1.3 Newtrinos?

Now let's start to consider puzzle-led speculations about potential unexpected BSM discoveries at the HL-LHC. Let's take the first example in our list of open questions and see, with a relatively superficial theory-led glance, what could happen at the HL-LHC. Working in two-component notation for fermions, let us add a single left and single right-handed fermion N and N^c to the SM, where we will only consider one generation, for simplicity. Our theory is

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_{\text{Kin}}(N, N^c) + \lambda L H N^c + M N N^c \quad , \tag{1.4}$$

where $\mathcal{L}_{\text{Kin}}(N, N^c)$ denotes the usual kinetic terms. Consider the symmetries. This theory respects a lepton-number symmetry $(L, N) \to e^{i\theta_L}(L, N)$, $N^c \to e^{-i\theta_L} N^c$. Furthermore M is the only parameter which explicitly breaks an additional symmetry $N \to e^{i\theta_N} N$. As a result, this mass scale can be naturally at any value. There is no particular reason for it to be large or small. For further explanation of this 'spurion' reasoning see App. (A).

Now consider the spectrum of states. Below the scale of EW symmetry breaking one has, in the neutrino sector,

$$\mathcal{L} = \mathcal{L}_{\text{Kin}} + m\nu N^c + M N N^c \quad . \tag{1.5}$$

Thus we have two left-handed fermions and one right-handed fermion, as well as a 'chiral' $\text{U}(1)_L$ symmetry. As the theory is chiral there must exist one massless fermion. Thus far this theory has nothing to do with neutrino masses since it predicts vanishing neutrino mass by dint of the symmetry. We will return to this point later. Let us make the rotation

$$\begin{pmatrix} \nu \\ N \end{pmatrix} = \begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix} \begin{pmatrix} \tilde{\nu} \\ \widetilde{N} \end{pmatrix} \tag{1.6}$$

where

$$\tan\phi = \frac{m}{M} \quad . \tag{1.7}$$

This leads to the diagonalised mass basis

$$\mathcal{L} = \mathcal{L}_{\text{Kin}} + \sqrt{M^2 + m^2}\, \widetilde{N} N^c \quad , \tag{1.8}$$

where we see that $\tilde{\nu}$ is massless, as expected. Thus we have a massless, mostly-active, neutrino state admixed with a sterile fermionic state with mixing angle ϕ. However, we also have a massive Dirac sterile fermion, of mass $\widetilde{M} = \sqrt{M^2 + m^2}$ whose left-handed component is admixed with an active neutrino. In accordance with PDG convention, we will henceforth refer to this state as a 'Heavy Neutral Lepton' (HNL) since it carries lepton number.

At the LHC every time a W^- boson is created, which is often, it can decay to a lepton and anti-neutrino, conserving lepton number. However, if $\widetilde{M} \ll M_W$ then it can also decay to an HNL with a rate which scales as $\sin^2\phi$ relative to the SM rate. Following this the HNL can decay through an off-shell W^- to an anti-lepton and the products of the W^-, which could be a pair of jets or a lepton-anti-neutrino pair. In fig. 2 the former possibility is shown. What is intriguing is that the rate for this decay can be very suppressed by the small mixing angle. For a relativistic HNL the displacement is of the order

$$d_N \sim 15 \left(\frac{\text{GeV}}{M}\right)^5 \left(\frac{10^{-4}}{\sin^2\phi}\right) \ \text{m} \quad . \tag{1.9}$$

