection. Note that for large values of α, heavy sterile neutrinos with relatively large mixing, $|U_{e4}|^2 \approx 10^{-2} - 10^{-4}$, are no longer bounded by beam dump experiments.

of the minimal-model. Although an enhanced total decay rate could produce a larger visible signal in another channel, or indeed in another experiment, if the decay is predominantly to three neutrinos or dark sector particles many existing bounds may not apply. We note that enhancements on the scale of $\alpha = 10^4$ could be expected if the novel decay proceeds without mixing suppression. Every bound presented in Figure (4.15) can be reinterpreted in terms of these non-minimal models using the scalings as discussed in Section (4.2.3), and highlights why searching across the whole parameter space is necessary in all kinematically allowed decay channels.

4.4.2 Timing information to study an observed signal

In addition to being able to reduce beam-related backgrounds, a precise knowledge of the timing of any observed events can also be used to discriminate between

potential models and aid parameter estimation. If a potential signal is observed, it would be highly desirable to establish whether the excess is associated with a heavy particle travelling from source to detector. An analysis based on the energy spectrum alone would struggle with this determination — we could not discount mis-understood beam-related backgrounds, unknown nuclear effects, or other models that mimic the low-energy spectrum. The angular distribution of events would also be highly informative, we have seen that heavy particle decays are likely to be associated with collimated decay products, but this would be only indirect evidence of a heavy particle, and could be associated with other models. For example, active neutrino scattering via a light mediator could also mimic this behaviour. However, as all beam-related backgrounds will be correlated with the Booster proton buckets, the observation of events with times outside of the BNB beam bucket window (and travelling in a forward direction) would be a smoking gun signal of a sub-luminal propagating parent.

We estimate the required timing resolution by simulating the distribution of arrival times for a given sterile neutrino mass. We then compute the compatibility of this data with a beam-bucket hypothesis, where all event timing is consistent (within errors) with being within the beam-bucket. We only study the shapes of these timing distributions, allowing the normalization to float, and in this sense the beam-bucket hypothesis encompasses all sources of particles which would appear beam-correlated. The beam-bucket hypothesis is defined as the assumption that all events originate in a 6 ns window surrounding the BNB beam spill, smeared by a Gaussian with a width of the assumed time resolution. We define our test statistic as [23]

$$t_m = -2\ln\left(\mathcal{L}\right) = 2\sum_{i=1}^{N}\left\{\mu_i(m) - n_i + n_i \ln\left[\frac{n_i}{\mu_i(m)}\right]\right\},$$

where $\mu_i(m)$ is the expected number of events in bin i if the sterile neutrino is of mass m. Using this statistic we have run a binned Maximum Likelihood analysis of the reconstructed time of arrival ΔT, assuming events are Poisson distributed. We compute the distribution for t_m by Monte Carlo to ensure that we account for all non-gaussianity in the likelihood function.

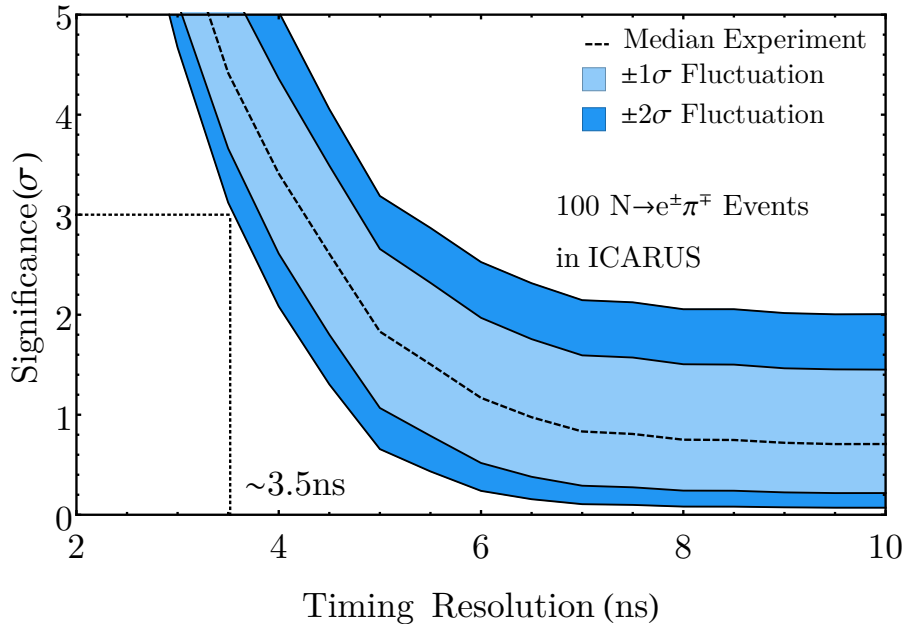

Figure 4.17: Expected significance at which ICARUS can exclude a beam-correlated origin from timing information alone, as a function of assumed timing resolution. This assumes a hypothetical signal of 100 $e^{\pm}\pi^{\mp}$ events consistent with $N \to e^{\pm}\pi^{\mp}$. For larger time resolutions, $\geqslant$ 6ns, although one could achieve a 2σ measurement, ICARUS would not be able to confirm that the events came from the inter-bucket window.

As the timing is solely a function of the initial sterile neutrino energy and mass, these results hold for all channels studied. Without loss of generality, we restrict our discussion to the semi-leptonic channel $N \to e^{\pm}\pi^{\mp}$. In Figure (4.17), we show the timing resolution required to exclude the beam-bucket hypothesis at a given statistical significance. This plot assumes that ICARUS has observed an excess of 100 events due to a 300 MeV sterile neutrino. To guarantee that ICARUS can reject the beam-bucket hypothesis at least 3σ significance in 95% of pseudo-experiments, we require a timing resolution of $\leqslant$ 3.5 ns.

If we relax this simulation and vary the signal strengths we see that as the number of observed events increases, the timing resolution required to rule out a beam-correlated origin decreases, as we show in Figure (4.18). Thus, even if only

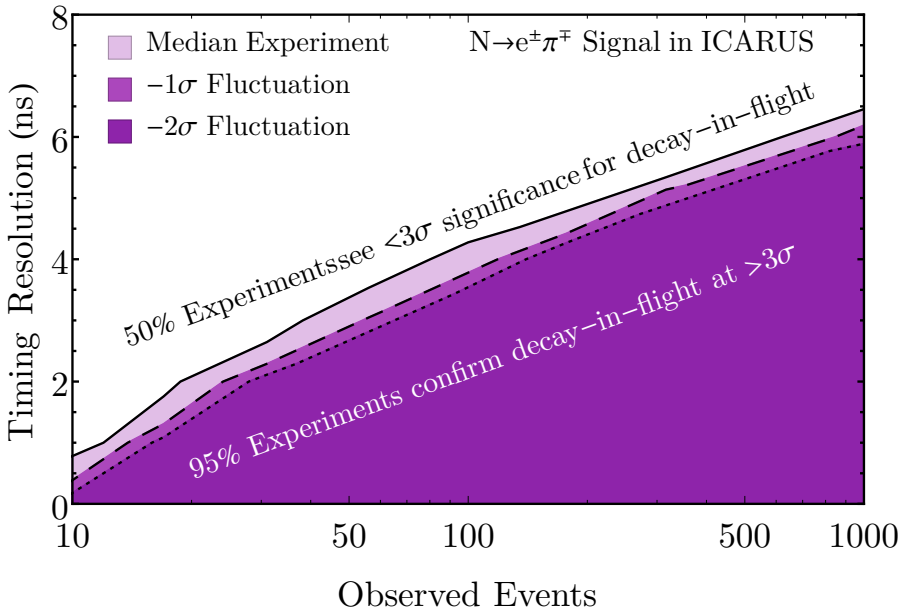

Figure 4.18: As the number of observed events goes down, it is significantly harder to establish a time delay. We show the timing resolution required for a given number of events for the median experiment (solid line) as well as for 1σ (dashed line) and 2σ (dotted line) downward fluctuations.

$\approx$ 6 ns timing resolution is achieved, any anomalous signal would prompt ICARUS to continue running until it can achieve 1000 events, thereby granting the same sensitivity to the temporal nature of the signal.

Although establishing that a signal arrived outside of the beam window would be an exciting sign of new physics, it would not necessarily establish a heavy propagating parent. For example, if an unaccounted for process had a fixed time delay with respect to the neutrino beam, Δ_t, such as the relaxation time of an excited atom, it could produce events in the inter-bucket region for $\Delta_t \approx \mathcal{O}(\text{ns})$. Similarly, other exotic BSM physics could be the source a fixed time decay signature without relying on a heavy propagating sterile neutrino. The scenario described in Ref. [209, 210] is one such case, it considers a sterile neutrino produced inside the detector through neutral current scattering of an active neutrino. The heavy particle promptly decays, with a decay length of the order $\mathcal{O}(1)$ m, producing the visible signal. Although the sterile neutrinos are produced inside the detector with no timing delay from active-neutrino scattering, the finite lifetime of these particles

could lead to events falling in the inter-bucket window.

