How broad is a neutrino?

Hannah Banks,[a,b] Kevin J. Kelly[b,c] and Matthew McCullough[b]

[a] *DAMTP, University of Cambridge,*
Wilberforce Road, Cambridge, U.K.

[b] *CERN, Theoretical Physics Department,*
Geneva, Switzerland

[c] *Department of Physics and Astronomy,*
Mitchell Institute for Fundamental Physics and Astronomy, Texas A&M University,
College Station, TX 77843, U.S.A.

E-mail: hmb61@cam.ac.uk, kjkelly@tamu.edu, matthew.mccullough@cern.ch

ABSTRACT: Canonical neutrino oscillations arise due to the propagation of three mass eigenstates from production to detection. We aspire to capture, in one simple framework, a broad range of new physics effects on neutrino propagation beyond this canonical picture — this can be done by promoting the neutrino propagators to the general Källén-Lehmann form. In this work we demonstrate how models predicting additional light propagating species of neutrino are naturally accommodated in this language and propose a simple model spectrum composed of just three 'broadened' states as a flexible ansatz by which to explore the phenomenology of new physics in neutrino propagation. Reinterpreting existing neutrino oscillation measurements, we illustrate how this framework provides the capacity to probe deviations from the standard three-neutrino scenario systematically and generally. Whilst current data allows for relatively strong constraints on broadened neutrinos, we find the upcoming JUNO experiment will yield significant improvements, particularly for the heaviest neutrino, paving the way to a clearer understanding of how neutrinos propagate in vacuum.

Contents

1 Introduction

Since their discovery, neutrinos have been an invaluable probe of physics 'beyond the Standard Model' (BSM). Indeed the observation of flavour changes in neutrino oscillations, which require non-zero neutrino masses, is arguably one of the strongest pieces of evidence for new BSM physics. This observation behooves particle physicists to probe, to the greatest degree possible, all aspects of the neutrino sector of the SM.

As neutrino oscillations arise as a result of massive neutrino propagation from the point of production to the point of detection, one may ask if these massive fermions are propagating from one point to another as expected. In this work we aim to develop a coherent framework in which this question may be posed theoretically and answered experimentally.[1]

In any free or interacting quantum field theory the propagator (two-point function) for a fermion may be captured by a Källén-Lehmann representation. Note that this is a non-perturbative representation, not reliant on any perturbative expansion, but only on

---

[1]Note that very recently ref. [1] appeared, with broadly similar motivations although very different specific considerations.
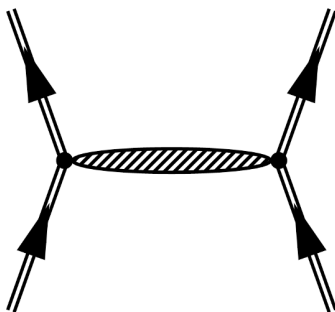
Figure 1. A schematic Feynman diagram for a general neutrino propagation process. The double lines show the external states at the production and detection vertices. The elliptical blob symbolises the Källén-Lehmann propagator. In the case of canonical oscillations, this reduces to the usual Feynman propagator.

the basic axioms of quantum mechanics. As a result, this representation may be included in any phenomenological description of neutrino oscillations, whether QM-like or QFT-like, and it will capture any QFT-compatible BSM modifications of neutrino propagation from production to detection, as illustrated schematically in figure 1. To this end, in section 2 we generalise the usual neutrino oscillation framework to include the more general Källén-Lehmann form, propagating this general form all the way through to generalised formulæ for oscillation probabilities of neutrino appearance and disappearance.

To illustrate the utility of this approach more concretely we provide two simple BSM scenarios in which the neutrino propagator becomes 'broad'. The first is the previously studied 'Pseudo-Dirac' neutrino model [2–9], where neutrino masses are effectively Dirac and a small amount of lepton number violation leads to a small splitting of mass eigenstates. A second model generalises this setup further, through a fermionic clockwork-inspired deconstruction of fermions on a circle, to give a broad band of states in lieu of a single neutrino mass eigenstate. The mapping of both models into the Källén-Lehmann representation of the propagator is developed and their impact on observable neutrino oscillation data is investigated. Motivated by these models and the fact that an almost limitless zoology of models is in principle realisable, we then present a simple 'top-hat' phenomenological ansatz which captures the dominant oscillation features of more complete microscopic models that lead to band-like neutrino spectral functions.

Armed with this formalism and utilitarian phenomenological ansatz, we first discuss how observations of neutrinos from distant sources can aide in testing these scenarios in section 3. Then, in section 4 we take the next step to see how well we currently understand BSM effects in neutrino propagation as it pertains to Earth-based experiments. In practical terms we do this by determining how well KamLAND, Daya Bay, T2K and NOvA measurements constrain non-SM contributions to neutrino propagation, finding that KamLAND has been particularly powerful. On the other hand, the upcoming JUNO experiment will provide an even more powerful and complementary probe, breaking flat directions that presently exist. Finally, section 6 offers some concluding remarks.

2 Neutrino propagation from Källén-Lehmann

There is a broad literature concerning neutrino oscillation amplitudes, ranging from textbook quantum-mechanical derivations to involved quantum field theory treatments. There is little to be gained in repeating these analyses here, thus we focus on the core novel ingredient of this work, wherein the key addition to typical treatments, reviewed for example in [10, 11], is to replace the free neutrino propagator and matrix elements found in QFT treatments

$$\sum_j U_{\alpha j}^* U_{\beta j} \frac{\slashed{p} + m_j}{p^2 - m_j^2 + i\epsilon} \ , \tag{2.1}$$

by the more general Källén-Lehmann propagator for spin-1/2 particles

$$G_{\alpha\beta}^{\mathrm{KL}}(p^2) = i \int_0^\infty d\mu^2 \frac{\slashed{p}\tilde{\rho}_{\alpha\beta}(\mu^2) + \rho_{\alpha\beta}(\mu^2)}{p^2 - \mu^2 + i\epsilon} \ , \tag{2.2}$$

where α and β denote the flavour eigenstate at the point of production and detection respectively. This captures any new-physics effects on the propagation of neutrinos consistent with the axioms of QFT and, essentially, promotes the sum over mass eigenstates to a continuum integral. The assumptions employed in deriving this general form of the propagator are broad but simple. They boil down to Lorentz-invariance, causality, and the usual assumptions of quantum mechanics including positive-definite norm states and boundedness of the Hamiltonian. Any BSM scenarios outside of these assumptions will not, generally, be captured. For instance, Lorentz-violating neutrino interactions, theories including ghost states, or propagation in Lorentz-violating backgrounds, cannot be captured. On the other hand, scenarios in which neutrinos have BSM decays modes will be captured. This can be understood through the optical theorem, where a particle decay width is related to the imaginary part of its propagator which is, in turn, given by $\tilde{\rho}(\mu^2)$ and $\rho(\mu^2)$. This would apply to any hidden QFT sector interactions that influence neutrino propagation including, for example, if the hidden sector is of an 'unparticle' type [12], with a continuum of states.

The ρ functions here encode both the density of states, and the matrix elements describing the overlap between the mass and interaction eigenstates. One recovers the usual Feynman propagator if $\tilde{\rho}(\mu^2)$ and $\rho(\mu^2)/\mu$ are identical and comprise delta functions $\delta(\mu^2 - m_j^2)$. More general scenarios such as neutrino mixing with hidden sector states, including for strongly-coupled hidden sectors, can be described by functions exhibiting a discrete or continuous density of states.

Since the propagator above and the spectral functions $\rho_{\alpha\beta}(\mu^2)$ and $\tilde{\rho}_{\alpha\beta}(\mu^2)$ are expressed in the interaction basis, whose eigenstates need not coincide with those of the vacuum Hamiltonion, the spectral functions need not be real nor obey the usual positivity requirements satisfied by Källén-Lehmann functions. An exception is in the case $\alpha = \beta$, when the mapping between mass- and interaction-bases depends only on real, positive factors such as $|U_{\alpha j}|^2$. In this case, the spectral functions are required to be real, and both $\tilde{\rho}_{\alpha\alpha}(\mu^2)$ and $(\mu\tilde{\rho}_{\alpha\alpha}(\mu^2) - \rho_{\alpha\alpha}(\mu^2))$ are positive-definite.

In all scenarios of interest, the neutrinos are ultra-relativistic. In this limit, chirality and helicity eigenstates coincide and the spin structure can be factored out of the two-point function. All the relevant oscillation physics is fully captured by the scalar amplitude which itself can be represented via the scalar Källén-Lehmann propagator. In practice this amounts to setting $\tilde{\rho}(\mu^2) = 0$ in eq. 2.2. Cataloguing the BSM possibilities for the neutrino sector thus fundamentally reduces to the study of a single scalar function, $\rho(\mu^2)$, which will be the object of interest in what follows.

Following the standard QFT calculations with this modification, the probability for neutrino flavour transitions for propagating neutrinos of momentum $\mathbf{p}$ is given by

$$P_{\alpha\beta} = \left| \int_0^\infty d\mu^2 e^{-i\sqrt{\mu^2 + \mathbf{p}^2}L} \rho_{\alpha\beta}(\mu^2) \right|^2 . \tag{2.3}$$

In this statement we have applied the approximation $L \approx T$, and are explicitly working in the limit in which the full QFT calculation reproduces that obtained in QM. This requires the additional assumption that the neutrino at the production and detection vertices can be approximated as a plane wave, or equivalently that the neutrino wavepackets maintain coherence over the entire distance of interest, L.

In order to expand $\sqrt{\mu^2 + \mathbf{p}^2}$, we must assume that all coherently-propagating neutrinos are relativistic. This requires that $\rho_{\alpha\beta}(\mu^2)$ does not have support for large μ^2, which is expected for light, oscillating neutrinos with large (above $\sim$MeV) energies. We thus have that $\sqrt{\mu^2 + \mathbf{p}^2} \approx E + \mu^2/(2E)$, and can express the transition probability as

$$P_{\alpha\beta} = \left| \int_0^\infty d\mu^2 e^{-i\frac{\mu^2 L}{2E}} \rho_{\alpha\beta}(\mu^2) \right|^2 . \tag{2.4}$$

This expression for the transition amplitude is simply the Fourier transform of the scalar spectral function.