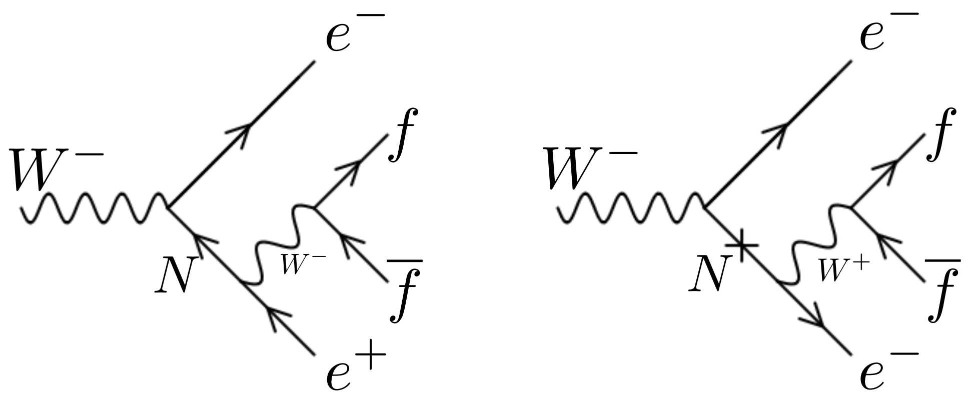

Figure 2: Left: Production of an HNL from W^- decay followed by HNL decay, which may be displaced. The final state respects lepton number and includes an opposite-sign, same-flavour pair of leptons alongside a pair of jets or lepton-antineutrino pair. Right: Production of a HNL from W^- decay followed by lepton-number violating oscillation into an anti-HNL and decay, which may also be displaced. The final state violates lepton number and includes a same-sign same-flavour pair of leptons alongside a pair of jets or anti-lepton - neutrino pair.

Sufficiently displaced, in fact, that the HNL can have propagated a finite distance before decay! This would be a truly spectacular signature of the appearance of a lepton and two jets, or two leptons, way out in the detector, seemingly from nothing!

Let's estimate the ultimate sensitivity one might have to such decays. Let us suppose that there is no SM background. This isn't strictly correct since, although the SM does not predict any processes with such spectacular displaced final states, detectors effects could, for instance, fake such processes. But, for the sake of illustration, let's consider that there is no background, such that any event would be a discovery. The production cross section for W boson production at 13 TeV is

$$\sigma^{13\text{ TeV}}_{W^-+W^+} \approx 2 \times 10^7 \text{ fb} \quad . \tag{1.10}$$

Note that, due to the composition of protons, W^+ and W^- cross sections are not the same. The branching ratio of a W boson into a single lepton-antineutrino flavour is 10%. Rescaling this by the mixing angle for a HNL we thus expect

$$N_{\text{Displaced}} \approx 2 \times 10^6 \sin^2\phi \times \mathcal{L}_{\text{Int}}(\text{fb}^{-1}) \quad , \tag{1.11}$$

spectacular displaced events. Inverting this to estimate the limit under the assumption that nothing has been seen we find the strict best-case-scenario limit

$$\sin^2\phi \lesssim 2 \times 10^{-10} \frac{3000 \text{ fb}^{-1}}{\mathcal{L}_{\text{Int}}} \quad . \tag{1.12}$$

We may compare this to the present limits, for instance from CMS as shown in fig. 3. We would have expected the strongest upper bound to be around $\sin^2\phi \sim 4 \times 10^{-9}$. However, we see that in fact the limits are far weaker. This leads me to ask you:

- What dictates the shape of the limit region?

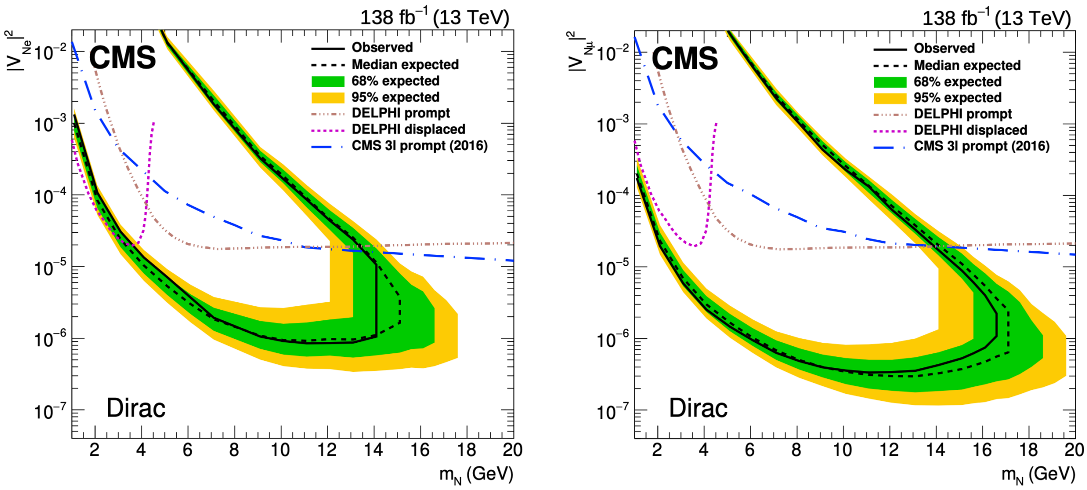

Figure 3: CMS search for displaced Dirac HNL decays [12]. ATLAS have also performed similar searches [13].