However, in both the BSM scenarios as well as generic backgrounds with a fixed timing delay, the temporal spectra of event arrival time would be expected to be constant across all three detectors. The SBN program is perfectly designed to account for this, however, through its multiple detectors at different baselines, as if the excess is indeed due to heavy particle propagation, then the sterile neutrino would have to travel further to reach each subsequent detector. This leads to observable shifts in the arrival timing spectra at each experiment. In particular ICARUS, would be most suited to studying heavy particle propagation, as particles must travel approximately 6 times further before detection.

We show the consequences of this effect in Figure (4.19) where the ratio of events at SBND and ICARUS are plotted as a function of time delay. A constant time delay would produce a ratio of unity, and curves that lie on the grey circle. We see a clear distortion in this ratio, with a generally low value inside the beam-bucket window and a larger value outside. Measuring this distortion would be definitive proof of the heavy particle having propagated the distance from target to detector and not merely being produced *in situ*. In figure Figure (4.20), we show how the attainable timing resolution affects this measurement. For a resolution of 10 ns, there is no spectral difference, but distortion starts to be apparent for resolutions better than 10 ns .

Assuming a positive signal is found and is identified as a heavy sterile neutrino decay thanks to the time delay, the temporal and energy analyses could be used to measure the heavy sterile neutrino mass. For an arrival time delay (behind a luminal or near luminal particle) over a distance L denoted by ΔT, the mass of a sterile neutrino with an energy E can be reconstructed as

$$m_N = E\sqrt{1 - \frac{1}{\left(1 + \frac{c\Delta T}{L}\right)^2}}.$$

Exact knowledge of the deposited energy and time of flight would be sufficient to

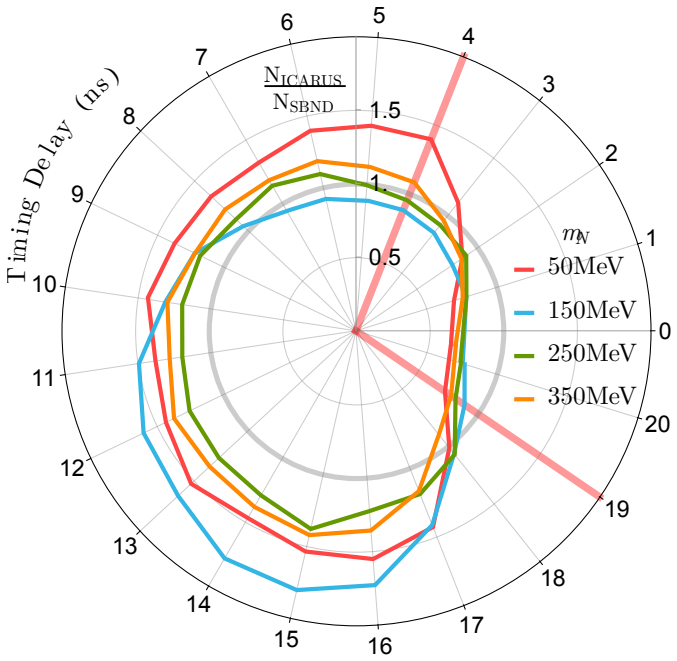

Figure 4.19: Ratio of arrival time spectra of $N \to \nu_\alpha e^+ e^-$ events in ICARUS to SBND after scaling out $1/R^2$ flux dependence. If events were due to an unknown background with a fixed time delay after the neutrino beam, one would expect the ratio to be a constant value of 1 (shaded grey ring). As the sterile neutrinos have to travel approximately 6 times further to ICARUS than SBND, increasingly higher energy sterile neutrinos can leave the beam-bucket (red arc) and populate the inter-bucket region leading to the distinct signature. This assumes a timing resolution of 1 ns.

establish the mass, but of course these data are in most cases not available: energy and timing resolution impair the reconstruction, and many channels have missing energy from active neutrinos in the final state. Moreover, due to the cyclic nature of the BNB beam buckets, an observed event could have originated from any of the previous buckets in the current spill, and not just the one closest to the tagged event timing. As such the absolute time of flight is not known. Only the relative timing since the last bucket, ΔT, is known and from this one can obtain up to 81 degenerate solutions for the sterile neutrino mass. Although absolute timing information could be found by studying the first few buckets for the onset of a signal, this would rely

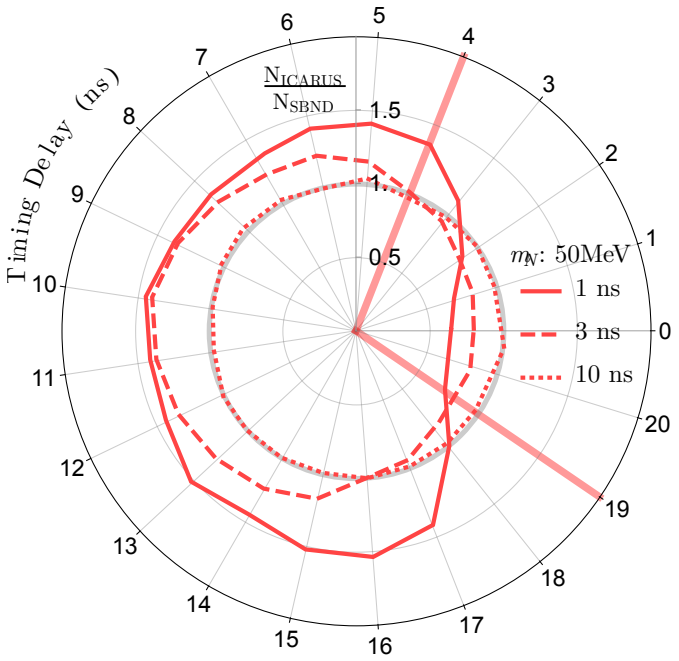

Figure 4.20: The same as left for fixed 50 MeV sterile neutrino with 1 ns (solid), 3 ns (dashed) and 10 ns (dotted) timing resolution showing the decreasing effect on the ratio. As can be seen for the 10ns resolution curve, once the timing resolution is comparable to the BNB frequency then naturally no affect can be seen, and the temporal ratio between SBND and ICARUS is flat

on precise absolute timing measurements between source and detector, and would also reduce the signal statistics by $\mathcal{O}(0.01)$ and we do not consider this information in the analysis. Given these limitations, we have studied how well m_N could be reconstructed, using then energy and periodic time since last bucket ΔT. We have generated Monte Carlo event data tagging each event by an arrival time, accounting for a systematic uncertainties on the time and energy measurement. We smear the true energy to represent detector effects as described above, and additionally smear the time of each event with a Gaussian of width $\sigma_T \approx 1$ ns for SBND and ICARUS. We use the same Monte Carlo analysis and test statistic as in the temporal analysis above, expanded to include a binned energy spectra. The reconstructed mass is defined as the mass which minimises the test statistic t_m.