2.1 Two-flavour example

To make concrete headway we now turn to two-flavour mixing as characterised by the unitary matrix

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} , \tag{2.5}$$

which rotates flavour states (ν_α, ν_β) into mass eigenstates (ν_1, ν_2). For example, the ν_α survival probability thus follows from the spectral density

$$\rho_{\alpha\alpha} = \delta(\mu^2 - m_1^2)\cos^2\theta - \delta(\mu^2 - m_2^2)\sin^2\theta , \tag{2.6}$$

giving

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2 2\theta \sin^2\left(\frac{\Delta m_{21}^2 L}{4E}\right) , \tag{2.7}$$

where $\Delta m_{21}^2 \equiv m_2^2 - m_1^2$. Similarly, the $\nu_\alpha \to \nu_\beta$ ($\alpha \neq \beta$) transition probability follows from

$$\rho_{\alpha\beta} = \sin\theta\cos\theta\left(\delta(\mu^2 - m_1^2) - \delta(\mu^2 - m_2^2)\right) , \tag{2.8}$$

giving

$$P\left(\nu_\alpha \to \nu_\beta\right) = \sin^2 2\theta \sin^2\left(\frac{\Delta m_{21}^2 L}{4E}\right). \tag{2.9}$$

We see that the predictions from this formalism map to textbook expressions in this standard case.

2.2 Pseudo-Dirac neutrinos

A simple and well-known possibility that converts a single neutrino mass eigenstate into multiple states is that of pseudo-Dirac neutrinos [2–9]. Structurally, the model consists of Dirac neutrinos, preserving a lepton-number symmetry, supplemented by a small source of explicit lepton-number violation which can be naturally small. This symmetry breaking thus splits the Dirac neutrinos into two Majorana mass eigenstates. Explicitly, the usual Lagrangian is

$$\lambda \nu_R L H + h.c., \tag{2.10}$$

which (after electroweak symmetry breaking) generates a Dirac mass $m_D = \lambda v/\sqrt{2}$. The small lepton-number violating Majorana mass term for the right-handed neutrino is

$$\frac{1}{2} M \nu_R^2 + h.c. \tag{2.11}$$

In the limit $m_D \gg M$, this has physical mass eigenvalues

$$M_\pm = m_D \pm \frac{M}{2} \ , \tag{2.12}$$

and the mixing angle between the EW gauge eigenstate and the mass eigenstates is

$$\tan(2\theta) = \frac{2m_D}{M}. \tag{2.13}$$

When $m_D \gg M$ the mixing angle is approximately maximal, $\theta \approx \pi/4$. This furnishes a basic example where the mass eigenstate that propagates from production to detection is, microscopically, multiple states of different mass.

This can be generalised to three flavours by the introduction of 3 right-handed neutrinos, ν_{Ri}, for $i \in \{1, 2, 3\}$. For simplicity, we take the alignment limit where the 3×3 Dirac and Majorana mass matrices can be simultaneously diagonalised. We shall henceforth refer to the diagonal entries of these as m_{D_i} and M_i respectively. The 6×6 unitary matrix describing the rotation of the $\{\nu_{eL}, \nu_{eR}, \nu_{\mu L}, \nu_{\mu R}, \nu_{\tau L}, \nu_{\tau R}\}$ flavour states into the mass eigenstates may be conveniently written as a product of 3×3 and 2×2 matrices, defined as

$$O_{\alpha\beta ab} = U_{\alpha\beta} R_{ab}(\theta_i) \ , \tag{2.14}$$

where $U_{\alpha\beta}$ is the usual leptonic mixing matrix and R_{ab} is a 2×2 rotation matrix with a rotation angle $\tan 2\theta_i = 2m_{D_i}/M_i$.

The spectral density for α to β flavour transitions is

$$\rho_{\alpha\beta}(\mu^2) = \sum_{i=1}^3 U_{\alpha i}^* U_{\beta i} \left[\cos^2(\theta_i)\delta(\mu^2 - M_{i,-}^2) + \sin^2(\theta_i)\delta(\mu^2 - M_{i,+}^2)\right] \ , \tag{2.15}$$

where $M_{i,+}, M_{i,-}$ refer to the masses of the physical Majorana mass eigenstates into which the i^{th} Dirac neutrino splits. Motivated by this, we now generalise further to a band of states.

2.3 A band of neutrinos

One way a band of neutrino states can be realised is through a 'clockwork'-inspired [13–15] fermion ring. Consider N identical copies of the Standard Model enjoying a large translation symmetry in theory space. These different sectors could have renormalisable (hence relevant at low energies) couplings to one another through one of three portals: Higgs, kinetic mixing and neutrino. We will focus on the latter and suppose the following ring of couplings for a single flavour of neutrinos

$$\mathcal{L}_\lambda = \frac{\lambda}{q} \left[\sum_{j=1}^{j=N-1} L_j H_j (\psi_j - q\psi_{j+1}) + L_N H_N (\psi_N - q\psi_1) \right] \, , \tag{2.16}$$

where the ψ are SM-neutral Weyl fermions, λ a small Yukawa coupling, and we assume $q > 1$. Upon electroweak symmetry breaking in all of the sectors the collective neutrino mass matrix, for one active flavour, becomes

$$M_\nu = \frac{m_D}{q} \begin{pmatrix} 1 & -q & 0 & \cdots & & 0 \\ 0 & 1 & -q & \cdots & & 0 \\ 0 & 0 & 1 & \cdots & & 0 \\ \vdots & \vdots & \vdots & \ddots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -q \\ -q & 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \, , \tag{2.17}$$

where $m_D = \lambda v / \sqrt{2}$. The physical eigenvalues of this mass matrix are

$$m_j^2 = \frac{m_D^2}{q^2} \left(1 + q^2 - 2q \cos\left(\frac{2\pi j}{N} \right) \right). \tag{2.18}$$

We thus have a band of N mass eigenstates centred at $\sim m_D$ with a breadth of $\sim m_D/q$. The overlap between these states and the flavour eigenstates is given by the elements of the rotation matrix

$$R_{jk} = \frac{\cos\left(\frac{2\pi jk}{N} \right) + \sin\left(\frac{2\pi jk}{N} \right)}{\sqrt{N}}. \tag{2.19}$$

Thus, for instance, the overlap between the j^{th} mass eigenstate of a given generation and the interaction eigenstate in that sector of the ring is given by the elements R_{1j}.

This can be further generalised to include 3 active generations of neutrino in each copy of the SM, with (site-independent) couplings λ_α where $\alpha \in \{1, 2, 3\}$. We take q to be the same for each generation. In our conventions, the spectral density describing transitions from the α to β flavour eigenstates is then

$$\rho_{\alpha\beta} = \sum_{i=1}^{3} U_{\alpha i}^* U_{\beta i} \left(\sum_{j=1}^{N} R_{1j}^2 \delta(\mu^2 - m_{ij}^2) \right) \, , \tag{2.20}$$

where the m_{ij} refer to the masses of the N physical mass eigenstates of the i^{th} generation. This model essentially replaces any would-be SM Dirac neutrino mass eigenstate, of mass $\sim m_D$, by a band of states spread about this mass scale.

Some comments are in order. This model serves only to illustrate that such a scenario is possible, but is not intended to advertise the model as being particularly strongly motivated in its own right. Furthermore, many details of the model are not necessary in order to realise the same qualitative scenario including, for instance, the translation symmetry that enforces equal q at each site which could be softly broken.

Finally, a brief comment on cosmology. If each sector were truly identical then it would be necessary that only the SM sector is reheated at the end of inflation, otherwise the neutrinos and photons of the hidden sectors would presumably lead to inconsistencies with cosmological observations. It may also be the case that the reheating temperature is necessarily low to avoid over-populating the other neutrino states.

2.4 Theory-space perspective

To understand the phenomenology of these more exotic scenarios it is instructive to consider the process in terms of the 'theory-space' sites, borrowing terminology from dimensional deconstruction [16, 17]. Consider an N-site band model. Only one of the sites corresponds to an active, and potentially measurable neutrino. The remaining $N-1$ sites are sterile. When we talk of physically measurable neutrino oscillations, we refer to the probability of starting in the active flavour eigenstate of the (α) generation and being measured in the active flavour eigenstate of the (β) generation at some later time, however here we will focus on a single-generation case.

Since the starting active flavour eigenstate is some superposition of the N mass eigenstates through which the system evolves, there will be a generally non-zero probability of being in one of the other $N-1$ sites when a neutrino is detected. In general, the higher the value of N, the lower the detection probabilities become due to the greater number of available sites to which the system has evolved during propagation.

To illustrate this, we consider the specific case of $N=6$, corresponding to one active neutrino site and 5 sterile sites. In the main panel of figure 2 we consider initialising the system in the active flavour state, which we number as state 1, and plot the overlap of the time-dependent state with the i^{th} site eigenstate as a function of the measurement time. Due to the discrete rotational symmetry in theory space the curves for $i=5$ and $i=6$ overlap exactly with those of states $i=3$ and $i=2$ respectively, so they are not visible on the plot. The inset images show the overlap with each site at a time t_j and thus how the measurement probability is distributed between the 6 sites. When viewed chronologically they illustrate the flow of probability in theory-space, which repeats periodically. As N is increased, the time taken to return to the starting distribution in which the system is in the active state increases, as there are a greater number of sites for the probability to flow through first.
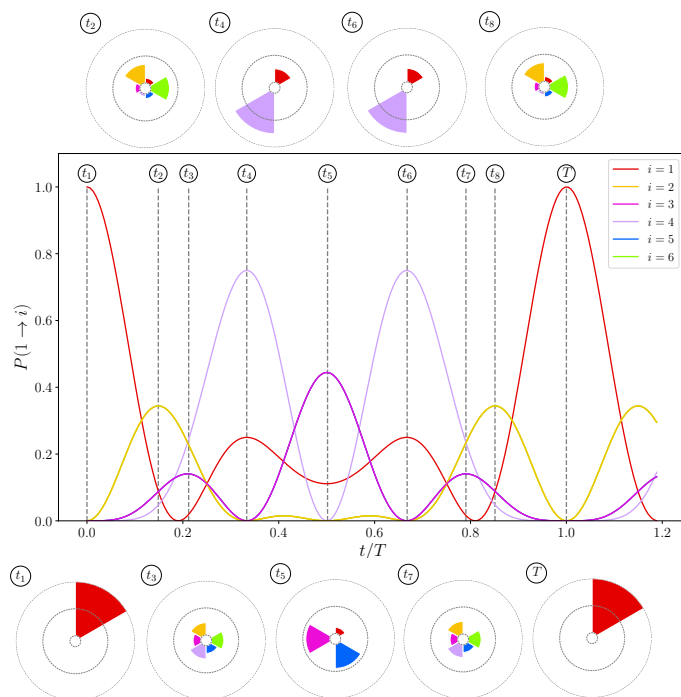
Figure 2. The flow of probability around the ring of fermions for $N = 6$ sites and $q = 10^3$. The neutrino is initialised at $t = 0$ in the active, $i = 1$ state and, as time evolves, may overlap with the other $N - 1$ sites, before returning to a maximum of being detected in the $n = 1$ state again after a time T. The probability that the system is in the i^{th} site on measurement after a time t is shown in the main figure, in units of the cycle period T. The circular bar charts illustrate, to scale, the flow of probability around the sites, plotted in the same colours as in the main figure, at a number of instances. The inner and outer concentric dashed circles denote measurement probabilities of 0.5 and 1 respectively.

