UNIVERSITY
OF OSLO

Master's thesis

A phenomenological study of heavy neutrinos and the seesaw mechanism

Tobias Opdalshei Baumann

60 ECTS study points

Department of Physics
Faculty of Mathematics and Natural Sciences

Spring 2023

Tobias Opdalshei Baumann

A phenomenological study of heavy neutrinos and the seesaw mechanism

Supervisors:
Prof. Farid Ould-Saada
Dr. Eirik Gramstad

Abstract

In this thesis, we review the motivations for proposing the existence of heavy neutrinos and the so-called seesaw mechanism. In particular, we show how this mechanism provides a natural explanation for the small masses of the Standard Model (SM) neutrinos.

We study three specific realisations of the seesaw mechanism; two with heavy neutrinos that couple to SM bosons with the neutrinos being either Majorana or pseudo-Dirac fermions, and one version of the Left-Right Symmetric Model (LRSM) with heavy Majorana neutrinos. For each of the models we perform simulations of heavy neutrino production at the LHC, considering the three lepton and missing energy final state for the Majorana and pseudo-Dirac models and the two-lepton plus two jets (2l+2j) final state for the LRSM. For the Majorana and pseudo-Dirac models, we study neutrino masses m_N in the range 200 GeV to 1200 GeV, while for the LRSM we look at models with W_R boson masses M_{W_R} from 3 TeV to 6 TeV, and neutrino masses that are between $M_{W_R}/2$ and M_{W_R}.

We study the kinematics of the simulated events for the different masses, and for the three lepton final state we develop a method of identifying the lepton pair from the production and decay of the heavy neutrino that is up to 50% more efficient than simply ordering the leptons by their transverse momenta p_T. For the 2l+2j final state, we find that the invariant mass of the four objects in the final state can be used to reconstruct M_{W_R}, and that the p_T of the leptons can be used to estimate the mass difference between W_R and the heavy neutrino.

We then perform a sensitivity study for the ATLAS detector at the LHC for the three signal processes after the current Run 3 and the future HL-LHC. For the Majorana and pseudo-Dirac signals we find that they are very difficult to discover and impossible to exclude even after the entire HL-LHC run, due to the strict constraints on the mixing between the heavy neutrinos and the SM neutrinos. For the LRSM models, we find that we can improve on results from previous studies. For the case $M_{W_R} = 3$ TeV we are able to put a 95% CL expected upper limit on the relative right-handed coupling strength κ_R of $\kappa_R \approx 0.2$ if no signal is seen after Run 3, and an upper limit of less than 0.1 after HL-LHC. We also find that models with masses as high as $M_{W_R} = 5$ TeV and $m_N = 4$ TeV can be discovered at 5σ during HL-LHC. If no signal is seen after HL-LHC, it will be possible to expand the current 95% CL excluded region in the (M_{W_R}, m_N) parameter space up to at least (5 TeV, 4.5 TeV) and the 90% CL excluded region will extend beyond (6 TeV, 4 TeV).

Acknowledgements

First of all, I want to thank my supervisors Farid and Eirik, without whom I could not have completed this project. I am especially thankful to Farid for generously offering to be my supervisor and presenting me with an exciting thesis topic when I had no clue what I wanted to do for my master's. I also need to thank you for always showing me the right way forward when I felt lost in the vastness of this project over the last year. Eirik deserves a huge thank you for the countless times he has helped me with all sorts of practical problems. Without you I would still be trying to get MadGraph to work, or just trying to remember what that magical Linux command was again (I will probably still be messaging you whenever I have a Linux problem). You are also a big personal role model for me, in particular in the way you are combining work and family life with a close-to professional running career. I will always have fond memories of the time I managed to (at least almost) keep up with you on our ski trip around Gålåvatnet.

I also want to thank all my friends and colleagues at the HEP section at UiO for the great year here. A special shout out goes to my row mates Martin and Ruben, and the foosball table, who have kept my spirits up even in these last most stressful months. I also want to thank my friends in my second home section, the theory section at UiO. Without spending a year with you I would probably not have been able put so many equations in my thesis.

Finally, I want to thank all my non-physics friends (normal people, some would say) and family. I haven't talked to you much over the last few weeks, but I still like you, I promise. I want to give a special thanks to the amazing OSI Friidrett/SK Vidar group, who make endless threshold intervals feel like the best activity one could do. I look forward to running more with you very soon.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Although the neutrinos have been known to exist for over 60 years [1], they are arguably the particles in the Standard Model of particle physics (SM) that we know the least about. The reason for this is their very weak interaction with other particles, making them difficult to study experimentally. Among the most important mysteries of the neutrinos are the origin of their masses and what type of fermion they are; Dirac or Majorana. The answers to these questions will open doors to new and unknown physics. In fact, the first decisive evidence for physics beyond the SM came from measurements of neutrinos, when the observation of neutrino oscillations in 1998 showed that the neutrinos have non-zero masses [2]. Even before this was shown experimentally, many theoretical models for the mass generation of the neutrinos had been conceived. Among these is the so-called seesaw mechanism, which attempts to explain the small non-zero masses of the SM neutrinos by introducing new heavy neutrinos.

In this thesis we will perform a phenomenological study of a few different versions of the seesaw mechanism. We will simulate particle collision events that can produce the associated heavy neutrinos within the framework of each model, and analyse the results with the goal of learning how to distinguish between the models. We also want to use our simulated events along with simulated background events made for the current Run of the ATLAS detector at the LHC to perform a sensitivity study for this detector to these models. With this analysis, we hope to say something about whether the heavy neutrinos are possible to detect currently or in the future, and which models or parts of parameter spaces can be excluded based on not having observed the expected signals yet.

The thesis is divided into two parts. In the first part, we give the theoretical background needed for our study. In chapter 2, we give a short review of the SM, and in particular how the fermions in the SM get their masses. We also present the evidence for the masses of the SM neutrinos being non-zero, and why this leads to the so-called neutrino mass problem. In chapter 3, we introduce the three versions of the seesaw mechanism that are the focus of our study. Finally, in chapter 4, we give a brief introduction to the LHC, particle collider kinematics and the ATLAS detector.