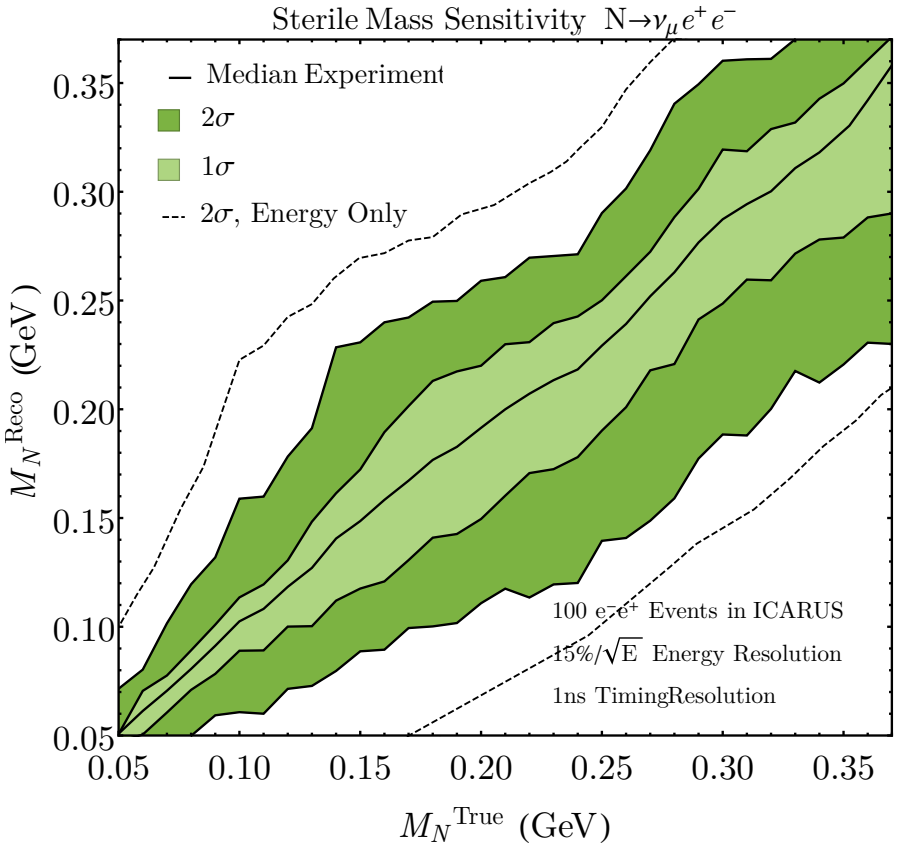

Figure 4.21: The reconstructed sterile neutrino mass as a function of true mass for energy only (dashed) and energy with timing information (solid). This assumes the heavy sterile neutrino decays to $N \to \nu e^+ e^-$ and that 100 events consistent with a decay-in-flight are seen at ICARUS. A 1 ns timing resolution has been assumed to have been is achieved.

The results of our analysis are shown in Figures 4.21 and 4.22. In both figures, we show the allowed region in reconstructed mass as a function of true sterile neutrino mass for an energy only analysis (dashed black lines), as well as for an energy and time-of-flight analysis (coloured bands). Figure (4.21) we show the results for the fully leptonic decay $N \to \nu e^+ e^-$ while Figure (4.22) shows our results for the semi-leptonic $N \to e^\pm \pi^\mp$ channel. In the case of the 2-body $N \to e^\pm \pi^\mp$ channel, resolution of approx 45 MeV at 2σ level is achievable for the entire range of sterile neutrino mass allowed. We estimate the $N \to \mu^- \pi^+$ channel would be approximately 10% better due to the improved energy resolution possible when reconstructing muons in LAr. For these semi-leptonic decays the energy spectrum

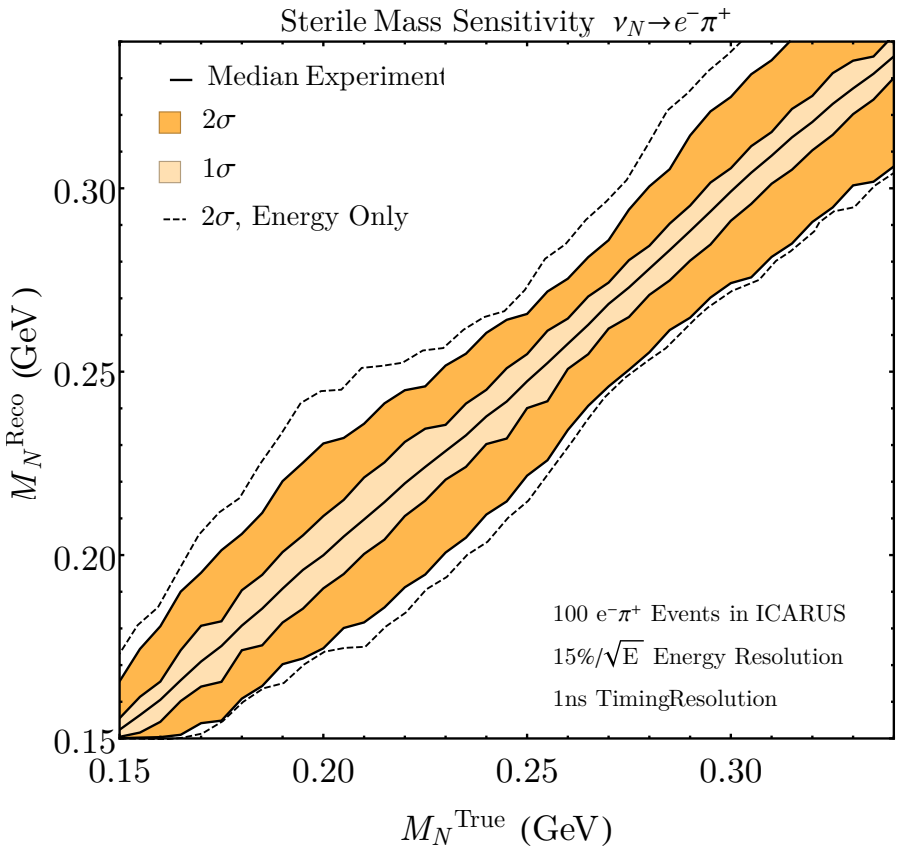

Figure 4.22: The same as Figure (4.21) but for heavy sterile decays to $e^{\pm}\pi^{\mp}$. Note that as there is no missing energy in this channel, the mass resolution is significantly better and less is gained from the addition of timing information.

is very informative, as the parent particle's energy can be reconstructed from the invariant mass of the decay products' four-momenta. As such, we see temporal information only trivially improves the reconstruction of parent mass. In contrast, for the 3-body $N \to \nu e^+ e^-$ channel, there is significant missing energy taken away by the active neutrino. In this case, timing information is much more valuable, almost halving the 2σ mass range from around 300 MeV to 150 MeV for widest region of parameter space.

4.5 Conclusions

In this Chapter, we have studied the prospects for the measurement of MeV-scale sterile neutrinos at the Fermilab Short-Baseline Neutrino program. MeV-scale neu-

tral states would naturally be produced in the Booster Neutrino Beam from mixing-suppressed meson decays. To assess SBN's potential to constrain these models, we have estimated the dominant backgrounds and signals. Thanks to excellent particle identification and the distinctive kinematic properties of our signal, high levels of background suppression can be expected, allowing SBN to improve on the current bounds on decaying sterile neutrinos over most of the parameter space.

We have shown that, in the absence of signal, SBN can place bounds on the active-sterile mixing-matrix elements of $|U_{e4}| \leqslant 10^{-6}$ for $m_N \leqslant 33$ MeV and $|U_{\mu 4}|^2 \leqslant 2 \times 10^{-8}$ for $m_N \leqslant 138$ MeV in the $N \to \nu e^+ e^-$ channels. For semi-leptonic decays, these bounds increase up to $|U_{e4}| \leqslant 8 \times 10^{-10}$ for $m_N \leqslant 388$ MeV and up to $|U_{\mu 4}|^2 \leqslant 7 \times 10^{-10}$ for $m_N \leqslant 493$ MeV. The neutral pion decay channel, $N \to \nu\pi^0$, which may be the dominant decay mode for masses in the range $m_\pi^0 \leqslant m_N \leqslant m_{\pi^\pm} + m_\mu$, can be used to place bounds of around $|U_{\alpha 4}|^2 \leqslant 3 \times 10^{-9}$. These can be expected to be the most stringent bounds placed by upcoming experiments in the region of parameter space we study here. In particular, although significantly more sensitive for sterile neutrino masses above the kaon mass, SHiP produces weaker constraints than PS-191 for the case of a single sterile neutrino produced via kaon decay [198]. Any improvements in the bounds by SBN at lower sterile neutrino masses, $m_N \leqslant m_K$, will therefore complement the strong bounds placed by SHiP at higher masses.[6]

We have also discussed searches for non-minimal models of heavy sterile neutrino decay, which could lead to observable decays over a wide range of parameter space which is conventionally excluded if the branching ratios are assumed to arise from

[6] A purpose built SHiP-style experiment situated near the beam dump of a lower energy beam such as BNB could improve on this. Simple scaling arguments like those in Table (4.1) predict event rates around 10^4 times larger than PS-191, even for steriles $m_N \leqslant m_K$. This is not unexpected as SHiP consists of a huge decay volume designed to search for such particle decays. Such a simple scaling, however, does not take into consideration the complexity of working near an active beam dump.

the minimal model. We have shown how to map existing minimal-model bounds onto non-minimal models and how bounds could be weakened in the case of specific enhancements to a decay channel. This motivates the search for particle decays in all channels over the full parameter space. We argue that some of these decay channels considered in this work are in fact poorly constrained by similar experiments, and show that SBN could place the first direct bounds on these processes.