2.5 A phenomenological ansatz

The two scenarios considered above are just examples of the rich landscape of BSM possibilities for the neutrino sector. Whilst it is of course possible to construct the relevant spectral function for any given model, calculate the transition probabilities, and extract the bounds from experimental data, this process would need to be undertaken separately for each individual model under consideration. Given the number and range of theoretical possibilities, a comprehensive survey is not only cumbersome, but fundamentally not feasible. Experimental analyses are thus typically limited to a handful of the simplest models.

We find that over the energies probed by existing oscillation experiments, the n-flavour oscillation probability distributions arising from both the pseudo-Dirac and band models can be sufficiently mimicked by a spectral function consisting of n top-hat functions. With respectively discrete and quasi-continuous spectral functions, the pseudo-Dirac and band cases span a broad landscape of plausible spectral functions and it is thus reasonable to expect the top-hat set-up to be capable of reproducing neutrino oscillation probability distributions for a broad range of microscopic scenarios. In this section, we will thus derive the general form of the transition probability generated by top-hat functions.
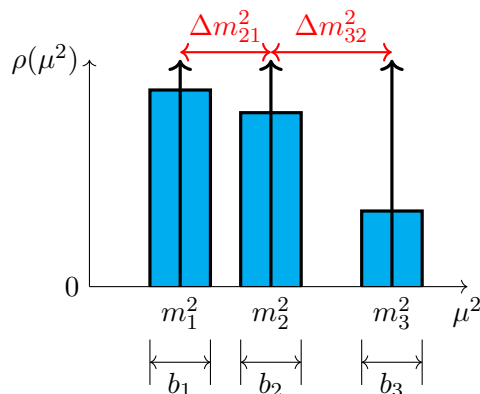
Figure 3. Sketch of the form of the top-hat $\rho_s(\mu^2)$ for 3 flavour oscillations as defined in eq. 2.21. The bold arrows represent the δ-function eigenstates of the standard scenario.

To this end, we begin by focusing on three-neutrino oscillations, replacing the delta functions (three at m_i^2) of the canonical spectral function with top-hat functions of (generally different) breadths b_i centred at these values. Explicitly, we parametrise the spectral function as

$$\rho_{ee}(\mu^2) = \begin{cases} \frac{1}{b_1}|U_{e1}|^2 = \frac{1}{b_1}\cos^2\theta_{12}\cos^2\theta_{13}, & m_1^2 - \frac{b_1}{2} \leq \mu^2 \leq m_1^2 + \frac{b_1}{2} \\ \frac{1}{b_2}|U_{e2}|^2 = \frac{1}{b_2}\cos^2\theta_{13}\sin^2\theta_{12}, & m_2^2 - \frac{b_2}{2} \leq \mu^2 \leq m_2^2 + \frac{b_2}{2} \\ \frac{1}{b_3}|U_{e3}|^2 = \frac{1}{b_3}\sin^2\theta_{13}, & m_3^2 - \frac{b_3}{2} \leq \mu^2 \leq m_3^2 + \frac{b_3}{2} \\ 0, & \text{otherwise} \end{cases} . \quad (2.21)$$

This is illustrated schematically in figure 3. We note that current measurements of the leptonic mixing matrix indicate $|U_{e3}|^2 \ll |U_{e1}|^2, |U_{e2}|^2$.

Upon taking the Fourier transform, we arrive at the amplitude

$$iA_{ee} = \sum_{i=1}^{3} \text{sinc}\left(\frac{Lb_i}{4E}\right) |U_{ei}|^2 e^{\frac{-iLm_i^2}{2E}} , \quad (2.22)$$
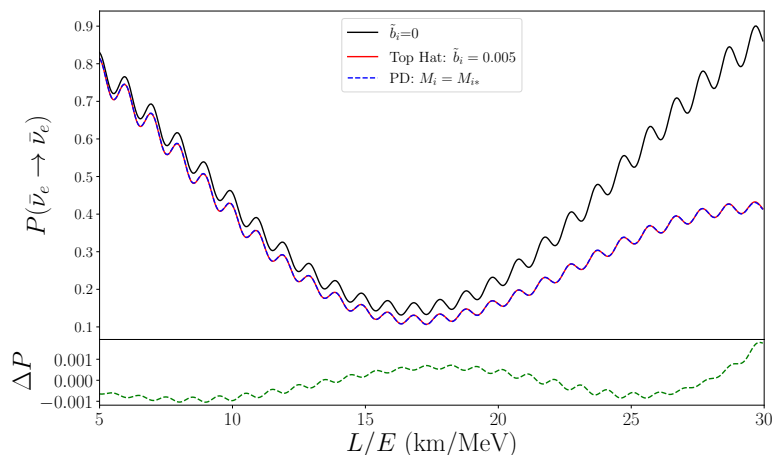
and probability

$$P(\nu_e \to \nu_e) = \left(\sum_i \text{sinc}(\alpha_i)|U_{ei}|^2\right)^2 - 4\sum_{i<j}|U_{ei}|^2|U_{ej}|^2\text{sinc}(\alpha_i)\,\text{sinc}(\alpha_j)\sin^2(\Delta_{ji}) , \quad (2.23)$$

where $\Delta_{ji} \equiv \Delta m_{ji}^2 L/(4E)$ and $\alpha_i \equiv b_i L/(4E)$. We note that since we are treating neutrino propagation in vacuum, the probability for antineutrino oscillations, $\bar{\nu}_e \to \bar{\nu}_e$, is identical.

For the case of equal breadths, $b_1 = b_2 = b_3 = b$, this simplifies to

$$P(\nu_e \to \nu_e) = \text{sinc}^2\left(\frac{Lb}{4E}\right)\left[1 - \sin^2 2\theta_{12}\cos^4\theta_{13}\sin^2\Delta_{21}\right.$$

$$\left. - \sin^2 2\theta_{13}\left(\cos^2\theta_{12}\sin^2\Delta_{31} + \sin^2\theta_{12}\sin^2\Delta_{32}\right)\right] , \quad (2.24)$$
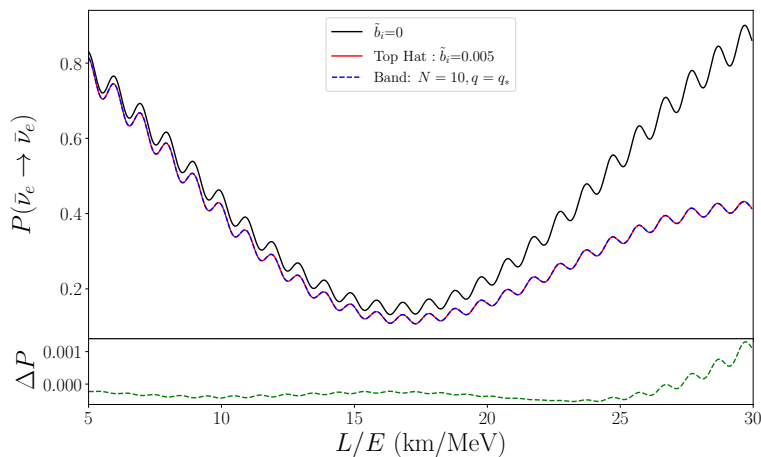
(a) Comparison of 3-flavour oscillation probabilities generated by a density of states comprised of 3 top-hat functions with fractional breadths $\tilde{b}_i = 0.005$, and a pseudo-Dirac density of states with masses $(M_1, M_2, M_3) = (1.23, 1.21, 1.39) \times 10^{-4}$eV, selected by performing a least squares fit to the top-hat probability distribution. The lower panel shows the fit residuals.
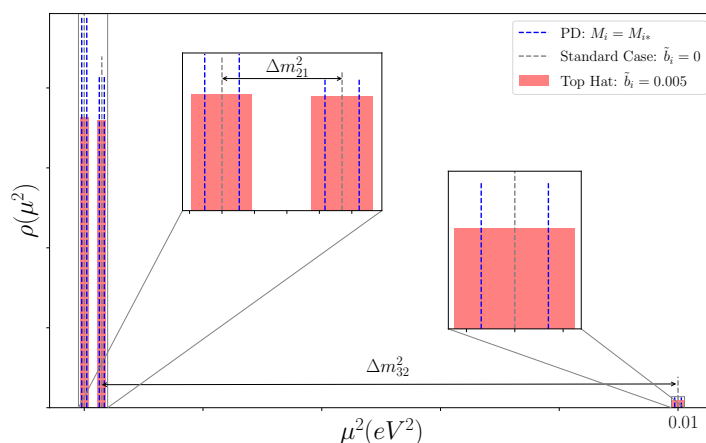
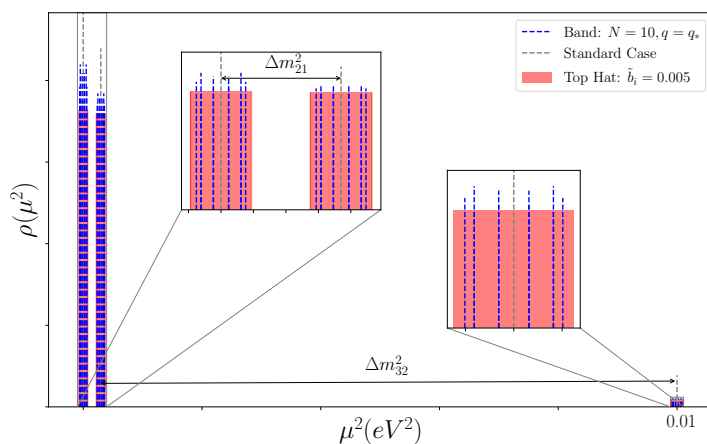**(b)** A plot of the spectral function $\rho(\mu^2)$ for these two models with the same parameters as in (a). Also shown is the triple δ-function density of states corresponding to the canonical scenario. The heights of the top-hat functions relative to each other are plotted to scale but the vertical extent of the (formally infinite) δ-functions for the pseudo-Dirac and conventional models are intended for illustrative purposes only.