In the second part of the thesis, we present the methods we used in our analyses and the results of our work. In chapter 5, we present the specific signal processes we will study within each of the three versions of the seesaw mechanism that we cover, and how we simulated these processes. We also give an overview of some important SM background processes that are similar to the signal processes. In chapter 6, we study the kinematics of the signal processes in detail to learn about their characteristics. In chapter

7, we perform the aforementioned sensitivity study for the ATLAS detector during Run 3 of the LHC and the High Luminosity LHC to each of our simulated signals. Finally, in chapter 8, we summarise our main results and present possibilities for future research.

Part I

Background

Chapter 2

The neutrino mass problem

In this thesis we will use natural units, setting $c = \hbar = 1$, meaning that mass, momentum and energy are all given units of energy. We will use the energy unit "electronvolt", or eV, which is defined by 1 eV $\approx 1.6 \times 10^{-19}$ J. We use the metric sign convention $g^{\mu\nu} = \text{diag}\,(+1, -1, -1, -1)$.

In this chapter we will introduce the main motivations for the search for heavy neutrinos and the study of the seesaw mechanism, from both theoretical and experimental points of view. When no other references are given, the information in this chapter is largely based on the textbooks by Peskin and Schroeder [3] and Thomson [4].

2.1 Introduction to the Standard Model

The Standard Model of particle physics (SM) is the model that describes all the currently known elementary particles and their interactions. The SM is a quantum field theory, meaning that the dynamical variables of the theory are fields, and we think of each field as having an associated particle. The "equations of motion" for the fields in a quantum field theory can be derived from the Lagrangian $\mathcal{L}$ of the theory, which contains the space-time dependent fields and their derivatives. The Lagrangian of the SM therefore gives all the information we need to predict the behaviour of the elementary particles. The SM Lagrangian is built on the principle of gauge invariance. This means that the Lagrangian is made to be invariant under certain symmetry transformations. The symmetries of the SM are defined by the so-called SM gauge group

$$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y. \tag{2.1}$$

Here $U(n)$ and $SU(n)$ are the unitary and special unitary groups of order n respectively, which are continuous groups with elements that can be represented as n by n unitary matrices. The multiplication symbol between each subgroup of the SM gauge group represents a direct product, while the subscripts C, L and Y are labels that remind us that only certain types of SM fields are charged under each group. Any element of one of the three subgroups can be written using the exponential map and the generators T^a, where a is an index labelling the different generators, of the group as

$$g = e^{i\theta_a T^a}. \tag{2.2}$$

Here θ_a are parameters that can be chosen to produce the specific group element g, and there is a sum over all values of a. Each subgroup of the SM gauge group then acts on the SM Lagrangian by transforming some fermion field ψ in a so-called gauge transformation

$$\psi \to \psi' = e^{-iq\theta_a(x)T^a}\psi, \tag{2.3}$$

where $\theta_a(x)$ are parameters which are now space-time dependent, and q is the charge of the field ψ under the specific subgroup. We therefore see that only fields with non-zero charges under a group will transform under a gauge transformation.

The Abelian group $U(1)$ only has one generator, which can be taken to be 1. The groups $SU(n)$ in general have $n^2 - 1$ generators, which in the so-called fundamental representation can be written as Hermitian n by n matrices. For $SU(2)$, the generators in the fundamental representation are the three Pauli matrices τ^a, $a = 1, \ldots, 3$, while for $SU(3)$ they are the eight Gell-Mann matrices λ^a, $a = 1, \ldots, 8$. This means that fields that transform under $SU(2)_L$ and $SU(3)_C$ need to be placed in two- and three component vectors respectively in order for equation (2.3) to make sense. These vectors are called $SU(2)_L$ doublets and $SU(3)_C$ triplets.

Now we can finally introduce the particle/field content of the SM. First, there are the fermions, which are characterised by their spin of 1/2 and their fields taking the form of four-component "spinors". The fermions of the SM include the leptons, which come in three generations and electrically charged and neutral types. The charged leptons are the electron $e^\pm$, the muon $\mu^\pm$ and the tau $\tau^\pm$. The neutral leptons are the neutrinos with flavours of the same name; the electron-, muon- and tau-neutrinos ν_e, ν_μ and ν_τ. The remaining fermions are the quarks, which also come in three generations and two types with different electric charges. There are the down-type quarks down d, strange s and bottom b with charges of $-1/3$, and the up-type quarks up u, charm c and top t with charges of $+2/3$. All of these fermions have left-handed fields[1] that are charged under the $SU(2)_L$ group. These left-handed fermion fields are therefore placed in $SU(2)_L$ doublets, as follows

$$Q_L^i = \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix} \quad E_L^i = \begin{pmatrix} e_L^i \\ \nu_L^i \end{pmatrix}, \tag{2.4}$$

where i is a generational index, such that Q_L^i represents the doublets of the quark fields of each generation, with for example u_L^i being the left-handed field of the up-quark, charm-quark and top-quark for $i = 1, 2, 3$ respectively. Similarly, E_L^i are the three generations of $SU(2)_L$ lepton doublet, with e_L^i being the fields of $e^\pm$, $\mu^\pm$ and $\tau^\pm$ for $i = 1, 2, 3$.

All of the quark fields are also charged under $SU(3)_C$, which means that they are additionally placed in $SU(3)_L$ triplets

$$q = \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix}, \tag{2.5}$$

where q can be any of the six types d, s, b, u, c and t.

To make the SM Lagrangian invariant under the general gauge transformation of the fermion fields shown in equation (2.3), we need to add new fields for each of the gauge groups to cancel the fermions' transformations. These are the so-called gauge (vector)

[1] We will explain what left-handed fields are in the next section.

bosons. There is one gauge boson associated with each generator of the SM gauge groups, meaning that there are eight gluons G^a_μ, $a = 1, \ldots, 8$ connected to the $SU(3)_C$ group, and there are the four electroweak bosons W^+_μ, W^-_μ, Z_μ and the photon A_μ connected to the $SU(2)_L$ and $U(1)_Y$ groups. The connection between the last four gauge bosons and the associated groups is actually slightly more complicated, since the gauge symmetry of $SU(2)_L \otimes U(1)_Y$ is "broken", which we will come back to in subsequent sections. These gauge bosons, which have fields with indices $\mu = 0, \ldots, 3$ that indicate that they transform as vectors under Lorentz transformations, mediate the interactions between the fermions in the SM that are charged under the gauge group associated with the boson. We therefore get the group of interactions between quarks mediated by gluons, which is described by the theory of Quantum Chromodynamics (QCD). There is also the theory of Electroweak (EW) interactions between the leptons, quarks and the bosons of $SU(2)_L \otimes U(1)_Y$. As mentioned, this gauge symmetry is broken in the SM, which occurs at an energy scale of around 100 GeV, and has the consequence that at lower energies there is the theory of Quantum Electrodynamics (QED) which describes the interactions between the photon and all electrically charged fermions.