We have commented in detail on the phenomenological role of timing information in this analysis. As well as providing a means of background suppression, nanosecond scale timing resolution could allow SBN to make direct measurements of the kinematic properties of heavy particle propagation. We have shown that if 100 events are seen at ICARUS, a 3.5 ns timing resolution would allow an observable timing delay to be established at 3σ in 95% of experiments. We have seen that the unique design of SBN would allow for the distribution of event times to be compared between the nearest and farthest detectors, allowing for any model with a finite time delay to be excluded when compared to a propagating sterile neutrino model. We have also shown how timing information can be used in sterile neutrino mass reconstruction. For the decay $N \to \nu e^+ e^-$, the inclusion of event timing information (with an assumed 1 ns resolution) can lead to the 2σ allowed region being reduced from around ± 300 MeV to ± 150 MeV.

We point out that this analysis is complementary to the central physics programme of SBN — studying eV-scale oscillating sterile neutrinos — and requires no additional detector or beam modifications. We have shown that SBN could contribute valuably to the search for sterile neutrino decays-in-flight, and moreover, if an anomalous signal is discovered, would play a central role in determining its origin, and the necessary extension of the Standard Model.

Chapter 5

Summary

"There is no real ending.
It's just the place where you stop the story."
Frank Herbert

The history of the Electro-weak sector of the standard model, including the very existence of the neutrino, has been rife with controversial, unexpected and era-defining discoveries. In particular, in the last few decades the implications of non-zero neutrino masses have had a large impact on the neutrino community; from a theoretical point of view the mechanism by which they are generated is unknown, and experimentally, it has allowed us to deeply probe the neutrino sector through neutrino flavour oscillations. This trend of discoveries is set to continue forward, with the first tantalising hints of CP violation in the lepton sector appearing from accelerator and reactor measurements, along with plans well underway to discover the neutrino mass ordering and θ_{23} octant.

In this thesis we have explored three distinct signatures of new low-scale physics that would conclusively prove that we do not live in a 3ν paradigm. These were directly measuring the 3 U_{PMNS} neutrino mixing matrix for signs of non-unitarity, looking for additional frequencies (corresponding to new mass-splittings) in neutrino flavour oscillations, and finally, searching for the direct decay products of heavier sterile neutrinos decaying into Standard Model particles.

The success of the Standard Model alongside the 3ν paradigm can correctly ex-

plain the vast majority of experimental observations in the neutrino sector, and as such it is often claimed we are in the *precision era* of neutrino physics. However, it must be noted that it is necessary to invoke the assumption of unitarity in order to provide indirect bounds on many of the 3×3 U_{PMNS} mixing matrix elements. When this stringent assumption is relaxed, we have shown in Chapter 2, the bounds on individual U_{PMNS} elements increase dramatically. Although we find no evidence for non-unitarity we must note there is significant room for new physics in the neutrino sector, especially when one considers mixing with the tau neutrino. The canonical way in which one can introduce non-unitarity in the 3×3 neutrino mixing matrix is though the addition of neutral singlet states of the Standard Model that mix with the active neutrinos. From a purely theoretical perspective, the addition of multiple sterile degrees of freedom is very well motivated and a common feature in Beyond the Standard Model scenario, as they can generate masses for the active neutrinos, as well as naturally explain their smallness. Two crucial concepts that the Standard Model as of yet has no explanation for. In addition to this, several anomalous experimental measurements suggest the existence of a light, but still massive, sterile neutrino of $\mathcal{O}(\text{eV}^2)$, which is in strict disagreement with the 3ν paradigm.

Sterile neutrinos, whose mass can lie at a huge variety of scales, have thus attracted great attention and effort in the international neutrino community, and a plethora of successful past, present and future experiments have arisen to search for signs of them in the neutrino sector. To this end, Fermilab's Short-Baseline Neutrino program was proposed to study the possible existence of the $\mathcal{O}(\text{ eV}^2)$ light oscillating sterile neutrino that would explain the LSND anomaly. Consisting of three LArTPC detectors at different baselines, SBND, MicroBooNE and ICARUS, the SBN program will have unprecedented sensitivity to neutrino flavour oscillations at the $0.01 - 100$ eV^2 mass scales. SBN will be able to probe the (3+1), (3+2) and (3+3) sterile neutrino scenarios, ruling out the vast majority of parameter space in the event of observations consistent with the null hypothesis. On the contrary, if the existence of one or more light sterile neutrinos are indeed confirmed by SBN, attention will turn to precision measurements of these new sterile states. In the case

of multiple light sterile neutrinos SBN has the potential to be the first experiment to directly observe CP violation in the lepton sector, induced by the (3+2) sterile neutrino scenario parameter ϕ_{54}.

Although designed primarily with the observation of such light oscillating sterile neutrino in mind, the precision LArTPC detectors that make up SBN will also be sensitive to significantly heavier $\mathcal{O}(100$ MeV) sterile neutrinos that would be produced alongside the active neutrinos in standard meson decay. These heavy sterile neutrinos do not contribute to neutrino oscillations, but rather they are directly observed through completely different phenomenology; their subsequent decay to electromagnetically charged Standard Model particles. By utilising the dramatic kinematic differences between neutrino scattering events, now a background, and such decay-in-flight signals, we show in Chapter 4 that SBN can improve upon the previous best bounds for sterile neutrinos below ≈ 500 MeV, in all channels studied. This is directly complementary to the planned experiment SHiP which would place very strong constraints on heavier, $m_N \geqslant$ GeV, sterile neutrinos produced in D-meson decays. SBN is extremely suited to studying this class of sterile neutrinos by taking full advantage of the multi-baseline setup to study the sub-luminal time of flight of these very heavy neutrinos. If sufficient timing resolution can achieved, $\leqslant 3.5ns$, we point out that such a time of flight analysis could lead to a "smoking-gun" signature of decay-in-flight sterile neutrinos.

We are poised at the beginning of a fascinating time for neutrino physics, with either the conformation of the LSND anomaly as a sterile phenomena, the first particle beyond the Standard Model and at the same time ruling out the 3ν paradigm, or the exclusion of the light oscillating neutrino hypothesis. In this thesis we have used three approaches to inform us about low-scale sterile neutrinos, their effect on unitarity, their effect on neutrino flavour oscillations and their decay to Standard Model particles. Measuring any one of these would invalidate the 3ν paradigm completely. We emphasise that SBN is a perfect example of a multi-faceted approach to sterile neutrino phenomenology, allowing us to cover a wider range of potential

sterile neutrino parameter space within the same experimental arrangement. This is all the while undergoing intense R&D and commissioning of LArTPC detectors, pioneering the technology that the next generation long baseline experiments, DUNE, will need to have a full and comprehensive understanding of in order to probe the neutrino sector as in depth as possible. It is during this crucial time of R&D that phenomenological studies, such as those presented in this thesis, are of greatest use to the neutrino community, increasing the physics goals of experiments, highlighting possible modifications, enhancements or suggesting new search strategies entirely, all with the end goal of exploring just what other surprises in the neutrino sector our Universe has for us to discover.

Bibliography

[1] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934.

[2] F. L. Wilson. Fermi's theory of beta decay. *American Journal of Physics*, 36(12):1150–1160, 1968.

[3] K. C. Wang. A suggestion on the detection of the neutrino. *Phys. Rev.*, 61:97–97, Jan 1942.

[4] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.

[5] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104:254–258, Oct 1956.

[6] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.

[7] G. Danby, J-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962.

[8] S. L. Glashow. Partial Symmetries of Weak Interactions. *Nucl. Phys.*, 22:579–588, 1961.

[9] F. Englert and R. Brout. Broken symmetry and the mass of gauge vector mesons. *Phys. Rev. Lett.*, 13:321–323, Aug 1964.

[10] P.W. Higgs. Broken symmetries, massless particles and gauge fields. *Physics Letters*, 12(2):132 – 133, 1964.

[11] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global conservation laws and massless particles. *Phys. Rev. Lett.*, 13:585–587, Nov 1964.

[12] S. Weinberg. A mode l of leptons. *Phys. Rev. Lett.*, 19:1264–1266, Nov 1967.

[13] A. Salam. Weak and Electromagnetic Interactions. *Conf. Proc.*, C680519:367–377, 1968.

[14] F.J. Hasert et al. Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment. *Nuclear Physics B*, 73(1):1 – 22, 1974.

[15] M. Banner et al. Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the cern pp collider. *Physics Letters B*, 122(5):476 – 485, 1983.

[16] G. Arnison et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at s=540 gev. *Physics Letters B*, 122(1):103 – 116, 1983.

[17] G. Aad et al. Combined Measurement of the Higgs Boson Mass in *pp* Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments. *Phys. Rev. Lett.*, 114:191803, 2015.

[18] M. L. Perl et al. Evidence for anomalous lepton production in $e^+ - e^-$ annihilation. *Phys. Rev. Lett.*, 35:1489–1492, Dec 1975.

[19] K. Kodama et al. Observation of tau neutrino interactions. *Phys. Lett.*, B504:218–224, 2001.

