Exercises

1. We will start by looking into the seesaw mechanism. In a single-generation scenario, the neutrino mass matrix would look like

$$\mathcal{L} \supset -\frac{1}{2} \begin{pmatrix} \nu & \nu^c \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & \mu \end{pmatrix} \begin{pmatrix} \nu \\ \nu^c \end{pmatrix} + \text{h.c.} \tag{0.1}$$

   (a) Find the eigenvalues of this matrix. Expand them in $m_D/\mu \ll 1$.

   After diagonalization, we get $m_{2,1} = \frac{1}{2}\left(\mu \pm \sqrt{\mu^2 + m_D^2}\right)$. Expanding that, we get $m_1 \simeq -m_D^2/\mu$ and $m_2 = \mu$. The latter corresponds to the mass of the heavy neutrino ν_h and the former to the mass of the light neutrino ν_ℓ.

   (b) What is the rotation angle that diagonalizes this symmetric matrix? Call the physical basis ν_ℓ and ν_h, where ν_h is the heavy neutrino, sometimes called a heavy neutral lepton (HNL). Which will interact more strongly with the rest of the SM particles?

   You can always diagonalize a 2×2 symmetric matrix

$$M = \begin{pmatrix} a & c \\ c & b \end{pmatrix} \text{ as } \mathrm{diag}(m_1, m_2) = R(\theta)^T.M.R(\theta) \text{ with}$$

   where $\tan 2\theta = \frac{2c}{b-a}$ and $R(\theta)$ the rotation matrix. In this case, the mixing angle in this scenario is $\tan 2\theta = m_D/\mu$, which, to leading order, implies $\theta \sim m_D/\mu$. Multiplying things out, you should see that the light mass state is the one that most closely resembles ν (the interacting neutrino state), and therefore, is the one that interacts more with SM particles (through the Weak force).

   (c) Identify ν_ℓ state with the neutrinos we observe in the laboratory and take $m_{\nu_\ell} \sim 0.05$ eV. Recalling that m_D comes from the Higgs mechanism, $m_D = yv/\sqrt{2}$ with $v = 246$ GeV and y a coupling constant, we will take m_D to be in the rage $[m_e, m_t] = [511 \text{ keV}, 176 \text{ GeV}]$, corresponding to $y \sim [3 \times 10^{-6}, 1]$. What are the values of μ that you need in order to obtain the assumed value for m_{ν_ℓ}? What values of the mixing angle do these correspond to?

   Putting in the numbers, we get roughly, $\mu \sim [10^3 \text{ GeV}, 10^{14}\text{GeV}]$ and $\theta \sim [10^{-7}, 10^{-12}]$. You can see that even if the Dirac mass term is as small as the electron one, the required Majorana mass (and therefore the mass of the propagating state ν_h) will be very large and will be quickly untestable as you raise m_D. All relevant cross sections and decay rates for detecting ν_h, in fact, scale as θ^2, so that makes these particles very hard to observe in the laboratory, even if we had the energy to produce them. We say then, that ν_h has *weaker-than-Weak* interactions, meaning that the only part of ν_h that does interact (a fraction θ of its superposition) interacts only through the Weak force.

2. This exercise will explore low-scale variants of the seesaw mechanism. The idea will be to turn the seesaw on its head and explore symmetries to make neutrino masses small.

Low-scale seesaws (sometimes referred to as double, extended, linear, or inverse seesaws) introduce a pair of neutral lepton per neutrino family (call them ν^c and n, for example). In a single-generation scenario, the neutrino mass matrix becomes

$$\mathcal{L} \supset -\frac{1}{2}\begin{pmatrix}\nu & \nu^c & n\end{pmatrix}\begin{pmatrix}0 & m_D & \varepsilon \\ m_D & \mu' & M \\ \varepsilon & M & \mu\end{pmatrix}\begin{pmatrix}\nu \\ \nu^c \\ n\end{pmatrix} + \text{h.c.} \tag{0.2}$$

(a) The quantum numbers of the neutral leptons are $L(\nu) = 1$, $L(\nu^c) = -1$, and $L(n) = 1$. Out of all the parameters in the matrix above, which ones violate lepton number?

By adding the quantum numbers of each mass term, we can deduce that μ, μ', and ϵ violate Lepton number.

(b) Calculate the determinant of this matrix. What happens when all the lepton-number-violating (LNV) parameters go to zero? What does this imply for light neutrino masses?