- What types of backgrounds could there be?
- What could limit signal statistics?

Since these are very low-background processes, every additional integrated fb^{-1} counts and the HL-LHC will break important additional new territory in the search for new physics. Here the appearance of lepton number in the final state would show that an HNL had been discovered. That alone would be reason for Champagne! I show in fig. 4 some estimates for future sensitivity, which indeed reflect the naïve estimate of eq. (1.12) better. The discrepancy between fig. 3 and fig. 4, in the sense that the growth of sensitivity is not just a factor $3000/138 \approx 20$ reflects the importance of the cuts used to suppress fake backgrounds in displaced searches and also the importance of considering as many channels as possible, since the latter also considers hadronic decay channels. Nonetheless, I think it is fair to assume that we can expect to probe a factor 20 or so deeper into new physics territory during the HL-LHC phase of operation.

Newtrino Oscillations?

The model, as introduced thus far, gives rise to a large mixing between active SM neutrinos and HNL, but does not generate neutrino masses. Thus if nature is to accommodate neutrino masses it must either introduce some additional states for the massless fermion to pair up with to form a Dirac fermion, or explicitly break the lepton number symmetry which, acting

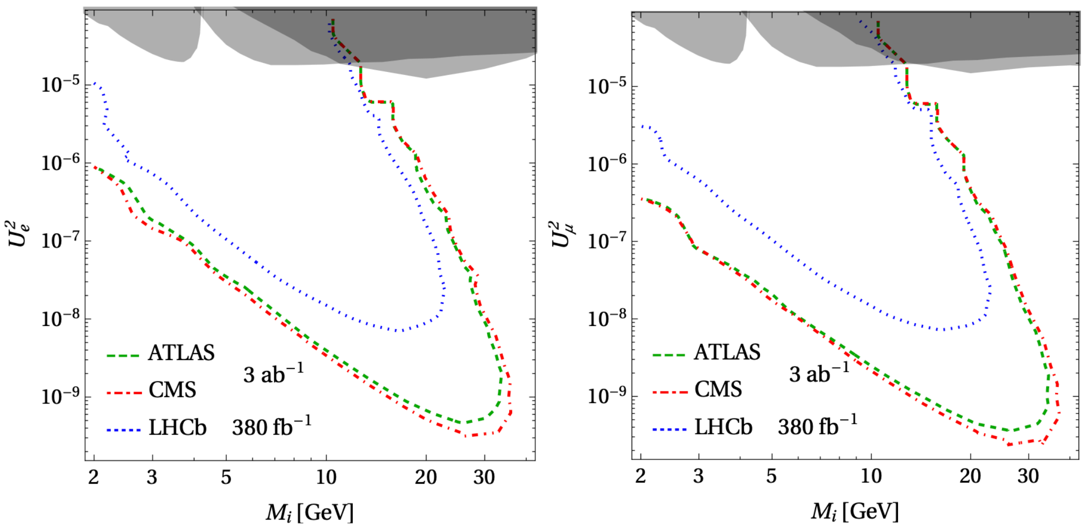

Figure 4: Projections for future HNL displaced decay searches at HL-LHC, taken from [14, 15].

as a type of chiral symmetry, was keeping one fermion massless. Let's consider the latter. There are a variety of ways in which one could break lepton number at the renormalisable level. For instance

$$\mathcal{L}_{\mathrm{U}(1)_L} = \frac{1}{2}\delta N^2 \quad , \quad \frac{1}{2}\delta^c N^{c2} \quad , \quad \delta_\lambda LHN \quad . \tag{1.13}$$

Whichever option is chosen, $\mathrm{U}(1)_L$ is explicitly broken and a Majorana neutrino mass will be generated. With the magnitude of breaking parameter chosen appropriately the observed active neutrino mass can be accommodated. This is all trivially extended to three generations.

However, in this instance something striking could arise at the HL-LHC. Now the massive Dirac HNL will also be split into two quasi-degenerate Majorana states. Importantly, the $\mathrm{U}(1)_L$ eigenstates are no longer aligned exactly with the mass eigenstates. Being so close in mass this means they can oscillate between $\mathrm{U}(1)_L$ eigenstates as they propagate, much in the same way that the known neutrinos oscillate amongst flavour eigenstates. See, for instance, [16]. This is depicted in the right panel of fig. 2. Whether or not such oscillations occur depends on whether they have completed before decay. The timescale for oscillation is proportional to the lepton-number violation scale, such as $1/\delta$ for the first example of eq. (1.13). This is bounded by not contributing too greatly to neutrino masses. In the present example the contribution to the mostly-active neutrino mass is[5]

$$\widetilde{m}_\nu \propto \delta^c \sin^2\phi \quad . \tag{1.14}$$

[5]Can you explain the scaling?