Figure 4. Evaluation of the top-hat phenomenological ansatz defined in section 2.5 to capture the oscillation behaviour of pseudo-Dirac models as detailed in section 2.2.

which we identify as the standard probability expression modulated by a factor of $\mathrm{sinc}^2\left(\frac{Lb}{4E}\right)$. We emphasise that these expressions only depend on the relative spacing of the states in μ^2 and not on their absolute value.

(a) Comparison of 3-flavour oscillation probabilities generated by a density of states comprised of 3 top-hat functions with fractional breadths $\tilde{b}_i = 0.005$, and a $N = 10$ band density of states. The value of $q = 987.2$ was selected using a least squares fit to the top-hat probability distribution.

(b) A plot of the spectral function $\rho(\mu^2)$ for these two models with the same parameters as in (a). Also shown is the triple δ-function density of states corresponding to the canonical scenario. The heights of the top-hat functions relative to each other are plotted to scale but the vertical extent of the (formally infinite) δ-functions for the band states and conventional model are intended for illustrative purposes only.

Figure 5. Evaluation of the top-hat phenomenological ansatz defined in section 2.5 to capture the oscillation behaviour of band models as defined in section 2.3.

For the purpose of this exercise we assume a normal neutrino mass ordering and set m_3 to have an absolute value of $0.1\,\mathrm{eV}$. We then fix m_1 and m_2 according to $m_i^2 = m_3^2 - \Delta m_{3i}^2$ for $i \in \{1, 2\}$. We use the current best fit values of the Δm_{ij}^2 and the mixing angles as determined by existing experiments.[2] We parametrise the breadths of the top-hat states according to a fractional value relative to the mass squared, $\tilde{b}_i \equiv b_i/m_i^2$. For reference, we highlight that at $\tilde{b}_1 = 0.01$, $b_1 \approx \Delta m_{21}^2$ and the spectral gap between the top-hats of state 1 and state 2 vanishes.

We now seek to demonstrate the capability of the top-hat to capture the phenomenology of the specific microscopic models introduced earlier. To approach this, we generate the (anti-)electron survival probability distribution as a function of baseline length over energy (L/E) for a given set of top-hats breadths (b_1, b_2, b_3), and perform a least squares fit of the pseudo-Dirac and band probability distributions over the range $L/E \in [5, 30]$ km/MeV, treating the set (M_1, M_2, M_3) as free parameters in the former and fixing N and fitting for q in the latter. As an illustration, we consider a top-hat set up with fractional breadths $\tilde{b}_i = 0.005$ which, for reference, corresponds to $b_1 \approx 0.5\Delta m_{21}^2$. Figure 4 (a) compares the probability distribution for this set up with that generated by the best fit of the pseudo-Dirac model. The best fit parameters are $(M_1, M_2, M_3) = (1.23, 1.21, 1.39) \times 10^{-4}\,\mathrm{eV}$. The residuals for the fit are shown in the lower panel and show a disagreement of less than 0.1% over the domain probed by JUNO. In figure 4 (b) we plot the density of states for these models with the parameters required for matched probability distributions as detailed above. Also shown is the triple δ-function spectrum of the canonical scenario. Whilst the heights of the top-hat states relative to each other are plotted to scale, the vertical extent of the (formally infinite) δ-functions for both the pseudo-Dirac model and the standard case are for illustrative purposes only. We observe that probability matching occurs when the breadth of the top-hat function for a given state is roughly twice the pseudo-Dirac mass splitting of that generation.

In figure 5 (a) we show the agreement between the same top-hat density of states, with the best fit of the band model for $N = 10$, which is achieved for $q = q_* = 987.2$. Once again we see an excellent agreement of within 0.1% over the JUNO energy range. As seen in the comparison of the spectral densities for the matched cases shown in figure 5 (b), the breadth of the top-hat of a given state should slightly exceed the breadth of the band. Note that whilst there are 10 mass eigenstates for each generation, degeneracies of the mass values of some states means that they do not appear as distinct lines on the plot.

Whilst we have merely shown matching to a single choice of top-hat breadths here, we have found that it is always possible to tune the model parameters to achieve an excellent agreement of the probability distributions for any set of top-hat breadths, and thus from the reverse perspective, to be able to find a choice of top-hat breadths which reproduce the probability distributions of these models for any given set of parameters.

_____

[2]Explicitly we take $\sin^2 \theta_{12} = 0.3$, $\sin^2 \theta_{13} = 2.2 \times 10^{-2}$, $\Delta m_{21}^2 = 7.4 \times 10^{-5}\,\mathrm{eV}^2$, and $\Delta m_{31}^2 = 2.5 \times 10^{-3}\,\mathrm{eV}^2$ [18].
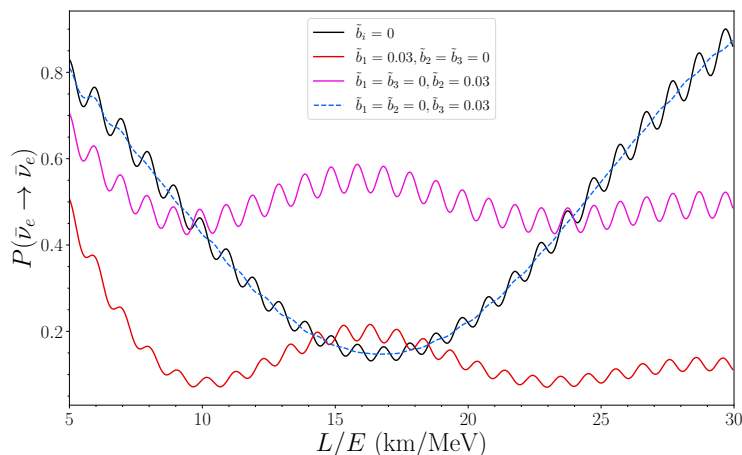
Figure 6. Plots of the anti-electron neutrino survival probability as a function of L/E for the 3-flavour top-hat density of states set up as defined in section 2.5 for the case where only one of the three states has a finite breadth, of fractional value $\tilde{b} = 0.03$ and the remaining two are δ-functions. Also shown for comparison is the probability distribution for the standard scenario in which the density of states comprises of 3 δ-functions ($b_i = 0$).

2.6 The top-hat landscape

Given their ability to capture the probability distributions of some specific theoretical models, it serves to explore the phenomenological space that can be spanned by top-hat spectral densities in this way. We will address this by exploring the effect on the transition probability of different top-hat breadths. We initially consider the case where only one of the three states has a finite breadth, and the remaining two are δ-functions. In figure 6 we plot the probability distributions generated by setting $\tilde{b} = 0.03$ for the non-zero breadth state. This comprises of $\mathcal{O}(1)$ amplitude oscillations driven by Δm_{21}^2 on which a smaller amplitude, higher frequency, oscillation driven by Δm_{13}^2 is superimposed. We note that the impact of broadening the third top-hat in the spectral function is to damp the amplitude of the Δm_{13}^2 oscillations. As apparent from eq. 2.23, modifications to the large amplitude Δm_{21}^2 oscillations arise from the sinc terms in b_1 and b_2, and thus occur on broadening of the breadths of the first and second top-hat states.

To gain a handle on the top-hat breadths required to produce measurable deviations from the standard case, we consider the impact on the probability distribution of setting $\tilde{b}_1 = \tilde{b}_2 = \tilde{b}_3 = \tilde{b}$, such that the breadths of the three states relative to their central mass-squared are equal. We plot the transition probability for increasing values of $\tilde{b}$ in figure 7. Very little deviation from the standard probability distribution occurs for $\tilde{b} \lesssim 0.002$, the point at which $b_1 \approx b_2$ becomes sizeable ($\sim 20\%$) relative to Δm_{21}^2. By nature of their larger amplitude, the greatest overall modifications to the probability distribution will arise from modifications to the Δm_{21}^2 oscillations and thus the degree of deviation from the standard probability distribution is largely controlled by the comparative sizes of $b_1 \approx b_2$ and Δm_{21}^2. Increasing $\tilde{b}_1$ (or $\tilde{b}_2$) both increases the frequency, and decreases the
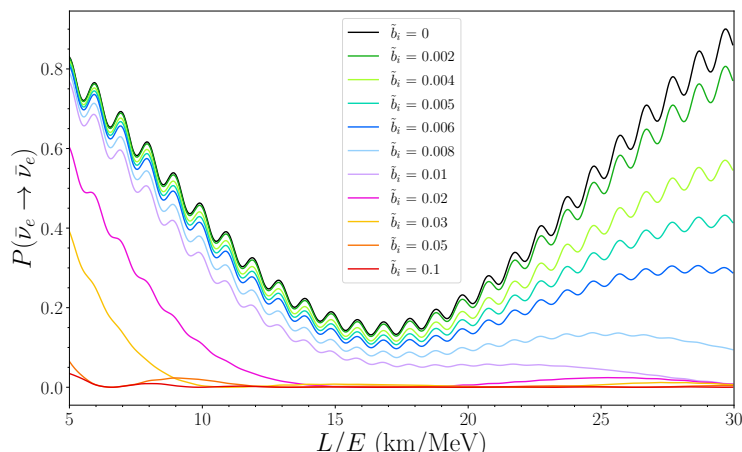
Figure 7. Plots of the anti-electron neutrino survival probability as a function of L/E for the 3-flavour top-hat density of states as defined in section 2.5 with $\tilde{b}_1 = \tilde{b}_2 = \tilde{b}_3 = \tilde{b}$, for various values of $\tilde{b}$. The case $\tilde{b}_i = 0$ corresponds to the standard scenario in which the density of states is comprised of 3 δ-functions.

extent of the central peak, of (one of) the sinc functions modulating the Δm_{21}^2 oscillations. If $b_1 \ll \Delta m_{21}^2$, the entire energy range probed by JUNO lies approximately at the central peak of the modulating sinc function. In the opposing regime, the JUNO energy range falls in the tails of the sinc function and probability is driven towards zero. If b_1 and Δm_{21}^2 are of the same order, the probed energy range lies on the falling edge of the sinc function central peak and we see sizeable corrections to the probability, which increase with L/E. For reference, we note that for the set of masses used in this figure, $b_1 \approx \Delta m_{21}^2$ when $\tilde{b}_1 \approx 0.01$.