$\det(M) = 2\epsilon M m_D - \epsilon^2\mu' - \mu m_D^2$. When all the LNV parameters go to zero, this determinant vanishes! This means that at least one of the eigenvalues is zero (recall that $\det(M) = \Pi_i\lambda_i$ and $\mathrm{Tr}(M) = \Sigma_i\lambda_i$). So, when lepton number is conserved, the model is no good for generating light neutrino masses. Note that in the standard seesaw (exercise 1), the LNV parameter (μ) is huge, so everything here is consistent.

(c) Keeping all these lepton-number-violating parameters zero and assuming the Dirac mass m_D is much smaller than all other parameters, diagonalize the matrix above and analyze what happens with the pair ν^c and n. What kind of particle are they?

By diagonalizing this matrix, we get $m_1 = 0$, $m_{2,3} = \mp\sqrt{M^2 + m_D^2}$. Note that this is a Dirac fermion! There are two degenerate (chiral/Weyl/2-component) fermions, which will make up one single Dirac particle, just like in the case of the electron or muons, etc. When m_D is small, this is easy to understand, the mass matrix looks very close to what we had for the electron case, (ee^c), except that there is an additional, unpaired fermion, namely the light neutrino.

(d) Turn on each LNV parameter at a time and diagonalize the matrix assuming $m_D \ll$ LNV parameters $\ll$ anything else. Mathematica, Wolfram Alpha, or Python are your friends here. You can assume all parameters are real and positive. What happens to the heavy neutral lepton masses? What kind of particle are they?

I won't bother with the full matrix, but the important part is the expression for the lightest mass,

$$m_\nu = \frac{\mu m_D^2 + \varepsilon\mu' - 2M\varepsilon m_D}{M^2 - \mu'\mu}. \tag{0.3}$$

Again, when all LNV parameters go to zero, this goes to zero. This gives us an additional way to make neutrino masses small! We can say that lepton number is *almost* conserved and that gives us an extra "knob" to control neutrino masses. The two other masses are roughly $m_{2,3} \sim \mp M + \mathcal{O}(\text{LNV parameters}/M)$ — this means that the mass eigenstates 2 and 3 combine into a Quasi-Dirac fermion. The two (Weyl) fermions have almost the same mass, but are split by the scale of LNV parameters. So, in the limit of lepton number (quasi-)preservation, the heavy neutrinos become more and more Dirac like.

(e) Estimate the mixing between the weak-interaction eigenstate ν and the heavy neutral leptons. What is the relationship between light neutrino masses and the mixing?

You can show that the (small) mixing angles are still given by things of the order of m_D/M. This is great because the small LNV parameters that can reduce the size of neutrino masses *do not* reduce the size of the mixing. The great advantage of low-scale seesaws is that the neutrino mass mechanism is a lot more testable. They are also "natural" in the technical sense: neutrino masses are small because there is a symmetry that is enhanced when they go to zero. This is the opposite of what happens in the usual seesaw: there the neutrino masses are small because the LNV scale is huge.

3. This exercise will show you how to estimate the signal event rate for heavy neutrinos produced by pion decay-at-rest sources and decaying inside a large-volume experiment. Consider a spallation source with $\mathcal{O}(\text{GeV})$ protons hitting a dense target. The number of pions produced at this source is related to the number of protons on target by $N_\pi \sim 10\% \times N_{\text{POT}}$.

(a) In terms of the mixing angle of the heavy neutrino with the muon neutrinos, $|U_{\mu 4}|^2$ [1], what is the flux of HNLs at a distance L from the target?

HNLs will be produced by pion decay. Since pions decay predominantly as $\pi^+ \to \mu^+ \nu_\mu$, we can just swap $\nu_\mu \to N$ by paying the price of $|U_{\mu 4}|^2$. For N_{HNLs} Given $N_\pi \sim 10\% \times N_{\text{POT}}$, we then have $N_{\text{HNLs}} \sim 10\% \times N_{\text{POT}} \times |U_{\mu 4}|^2$. Since we want the flux, we divide by the area of the sphere centered around the source (note the neutrino and HNL emission is isotropic!),

$$\Phi_{\text{HNLs}} = \frac{N_{\text{HNLs}}}{4\pi L^2}. \tag{0.4}$$

(b) Knowing that the muon lifetime is determined by the decay rate $\Gamma = G_F^2 m_\mu^5/192\pi^3$, estimate the analogous decay rate for $\nu_h \to \nu e^+ e^+$.