Therefore we have that

$$\delta^c \lesssim \frac{\text{eV}}{\sin^2\phi} \quad , \tag{1.15}$$

and the length scale associated with HNL oscillations satisfies

$$d_N^{\text{Osc}} \gtrsim 2\sin^2\phi \times 10^{-7}\ \text{m} \quad . \tag{1.16}$$

Comparison with eq. (1.9) reveals that there is ample distance for many oscillations of HNLs between lepton number eigenstates before decay. Thus, if one produces a HNL at the HL-LHC which is, for practical purposes, a $U(1)_L$ eigenstate carrying lepton number 1, then by the time it decays it can have oscillated into an admixture of 1 and -1 states. In other words, between production and decay lepton-number is violated, as depicted on the right hand side of fig. 2. There are no lepton-number violating processes in the SM thus such a final state would be very striking indeed. Furthermore, this is true even if the decay is prompt, without a displaced vertex. Such a signature could provide the first microscopic window onto the origin of neutrino masses!

2 Weighing up the Weak Scale

Now we turn to the second reason I expect BSM physics could show up at the HL-LHC. For me, it is the strongest motivation.

As I argued, since discovering the Higgs boson we have cornered the 'How?' of EW symmetry breaking. However, we don't yet have the 'Why?' I described the difference between the two in terms of the Ginzburg-Landau model. However, in Footnote 2 I also claimed there is also a strong analogy between the Higgs sector of the SM and the pions of QCD, that I would return to. Well, this is the moment.

As we know, in the 'UV' (i.e. short distance physics) the quarks and gluons of QCD have perturbative interactions amongst themselves. However, in the 'IR' (i.e. long distance physics) this coupling becomes strong and the quarks and gluons cannot be considered asymptotic states on long distance scales. Instead, they are bound into hadrons which become the true asymptotic states on distance scales greater than the size of the proton, for example. In this sense, the hadrons (especially the pions) and their associated 'IR' theory description known as the 'Chiral Lagrangian' are the 'How' of chiral symmetry breaking and confinement, and QCD is the 'Why'.

To study this more closely let's go back to life below 1 GeV. Working below this energy scale we observe that there are three pion degrees of freedom, packaged into a neutral pseudoscalar field π^0 and a charged field $\pi^\pm = \pi_1 \pm i\pi_2$, with masses $m_0 = 135$ MeV and $m_\pm = 140$ MeV. Clearly they are very close in mass, so one might assume there is some symmetry that enforces their mass to be equal. In fact, it is a good idea to think of these pions to be packaged into the adjoint representation of SU(2), as $\Pi = e^{\sum_i \pi_i \sigma_i / f_\pi}$, where the latter are simply the Pauli matrices, and an SU(2) transformation takes $\Pi \to U \Pi U$, where U is a unitary 2×2 matrix.

Now we may trivially write their mass in an explicitly symmetry-invariant manner

$$\mathcal{L}_{Mass} = \frac{1}{2} m_\pi^2 f_\pi^2 \text{Tr}\Pi \to \frac{1}{2} m_\pi^2 (\pi_0^2 + \pi_1^2 + \pi_2^2) + ... = \frac{1}{2} m_\pi^2 \pi_0^2 + m_\pi^2 \pi^+ \pi^- + ... \quad . \tag{2.17}$$

Lets do a spurion analysis. The parameter m_π is the only spurion that breaks a shift symmetry acting on the pions, thus if we think of this as an EFT then perturbative effects will not generate large corrections to their mass. Furthermore, this parameter respects the SU(2) symmetry, thus all quantum corrections will respect the symmetry and the pions will continue to have the same mass.