[20] S. Abachi et al. Search for high mass top quark production in $p\overline{p}$ collisions at $\sqrt{s} =$ 1.8 tev. *Phys. Rev. Lett.*, 74:2422–2426, Mar 1995.

[21] F. Abe et al. Observation of top quark production in $\overline{p}p$ collisions with the collider detector at fermilab. *Phys. Rev. Lett.*, 74:2626–2631, Apr 1995.

[22] S. W. Herb, D. C. Hom, L. M. Lederman, J. C. Sens, H. D. Snyder, J. K. Yoh, J. A. Appel, B. C. Brown, C. N. Brown, W. R. Innes, K. Ueno, T. Yamanouchi, A. S. Ito, H. Jöstlein, D. M. Kaplan, and R. D. Kephart. Observation of a dimuon resonance at 9.5 gev in 400-gev proton-nucleus collisions. *Phys. Rev. Lett.*, 39:252–255, Aug 1977.

[23] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014.

[24] J. A. Formaggio and G. P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev. Mod. Phys.*, 84:1307–1341, 2012.

[25] E. Majorana. Teoria simmetrica dell'elettrone e del positrone. *Il Nuovo Cimento (1924-1942)*, 14(4):171, 2008.

[26] S. Weinberg. Baryon- and lepton-nonconserving processes. *Phys. Rev. Lett.*, 43:1566–1570, Nov 1979.

[27] S. Weinberg. Varieties of baryon and lepton nonconservation. *Phys. Rev. D*, 22:1694–1700, Oct 1980.

[28] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958. [Zh. Eksp. Teor. Fiz.34,247(1957)].

[29] B. Pontecorvo. Mesonium and anti-mesonium. *Sov. Phys. JETP*, 6:429, 1957. [Zh. Eksp. Teor. Fiz.33,549(1957)].

[30] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870, 1962.

[31] C. Giunti and C. W. Kim. Coherence of neutrino oscillations in the wave packet approach. *Phys. Rev.*, D58:017301, 1998.

[32] D. Hernandez and A. Yu. Smirnov. Active to sterile neutrino oscillations: Coherence and MINOS results. *Phys. Lett.*, B706:360–366, 2012.

[33] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys.Rev.Lett.*, 81:1562–1567, 1998.

[34] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz. Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity. *JHEP*, 01:087, 2017.

[35] D. V. Forero, M. Tortola, and J. W. F. Valle. Neutrino oscillations refitted. *Phys. Rev.*, D90(9):093006, 2014.

[36] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

[37] G. Bellini, A. Ianni, L. Ludhova, F. Mantovani, and W. F. McDonough. Geo-neutrinos. *Progress in Particle and Nuclear Physics*, 73:1–34, November 2013.

[38] M. Schwartz. Feasibility of using high-energy neutrinos to study the weak interactions. *Phys. Rev. Lett.*, 4:306–307, Mar 1960.

[39] B. Pontecorvo.

[40] K. Lande, B. T. Cleveland, T. Daily, R. Davis, J. Distel, C. K. Lee, A. Weinberger, P. Wildenhain, and J. Ullman. Solar neutrino observations with the Homestake ^{37}Cl detector. *AIP Conf. Proc.*, 243:1122–1133, 1992.

[41] W Hampel et al. Final results of the 51cr neutrino source experiments in gallex. *Physics Letters B*, 420(1–2):114 – 126, 1998.

[42] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. iii. results for the 2002˘2007 data-taking period. *Phys. Rev. C*, 80:015807, Jul 2009.

[43] B. Aharmim et al. Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory. *Phys. Rev.*, C88:025501, 2013.

[44] B. Pontecorvo.

[45] S. P. Mikheev and A. Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917, 1985. [Yad. Fiz.42,1441(1985)].

[46] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, May 1978.

[47] P. Adamson et al. Measurement of the neutrino mixing angle θ_{23} in NOvA. *Submitted to: Phys. Rev. Lett.*, 2017.

[48] K. Abe et al. First combined analysis of neutrino and antineutrino oscillations at T2K. 2017.

[49] P. Adamson et al. Combined analysis of ν_μ disappearance and $\nu_\mu \rightarrow \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos. *Phys. Rev. Lett.*, 112:191801, 2014.

[50] M. Apollonio et al. Limits on neutrino oscillations from the CHOOZ experiment. *Phys. Lett.*, B466:415–430, 1999.

[51] P. F. Harrison, D. H. Perkins, and W. G. Scott. Tri-bimaximal mixing and the neutrino oscillation data. *Phys. Lett.*, B530:167, 2002.

[52] H. Fritzsch and Z. Z. Xing. Lepton mass hierarchy and neutrino oscillations. *Physics Letters B*, 372(3):265 – 270, 1996.

[53] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[54] X. Guo et al. A Precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya-Bay. 2007.

[55] Y. Abe et al. Reactor electron antineutrino disappearance in the Double Chooz experiment. *Phys. Rev.*, D86:052008, 2012.

[56] J. K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, 2012.

[57] F. P. An et al. Measurement of electron antineutrino oscillation based on 1230 days of operation of the Daya Bay experiment. 2016.

[58] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.

[59] K. Abe et al. Letter of Intent: The Hyper-Kamiokande Experiment; Detector Design and Physics Potential —. 2011.

[60] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2015.

[61] K. Abe et al. Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6e20 protons on target. *Phys.Rev.*, D91(7):072010, 2015.

[62] P. Adamson et al. Electron neutrino and antineutrino appearance in the full MINOS data sample. *Phys.Rev.Lett.*, 110(17):171801, 2013.

[63] R. Wendell et al. Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II, and III. *Phys. Rev.*, D81:092004, 2010.

[64] Alexandre B. Sousa. First MINOS+ Data and New Results from MINOS. *AIP Conf. Proc.*, 1666:110004, 2015.

[65] K. Abe et al. Precise Measurement of the Neutrino Mixing Parameter theta23 from Muon Neutrino Disappearance in an Off-Axis Beam. *Phys. Rev. Lett.*, 112(18):181801, 2014.

[66] X. Qian, C. Zhang, M. Diwan, and P. Vogel. Unitarity Tests of the Neutrino Mixing Matrix. 2013.

[67] M.C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz. Updated fit to three neutrino mixing: status of leptonic CP violation. 2014.

[68] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, A. Palazzo, and A. M. Rotunno. Global analysis of neutrino masses, mixings and phases: entering the era of leptonic CP violation searches. *Phys. Rev.*, D86:013012, 2012.

[69] S. Antusch, C. Biggio, E. Fernandez-Martinez, M.B. Gavela, and J. Lopez-Pavon. Unitarity of the Leptonic Mixing Matrix. *JHEP*, 0610:084, 2006.

[70] S. Antusch and O. Fischer. Non-unitarity of the leptonic mixing matrix: Present bounds and future sensitivities. 2014.

[71] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam. *Phys.Rev.*, D64:112007, 2001.

[72] A.A. Aguilar-Arevalo et al. Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment. *Phys.Rev.Lett.*, 110:161801, 2013.

[73] J.N. Abdurashitov et al. Measurement of the response of the Russian-American gallium experiment to neutrinos from a Cr-51 source. *Phys.Rev.*, C59:2246–2263, 1999.

[74] G. Mention, M. Fechner, Th. Lasserre, Th.A. Mueller, D. Lhuillier, et al. The Reactor Antineutrino Anomaly. *Phys.Rev.*, D83:073006, 2011.

[75] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz. Sterile Neutrino Oscillations: The Global Picture. *JHEP*, 1305:050, 2013.

[76] J. M. Conrad, C. M. Ignarra, G. Karagiorgi, M. H. Shaevitz, and J. Spitz. Sterile Neutrino Fits to Short Baseline Neutrino Oscillation Measurements. *Adv. High Energy Phys.*, 2013:163897, 2013.

[77] F. J. Escrihuela, D. V. Forero, O. G. Miranda, M. Tortola, and J. W. F. Valle. On the description of non-unitary neutrino mixing. *Phys. Rev.*, D92(5):053009, 2015.

[78] Y. Declais et al. Search for neutrino oscillations at 15-meters, 40-meters, and 95-meters from a nuclear power reactor at Bugey. *Nucl. Phys.*, B434:503–534, 1995.

[79] I. E. Stockdale et al. Search for ν_μ and $\bar{\nu}_\mu$ oscillations in the mass range $15 < \delta m^2 < 1000\,\mathrm{eV}^2/c^4$. *Z. Phys.*, C27:53, 1985.

[80] D. Naples et al. A High statistics search for neutrino(e) (anti-neutrino(e)) —¿ neutrino(tau) (anti-neutrino(tau)) oscillations. *Phys. Rev.*, D59:031101, 1999.