3 The decoherence limit

When propagating over long distances the neutrino wave packets will ultimately decohere, leading to asymptotic neutrino detection probabilities. One may recall this from the standard two-neutrino treatment, wherein for $\Delta m_{21}^2 L \gg 4E$, there are many oscillations such that

$$\sin^2 \left(\frac{\Delta m_{21}^2 L}{4E} \right) \to \frac{1}{2} \; , \tag{3.1}$$

and the survival probability asymptotes to

$$P\left(\nu_\alpha \to \nu_\alpha\right) = 1 - \frac{1}{2} \sin^2 2\theta. \tag{3.2}$$

Due to our simplifying ansatz of flavour alignment, these same limits factorise within the various scenarios considered above, when L/E is much larger than any mass-squared-separation scale of interest for the model at hand.

In this limit, coherence is lost among the mass eigenstates and simplified expressions for the various models may be found.

For the pseudo-Dirac case, assuming all mass eigenstates having splittings such that coherence is lost amongst them, the probabilities become

$$
\begin{aligned}
P_{\alpha\beta} &= \sum_{i=1}^{3} (U_{\alpha i}^* U_{\beta i})^2 \left[\cos^4(\theta_i) + \sin^4(\theta_i) \right] \\
&\rightarrow \frac{1}{2} \sum_{i=1}^{3} (U_{\alpha i}^* U_{\beta i})^2 \; ,
\end{aligned}
\tag{3.3}
$$

where in the final expression we have employed the limit $M \ll m_D$, thus $\theta_i \rightarrow \pi/4$. Importantly, note that this is a factor $1/2$ smaller than the corresponding standard three-neutrino result in the decoherence limit.

Now consider the band model, again with a band splitting for each mass eigenstate that is great enough for coherence to be lost amongst the band. In this case the probabilities become

$$
\begin{aligned}
P_{\alpha\beta} &= \sum_{i=1}^{3} (U_{\alpha i}^* U_{\beta i})^2 \sum_{j=1}^{N} R_{1j}^4 \\
&\rightarrow \frac{3}{2N} \sum_{i=1}^{3} (U_{\alpha i}^* U_{\beta i})^2.
\end{aligned}
\tag{3.4}
$$

The emerging pattern is physically intuitive. Given a long-enough propagation distance the individual mass wave packets separate. In these flavour-aligned models, what would have been one mass eigenstate is replaced by N individual separated states, of which an effective $N-1$ are sterile and undetectable. The overall numerical coefficient depends on the specific details of the model, however the suppression of the resulting signal universally scales inversely proportionally to the number of available states.

Since the top-hat model effectively corresponds to an infinite number of states, one would expect that in the decoherence limit the various detection probabilities would asymptotically vanish. In this case the electron neutrino survival probability asymptotes to

$$
P(\nu_e \rightarrow \nu_e) \rightarrow \sum_i \frac{1}{2} \left(\frac{4E}{L b_i} \right)^2 |U_{ei}|^4 \; ,
\tag{3.5}
$$

which indeed vanishes asymptotically as an inverse quadratic of the length scale. This will be important when we come to consider experimental constraints. Furthermore, the mixing angles are also important. For instance, one has

$$
|U_{e1}|^4 \approx 0.5 \, , \quad |U_{e2}|^4 \approx 0.09 \, , \quad |U_{e3}|^4 \approx 5 \times 10^{-4} \, ,
\tag{3.6}
$$

hence, depending on the situation, it may be that constraints on the survival probability of the total electron neutrino flux will only likely significantly constrain broadening in the lightest neutrinos.

In the following subsections, we discuss how observations of neutrinos from very distant sources can, in principle, place strong constraints on broadened neutrinos through nonzero b_i. In contrast, sections 4 and 5 will demonstrate how terrestrial, neutrino-oscillation focused measurements, are an interesting avenue to potentially discover a nonzero b_i.

3.1 Supernova constraints

Let us first consider the longest baseline constraints arising from the detection of SN1987A neutrinos by Kamiokande-II [19, 20], IMB [21, 22], and Baksan [23]. Constraints on the pseudo-Dirac case were considered in ref. [9], which informs our comments here. In ref. [9] it was found that constraints were not strong in the decoherence limit since the flux reduction by a factor of 2 could be accommodated by a doubling of the supernova energy. Stronger limits were found for mass splittings in the L/E-range corresponding to SN1987A as this modifies the energy-dependence of the spectrum. However, in this work we are interested in large mass splittings $\Delta m^2 \gg 10^{-19}$ eV^2, such that the full decoherence limit is reached. We would thus expect that in the band model there would be similar flexibility, such that $N = 3$ would be acceptable, and perhaps even larger. On the contrary, for the top-hat ansatz we would expect a significant suppression of the signal for a universal breadth of $b \gg 10^{-19}$ eV^2.

There is, however, the aforementioned caveat which concerns the universality of the breadth. We will illustrate this with the top-hat case, although similar aspects apply to the other two models. Due to the smallness of $|U_{e2}|^4$ and $|U_{e3}|^4$, from eq. 3.5 it appears there would be no strong constraint on b_2 and b_3, since a reduction of the flux due to decoherence in these modes would not be sufficient to generate tension with observations. On the contrary, we expect a limit in the ballpark of $b_1 \lesssim 10^{-19}$ eV^2 applies, otherwise the neutrino flux would be too greatly depleted. The analysis of ref. [9] assumed universal splittings, however it would be interesting to repeat this analysis under non-universal assumptions, especially with application to the models considered here.

3.2 Solar constraints

The case for solar neutrinos is somewhat more complex. While matter effects are important, it is still ultimately the element $|U_{ei}|^4$ which controls the magnitude of electron neutrino disappearance on broadening the i^{th} mass eigenstate. However, the overall fluxes are measured with greater precision and the physics of production understood with greater certainty. As a result, one again expects the most significant constraints on b_1, however constraints on b_2 may also be relevant. Estimating that strong constraints arise whenever $Lb_i/E \ll 1$, then for typical solar neutrino energies and the Earth-Sun baseline one expects the limits to be in the region of $b_1 \lesssim 10^{-12}$ eV^2.

Solar constraints on the pseudo-Dirac case were studied in detail in ref. [6], where indeed limits which would roughly correspond to $b_1 \lesssim 10^{-12}$ eV^2 were found, however in this case constraints in the region of $b_2 \lesssim 10^{-11}$ eV^2 also arise, due to the well-measured neutrino flux which gives sensitivity to effects at the $|U_{e2}|^4$-level. Which width is most strongly constrained is determined from an interplay between the precision of measurement and the

magnitude broadening effect which, since it is controlled by the parameter combination $\sim b_j L/E$, is stronger at lower energies for a given fixed width and baseline.

3.3 Atmospheric & astrophysical constraints

Due to the typical baseline and energies involved, one expects constraints from atmospheric neutrinos to be at the level of $b_i \lesssim 10^{-4}$ eV^2. However, as we will demonstrate in section 4, future long-baseline reactor antineutrino experiments will probe smaller breadths for all b_i, and thus we will not consider atmospheric constraints further here.

Finally, the observation of extragalactic neutrinos at neutrino observatories [24, 25] can provide an additional handle on neutrinos traveling great distances. Measuring the ratios of different flavours of the neutrinos upon arrival at Earth [26, 27] can, in principle, help in constraining many of the models discussed here, however, such constraints would be subject to uncertainties on the overall neutrino flux, among others. Nevertheless, as precision improves [28], the BSM power of these measurements should be considered in more detail.

3.4 Summary

It is clear that astrophysical probes of neutrino oscillations allow for baselines that go deep into the decoherence regime. Indeed, due to the form of the PMNS matrix they lead to very strong constraints on the breadth of the lightest mass eigenstate, at the scale of $b_1 \lesssim 10^{-19}$ eV^2 and slightly weaker constraints at the level of $b_2 \lesssim 10^{-11}$ eV^2. However, they do not probe b_3 with the same power due to the smallness of $|U_{e3}|^2$. Moreover, these observations all probe the physics of broadened neutrinos in the decoherence (classical) regime as opposed to making measurements where the propagating neutrinos maintain coherence.

As a result, in section 4, we turn to terrestrial probes of neutrino breadths for all mass eigenstates, for their novelty in probing the quantum interference effects of broadened neutrinos, and also as a complementary probe to the methods discussed in this section, subject to a very different set of measurement techniques and assumptions.

4 Terrestrial experiments for constraining spectral functions

Having discussed the strengths and weaknesses of very long-baseline, astrophysical constraints on neutrino breadths in section 3, we now shift our focus to terrestrial neutrino oscillation experiments. Section 2 established our phenomenological description of the neutrino spectral functions; now we explore these experiments and their ability to test this scenario. In developing this phenomenological approach, we have focused on the case where neutrinos are propagating for long proper times (large L/E_ν) in vacuum, where their interactions with any matter along the path of propagation can be neglected. To date, the experiments consistent with this assumption are those measuring oscillations of electron antineutrinos $\overline{\nu}_e$ produced in nuclear reactors. These oscillations, with $E_\nu \sim 1-10$ MeV, have been measured at a variety of baseline lengths, allowing for world-leading measurements of both mass-squared splittings Δm_{21}^2 and Δm_{31}^2 in the standard framework. A

second class of experiments measure the oscillations of ν_μ produced at $\sim$ GeV energies in proton-induced neutrino beams, travelling hundreds of kilometers. These ν_μ disappearance experiments offer comparable sensitivity to Δm_{31}^2 and operate in a similar L/E_ν regime to the reactor antineutrino experiments.

Reactor antineutrino $\overline{\nu}_e$ experiments. The most precise measurements of reactor antineutrinos are those from KamLAND [29–31], with baselines of $L \approx 200$ km, and Daya Bay [32–36], with baselines $L \approx 1$ km. These two correspond to measurements at $L/E \approx$ 40 km/MeV and 0.5 km/MeV, respectively. Due to the hierarchicy between $\Delta m_{21}^2 \approx$ 7.5×10^{-5} eV^2 and $\Delta m_{31}^2 \approx 2.5 \times 10^{-3}$ eV^2, KamLAND is sensitive to oscillations driven by Δm_{21}^2 (where oscillations due to Δm_{31}^2 have averaged out over the energy uncertainty of its detector) and Daya Bay is sensitive to oscillations driven by Δm_{31}^2 (where the oscillations due to Δm_{21}^2 have yet to develop significantly at the Daya Bay L/E ranges). Because of this, we expect that when studying the generalised spectral functions described in section 2.5, KamLAND will be sensitive to nonzero b_1 and b_2 on the order of Δm_{21}^2 and that Daya Bay will be sensitive to b_3 on the order of Δm_{31}^2.

