Neutrino Oscillations

Moral: Just because a calculation achieves the correct answer, that does not necessarily serve as evidence that its methods or assumptions are correct.

The Particle Data Group (PDG) produces a giant tome, amassing the important measurements and results in particle physics. It also contains review articles, which are supposed to summarize the state-of-the-art theoretical aspects of a subfield of research. Its section on Neutrino Masses, Mixing, and Oscillations has a typical derivation of neutrino flavor mixing in vacuum that is presented to most graduate students in physics. In particular, the PDG derivation assumes that the neutrino mass eigenstates (1) all have equal momenta, and (2) all travel the same distance over the same amount of time. These assumptions are often stated without clear justification. In fact, we'll see they're unjustified!

First, some background information. Neutrinos used in experiments are produced in weak processes, where a source particle (P) decays, producing a neutrino ν along with one or many other particles (call the set of these particles X): $P \to \nu_\ell + X$. The neutrino produced in such a process will be of one of three flavors (denoted by the index $\ell = e, \mu, \tau$): electron type (ν_e), muon type (ν_μ), or tau type (ν_τ), and which type depends on the details of the process. These three ν_ℓ are called flavor eigenstates. Now for the weird part: each flavor eigenstate is a linear superposition of three mass eigenstates ν_1, ν_2, and ν_3. These mass eigenstates are aptly named, as they all have well-defined masses. Neutrino masses are quite small compared to every other mass scale, so they will be ultrarelativistic in nearly all experimental setups. These mass eigenstates then propagate through space as time, and are detected a certain distance away on the other end of the experiment in the process $D + \nu_{\ell'} \to$ (stuff that's detected), where D is the detector particle, and the stuff that's detected depends on the flavor ℓ'. This problem asks you to calculate the probability that a neutrino of flavor ℓ will be produced at one end of the experiment and will be detected as flavor ℓ' at the other end. We'll get the same neutrino oscillation formula as the one used by experiments, but the way we calculate it should shouldn't include any theoretical funny business.

(a) We'll begin by replicating the PDG calculation. This method will use a theoretical widget that represents a single neutrino at a certain point in spacetime: $|\nu(t, x)\rangle$. We're only going to consider one dimension of space, which we will take to be the ray between where the neutrino is produced and where it is detected.

First, the fact that a flavor eigenstate ν_ℓ is a linear superposition of mass eigenstates is represented as:

$$|\nu_\ell(t,x)\rangle = \sum_i U^*_{\ell i} \, |\nu_i(t,x)\rangle \, , \qquad \sum_i U^*_{\ell i} U_{i\ell'} = \delta_{\ell\ell'} \tag{1}$$

Second, this single-particle quantum state is translated through spacetime in the usual way for a free particle:

$$e^{-i(\hat{H}t-\hat{P}x)}\left|\nu_\alpha(0,0)\right\rangle = \left|\nu_\beta(t,x)\right\rangle \tag{2}$$

where $\hat{H}$ and $\hat{P}$ are the Hamiltonian and momentum operators, respectively, and α and β symbolize indices indicating either flavor eigenstates or mass eigenstates. However, note that the action of operators $\hat{H}$ and $\hat{P}$ are only defined on the mass eigenstates, so:

$$\hat{H}\left|\nu_i(t,x)\right\rangle = E_i\left|\nu_i(t,x)\right\rangle, \qquad \hat{P}\left|\nu_i(t,x)\right\rangle = p_i\left|\nu_i(t,x)\right\rangle \tag{3}$$

where E_i and p_i are the energy and momentum of the mass eigenstate i. Finally, the inner products are defined as:

$$\langle\nu_\alpha(t,x)|\nu_\beta(t,x)\rangle = \delta_{\alpha\beta} \tag{4}$$

where α and β are indices that can either indicate a flavor eigenstate or mass eigenstate label. Note that the above inner product is only defined at the same position in spacetime, so if you want to compare states that are not in the same position in spacetime, you'll need to translate using Eq. (2).