Finally, the last field of the SM is the scalar field called the Brout-Englert-Higgs (BEH) field ϕ, with the associated scalar boson called the Higgs h. This field is also charged under $SU(2)_L$, and therefore forms a doublet

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}. \tag{2.6}$$

2.2 A problem with fermion masses in the Standard Model

To acquire mass for the fermions in the Standard Model is surprisingly non-trivial. For any SM fermion with a Dirac spinor ψ, the most natural attempt would be to, analogous to the Dirac Lagrangian, add the term

$$\mathcal{L} \sim -m\overline{\psi}\psi \tag{2.7}$$

to the SM Lagrangian. Here m would be a constant with mass dimension one, as is expected for a mass, and $\overline{\psi}$ denotes the Dirac adjoint of ψ, defined as

$$\overline{\psi} = \psi^\dagger \gamma^0, \tag{2.8}$$

where γ^0 is the time-like gamma matrix. However, terms like the one in (2.7) are not allowed in the SM Lagrangian due to the requirement that the Lagrangian must be invariant under the transformations of the SM gauge group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. To show that the term in (2.7) is not invariant under this SM gauge group, and specifically not under the transformations of $SU(2)_L$, we decompose the spinors into their left- and right-handed chiral components using the projection operators

$$\begin{aligned} P_L &= \frac{1}{2}\left(1 - \gamma^5\right), \\ P_R &= \frac{1}{2}\left(1 + \gamma^5\right), \end{aligned} \tag{2.9}$$

where γ^5 is the chirality operator, also referred to as "the fifth gamma matrix", and is defined by

$$\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3. \tag{2.10}$$

The four gamma matrices γ^μ, $\mu = 0, \ldots, 3$ are defined by the anticommutation relations

$$\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}\mathbb{I}_4, \tag{2.11}$$

where $\mathbb{I}_4$ is the four by four identity matrix. Using the properties of the four gamma matrices, one can show that the operators P_L and P_R have the standard projection operator properties

$$P_L + P_R = 1, \quad P_R P_R = P_R, \quad P_L P_L = P_L \quad \text{and} \quad P_L P_R = P_R P_L = 0. \tag{2.12}$$

Using this we can write

$$\psi = 1\psi = (P_L + P_R)\,\psi = \psi_L + \psi_R, \tag{2.13}$$

where ψ_L and ψ_R are now the left-handed and right-handed chiral components of ψ. Since the Dirac adjoint is linear, it follows that

$$\overline{\psi} = \overline{\psi}_L + \overline{\psi}_R \tag{2.14}$$

Using the fact that γ^5 is hermitian and anti-commutes with all four gamma matrices, we can also write these as

$$\begin{aligned} \overline{\psi}_L &= (P_L\psi)^\dagger\,\gamma^0 = \psi^\dagger P_L^\dagger \gamma^0 = \psi^\dagger P_L \gamma^0 = \overline{\psi} P_R, \\ \overline{\psi}_R &= (P_R\psi)^\dagger\,\gamma^0 = \psi^\dagger P_R^\dagger \gamma^0 = \psi^\dagger P_R \gamma^0 = \overline{\psi} P_L. \end{aligned} \tag{2.15}$$

Now we can expand the Dirac Lagrangian mass term from equation (2.7) in terms of chiral spinors

$$\begin{aligned} \mathcal{L} &\sim -m\left(\overline{\psi}_L + \overline{\psi}_R\right)(\psi_L + \psi_R) = -m\left(\overline{\psi} P_R + \overline{\psi} P_L\right)(P_L\psi + P_R\psi) \\ &= -m\overline{\psi}\,(P_R + P_L)\,(P_L + P_R)\,\psi = -m\overline{\psi}\,(P_R P_R + P_L P_L)\,\psi \\ &= -m\left(\overline{\psi}_L \psi_R + \overline{\psi}_R \psi_L\right). \end{aligned} \tag{2.16}$$

Now we can see the problem with this term in the SM Lagrangian. The left-handed chiral spinors of fermions in the SM are part of $SU(2)_L$ doublets, that are transformed under the $SU(2)_L$ group as

$$L \to L' = e^{ig_W \theta_a(x)\tau^a} L \quad \text{and} \quad \overline{L} \to \overline{L}' = \overline{L} e^{-ig_W \theta_a(x)\tau^a}, \tag{2.17}$$

where L is any SM $SU(2)_L$ doublet containing two left-handed fermion fields, g_W is the weak coupling constant, a is an index that is implicitly summed over, θ_a are space-time dependent parameters of the transformation and τ^a, $a = 1, \ldots, 3$ are the hermitian Pauli matrices. The right-handed spinors, on the other hand, are singlets under $SU(2)_L$, meaning they have no charge and do not change at all under $SU(2)_L$ transformations. This means that the Lagrangian term in equation (2.16), which came from the proposed mass term in equation (2.7), cannot be invariant under $SU(2)_L$ transformations, as it has a non-zero charge. In other words, the right-handed fields are unable to cancel the transformation of the left-handed fields, since they do not transform at all.

This is one of the main reasons why fermions must be massless as long as the symmetries of the SM are upheld. But there is no doubt that we observe massive fermions in nature. So how can the SM explain this? The answer is to break the symmetries of the model.

2.3 Spontaneous symmetry breaking, the Brout-Englert-Higgs mechanism and electroweak unification

In addition to the difficulties we have demonstrated associated with writing down fermion mass terms in the SM, there is also the problem that all gauge boson fields must be massless in any gauge theory. The reason for this is that the natural gauge boson mass term[2]

$$\mathcal{L} \sim -\frac{1}{2} M^2 V_\mu^\dagger V^\mu, \tag{2.18}$$

where M is a constant of mass dimension one and V_μ is a gauge boson field, is not gauge invariant. However, we know that the mediators of the weak interaction, the $W^\pm$- and Z bosons, have masses. This is another indication that the symmetries of the SM need to be broken, and it gives us a reason to suspect that the symmetry breaking mechanism is related to the weak interaction in particular.