In simulating KamLAND, we develop our analysis to match the results of the most recent collaboration results in ref. [31]. This is modified from analyses from refs. [37, 38] built off the collaboration results of ref. [30] to include a broader range of L/E for the greatest possible sensitivity to the b_i. We have verified that with b_i set to zero, our analysis roughly reproduces the results of ref. [31] by weighting the different reactor contributions to the final result properly and averaging over the L/E ranges considered. We also include a correlated 4.1% systematic normalization uncertainty as done in ref. [30] to arrive at our results.

Our simulation of Daya Bay is modified from the analysis of ref. [39] (see ref. [40] for further discussion), developed to match the official results from ref. [35]. This includes nuisance parameters reflecting the systematic uncertainty associated with the neutrino flux for each bin in the analysis, correlated across the three experimental halls of Daya Bay, allowing for reduced dependence on the neutrino flux prediction for the results. Refs. [39, 41] provide further details for the simulation of neutrino event-rate expectations at Daya Bay in the context of additional sterile neutrino species.

We also consider the possibility of testing these phenomenological spectral functions in the future, namely by the JUNO [42, 43] experiment. JUNO is a medium-baseline reactor experiment that will measure antineutrino oscillations with $L \approx 50$ km, in the L/E regime between that tested by Daya Bay and KamLAND. This will allow JUNO to simultaneously measure oscillations driven by the two mass-squared splittings in a precise way. Previous analyses, including refs. [44–49], have demonstrated that JUNO is an impressive discovery ground for BSM physics — here, we demonstrate that in our broad-neutrino framework, JUNO will be sensitive to all three spectral-function breadths b_i simultaneously, and should exhibit impressive capability in searching for nonzero b_i's. To simulate JUNO we use the same analysis described in refs. [37, 38], modified to accommodate our scenario. We refer the reader to refs. [37, 38] and references therein for more detail. Notably, the simulation that we use here includes a correlated flux systematic uncertainty of 2%, an uncorrelated uncertainty of 0.8% for each reactor considered, and a spectral shape uncertainty of 1%.
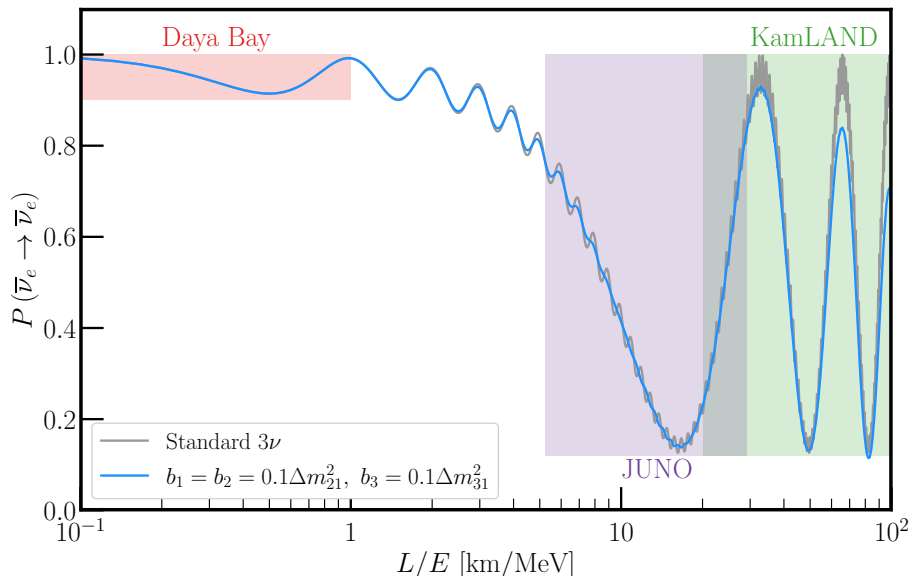
Figure 8. Oscillation probability for reactor antineutrinos as a function of L/E for the standard three-neutrino case (grey) and including nonzero spectral-function breadths as indicated in the legend (blue). We shade the regions of L/E probed by existing/future experiments Daya Bay (red), JUNO (purple), and KamLAND (green).

The relevant L/E ranges of these three reactor antineutrino experiments are displayed in figure 8. Each experiment's L/E range is shown as a shaded box, with Daya Bay, JUNO, and KamLAND in red, purple, and green, respectively. We also show the oscillation probability $P(\overline{\nu}_e \rightarrow \overline{\nu}_e)$ as a function of L/E that is/can be measured by these three experiments. We display oscillation probabilities for two cases: grey for the standard three-neutrino scenario[3] and in blue where we additionally include nonzero spectral-function breadths, $b_1 = b_2 = 0.1\Delta m_{21}^2$ and $b_3 = 0.1\Delta m_{31}^2$.

Here, the advantage of exploring these effects at large L/E is clear, as in figures 4–7. Because KamLAND operated at such large L/E, we would expect powerful sensitivity. However, not present in this figure (but present in our simulations) are the effects of finite energy resolution by the respective detectors. For instance, KamLAND has larger energy uncertainty than JUNO will and therefore is not as sensitive to fast, Δm_{31}^2-driven oscillations in its range of L/E (and does not have sensitivity to b_3). Thus, despite its lower L/E, we expect JUNO to be the most powerful of these three.

Long-baseline ν_μ disappearance experiments. Throughout this work, we are interested in scenarios where neutrino propagation in vacuum is a suitable description. While modern-day (and future) long-baseline experiments measuring $P(\nu_\mu \rightarrow \nu_e)$ require the consideration of neutrino interactions with matter for an accurate calculation of oscillation probabilities, the disappearance probability $P(\nu_\mu \rightarrow \nu_\mu)$ and its CP-conjugate are insensitive to standard matter effects for the baselines/energies of interest.

---

[3]Assuming the same values for the mass splittings and mixing angles as stated in section 2.5.

For that reason, and for complete comparison against the tests from reactor antineutrino measurements, we include adapted simulations of the T2K [50] and NOvA [51] experiments from ref. [52] to account for these effects in long-baseline $\nu_\mu \to \nu_\mu$ and $\overline{\nu}_\mu \to \overline{\nu}_\mu$ oscillations. The oscillation probabilities follow from eq. (2.23) with the substitution $|U_{ei}|^2 \to |U_{\mu i}|^2$. When considering present T2K and NOvA results, we consider only statistical uncertainties when determining constraints on b_i — such an approach has been demonstrated to be sufficient, even in the context of new-physics searches, in refs. [53, 54] among others.

We also consider future long-baseline experiments DUNE [55, 56] and Tokai to Hyper-Kamiokande (T2HK) [57, 58], which we will comment on in section 5. For these future projections, flux normalization uncertainties of 5% are incorporated to better reproduce the experiments' expected sensitivities.

In the context of figure 8, all of these experiments are situated at a similar L/E to Daya Bay. Therefore, we expect sensitivity to b_3 but not competitive with what JUNO will have to offer in the coming decade, due to JUNO's larger L/E and powerful energy resolution.

5   Current & future constraints on spectral functions

In this section, we provide the current constraints on the breadths b_1, b_2, and b_3. For simplicity, we will focus on the scenario in which the neutrino masses follow the normal ordering $m_3 > m_2 > m_1$. We will present results in terms of the dimensionless $\tilde{b}_i \equiv b_i/m_i^2$. We choose $m_1 = 10^{-2}\,\mathrm{eV}$ as a benchmark for this presentation. Given the discussion in section 2, we expect that the constraints on b_i from these experiments would be largely unchanged if we considered the inverted mass ordering $m_2 > m_1 > m_3$, however the dimensionless $\tilde{b}_i$ would given the change in the overall m_i^2.

When analysing current data, to estimate the constraints on $\tilde{b}_i$, we marginalise over the relevant "standard" oscillation parameters in a given analysis. For KamLAND, this means marginalising over $\sin^2\theta_{12}$ and Δm_{21}^2, for Daya Bay, $\sin^2\theta_{13}$ and Δm_{31}^2, and for our analysis of T2K and NOvA's ν_μ disappearance channels, Δm_{31}^2 and $\sin^2\theta_{23}$. This additional parameter space, relative to fixing the standard oscillation parameters to their best-fit values, leads to a factor of ∼two reduction in constraining power of the $\tilde{b}_i$ in some cases, similar to effects observed in ref. [46] in the context of neutrino decoherence searches.

We present current constraints on the three $\tilde{b}_i$ in figure 9. We find that current data from KamLAND, Daya Bay, T2K, and NOvA are consistent with $b_1 = b_2 = b_3 = 0$. Here, we compare 1σ (dotted), 2σ (dashed), and 3σ (dashed) constraints for the different experiments, KamLAND in blue, Daya Bay in green, and a combined analysis of T2K and NOvA in orange. As expected from our discussion in section 4, we find that KamLAND has strong sensitivity to $\tilde{b}_1$ and $\tilde{b}_2$, but no sensitivity to $\tilde{b}_3$. In contrast, Daya Bay and the long-baseline ν_μ disappearance measurements are able to constrain $\tilde{b}_3$. The 2σ and 3σ Daya Bay constraints do not appear on this space as they are cannot constrain $\tilde{b}_3 < 1$ at this level.
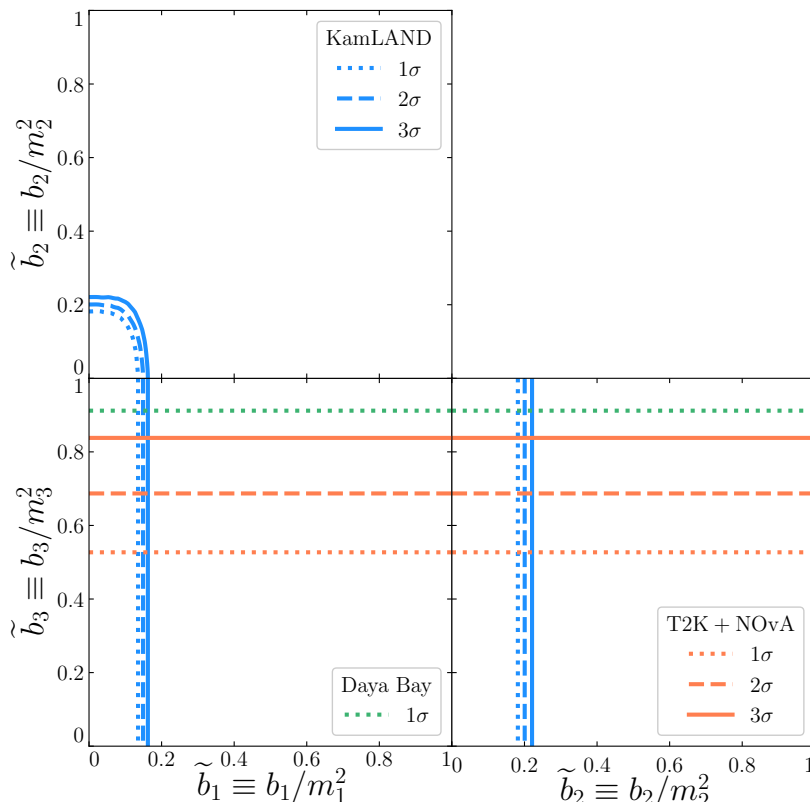
Figure 9. Current constraints from KamLAND (blue), Daya Bay (1σ only, green), and long-baseline ν_μ disappearance measurements from T2K and NOvA (orange) at 1, 2, and 3σ CL (dotted, dashed and solid respectively) on the reduced breadths b_i, relative to the overall neutrino-masses-squared when we assume $m_1 = 10^{-2}\,\text{eV}$.