So far so good. We have a pretty decent theory for the pions. However, there is an elephant in the room. The charged pions interact with the photon through the kinetic terms

$$\mathcal{L}_{Kin} = \frac{1}{2}(\partial_\mu \pi_0)^2 + |(\partial_\mu + ieA_\mu)\pi^+|^2 \quad . \tag{2.18}$$

This interaction not only breaks the SU(2) symmetry, since it only affects the charged pions, but it also breaks their shift symmetry! Although it may look innocuous, this is not some minor modification of the theory. In fact, it *completely* destablises the entire setup. Even without performing any calculations we know we now have a spurion parameter e that breaks these symmetries, thus if we consider this as an effective field theory, which we should, then there is absolutely nothing to forbid corrections arising at the quantum level that scale as

$$\delta\mathcal{L}_{Mass} \sim \frac{e^2}{(4\pi)^2}\Lambda^2 \pi^+\pi^- \quad , \tag{2.19}$$

where the 4π factor is typical for a quantum correction. Now we have a hierarchy problem, since if $\Lambda \gtrsim 750$ MeV then we would have a huge puzzle, as these corrections would be greater than the observed mass splitting. How can we address this puzzle? The most obvious answer is that it must be the case that $\Lambda \lesssim 750$ MeV. In other words there *must* be new fields and interactions that become relevant at a scale of $E \sim 750$ MeV that will somehow tame these corrections. It turns out that nature did indeed choose this route, and in fact the ρ-meson shows up, alongside all the rest of the fields associated with QCD, and then eventually at higher energies the quarks and gluons themselves. All of this physics at the cutoff and above then explains why the pions mass splitting is what it is (see [17]). The actual correction is

$$m_{\pi^\pm}^2 - m_{\pi^0}^2 \approx \frac{3e^2}{(4\pi)^2}\frac{m_\rho^2 m_{a_1}^2}{m_\rho^2 + m_{a_1}^2}\log\left(\frac{m_{a_1}^2}{m_\rho^2}\right) \tag{2.20}$$

where ρ and a_1 are the lightest vector and axial vector resonances. So this hierarchy problem is resolved very clearly in QCD. The quadratic correction from electromagnetism very much exists and is calculable. New composite resonances kick in to tame these quadratic corrections, and soon after that, above the QCD scale, the pion itself is no longer a physical state as it is a composite made up of fermions. Fermions do not receive quadratic corrections to their mass, so we can understand why the pion mass splitting is not sensitive to physics at, for example, the Planck scale!

Imagine, however, that the expected new physics had not shown up at $E \sim 750$ MeV. We would have a huge puzzle and we would have to try and understand what is going on. We could simply add an additional parameter to our action

$$\delta\mathcal{L}_{Tune} \sim \delta_m^2 \pi^+\pi^- \quad , \tag{2.21}$$

and then fine-tune this against the other corrections to keep the sum small, however this would seem very ad. hoc. Nature did not choose this route. Instead, nature chose for the mass splitting to be natural (= not fine-tuned). In essence, the requirement of naturalness is satisfied precisely as we would expect from taking the measured mass splitting and turning it around to predict new fields at some energy scale!

Nowadays with the Higgs boson we are in a similar situation, except we are just working at higher energies. We have a scalar field, the Higgs. If the Lagrangian were simply the kinetic terms and its mass, then we would have no problem at all, because the mass would be the only parameter that breaks a shift symmetry for the Higgs, hence it would be stable against quantum corrections. However, we also have the gauge interactions that break any shift symmetry, just like the pions, but also more importantly the Yukawa interactions

$$\mathcal{L}_{Yukawa} = \lambda H Q U^c + h.c.... \quad . \tag{2.22}$$

These interactions also break the shift symmetry, where the top Yukawa is the most significant breaking term. We may thus pursue exactly the same reasoning as for the pions. Whatever the UV-completion of the Higgs sector, at the quantum level there should arise corrections to the Higgs mass that scale as

$$\delta\mathcal{L}_{Mass} \sim 6\frac{\lambda_t^2}{(4\pi)^2}\Lambda^2|H|^2 \quad . \tag{2.23}$$

In natural units $\lambda_t \approx 1$, thus for these mass-squared corrections to remain below the EW scale we require $\Lambda \lesssim 500$ GeV. Just as for the pions, unless some new physics kicks in around this scale we have an issue, which is that if the cutoff of the SM exceeds 500 GeV, then there must be some sort of fine-tuning taking place.

So, we see that the hierarchy problem is not some wishy-washy notion, but is in fact very crisp and familiar and it points directly to the $\sim$ TeV scale as somewhere where *something* ought to be going on. For the pions the reasoning of EFT worked beautifully, so what is going on with the Higgs? This question has pestered theorists for decades. In fact, a significant portion of all BSM theories so far proposed are either concerned with the hierarchy problem, or in some way framed within the context of a solution. I will now sketch some of the ideas that have been considered and how they inform possibilities for BSM physics at the HL-LHC.