There turns out to be a convenient way to give mass to the gauge bosons of a gauge theory, by coupling a scalar field with certain properties to the gauge bosons. This can allow the symmetry of the theory to be broken spontaneously by the scalar field, while keeping the explicit symmetry of the Lagrangian. This method of giving mass to gauge bosons was discovered almost simultaneously by Robert Brout and François Englert [5], and Peter Higgs [6] in 1964, and is therefore called the Brout-Englert-Higgs (BEH) mechanism. In order to give the desired results, the scalar field must have a potential with a set of degenerate minima. The minima of the potential correspond to the lowest energy states of the field, which are often called its vacuum states. However, at any given space-time point x the field can only be in one of these vacuum states. This "choice" of the physical vacuum state spontaneously breaks the symmetry of the Lagrangian.

A few years before the BEH mechanism was conceived, Glashow proposed a unification of the weak and electromagnetic interactions [7], forming an electroweak interaction with gauge group $SU(2)_L \otimes U(1)_Y$. In 1967, this unified gauge theory was combined with the spontaneous symmetry breaking of the BEH mechanism in what is now known as the Glashow-Weinberg-Salam theory. This theory completed the main theoretical foundation of the Standard Model as we know it today. In this theory, the scalar field that is introduced to spontaneously break the gauge symmetry is called the BEH field. The BEH field is coupled to the electroweak gauge group to give the weak gauge bosons masses. In order to give the W bosons and the Z boson masses, the BEH field needs to "give" one degree of freedom to each of their fields[3]. There needs to be at least one degree of freedom remaining in the BEH field corresponding to the massive Higgs boson, so the BEH field needs at least four degrees of freedom in total. The simplest way to do this is to make a $SU(2)_L$ doublet of complex scalar fields, meaning the field has a weak isospin charge of $I_W = \frac{1}{2}$. The doublet needs to contain one electrically charged field and one electrically neutral field to give mass to both the electrically charged W bosons and the electrically neutral Z boson. We therefore get the following form of the Standard Model BEH field

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}, \tag{2.19}$$

[2]This is the form of the mass term for neutral bosons. For charged bosons the factor $\frac{1}{2}$ disappears.

[3]Massive vector bosons have three degrees of freedom, while massless only have two.

where $\phi_1, \ldots, \phi_4$ are the real valued component fields of the two complex fields ϕ^+ and ϕ^0. In order for the electroweak unification to give interactions that agree with observations, the following relation between the weak and electromagnetic charges must be fulfilled for all fields

$$Y = 2\left(Q - I_W^{(3)}\right), \tag{2.20}$$

where Q is the electric charge of the field, $I_W^{(3)}$ is called the third component of weak isospin and it is $+\frac{1}{2}$ and $-\frac{1}{2}$ for fields in the upper and lower part of a $SU(2)_L$ doublet respectively, and Y is the weak hypercharge of $U(1)_Y$ and must be the same for both component fields of a $SU(2)_L$ doublet. Using this relation, we see that the BEH field must have a weak hypercharge of $Y = 1$.

The BEH field is then given the potential

$$V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2, \tag{2.21}$$

where μ^2 and λ are real constants with mass dimensions two and zero respectively. If μ^2 is negative while λ is positive, the potential has an infinite set of minima for

$$\phi^\dagger\phi = \frac{1}{2}\left(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2\right) = -\frac{\mu^2}{2\lambda} = \frac{v^2}{2}, \tag{2.22}$$

where

$$v = \sqrt{\frac{-\mu^2}{\lambda}} \tag{2.23}$$

is called the vacuum expectation value, or vev, of the BEH field[4]. In the four dimensional space where each coordinate axis corresponds to one of the component fields of ϕ, such that any point in the space corresponds to a particular value of the BEH field, these minima form the three dimensional surface of a four dimensional hyper-sphere with radius v. This symmetry of the minima is not accidental, but a consequence of the fact that the BEH potential in (2.21), and in fact the whole Lagrangian, is invariant under $SU(2)_L$ transformations

$$\phi \to \phi' = e^{ig_W\theta_a(x)\tau^a}\phi, \tag{2.24}$$

which are exactly the transformations that keep the inner product $\phi^\dagger\phi$ constant. We can see this by computing

$$\phi'^\dagger\phi' = \phi^\dagger e^{-ig_W\theta_a(x)\tau^a} e^{ig_W\theta_a(x)\tau^a}\phi = \phi^\dagger\phi, \tag{2.25}$$

where we used that the generators τ^a of $SU(2)_L$ are hermitian. This shows that the symmetry of the hyper-sphere surface is simply a representation of the symmetry of $SU(2)_L$, and that a $SU(2)_L$ transformation like the one in (2.24) will move the state of the field along this three dimensional surface.

Since the theory is supposed to be invariant under local $SU(2)_L$ transformations, this means that we can choose the vacuum state of the field at all space-time points to be any point on the hyper-sphere surface without any physical consequences. In this case, we choose the vacuum expectation value to be only in the ϕ_3-direction. This choice of a direction in the four dimensional space represents a breaking of the $SU(2)_L$ symmetry of

[4]Another common convention puts a negative sign in front of μ^2 in the definition of the potential, making μ^2 itself positive, in which case the negative sign in the expression for v disappears. In any case, v is a positive real number of mass dimension one.

the system, which is exactly what we wanted. With this choice of the vev, the BEH field can be re-parameterised around the vacuum value as $\phi_3(x) = v + \eta(x)$, where η is a field that represents the deviation of ϕ_3 from the vev. The BEH doublet can then be written in terms of real scalar fields that all have a vacuum expectation value of zero

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ v + \eta + i\phi_4 \end{pmatrix}. \tag{2.26}$$

Writing out the Lagrangian for the BEH field in this form and coupling it to the electroweak gauge group $SU(2)_L \otimes U(1)_Y$ gives the desired mass terms for the gauge bosons. However, the component fields ϕ_1, ϕ_2 and ϕ_4 also enter the Lagrangian as un-physical massless scalar fields called Goldstone bosons. To remove these from the theory, we can perform another gauge transformation (still not changing the physics described by the Lagrangian!) to transform the BEH doublet into the form

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + \eta(x) \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \tag{2.27}$$

where we have renamed η as h to indicate that this is the physical BEH field. The gauge in which the Goldstone bosons are removed from the Lagrangian and the BEH doublet takes this form is called unitary gauge.