	$\tilde{b}_1$	b_1 [eV^2]	$\tilde{b}_2$	b_2 [eV^2]	$\tilde{b}_3$	b_3 [eV^2]
Current	5.7×10^{-2}	5.7×10^{-6}	1.6×10^{-1}	2.8×10^{-5}	5.3×10^{-1}	1.4×10^{-3}
Future	5.1×10^{-2}	5.1×10^{-6}	4.0×10^{-2}	7.0×10^{-6}	1.3×10^{-2}	3.3×10^{-5}

Table 1. Current and future (expected) terrestrial constraints on $\tilde{b}_i \equiv b_i/m_i^2$ and b_i at 1σ confidence.

In contrast, future projections on $\tilde{b}_i$ from JUNO are displayed in figure 10, compared against KamLAND's constraints. Note that we have changed the axes ranges here such that the Daya Bay, T2K, and NOvA constraints are no longer visible. Additionally, we have explored the capability of the future T2HK and DUNE measurements of ν_μ disappearance in this context; their sensitivities are both also outside the range shown in figure 10 — they have the power to constrain $\tilde{b}_3 \lesssim 7 \times 10^{-2}$ and have no sensitivity to $\tilde{b}_1$ nor $\tilde{b}_2$.

Taking this set of constraints, we can derive 1σ upper limits on $\tilde{b}_i$ as well as b_i given our benchmark $m_1 = 10^{-2}\,\text{eV}$, which we present in table 1. While the $\tilde{b}_i$ are useful for dimensionless comparisons, the future sensitivity on the absolute b_i are notable in their own right, demonstrating sensitivity to meV-scale phenomena.
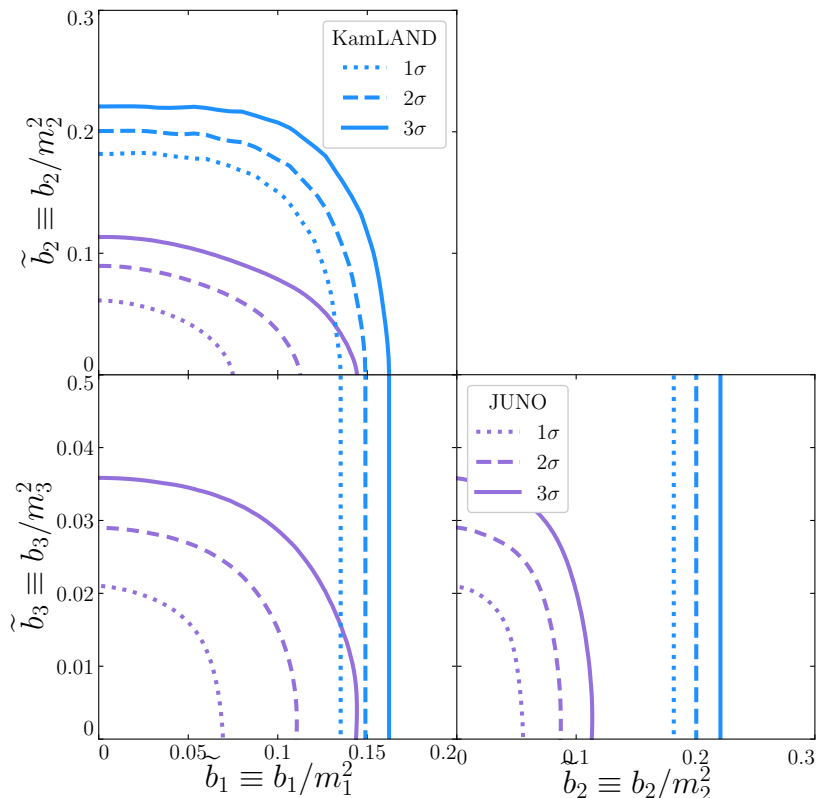
Figure 10. Expected future constraints by JUNO (purple) in the absence of a new-physics signal in the parameter space of $\tilde{b}_i$, compared against the current constraints from KamLAND (blue). Compared to figure 9, the data range here is so narrow that the constraints from Daya Bay, T2K, and NOvA do not appear.

6 Conclusions

For decades the neutrino sector has provided us with a unique window through which to study the hidden world of matter. Indeed it remains a theoretically well-motivated location to hunt for new physics. With new neutrino oscillation observatories planned for the near future, it is timely to explore the diverse theoretical and experimental BSM landscape of neutrino oscillations.

In light of this, we have proposed a new general framework to organise future explorations of the neutrino sector, capturing new physics effects on neutrino propagation in a single spectral function. We demonstrated how this language both reproduces conventional neutrino flavour oscillation calculations, and is simultaneously capable of describing the phenomenology of more exotic theoretical models including those with discrete and continuous mass spectra. The relevant phenomenological features can in both cases be mimicked by a 'toy' mass-spectrum comprising three top-hat functions with a model-specific choice of breadths. We emphasise that this top-hat set up should not be considered a concrete theoretical model, but moreover as a convenient phenomenological ansatz offering the possibility of model-independent analyses of experimental data. Instead of having

to sequentially re-interpret searches for different theoretical models, one can equivalently constrain the 'breadth' of neutrino spectral functions and thus probe the nature of neutrino propagation effects directly. In the instance of a positive hint for a non-zero neutrino breadth, whilst there may not be a unique invertible mapping from a given set of top-hat breadths to the true mass spectrum, a preference for top-hat functions of non zero breadth is likely to indicate the presence of extra states in the mass-spectrum, and thus the existence of additional sterile neutrinos.

After discussing how long-distance neutrino measurements can test this non-zero breadth in section 3, we demonstrated the utility of this approach with terrestrial experiments in sections 5, where we explored the landscape of existing and future oscillation experiments and compared their capacity to probe new physics in neutrino propagation. We found that the long-baseline anti-electron neutrino oscillation experiment KamLAND constrains the breadths of the two lower mass states to remarkable sensitivity but provides no information on the breadth of the third state. Daya Bay and current long-baseline ν_μ disappearance searches close this gap somewhat, but current constraints are comparatively weak. The near-future mid-baseline anti-electron neutrino experiment JUNO is projected to significantly improve upon the sensitivities of existing searches, most noticeably for that of the third, highest mass, state.

We highlight that whilst current data is consistent with the predictions of the conventional 3-neutrino model, the possibility of finite breadths are by no means excluded, particularly for the highest mass state, m_3. Finding new physics in the neutrino sector is possible, as testified by the rich and diverse range of BSM scenarios considered in literature. Armed with the tools to probe the neutrino mass spectrum in a general manner, this framework offers a useful and previously unexploited manner by which to harness the capabilities of future experiments to answer the question: how broad is a neutrino?

Acknowledgments

References

[1] T. Gherghetta and A. Shkerin, *Out of this world neutrino oscillations*, FTPI-MINN-22-22 (2022) [INSPIRE].

[2] L. Wolfenstein, *Different Varieties of Massive Dirac Neutrinos*, *Nucl. Phys. B* **186** (1981) 147 [INSPIRE].

[3] S.T. Petcov, *On Pseudodirac Neutrinos, Neutrino Oscillations and Neutrinoless Double beta Decay*, *Phys. Lett. B* **110** (1982) 245 [InSPIRE].

[4] S.M. Bilenky and B. Pontecorvo, *Neutrino Oscillations With Large Oscillation Length in Spite of Large (Majorana) Neutrino Masses?*, *Sov. J. Nucl. Phys.* **38** (1983) 248 [InSPIRE].

[5] M. Kobayashi and C.S. Lim, *Pseudo Dirac scenario for neutrino oscillations*, *Phys. Rev. D* **64** (2001) 013003 [hep-ph/0012266] [InSPIRE].

[6] A. de Gouvea, W.-C. Huang and J. Jenkins, *Pseudo-Dirac Neutrinos in the New Standard Model*, *Phys. Rev. D* **80** (2009) 073007 [arXiv:0906.1611] [InSPIRE].

[7] G. Anamiati, R.M. Fonseca and M. Hirsch, *Quasi Dirac neutrino oscillations*, *Phys. Rev. D* **97** (2018) 095008 [arXiv:1710.06249] [InSPIRE].

[8] G. Anamiati, V. De Romeri, M. Hirsch, C.A. Ternes and M. Tórtola, *Quasi-Dirac neutrino oscillations at DUNE and JUNO*, *Phys. Rev. D* **100** (2019) 035032 [arXiv:1907.00980] [InSPIRE].

[9] I. Martinez-Soler, Y.F. Perez-Gonzalez and M. Sen, *Signs of pseudo-Dirac neutrinos in SN1987A data*, *Phys. Rev. D* **105** (2022) 095019 [arXiv:2105.12736] [InSPIRE].

[10] M. Beuthe, *Oscillations of neutrinos and mesons in quantum field theory*, *Phys. Rept.* **375** (2003) 105 [hep-ph/0109119] [InSPIRE].

[11] E.K. Akhmedov and J. Kopp, *Neutrino Oscillations: Quantum Mechanics vs. Quantum Field Theory*, *JHEP* **04** (2010) 008 [arXiv:1001.4815] [InSPIRE].