[81] A. Romosan et al. A High statistics search for muon-neutrino (anti-muon-neutrino) –¿ electron-neutrino (anti-electron-neutrino) oscillations in the small mixing angle regime. *Phys. Rev. Lett.*, 78:2912–2915, 1997.

[82] K. S. McFarland et al. A Limit on ν_μ (anti-muon-neutrino) $\rightarrow$ ν_τ (anti-tau-neutrino) oscillations from a precision measurement of neutrino - nucleon neutral current interactions. *Phys. Rev. Lett.*, 75:3993–3996, 1995.

[83] F. Dydak et al. A search for ν_μ oscillations in the δm^2 range $0.3 - 90\,\mathrm{eV}^2$. *Phys. Lett.*, B134:281, 1984.

[84] E. Eskut et al. Final results on nu(mu) —¿ nu(tau) oscillation from the CHORUS experiment. *Nucl. Phys.*, B793:326–343, 2008.

[85] F. P. An et al. Improved Measurement of Electron Antineutrino Disappearance at Daya Bay. *Chin. Phys.*, C37:011001, 2013.

[86] F. P. An et al. Search for a Light Sterile Neutrino at Daya Bay. *Phys. Rev. Lett.*, 113:141802, 2014.

[87] M. Antonello et al. Search for anomalies in the ν_e appearance from a ν_μ beam. *Eur.Phys.J.*, C73:2599, 2013.

[88] B. Armbruster et al. Upper limits for neutrino oscillations muon-anti-neutrino —¿ electron-anti-neutrino from muon decay at rest. *Phys. Rev.*, D65:112001, 2002.

[89] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.

[90] A. Gando et al. Constraints on theta13 from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND. *Phys. Rev.*, D83:052002, 2011.

[91] P. Adamson et al. Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys. Rev. Lett.*, 110(25):251801, 2013.

[92] P. Adamson et al. Search for sterile neutrino mixing in the MINOS long baseline experiment. *Phys. Rev.*, D81:052004, 2010.

[93] P. Astier et al. Search for nu(mu) —¿ nu(e) oscillations in the NOMAD experiment. *Phys. Lett.*, B570:19–31, 2003.

[94] P. Astier et al. Final NOMAD results on muon-neutrino —¿ tau-neutrino and electron-neutrino —¿ tau-neutrino oscillations including a new search for tau-neutrino appearance using hadronic tau decays. *Nucl. Phys.*, B611:3–39, 2001.

[95] R Patterson. First oscillation resultsfrom NOvA. Joint Experimental-Theoretical Seminar, Fermilab, August 6th, 2015, 2015.

[96] S. Avvakumov et al. A Search for muon-neutrino —¿ electron-neutrino and muon-anti-neutrino —¿ electron-anti-neutrino oscillations at NuTeV. *Phys. Rev. Lett.*, 89:011804, 2002.

[97] N. Agafonova et al. Discovery of tau neutrino appearance in the CNGS neutrino beam with the OPERA experiment. 2015.

[98] N. Agafonova et al. Limits on muon-neutrino to tau-neutrino oscillations induced by a sterile neutrino state obtained by OPERA at the CNGS beam. *JHEP*, 06:069, 2015.

[99] G. Cheng et al. Dual baseline search for muon antineutrino disappearance at $0.1\text{eV}^2 < \Delta m^2 < 100\text{eV}^2$. *Phys. Rev.*, D86:052009, 2012.

[100] K. Abe et al. Limits on sterile neutrino mixing using atmospheric neutrinos in Super-Kamiokande. *Phys. Rev.*, D91:052019, 2015.

[101] T. Toshito. Super-Kamiokande atmospheric neutrino results. In *36th Rencontres de Moriond on Electroweak Interactions and Unified Theories Les Arcs, France, March 10-17, 2001*, 2001.

[102] Y. Koshio. Solar results from Super-Kamiokande. *AIP Conf. Proc.*, 1666:090001, 2015.

[103] K. Abe et al. Observation of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys. Rev. Lett.*, 112:061802, 2014.

[104] A.C. Hayes, J.L. Friar, G.T. Garvey, G. Jungman, and G. Jonkmans. Systematic Uncertainties in the Analysis of the Reactor Neutrino Anomaly. *Phys.Rev.Lett.*, 112:202501, 2014.

[105] Y. F. Li, J. Cao, Y. Wang, and L. Zhan. Unambiguous Determination of the Neutrino Mass Hierarchy Using Reactor Neutrinos. *Phys.Rev.*, D88:013008, 2013.

[106] S. B. Kim. Talk at International Workshop on RENO-50 toward Neutrino Mass Hierarchy, Seoul, South Korea, 2013.

[107] S. Geer. Neutrino beams from muon storage rings: Characteristics and physics potential. *Phys. Rev. D*, 57:6989–6997, Jun 1998.

[108] M. Antonello et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. 2015.

[109] C. Adams et al. The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe. 2013.

[110] A Himmel. HK Sensitivity to Sterile Neutrinos and Lorentz Violation. Talk at 6th Open Meeting for the HyperK Project, Kavli IPMU, The University of Tokyo, 2015.

[111] S. Schael et al. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006.

[112] W. de Boer Ugo A. and H. Fürstenau. Comparison of grand unified theories with electroweak and strong coupling constants measured at lep. *Physics Letters B*, 260(3):447 – 455, 1991.

[113] C. Giunti, C. W. Kim, and U. W. Lee. Running coupling constants and grand unification models. *Mod. Phys. Lett.*, A6:1745–1755, 1991.

[114] P. Minkowski. mu-¿e gamma at a rate of one out of 109 muon decays? *Physics Letters B*, 67(4):421 – 428, 1977.

[115] T. Yanagida. Workshop on the baryon number of the universe and unified theories, tsukuba, japan.

[116] P. Ramond M. Gell-Mann and R. Slansky. in supergravity, editors f. van nieuwenhuizen and d. freedman. 1979.

[117] R. N. Mohapatra and G. Senjanović. Neutrino mass and spontaneous parity nonconservation. *Phys. Rev. Lett.*, 44:912–915, Apr 1980.

[118] P. C. de Holanda and A. Yu. Smirnov. Homestake result, sterile neutrinos and low-energy solar neutrino experiments. *Phys. Rev.*, D69:113002, 2004.

[119] P. C. de Holanda and A. Yu. Smirnov. Solar neutrino spectrum, sterile neutrinos and additional radiation in the Universe. *Phys. Rev.*, D83:113011, 2011.

[120] P. S. Bhupal Dev and A. Pilaftsis. Light and Superlight Sterile Neutrinos in the Minimal Radiative Inverse Seesaw Model. *Phys. Rev.*, D87(5):053007, 2013.

[121] A. Palazzo. Constraints on very light sterile neutrinos from θ_{13}-sensitive reactor experiments. *JHEP*, 10:172, 2013.

[122] M. Galeazzi, F. Fontanelli, F. Gatti, and S. Vitale. Limits on the existence of heavy neutrinos in the range 50˘1000 ev from the study of the ^{187}Re beta decay. *Phys. Rev. Lett.*, 86:1978–1981, Mar 2001.

[123] E. Holzschuh, L. Palermo, H. Stussi, and P. Wenk. The beta-spectrum of S-35 and search for the admixture of heavy neutrinos. *Phys. Lett.*, B482:1–9, 2000.

[124] J. Deutsch, M. Lebrun, and R. Prieels. Searches for admixture of massive neutrinos into the electron flavour. *Nuclear Physics A*, 518(1–2):149 – 155, 1990.

[125] E Holzschuh, W Kundig, L Palermo, H Stüssi, and P Wenk. Search for heavy neutrinos in the beta-spectrum of 63ni. *Physics Letters B*, 451(1–2):247 – 255, 1999.

[126] D.N. Abdurashitov et al. The current status of "troitsk nu-mass" experiment in search for sterile neutrino. *Journal of Instrumentation*, 10(10):T10005, 2015.

[127] R. Adhikari et al. A White Paper on keV Sterile Neutrino Dark Matter. 2016.

[128] R.E. Shrock. New tests for and bounds on neutrino masses and lepton mixing. *Physics Letters B*, 96(1):159 – 164, 1980.

[129] W Y. Keung and G. Senjanović. Majorana neutrinos and the production of the right-handed charged gauge boson. *Phys. Rev. Lett.*, 50:1427–1430, May 1983.

[130] A. Datta, M. Guchait, and A. Pilaftsis. Probing lepton number violation via majorana neutrinos at hadron supercolliders. *Phys. Rev.*, D50:3195–3203, 1994.

[131] J. C. Helo, M. Hirsch, and S. Kovalenko. Heavy neutrino searches at the LHC with displaced vertices. *Phys. Rev.*, D89:073005, 2014. [Erratum: Phys. Rev.D93,no.9,099902(2016)].