In unitary gauge, the Lagrangian for the BEH field coupled to the electroweak interaction has almost all the desired properties. Due to the vev of the BEH field, the kinetic part of the Lagrangian gives terms that are quadratic in the boson fields, just like the mass term in equation (2.18). From the potential part of the Lagrangian, the BEH field itself also gets a mass due to the non-zero vev. The only thing that remains is to take linear combinations of the four bosons of $SU(2)_L \otimes U(1)_Y$ to get the correct interactions of the physical photon, $W^\pm$- and Z bosons, and make sure that the photon remains massless. These requirements give relations between the weak and electric coupling constants, and the masses of the W and Z bosons, usually parameterised by the weak mixing angle θ_W. These are some of the main predictions of the theory of the BEH mechanism and electroweak unification, and they are verified experimentally to high precision.

2.4 Fermion masses in the Standard Model, revisited

Armed with the knowledge of the BEH mechanism, we can now explain how, at least most of, the fermions in the SM get their mass. We have introduced the BEH field in the form of a $SU(2)_L$ doublet and with weak hypercharge $Y = 1$ to the SM Lagrangian. This allows us to write new gauge invariant interaction terms between fermions and the BEH field, called Yukawa terms. They are

$$\mathcal{L}_{\text{Yuk}} = -\lambda_d^{ij} \overline{Q}_L^i \phi d_R^j + \lambda_u^{ij} \overline{Q}_L^i \phi^c u_R^j - \lambda_l^{ij} \overline{E}_L^i \phi e_R^j + h.c. \tag{2.28}$$

where i and j are generational indices that are summed over in each term, λ_d^{ij}, λ_u^{ij} and λ_l^{ij} are complex constants with mass dimension zero, Q_L^i is the ith generation of $SU(2)_L$ Quark doublet, d_R^i and u_R^i are the $SU(2)_L$ singlet fields of the quarks in the lower and upper part of Q_L^i, E_L^i is the ith generation of $SU(2)_L$ lepton doublet, e_R^i is the singlet

Table 2.1: The transformation properties of the different types of fermion fields and the BEH field under each of the three SM gauge groups. The first number in the parenthesis gives the size of the $SU(3)_C$ multiplet of the fermion (3=triplet, 2=doublet and 1=singlet). The second number gives the same for $SU(2)_L$. The last number is the weak hypercharge of the field, as given by the relation (2.20).

Field	Transformation ($SU(3)_C$, $SU(2)_L$, $U(1)_Y$)
Q_L	(3, 2, $\frac{1}{3}$)
E_L	(1, 2, -1)
d_R	(3, 1, $-\frac{2}{3}$)
u_R	(3, 1, $\frac{4}{3}$)
e_R	(1, 1, -2)
ν_R	(1, 1, 0)
ϕ	(1, 2, 1)

field of the charged lepton in E_L^i, ϕ is the BEH doublet and the field ϕ^c is called the conjugate BEH doublet and is defined as

$$\phi^c = -i\sigma_2\phi^*. \tag{2.29}$$

The "$h.c.$" at the end of equation (2.28) is short for the hermitian conjugate of all the previous terms in the equation. To show that the Yukawa terms in equation (2.28) are gauge invariant, we use the transformation properties of the fields that are listed in table 2.1. We also know that any contraction of two $SU(2)_L$ doublets or two $SU(3)_C$ triplets is invariant, as was shown in equation (2.25) for the case of $SU(2)_L$. The exact same argument shows that contractions of $SU(3)_C$ triplets are invariant, since this group also has hermitian generators. Then we see that the terms containing ϕ are invariant under $SU(2)_L$ since there is a contraction between the fermion doublet and the BEH doublet. In the last term, all of the fields are singlets under the strong interaction, so they are obviously $SU(3)_C$ invariant. In the first two terms, there is a contraction between the $SU(3)_C$ triplets of the two quark fields, so the terms are also invariant. For invariance under $U(1)_Y$ it suffices to show that the sum of hypercharges in each term is zero. Noting that the sign of the hypercharge is flipped when taking the Dirac adjoint of a field, we see that the sum of hypercharges in the first term is $-\frac{1}{3}+1-\frac{2}{3}=0$ and in the last term it is $1+1-2=0$. To find the sum of hypercharges of the second term we need to look at how the conjugate BEH doublet transforms under $U(1)_Y$

$$\phi^{c\prime} = -i\sigma_2\phi^{\prime *} = -i\sigma_2 e^{-ig'Y\theta}\phi^* = e^{ig'(-Y)\theta}\left(-i\sigma_2\phi^*\right) = e^{ig'(-Y)\theta}\phi^c. \tag{2.30}$$

This shows that ϕ^c transforms just like the un-conjugated field ϕ, only with the hypercharge multiplied by -1. Thus ϕ^c has $Y=-1$. Then the sum of hypercharges of the second term is $-\frac{1}{3}-1+\frac{4}{3}=0$. Finally, we need to show that the second term is invariant under $SU(2)_L$. To do this we look at how the conjugate BEH doublet transforms under a general $SU(2)_L$ transformation, for which we will need the following properties of the

Pauli matrices

$$\sigma_1^2 = \sigma_2^2 = \sigma_3^2 = I, \tag{2.31a}$$
$$\sigma_a\sigma_b = -\sigma_b\sigma_a \quad (a \neq b)\,, \tag{2.31b}$$
$$\sigma_1^* = \sigma_1, \quad \sigma_2^* = -\sigma_2, \quad \sigma_3^* = \sigma_3. \tag{2.31c}$$