Consider a source particle P at the origin of spacetime $(t,x)=(0,0)$. It then decays, producing a neutrino ν_ℓ, which in turn can be expressed as a linear superposition of mass eigenstates:

$$\left|\nu_\ell(0,0)\right\rangle = \sum_i U^*_{\ell i}\left|\nu_i(0,0)\right\rangle \tag{5}$$

Then the mass eigenstates travel from one end of the experiment to the other, a distance L away. The probability that a neutrino of flavor ℓ' is detected at spacetime point (T,L), where T is the transit time, is:

$$P_{\nu_\ell\to\nu_{\ell'}} = \left|\langle\nu_{\ell'}(T,L)|\nu_\ell(0,0)\rangle\right|^2 \tag{6}$$

Since the neutrinos are ultrarelativistic, we can make the following three approximations:

(i) The transit time and the distance traveled are extremely similar in value $T \simeq L$.

(ii) Momenta of the mass eigenstates are extremely similar in value $p_i \simeq p_j \equiv p$.

(iii) The energy of the mass eigenstates can be approximated by

$$E_i = \sqrt{p_i^2+m_i^2} \simeq p + \frac{m_i^2}{2p} + \mathcal{O}(m_i^4) \tag{7}$$

where m_i is the mass of the mass eigenstate ν_i.

Given this setup, what is the transition probability $P_{\nu_\ell\to\nu_{\ell'}}$, in terms of L, m_i, and p?

(b) Do the three approximations in part (a) seem peculiar? If the neutrinos are produced at the same time, wouldn't they have different transit times, reaching the detector at slightly different times? Different mass eigenstates have different masses, so wouldn't they have different momenta? The PDG and other textbooks with this derivation do not fully explain why these assumptions are made!

Calculate the on-shell energy and momenta of all particles participating in the neutrino production. Let the source particle has mass M and the invariant mass of the set of particles produced along with the neutrino to be M_X. Find a the perturbative expansion in powers of m_i for the neutrino momentum, energy, and transit time. For simplicity, assume the source particle is at rest. Express yours answer in terms of M, M_X, and m_i. Are the three approximations in part (a) correct?

(c) Calculate the neutrino oscillation probability $P_{\nu_\ell \to \nu_{\ell'}}$ again, but this time assume the neutrinos are produced at the same time but reach the detector at slightly different times, and use the on-shell calculations you did in part (b), with self-consistent perturbative expansion in powers of m_i. Express your answer in terms of M, M_X, m_i, and L.

(d) Now the punch to the gut: you should have found that the neutrino oscillation phases in parts (a) and (c) differ by a factor of 2. Even though the approximations made in part (a) seemed plucked from the sky, and the approximations leading to part (c) seemed more reasonable (since they're based on energy and momentum conservation and basic kinematics), the answer to part (a) is actually the oscillation equation that is universally excepted as correct. What's going on here?

"I'd just like to know what in hell is happening, that's all! I'd like to know what in hell is happening! Do you know what in hell is happening?"

Let's make an attempt to resolve this paradox. Just as a shot in the dark, maybe you think that energy and momentum conservation is a safe thing to assume, but you don't think that it's safe to assume that neutrinos constructively interfere at the detector if they arrive at slightly different times. Calculate the oscillation probability $P_{\nu_\ell \to \nu_{\ell'}}$ again, but assuming that all a neutrinos have the same transit time $T_i = T \simeq L$. You should find that your answer results in the same answer at part (a). Does this actually resolve the paradox?

(e) We have two expressions for the neutrino oscillation probability, differing by a factor of 2 in the phase. Which one is right? It's not totally clear, given only the ingredients in this problem so far. There has been much ink spilled on this topic since neutrino oscillations were first discussed by Pontecorvo in the 1950's, but usually the conclusion in the literature is in favor of the answer to part (a). These papers often invoke the concept of a "neutrino wave packet" for each mass eigenstate, which spread away from each other on their journey to the detector (since they have different masses). However, even though this picture has different neutrino transit times, this wave packet approach does not result in the oscillation phase in part (c), and instead gives the same answer as part (a), which is usually explained with the idea of "smearing." But so many new ingredients! And where do neutrino wave packets come from anyways? Shouldn't one be able to derive all of this from first principles in quantum mechanics? Is there a semi-classical approach that can be extracted from this more general method that we can use to resolve this paradox once and for all?