[132] P. J. Fox, R. Harnik, J. Kopp, and Y. Tsai. Missing Energy Signatures of Dark Matter at the LHC. *Phys. Rev.*, D85:056011, 2012.

[133] L. Basso, A. Belyaev, S. Moretti, and C. H. Shepherd-Themistocleous. Phenomenology of the minimal B-L extension of the Standard model: Z' and neutrinos. *Phys. Rev.*, D80:055030, 2009.

[134] A. Pilaftsis. The Little Review on Leptogenesis. *J. Phys. Conf. Ser.*, 171:012017, 2009.

[135] P. Minkowski. mu gt; e gamma at a Rate of One Out of 1-Billion Muon Decays? *Phys.Lett.*, B67:421, 1977.

[136] M. Gell-Mann, P. Ramond, and R. Slansky. Complex Spinors and Unified Theories. *Conf. Proc.*, C790927:315–321, 1979.

[137] R. N. Mohapatra and G. Senjanovic. Neutrino Mass and Spontaneous Parity Violation. *Phys. Rev. Lett.*, 44:912, 1980.

[138] M. Fukugita and T. Yanagida. Barygenesis without grand unification. *Physics Letters B*, 174(1):45 – 47, 1986.

[139] K. N. Abazajian et al. Light Sterile Neutrinos: A White Paper. 2012.

[140] S. Parke and M. Ross-Lonergan. Unitarity and the three flavor neutrino mixing matrix. *Phys. Rev.*, D93(11):113009, 2016.

[141] G. H. Collin, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz. First Constraints on the Complete Neutrino Mixing Matrix with a Sterile Neutrino. *Phys. Rev. Lett.*, 117(22):221801, 2016.

[142] C. Athanassopoulos et al. Evidence for anti-muon-neutrino —¿ anti-electron-neutrino oscillations from the LSND experiment at LAMPF. *Phys. Rev. Lett.*, 77:3082–3085, 1996.

[143] A. Bolshakova et al. Revisiting the 'LSND anomaly' I: impact of new data. *Phys. Rev.*, D85:092008, 2012.

[144] A. Bolshakova et al. Revisiting the 'LSND anomaly' II: critique of the data analysis. *Phys. Rev.*, D85:092009, 2012.

[145] A. A. Aguilar-Arevalo et al. The Neutrino Flux prediction at MiniBooNE. *Phys. Rev.*, D79:072002, 2009.

[146] A. A. Aguilar-Arevalo et al. Unexplained Excess of Electron-Like Events From a 1-GeV Neutrino Beam. *Phys. Rev. Lett.*, 102:101802, 2009.

[147] A. A. Aguilar-Arevalo et al. A Combined $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillation Analysis of the MiniBooNE Excesses. 2012.

[148] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten. Reanalysis of the gallex solar neutrino flux and source experiments. *Physics Letters B*, 685(1):47 – 54, 2010.

[149] D. Dwyer. The energy spectrum of reactor antineutrinos. *PoS NEUTEL* **2015**, *023 (2015)*.

[150] A. C. Hayes. Uncertainties in reactor neutrino fluxes and in the anomaly. In *Proceedings, 50th Rencontres de Moriond Electroweak Interactions and Unified Theories: La Thuile, Italy, March 14-21, 2015*, page 241, 2015.

[151] I. Alekseev et al. DANSS: Detector of the reactor AntiNeutrino based on Solid Scintillator. *JINST*, 11(11):P11011, 2016.

[152] Y. J. Ko et al. Sterile neutrino search at NEOS Experiment. 2016.

[153] A. P. Serebrov et al. NEUTRINO4 experiment: preparations for search for sterile neutrino at 100 MW Reactor SM-3 at 6-12 Meters. 2012.

[154] C. Lane et al. A new type of Neutrino Detector for Sterile Neutrino Search at Nuclear Reactors and Nuclear Nonproliferation Applications. 2015.

[155] J. Ashenfelter et al. PROSPECT - A Precision Reactor Oscillation and Spectrum Experiment at Short Baselines. In *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013.

[156] V. Hélaine. Sterile neutrino search at the ILL nuclear reactor: the STEREO experiment. In *NuPhys2015: Prospects in Neutrino Physics (NuPhys) London, UK, December 16-18, 2015*, 2016.

[157] N. Ryder. First results of the deployment of a SoLid detector module at the SCK-CEN BR2 reactor. *PoS*, EPS-HEP2015:071, 2015.

[158] A. A. Aguilar-Arevalo et al. A Search for muon neutrino and antineutrino disappearance in MiniBooNE. *Phys. Rev. Lett.*, 103:061802, 2009.

[159] C. Giunti. Status of Light Sterile Neutrinos. In *PoS EPS* **-HEP2015**, *067 (2015).*

[160] M. G. Aartsen et al. Searches for Sterile Neutrinos with the IceCube Detector. *Phys. Rev. Lett.*, 117(7):071801, 2016.

[161] P. Adamson et al. Search for Sterile Neutrinos Mixing with Muon Neutrinos in MINOS. *Phys. Rev. Lett.*, 117(15):151803, 2016.

[162] G. H. Collin, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz. Sterile Neutrino Fits to Short Baseline Data. *Nucl. Phys.*, B908:354–365, 2016.

[163] A. Szelc. Recent Results from ArgoNeuT and Status of MicroBooNE. *Talk at Neutrino, 2014*, 2014.

[164] R. Acciarri et al. Detection of back-to-back proton pairs in charged-current neutrino interactions with the ArgoNeuT detector in the NuMI low energy beam line. *Phys. Rev.*, D90(1):012008, 2014.

[165] R Gandhi, B. Kayser, M. Masud, and S. Prakash. The impact of sterile neutrinos on CP measurements at long baselines. *JHEP*, 11:039, 2015.

[166] H. W. Zaglauer and K. H. Schwarzer. The Mixing Angles in Matter for Three Generations of Neutrinos and the Msw Mechanism. *Z. Phys.*, C40:273, 1988.

[167] K. Kimura, A. Takamura, and H. Yokomakura. Exact formulas and simple CP dependence of neutrino oscillation probabilities in matter with constant density. *Phys. Rev.*, D66:073005, 2002.

[168] A. A. Aguilar-Arevalo et al. A Search for electron neutrino appearance at the $\Delta m^2 \sim 1\text{eV}^2$ scale. *Phys. Rev. Lett.*, 98:231801, 2007.

[169] A. A. Aguilar-Arevalo et al. Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations. *Phys. Rev. Lett.*, 105:181801, 2010.

[170] P. Adamson et al. First Measurement of ν_μ and ν_e Events in an Off-Axis Horn-Focused Neutrino Beam. *Phys. Rev. Lett.*, 102:211801, 2009.

[171] J. M. Conrad and M. H. Shaevitz. Limits on electron neutrino disappearance from the karmen and lsnd ν_e-carbon cross section data. *Phys. Rev. D*, 85:013017, Jan 2012.

[172] B. Achkar et al. Search for neutrino oscillations at 15, 40 and 95 meters from a nuclear power reactor at bugey. *Nuclear Physics B*, 434(3):503 – 532, 1995.

[173] K. B. M. Mahn et al. Dual baseline search for muon neutrino disappearance at $0.5\text{eV}^2 < \Delta m^2 < 40\text{eV}^2$. *Phys. Rev.*, D85:032007, 2012.

[174] P. Adamson et al. Observation of muon neutrino disappearance with the minos detectors in the numi neutrino beam. *Phys. Rev. Lett.*, 97:191801, Nov 2006.

[175] P. Adamson et al. A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam. *Phys. Rev.*, D77:072002, 2008.

[176] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa. A New calculation of the atmospheric neutrino flux in a 3-dimensional scheme. *Phys. Rev.*, D70:043008, 2004.

[177] M. C. Gonzalez-Garcia and M. Maltoni. Atmospheric neutrino oscillations and new physics. *Phys. Rev.*, D70:033010, 2004.

[178] S. H. Ahn et al. Detection of accelerator produced neutrinos at a distance of 250-km. *Phys. Lett.*, B511:178–184, 2001.

[179] M. H. Ahn et al. Indications of neutrino oscillation in a 250 km long baseline experiment. *Phys. Rev. Lett.*, 90:041801, 2003.

[180] M. H. Ahn et al. Measurement of Neutrino Oscillation by the K2K Experiment. *Phys. Rev.*, D74:072003, 2006.

[181] C. Adams and others. Lar1nd proposal, fermilab proposal p-1053 (2013).

[182] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.