Combining these properties, we find that

$$\sigma_2\sigma_a^* = -\sigma_a\sigma_2 \tag{2.32}$$

for all three values of a, and

$$\begin{aligned}(\theta_a\sigma^{a*})^2 &= \theta_1^2 + \theta_2^2 + \theta_3^2 \\ &\quad - \theta_1\theta_2\left(\sigma_1\sigma_2 + \sigma_2\sigma_1\right) + \theta_1\theta_3\left(\sigma_1\sigma_3 + \sigma_3\sigma_1\right) - \theta_2\theta_3\left(\sigma_2\sigma_3 + \sigma_3\sigma_2\right) \\ &= \theta_1^2 + \theta_2^2 + \theta_3^2 = \theta_a\theta^a.\end{aligned} \tag{2.33}$$

Using this, the transformation of the conjugate BEH field can be written as

$$\begin{aligned}\phi^{c\prime} &= -i\sigma_2\phi^{\prime *} = -i\sigma_2 e^{-\frac{1}{2}ig_W\theta_a\sigma^{a*}}\phi^* \\ &= \sigma_2\left(1 - \frac{1}{2}ig_w\theta_a\sigma^{a*} + \frac{1}{2}\left[\frac{1}{2}ig_w\theta_a\sigma^{a*}\right]^2 - \frac{1}{6}\left[\frac{1}{2}ig_w\theta_a\sigma^{a*}\right]^3 + \dots\right)(-i\phi^*) \\ &= \sigma_2\left(1 - \frac{1}{2}ig_w\theta_a\sigma^{a*} + \frac{1}{2}\left[\frac{1}{2}ig_w\right]^2\theta_a\theta^a - \frac{1}{6}\left[\frac{1}{2}ig_w\right]^3\theta_b\theta^b\theta_a\sigma^{a*} + \dots\right)(-i\phi^*) \\ &= \left(1 + \frac{1}{2}ig_w\theta_a\sigma^a + \frac{1}{2}\left[\frac{1}{2}ig_w\right]^2\theta_a\theta^a + \frac{1}{6}\left[\frac{1}{2}ig_w\right]^3\theta_b\theta^b\theta_a\sigma^a + \dots\right)(-i\sigma_2\phi^*) \\ &= e^{\frac{1}{2}ig_W\theta_a\sigma^a}\phi^c,\end{aligned} \tag{2.34}$$

which has exactly the same form as the transformation of the regular BEH doublet ϕ, as shown in equation (2.24). We also note that the hermitian conjugates of terms that are invariant under a gauge transformation must also be invariant, and thus the the entire Lagrangian in equation (2.28) is gauge invariant.

2.4.1 Right-handed neutrinos

Interestingly, there is no Yukawa term for leptons involving the conjugate BEH doublet in equation (2.28). The reason for this is that this term requires the presence of the right-handed neutrino fields ν_R^i in order for the term to be gauge invariant, and this is a problem in the Standard Model.

The neutrinos are the only SM fermions that do not interact with either the gluons or the photon, since they have no colour or electric charge. The left-handed neutrinos (and their CP conjugate counterparts, the right-handed antineutrinos) only couple to the W and Z bosons due to their non-zero weak isospin and hypercharge. However, right-handed neutrinos, which would be singlet under $SU(2)_L$, would have no weak isospin or hypercharge either. They would therefore not interact with any other SM particle, except to a completely negligible extent through gravity[5], or potentially through mixing. It is therefore not surprising that we have not seen any evidence of the existence of right-handed neutrinos. For this reason they are not included in the Standard Model. Thus neither is the Yukawa term for the neutrinos.

2.4.2 Fermion mass terms

When the BEH field now acquires a vacuum expectation value and the Lagrangian is written in unitary gauge, such that the BEH field is on the form of equation (2.27), the Yukawa terms become

$$\mathcal{L}_{\text{Yuk}} = -\lambda_d^{ij}\frac{1}{\sqrt{2}}(v+h)\,\overline{d}_L^i d_R^j - \lambda_u^{ij}\frac{1}{\sqrt{2}}(v+h)\,\overline{u}_L^i u_R^j - \lambda_l^{ij}\frac{1}{\sqrt{2}}(v+h)\,\overline{e}_L^i e_R^j + h.c. \tag{2.35}$$

where d_L^i and u_L^i are the down-type and up-type quark fields the component fields of the doublet Q_L^i, and e_L^i is the lower component field of E_L^i.

There are two types of terms in equation (2.35). The first type is for the down-type quarks

$$\mathcal{L}_{d,h} = -\lambda_d^{ij}\frac{1}{\sqrt{2}}\overline{d}_L^i d_R^j h + h.c., \tag{2.36}$$

which has the form of an interaction term between the quarks and the Higgs boson, with λ_d^{ij} as coupling constants. The other type is for the same quarks

$$\mathcal{L}_{\text{mass},d} = -\lambda_d^{ij}\frac{1}{\sqrt{2}}v\overline{d}_L^i d_R^j + h.c., \tag{2.37}$$

which has a form that is very similar to the infamous fermion mass term shown in equation (2.16). However, the Lagrangian in equation (2.37) has some issues in its current state. Since there are a priori no restrictions on the parameters λ^{ij}, it seems to allow "off-diagonal" terms of the form $\overline{d}^i d^j$ where i and j are different. But these terms can not be physical, since they would allow quarks of one generation to turn into quarks of another generation without interacting with any other particles. To find the mass terms of the physical fermion fields, we therefore need to mix the different generations to make fields for which the matrices λ_d, λ_u and λ_l are diagonal. The fields for which this is the case are called the mass eigenstates of the down-type quarks, the up-type quarks and the charged leptons respectively.

2.4.3 Mixing in the quark sector

For the quarks, we need to diagonalise both λ_d and λ_u. It can be shown[6] that any complex 3 by 3 matrix M can be written as

$$M = UDW, \tag{2.38}$$

[5]The gravitational effect of right-handed neutrinos could only be detectable on a cosmic scale. In particular, right-handed neutrinos could change the effective number of neutrino species, which again is related to the energy density and expansion rate of the universe, as explained in section 4 of *The Phenomenology of Right Handed Neutrinos* [8].