The resolution is that the flavor eigenstate ν_ℓ is not actually a particle: it is not an irreducible representation of the Lorentz group, it does not have a mass, it does not have a position in spacetime. The object $|\nu_\ell(t,x)\rangle$ makes zero sense; it's an artifact of cultural baggage – a bookkeeping device to keep track of what linear superposition of mass eigenstates you need to consider. So, the error, the reason for all this confusion, was the very first step. The calculations in parts (a), (c), and (d) cannot be trusted, even if one of them may give the right answer. How do we fix this up?

You fix it up by doing the calculation for $P_{\nu_\ell \to \nu_{\ell'}}$ from scratch in quantum field theory, which is designed to handle these question about spacetime. These calculation are beyond the scope of this current problem. See, for example, this paper. But after doing the calculation, you can extract what the semi-classical assumptions should be:

(i) All particles travel in straight lines in spacetime.

(ii) Energy and momentum are conserved at the source's decay vertex, i.e., $p_S^\mu = p_X^\mu + p_i^\mu$, where p_S^μ, p_X^μ, and p_i^μ are the on-shell 4-momenta of the source, X state, and neutrino, respectively.

(iii) Different neutrino mass eigenstates are emitted from the source particle at different points in spacetime x_i^μ, but are detected at the same spacetime location x_D^μ.

(iv) The total oscillation phase for diagram includes contributions from the phase of the source, the phase of the X state, and the phase of the neutrino:

$$-ip_S \cdot x_i + ip_X \cdot x_i - ip_i \cdot (x_D - x_i) \tag{8}$$

(Hint: can you simplify this phase using energy and momentum conservation?)

(v) The 4-momentum of the neutrino is the following:

$$p_i^\mu = \left(\gamma E + \gamma v\sqrt{E^2 - m_i^2}, \gamma\sqrt{E^2 - m_i^2} + \gamma v E\right) \tag{9}$$

where v is the velocity of the source in the lab frame, and E is the energy of the neutrino in the CM frame of the source, as if the neutrinos were massless, i.e., $E = (M^2 - M_X^2)/2M$.

To illustrate why assumption (iii) is the case, consider Fig. 1. In quantum mechanics, you have to sum over variables that you don't measure. So, since you don't know when the source decayed, you have to sum over it at the amplitude level.

Using this picture, calculate again $P_{\nu_\ell \to \nu_{\ell'}}$. You should get the same answer as part (a). Should you trust this result?

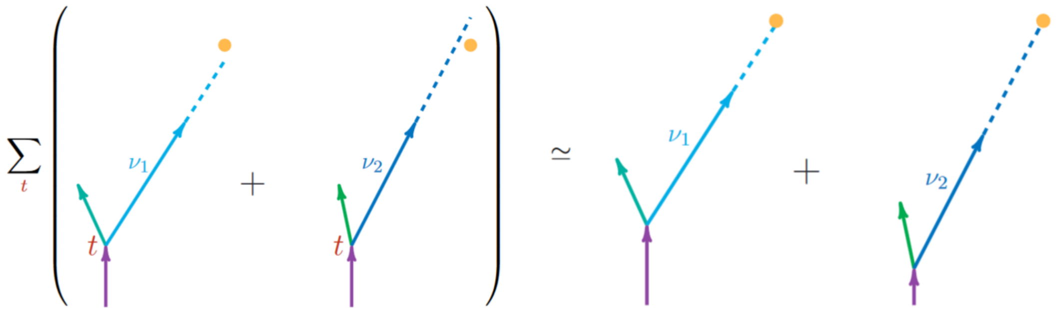

Figure 1: A cartoon of how neutrinos ν_1 and ν_2, where $m_2 > m_1$, are emitted at different times so they reach the detector (yellow dot) at the same point in spacetime. Here, the horizontal direction is space, and the vertical direction is time.