[183] A. Zee. A Theory of Lepton Number Violation, Neutrino Majorana Mass, and Oscillation. *Phys. Lett.*, B93:389, 1980. [Erratum: Phys. Lett.B95,461(1980)].

[184] K. S. Babu. Model of 'Calculable' Majorana Neutrino Masses. *Phys. Lett.*, B203:132–136, 1988.

[185] N. Arkani-Hamed, S. Dimopoulos, G. R. Dvali, and J. March-Russell. Neutrino masses from large extra dimensions. *Phys. Rev.*, D65:024032, 2002.

[186] E. K. Akhmedov and A. Y. Smirnov. Paradoxes of neutrino oscillations. *Phys. Atom. Nucl.*, 72:1363–1381, 2009.

[187] E. D. Church. LArSoft: A Software Package for Liquid Argon Time Projection Drift Chambers. 2013.

[188] J. S. Marshall and M. A. Thomson. The Pandora Software Development Kit for Pattern Recognition. *Eur. Phys. J.*, C75(9):439, 2015.

[189] M. Sorel. Expected performance of an ideal liquid argon neutrino detector with enhanced sensitivity to scintillation light. *JINST*, 9:P10002, 2014.

[190] M. Antonello et al. Precise 3D track reconstruction algorithm for the ICARUS T600 liquid argon time projection chamber detector. *Adv. High Energy Phys.*, 2013:260820, 2013.

[191] G. Bernardi et al. Search for Neutrino Decay. *Phys. Lett.*, B166:479, 1986.

[192] G. Bernardi et al. FURTHER LIMITS ON HEAVY NEUTRINO COUPLINGS. *Phys. Lett.*, B203:332, 1988.

[193] T. Asaka and M. Shaposhnikov. The nuMSM, dark matter and baryon asymmetry of the universe. *Phys. Lett.*, B620:17–26, 2005.

[194] A. M. Cooper-Sarkar et al. Search for Heavy Neutrino Decays in the BEBC Beam Dump Experiment. *Phys. Lett.*, B160:207, 1985.

[195] F. Bergsma et al. A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-GeV to 2.8-GeV. *Phys. Lett.*, B166:473, 1986.

[196] A. Vaitaitis et al. Search for neutral heavy leptons in a high-energy neutrino beam. *Phys. Rev. Lett.*, 83:4943–4946, 1999.

[197] M. Anelli et al. A facility to Search for Hidden Particles (SHiP) at the CERN SPS. 2015.

[198] S. Alekhin et al. A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case. 2015.

[199] D. Gorbunov and M. Shaposhnikov. How to find neutral leptons of the νMSM? *JHEP*, 10:015, 2007. [Erratum: JHEP11,101(2013)].

[200] T. Asaka, S. Eijima, and A. Watanabe. Heavy neutrino search in accelerator-based experiments. *JHEP*, 03:125, 2013.

[201] R. E. Shrock. General theory of weak processes involving neutrinos. i. leptonic pseudoscalar-meson decays, with associated tests for, and bounds on, neutrino masses and lepton mixing. *Phys. Rev. D*, 24:1232–1274, Sep 1981.

[202] A. Atre, T. Han, S. Pascoli, and B. Zhang. The Search for Heavy Majorana Neutrinos. *JHEP*, 05:030, 2009.

[203] R. E. Shrock. General theory of weak processes involving neutrinos. ii. pure leptonic decays. *Phys. Rev. D*, 24:1275–1309, Sep 1981.

[204] P. B. Pal and L. Wolfenstein. Radiative decays of massive neutrinos. *Phys. Rev. D*, 25:766–773, Feb 1982.

[205] F. del Aguila, S. Bar-Shalom, A. Soni, and J. Wudka. Heavy Majorana Neutrinos in the Effective Lagrangian Description: Application to Hadron Colliders. *Phys. Lett.*, B670:399–402, 2009.

[206] A. Aparici, K. Kim, A. Santamaria, and J. Wudka. Right-handed neutrino magnetic moments. *Phys. Rev.*, D80:013010, 2009.

[207] S. Bhattacharya and title = Wudka, J.

[208] S. Weinberg. Baryon and Lepton Nonconserving Processes. *Phys. Rev. Lett.*, 43:1566–1570, 1979.

[209] S. N. Gninenko. The MiniBooNE anomaly and heavy neutrino decay. *Phys. Rev. Lett.*, 103:241802, 2009.

[210] S. N. Gninenko. A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments. *Phys. Rev.*, D83:015015, 2011.

[211] L. Duarte, I. Romero, J. Peressutti, and O. A. Sampayo. Effective Majorana neutrino decay. 2016.

[212] B. Batell, M. Pospelov, and B. Shuve. Shedding Light on Neutrino Masses with Dark Forces. 2016.

[213] R. Foot, X. G. He, H. Lew, and R. R. Volkas. Model for a light Z-prime boson. *Phys. Rev.*, D50:4571–4580, 1994.

[214] D. I. Britton, S. Ahmad, D. A. Bryman, R. A. Burnham, E. T. H. Clifford, P. Kitching, Y. Kuno, J. A. Macdonald, T. Numao, A. Olin, J. M. Poutissou, and M. S. Dixit. Improved search for massive neutrinos in $\pi^+ \to e^+\nu$ decay. *Phys. Rev. D*, 46:R885–R887, Aug 1992.

[215] D. I. Britton, S. Ahmad, D. A. Bryman, R. A. Burnham, E. T. H. Clifford, P. Kitching, Y. Kuno, J. A. Macdonald, T. Numao, A. Olin, J-M. Poutissou, and M. S. Dixit. Measurement of the $\pi^+ \to e^+\nu$ branching ratio. *Phys. Rev. Lett.*, 68:3000–3003, May 1992.

[216] P. Coloma, B. A. Dobrescu, C. Frugiuele, and R. Harnik. Dark matter beams at LBNF. *JHEP*, 04:047, 2016.

[217] R. Dharmapalan et al. Low Mass WIMP Searches with a Neutrino Experiment: A Proposal for Further MiniBooNE Running. 2012.

[218] J. Badier et al. Mass and Lifetime Limits on New Longlived Particles in 300-GeV/$c\pi^-$ Interactions. *Z. Phys.*, C31:21, 1986.

[219] B. Fields and S. Sarkar. Big-Bang nucleosynthesis (2006 Particle Data Group mini-review). 2006.

[220] A. D. Dolgov, S. H. Hansen, G. Raffelt, and D. V. Semikoz. Heavy sterile neutrinos: Bounds from big bang nucleosynthesis and SN1987A. *Nucl. Phys.*, B590:562–574, 2000.

[221] A. D. Dolgov, S. H. Hansen, G. Raffelt, and D. V. Semikoz. Cosmological and astrophysical bounds on a heavy sterile neutrino and the KARMEN anomaly. *Nucl. Phys.*, B580:331–351, 2000.

[222] G. Bernardi. *Mesure du flux de ν_e dans le faisceau de neutrinos de basse energie du PS: intrepretation dans le cadre d'une oscillation $\nu_\mu \to \nu_e$.* PhD thesis, Unversité Pierre et Marie Curie, 1985.

[223] M. O. Wascko. Charged current single pion cross section measurement at MiniBooNE. *Nucl. Phys. Proc. Suppl.*, 159:50–55, 2006. [,50(2006)].

[224] A. Higuera et al. Measurement of Coherent Production of $\pi^{\pm}$ in Neutrino and Antineutrino Beams on Carbon from E_ν of 1.5 to 20 GeV. *Phys. Rev. Lett.*, 113(26):261802, 2014.

[225] D. Rein and L. M. Sehgal. PCAC and the Deficit of Forward Muons in pi+ Production by Neutrinos. *Phys. Lett.*, B657:207–209, 2007.

[226] Joshua B. Spitz. *Measuring Muon-Neutrino Charged-Current Differential Cross Sections with a Liquid Argon Time Projection Chamber.* PhD thesis, Yale U., 2011.

[227] R. Acciarri et al. Summary of the Second Workshop on Liquid Argon Time Projection Chamber Research and Development in the United States. *JINST*, 10(07):T07006, 2015.

[228] M. Toups on behalf of MicroBooNE Collaboration. First Results From MicroBooNE, Neutrino London, July, 2016. 2016.

[229] V. A. Duk et al. Search for Heavy Neutrino in $K^- -> \mu^- \nu_h (\nu_h -> \nu\gamma)$ Decay at ISTRA+ Setup. *Phys. Lett.*, B710:307–317, 2012.

[230] S. N. Gninenko and N. V. Krasnikov. Limits on the magnetic moment of sterile neutrino and two photon neutrino decay. *Phys. Lett.*, B450:165–172, 1999.

[231] G. J. Feldman and R. D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev.*, D57:3873–3889, 1998.