Well, no trick here, just again, take $\Gamma_{\text{N}} \sim |U_{\mu 4}|^2 G_F^2 m_\mu^5/192\pi^3$. In principle, since I did not specify what kind of outgoing neutrino we have in the decay, you could also write $\Gamma_{\text{N}} \sim (|U_{e4}|^2 + |U_{\mu 4}|^2 + |U_{\tau 4}|^2) G_F^2 m_\mu^5/192\pi^3$, corresponding to the sum of $\nu_h \to \nu_e e^+ e^-$, $\nu_h \to \nu_\mu e^+ e^-$, and $\nu_h \to \nu_\tau e^+ e^-$. This is not exactly correct since we are neglecting factors of order 1 and the fact that the decay

[1] Note this is analogous to θ^2 in exercise 1, except that it is specific to the muon neutrino flavor

$\nu_h \to \nu_e e^+ e^-$ can proceed through either W or Z, but it is good enough for our estimates.

(c) What is the probability for ν_h to survive the distance L without decaying? What is then the total probability for it to decay within a region of volume $V = \lambda^3$ located at a distance L from the source? You can assume $\lambda \ll L$.

A particle decay survival probability is given by $P = e^{-\tau/\tau_L}$ with τ_L the lab-frame lifetime. In terms of the decay with, $\tau_L = \gamma/\Gamma$, where γ is the boost of the HNL. Now, we want the probability to decay inside the detector after a distance L. First, the particle must survive a distance L, so $P_{\text{survival}} = e^{-\Gamma_{\mathrm{N}} L/\beta\gamma}$, where β is the velocity of the HNL, $\beta = v/c$. Now we want it to decay within a linear segment inside the detector. Let's assume it is of size λ, so $P_{\text{decay}} = 1 - e^{-\Gamma_N \lambda/\beta\gamma}$. Overall, the probability for a HNL to make it to the detector and decay inside of it is just the product of these two things: $P_{\text{decay}} \times P_{\text{survival}}$.

(d) Putting all of this together, write down a formula for the number of HNL decays within this detector of size V. You can assume the HNL decays only via $N \to \nu e^+ e^-$ with a rate Γ_N. Expand this formula assuming $\Gamma_N L$ and $\Gamma_N \lambda$ are very small. How does the signal event rate depend on the decay rate Γ_N? What if we add another decay channel for the HNL, like $N \to \nu\nu\bar{\nu}$?

First we want the number of HNLs that pass through our detector. We can estimate it assuming the area of the detector is roughly λ^2. In that case, the total number of HNL decays is

$$N_{\text{tot}} = \frac{N_{\text{HNLs}}}{4\pi L^2} \times \lambda^2 \times P_{\text{decay}} \times P_{\text{survival}} \sim \frac{N_{\text{HNLs}}}{4\pi L^2} \times \lambda^2 \times \frac{\Gamma_N \lambda}{\beta\gamma}, \tag{0.5}$$

where we expanded the exponentials to first order. Note that this expression is actually proportional to the volume of the detector, $V \sim \lambda^3$. It is also proportional to the signal rate Γ_N. If you consider another decay rate, then we would need to include an additional branching ratio factor in the whole formula. You can convince yourself that in the end, since both exponentials only depend on the total decay rate $\Gamma_N = \Gamma_{\nu e^+ e^-} + \Gamma_{\nu\nu\nu}$, in the long-lifetime regime ($\Gamma_N L/\gamma \ll 1$ and $\Gamma_N \lambda/\gamma \ll 1$), the total lifetime drops out of the formula and the signal is only proportional to the decay rate you are interested in ($\Gamma_{\nu e^+ e^-}$ in our case).

(e) In your preferred coding language, calculate this signal rate for fixed values of $|U_{\mu 4}|^2$ and m_4. Do this for a range of $|U_{\mu 4}|^2 \in [10^{-8}, 0.1]$ and $2\text{ MeV} < m_4 < 30\text{ MeV}$. Try drawing contours of constant event rate in this parameter space.

(f) Suppose the HNL decays via a stronger $G_X \gg G_F$ four-fermion interaction with SM particles. What happens to the contours in that case? Try, for instance, $G_X \sim 10^3 G_F$.