[12] H. Georgi, *Unparticle physics*, *Phys. Rev. Lett.* **98** (2007) 221601 [hep-ph/0703260] [InSPIRE].

[13] K. Choi and S.H. Im, *Realizing the relaxion from multiple axions and its UV completion with high scale supersymmetry*, *JHEP* **01** (2016) 149 [arXiv:1511.00132] [InSPIRE].

[14] D.E. Kaplan and R. Rattazzi, *Large field excursions and approximate discrete symmetries from a clockwork axion*, *Phys. Rev. D* **93** (2016) 085007 [arXiv:1511.01827] [InSPIRE].

[15] G.F. Giudice and M. McCullough, *A Clockwork Theory*, *JHEP* **02** (2017) 036 [arXiv:1610.07962] [InSPIRE].

[16] N. Arkani-Hamed, A.G. Cohen and H. Georgi, *(De)constructing dimensions*, *Phys. Rev. Lett.* **86** (2001) 4757 [hep-th/0104005] [InSPIRE].

[17] N. Arkani-Hamed, A.G. Cohen and H. Georgi, *Electroweak symmetry breaking from dimensional deconstruction*, *Phys. Lett. B* **513** (2001) 232 [hep-ph/0105239] [InSPIRE].

[18] I. Esteban, M.C. Gonzalez-Garcia, M. Maltoni, T. Schwetz and A. Zhou, *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, *JHEP* **09** (2020) 178 [arXiv:2007.14792] [InSPIRE].

[19] KAMIOKANDE-II collaboration, *Observation of a Neutrino Burst from the Supernova SN 1987a*, *Phys. Rev. Lett.* **58** (1987) 1490 [InSPIRE].

[20] K.S. Hirata et al., *Observation in the Kamiokande-II Detector of the Neutrino Burst from Supernova SN 1987a*, *Phys. Rev. D* **38** (1988) 448 [InSPIRE].

[21] R.M. Bionta et al., *Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud*, *Phys. Rev. Lett.* **58** (1987) 1494 [InSPIRE].

[22] IMB collaboration, *Angular Distribution of Events From Sn1987a*, *Phys. Rev. D* **37** (1988) 3361 [INSPIRE].

[23] E.N. Alekseev, L.N. Alekseeva, I.V. Krivosheina and V.I. Volchenko, *Detection of the Neutrino Signal From SN1987A in the LMC Using the Inr Baksan Underground Scintillation Telescope*, *Phys. Lett. B* **205** (1988) 209 [INSPIRE].

[24] IceCube collaboration, *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*, *Science* **342** (2013) 1242856 [arXiv:1311.5238] [INSPIRE].

[25] IceCube collaboration, *Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data*, *Phys. Rev. Lett.* **113** (2014) 101101 [arXiv:1405.5303] [INSPIRE].

[26] IceCube collaboration, *A combined maximum-likelihood analysis of the high-energy astrophysical neutrino flux measured with IceCube*, *Astrophys. J.* **809** (2015) 98 [arXiv:1507.03991] [INSPIRE].

[27] IceCube-Gen2 collaboration, *IceCube-Gen2: the window to the extreme Universe*, *J. Phys. G* **48** (2021) 060501 [arXiv:2008.04323] [INSPIRE].

[28] N. Song, S.W. Li, C.A. Argüelles, M. Bustamante and A.C. Vincent, *The Future of High-Energy Astrophysical Neutrino Flavor Measurements*, *JCAP* **04** (2021) 054 [arXiv:2012.12893] [INSPIRE].

[29] KamLAND collaboration, *Precision Measurement of Neutrino Oscillation Parameters with KamLAND*, *Phys. Rev. Lett.* **100** (2008) 221803 [arXiv:0801.4589] [INSPIRE].

[30] KamLAND collaboration, *Constraints on θ_{13} from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND*, *Phys. Rev. D* **83** (2011) 052002 [arXiv:1009.4771] [INSPIRE].

[31] KamLAND collaboration, *Reactor On-Off Antineutrino Measurement with KamLAND*, *Phys. Rev. D* **88** (2013) 033001 [arXiv:1303.4667] [INSPIRE].

[32] Daya Bay collaboration, *Observation of electron-antineutrino disappearance at Daya Bay*, *Phys. Rev. Lett.* **108** (2012) 171803 [arXiv:1203.1669] [INSPIRE].

[33] Daya Bay collaboration, *Spectral measurement of electron antineutrino oscillation amplitude and frequency at Daya Bay*, *Phys. Rev. Lett.* **112** (2014) 061801 [arXiv:1310.6732] [INSPIRE].

[34] Daya Bay collaboration, *New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay*, *Phys. Rev. Lett.* **115** (2015) 111802 [arXiv:1505.03456] [INSPIRE].

[35] Daya Bay collaboration, *Measurement of electron antineutrino oscillation based on 1230 days of operation of the Daya Bay experiment*, *Phys. Rev. D* **95** (2017) 072006 [arXiv:1610.04802] [INSPIRE].

[36] Daya Bay collaboration, *Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay*, *Phys. Rev. Lett.* **121** (2018) 241805 [arXiv:1809.02261] [INSPIRE].

[37] S.A.R. Ellis, K.J. Kelly and S.W. Li, *Leptonic Unitarity Triangles*, *Phys. Rev. D* **102** (2020) 115027 [arXiv:2004.13719] [INSPIRE].

[38] S.A.R. Ellis, K.J. Kelly and S.W. Li, *Current and Future Neutrino Oscillation Constraints on Leptonic Unitarity*, *JHEP* **12** (2020) 068 [arXiv:2008.01088] [INSPIRE].

[39] C.A. Argüelles, T. Bertólez-Martínez and J. Salvado, *Impact of Wave Packet Separation in Low-Energy Sterile Neutrino Searches*, `arXiv:2201.05108` [InSPIRE].

[40] E. Akhmedov and A.Y. Smirnov, *Damping of neutrino oscillations, decoherence and the lengths of neutrino wave packets*, *JHEP* **11** (2022) 082 [`arXiv:2208.03736`] [InSPIRE].

[41] M. Dentler, A. Hernández-Cabezudo, J. Kopp, M. Maltoni and T. Schwetz, *Sterile neutrinos or flux uncertainties? — Status of the reactor anti-neutrino anomaly*, *JHEP* **11** (2017) 099 [`arXiv:1709.04294`] [InSPIRE].

[42] JUNO collaboration, *Neutrino Physics with JUNO*, *J. Phys. G* **43** (2016) 030401 [`arXiv:1507.05613`] [InSPIRE].

[43] JUNO collaboration, *Sub-percent precision measurement of neutrino oscillation parameters with JUNO*, *Chin. Phys. C* **46** (2022) 123001 [`arXiv:2204.13249`] [InSPIRE].

[44] T. Abrahão, H. Minakata, H. Nunokawa and A.A. Quiroga, *Constraint on Neutrino Decay with Medium-Baseline Reactor Neutrino Oscillation Experiments*, *JHEP* **11** (2015) 001 [`arXiv:1506.02314`] [InSPIRE].

[45] Y.P. Porto-Silva, S. Prakash, O.L.G. Peres, H. Nunokawa and H. Minakata, *Constraining visible neutrino decay at KamLAND and JUNO*, *Eur. Phys. J. C* **80** (2020) 999 [`arXiv:2002.12134`] [InSPIRE].

[46] A. de Gouvea, V. de Romeri and C.A. Ternes, *Probing neutrino quantum decoherence at reactor experiments*, *JHEP* **08** (2020) 018 [`arXiv:2005.03022`] [InSPIRE].

[47] P. Huber, H. Minakata, D. Minic, R. Pestes and T. Takeuchi, *Neutrino oscillations at JUNO, the Born rule, and Sorkin's triple path interference*, *Phys. Rev. D* **105** (2022) 115013 [`arXiv:2105.14061`] [InSPIRE].

[48] V.S. Basto-Gonzalez, D.V. Forero, C. Giunti, A.A. Quiroga and C.A. Ternes, *Short-baseline oscillation scenarios at JUNO and TAO*, *Phys. Rev. D* **105** (2022) 075023 [`arXiv:2112.00379`] [InSPIRE].

[49] E. Marzec and J. Spitz, *Neutrino decoherence and the mass hierarchy in the JUNO experiment*, *Phys. Rev. D* **106** (2022) 053007 [`arXiv:2208.04277`] [InSPIRE].

[50] T2K collaboration, *Improved constraints on neutrino mixing from the T2K experiment with 3.13×10^{21} protons on target*, *Phys. Rev. D* **103** (2021) 112008 [`arXiv:2101.03779`] [InSPIRE].

[51] NOvA collaboration, *Improved measurement of neutrino oscillation parameters by the NOvA experiment*, *Phys. Rev. D* **106** (2022) 032004 [`arXiv:2108.08219`] [InSPIRE].

[52] A. de Gouvêa, G. Jusino Sánchez and K.J. Kelly, *Very light sterile neutrinos at NOvA and T2K*, *Phys. Rev. D* **106** (2022) 055025 [`arXiv:2204.09130`] [InSPIRE].

[53] P.B. Denton, J. Gehrlein and R. Pestes, *CP-Violating Neutrino Nonstandard Interactions in Long-Baseline-Accelerator Data*, *Phys. Rev. Lett.* **126** (2021) 051801 [`arXiv:2008.01110`] [InSPIRE].

[54] S.S. Chatterjee and A. Palazzo, *Nonstandard Neutrino Interactions as a Solution to the NOνA and T2K Discrepancy*, *Phys. Rev. Lett.* **126** (2021) 051802 [`arXiv:2008.04161`] [InSPIRE].

[55] DUNE collaboration, *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF*, FERMILAB-DESIGN-2016-02 (2015), `arXiv:1512.06148` [InSPIRE].

[56] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics*, FERMILAB-PUB-20-025-ND (2020), [arXiv:2002.03005](https://arxiv.org/abs/2002.03005) [ɪɴSPIRE].

[57] HYPER-KAMIOKANDE PROTO- collaboration, *Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande*, *PTEP* **2015** (2015) 053C02 [[arXiv:1502.05199](https://arxiv.org/abs/1502.05199)] [ɪɴSPIRE].

[58] HYPER-KAMIOKANDE collaboration, *Hyper-Kamiokande Design Report*, [arXiv:1805.04163](https://arxiv.org/abs/1805.04163) [ɪɴSPIRE].