[6]See for example section 11.3 of *Gauge Theory of Elementary Particle Physics* [9].

where U and W are 3 by 3 unitary matrices, and D is a 3 by 3 diagonal matrix with real and positive numbers on the diagonal. Therefore, we can without loss of generality write

$$\lambda_u = U_u D_u W_u^\dagger, \quad \lambda_d = U_d D_d W_d^\dagger, \tag{2.39}$$

where U_u, W_u, U_d and W_d are all unitary matrices and D_u and D_d are diagonal with real and positive eigenvalues. If we then transform the quark fields according to the linear transformations

$$u_L^i \to U_u^{ij} u_L^j, \quad u_R^i \to W_u^{ij} u_R^j, \quad d_L^i \to U_d^{ij} d_L^j, \quad d_R^i \to W_d^{ij} d_R^j, \tag{2.40}$$

where we keep the same names for the fields for simplicity, we see that all of the unitary matrices disappear from the fermion-Higgs interaction terms and the fermion mass terms. For example, the terms in equation (2.37) become

$$\begin{aligned}
\mathcal{L}_{\text{mass},d} &= -\lambda_d^{ij} \frac{1}{\sqrt{2}} v \overline{d}_L^i d_R^j + h.c. \\
&\to -\frac{1}{\sqrt{2}} v \left(U_d^{ik} d_L^k \right)^\dagger \gamma^0 \lambda_d^{ij} W_d^{jl} d_R^l + h.c. \\
&= -\frac{1}{\sqrt{2}} v d_L^{\dagger k} \gamma^0 U_d^{ik*} \left(U_d D_d W_d^\dagger \right)^{ij} W_d^{jl} d_R^l + h.c. \\
&= -\frac{1}{\sqrt{2}} v \overline{d}_L^k U_d^{\dagger ki} U_d^{im} D_d^{mn} W_d^{\dagger nj} W_d^{jl} d_R^l + h.c. \\
&= -\frac{1}{\sqrt{2}} v \overline{d}_L^k D_d^{kl} d_R^l + h.c. \\
&= -\frac{1}{\sqrt{2}} v D_d^{ii} \left(\overline{d}_L^i d_R^i + \overline{d}_R^i d_L^i \right).
\end{aligned} \tag{2.41}$$

In the last line, there are only terms containing fields of the same generation. In fact, looking at equation (2.16) from section 2.2, we see that we can write this last line using the complete fields of the quarks given by

$$d^i = d_L^i + d_R^i \tag{2.42}$$

as

$$\mathcal{L}_{\text{mass},d} = -\frac{1}{\sqrt{2}} v D_d^{ii} \left(\overline{d}^i d^i \right). \tag{2.43}$$

This shows that in the mass eigenstates of the quarks, we get terms that are exactly like the standard Dirac fermion mass terms shown in equation (2.7). Comparing equation (2.43) with equation (2.7), we see that the masses that the down-type quarks get from the BEH mechanism are

$$m_d^i = \frac{1}{\sqrt{2}} v D_d^{ii}. \tag{2.44}$$

The corresponding masses for the up-type quarks, which can be found in the same way by inserting the u_L^i and u_R^i fields in the mass eigenstate basis into the Lagrangian in equation (2.35), are

$$m_u^i = \frac{1}{\sqrt{2}} v D_u^{ii}. \tag{2.45}$$

An interesting thing to note about the quark masses is that they depend on both the BEH vev v and the diagonalised Yukawa couplings D_u and D_d. The value of v can be measured independently, for instance by measuring the Higgs boson mass or the weak gauge boson masses. The entries of D_u and D_d, however, come directly from the arbitrary Yukawa couplings λ_u and λ_d, which are not restricted by theory or any independent experiments[7]. This therefore serves more as a parametrisation of the quark masses, rather than an explanation for their values.

It is important to ask ourselves whether the transformation we did in equation (2.40) to write all the quark fields in their mass eigenstates affected any other parts of the SM Lagrangian. In any terms where a chiral quark field is multiplied by the Dirac adjoint of the same field, the unitary matrices will cancel as long as they commute with whatever is between the two quark fields. Since the unitary matrices are constant they commute with derivatives, and they also commute with the gamma matrices and the generators of the $SU(3)_C$ gauge group, since they work in different spaces. Therefore terms that are invariant under the change to mass eigenstate basis include the kinetic terms, strong interactions, and photon and Z boson interactions. However, in the charged current weak interactions involving the W bosons, the terms contain left-handed fields of one up-type quark and one down-type quark. Before transforming to the mass eigenstate basis, these terms have the form

$$\mathcal{L}_{CC} = \frac{g_W}{\sqrt{2}} W_\mu^+ \overline{u}_L^i \gamma^\mu d_L^i + h.c. \tag{2.46}$$

After the transformation, they become

$$\begin{aligned}\mathcal{L}_{CC} &= \frac{g_W}{\sqrt{2}} W_\mu^+ \overline{u}_L^j U_u^{ij*} \gamma^\mu U_d^{ik} d_L^k + h.c. \\ &= \frac{g_W}{\sqrt{2}} W_\mu^+ \gamma^\mu \overline{u}_L^j U_u^{\dagger ji} U_d^{ik} d_L^k + h.c. \\ &= \frac{g_W}{\sqrt{2}} W_\mu^+ \gamma^\mu \overline{u}_L^j \left(U_u^\dagger U_d\right)^{jk} d_L^k + h.c.\end{aligned} \tag{2.47}$$

We see that the unitary matrices do not disappear in these terms. Instead we get a new unitary matrix

$$V = U_u^\dagger U_d \tag{2.48}$$

which the down-type quarks are multiplied by. This matrix is known as the Cabibbo–Kobayashi–Maskawa (CKM) matrix. We now define the weak eigenstate basis of the quark fields as the ones that are diagonal in the charged current interaction. Marking the fields in the weak eigenstate basis with primes to distinguish them from the mass eigenstates, the charged current in terms of weak eigenstates are by definition

$$\mathcal{L}_{CC} = \frac{g_W}{\sqrt{2}} W_\mu^+ \overline{u}_L^{i\prime} \gamma^\mu d_L^{i\prime} + h.c. \tag{2.49}$$

Comparing this with the same terms in the mass eigenstate basis, we see that we can identify

$$u_L^{i\prime} = u_L^i, \quad d_L^{i\prime} = V^{ij} d_L^j. \tag{2.50}$$

[7] The elements of D_u and D_d can of course be measured from the cross sections of the decays of the Higgs boson to quarks, since the exact same matrix elements determine these cross sections (as can be seen from equations (2.36) and (2.37)). However, one could just as well measure the masses of the quarks using some other method, and use the mass values to determine the couplings of the Higgs boson to quarks. The measurement of D_u and D_d elements from Higgs decays can therefore not be called an independent measurement of the quark masses.

This shows that while we can make the weak eigenstate of the up-type quark equal to their mass eigenstates, the weak eigenstates of the down-type quarks are a rotation of the physical mass eigenstates. This is what we mean when we say that the quarks are mixing, and it has drastic consequences for the symmetries of the theory. While the combined quark number (where a quark is assigned quark number 1 and an antiquark quark number -1) is still conserved in the entire SM, the individual quark flavour numbers, where each generation is given its own separate number, can now be violated by the charged current weak interaction. Additionally, it can be shown[8] that if the CKM matrix has any complex elements, these interactions will violate CP symmetry, which is otherwise an exact symmetry of the SM. Experiments show that the CKM matrix is indeed complex, and CP violations have been observed in weak interactions [10].

2.4.4 Lepton masses

The change to mass eigenstate basis in the lepton sector is slightly different from the quark sector, due to the missing right-handed neutrinos. To find the mass eigenstates of the leptons we only need to diagonalise the matrix λ_l. This can again be done with two unitary matrices U_l and W_l such that

$$\lambda_l = U_l D_l W_l^\dagger, \tag{2.51}$$

where D_l is a diagonal matrix with positive real numbers on the diagonal. To eliminate the unitary matrices from the lepton mass terms, we need only to transform the charged lepton fields as follows

$$e_L^i \to U_l^{ij} e_L^j, \quad e_R^i \to W_l^{ij} e_R^j. \tag{2.52}$$

Then the mass terms for the charged leptons can be written as

$$\mathcal{L}_{\text{mass},l} = -\frac{1}{\sqrt{2}} v D_l^{ii} \left(\bar{e}_L^i e_R^i + \bar{e}_R^i e_L^i \right), \tag{2.53}$$

which means that similarly to the quark masses, the masses given to the charged leptons by the BEH mechanism are

$$m_l^i = \frac{1}{\sqrt{2}} v D_l^{ii}. \tag{2.54}$$

The masses again depend on the arbitrary matrix elements D_l^{ii}, an thus the theory does not predict the values of the charged lepton masses. The neutrinos, meanwhile, get neither mass terms or Higgs interactions since they have no Yukawa term. It therefore does not really make sense to speak of the mass eigenstates of the neutrinos in this framework, and there is no reason to change basis for the neutrino fields ν_L^i. However, if the neutrino fields are not transformed, the matrix U_l will appear in the charged current interactions of leptons in the same way as the CKM matrix did for the quarks. We therefore choose to transform the neutrinos with the same matrix as the charged leptons, giving

$$\nu_L^i \to U_l^{ij} \nu_L^j. \tag{2.55}$$

In this way, U_L can be factored out of the lepton $SU(2)_L$ doublets, and it therefore disappears from all the SM Lagrangian terms, including the charged current interactions. This choice of basis for the neutrino states is therefore equivalent with defining the weak

[8] See for example section 20.3 of Peskin and Schroeder [3].

eigenstates of the charged leptons to be the same as their mass eigenstates. It also means that without the right-handed neutrino fields and neutrino mass terms, there is no mixing between the lepton generations, and each individual lepton flavour number is conserved. Then any lepton interactions also have exact CP symmetry. However, as we will see in the following section, the observation of neutrino oscillation changes all of this.

2.5 Neutrino oscillations and evidence for non-zero neutrino masses

Because the SM neutrinos are so difficult to detect, no experiment to this day has been able to directly measure their mass. In section 2.4, we saw that the three neutrinos are all assumed to be massless in the SM. Nevertheless, the consequences of non-zero neutrino masses have been explored theoretically for almost as long as neutrinos have been known to exist, for example by Bruno Pontecorvo in 1967 [11]. One of the main consequences that were explored was the possibility of neutrino oscillations. These oscillations also turned out to be the phenomenon that provided the first decisive evidence for non-zero neutrino masses, so we will look at this theory in more detail.

2.5.1 Neutrino Oscillations

To see how non-zero neutrino masses can lead to oscillations, we assume that the right-handed neutrino fields do exist, and that the neutrinos have some complex mass matrix λ_ν which gives the mass terms[9]

$$\mathcal{L}_{\text{mass},\nu} = -\lambda_\nu^{ij} \overline{\nu}_L^i \nu_R^j + h.c. \tag{2.56}$$

Then the mass matrix λ_ν must be diagonalised in the neutrino mass eigenstate basis, and we do no longer have the freedom to transform the neutrino fields such that the charged current interaction of the leptons remains diagonal with respect to the lepton generations. Instead we get the same situation as in the quark sector. If the unitary matrices that are needed to diagonalise λ_ν are U_ν and $W_\nu^\dagger$, then the charged current interactions in the mass eigenstate basis are

$$\mathcal{L}_{CC} = \frac{g_W}{\sqrt{2}} W_\mu^+ \gamma^\mu \overline{\nu}_L^j \left(U_\nu^\dagger U_l\right)^{jk} e_L^k + h.c. \tag{2.57}$$

We see that the matrix $U_\nu^\dagger U_l$ takes the role of the CKM matrix for the quarks. By convention, we choose the charged lepton weak eigenstates to be the same as their mass eigenstates, which means that the mixing matrix for the neutrinos is

$$U = \left(U_\nu^\dagger U_l\right)^\dagger = U_l^\dagger U_\nu, \tag{2.58}$$

which is known as the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix. Then the weak eigenstates of the leptons are related to the mass eigenstates by

$$e_L^{i\prime} = e_L^i, \quad \nu_L^{i\prime} = U^{ij} \nu_L^j. \tag{2.59}$$

[9]The neutrino mass terms must not necessarily originate in the BEH mechanism, and we therefore write them in a more generic form.

This means that when a neutrino is produced in a charged current interaction, its quantum mechanical state can be described as a superposition of the three mass eigenstates. For oscillations to occur, this superposition must remain coherent over significant amounts of time, which can only happen if the mass differences between the states are small. This is one of the reasons why oscillations are not observed in charged leptons or quarks[10], and it is explained in more detail in *Do charged leptons oscillate?* [12]. Assuming that the mass differences between the neutrinos are small, we can write a weak eigenstate $\nu_L^{j'}$ as a coherent superposition of mass eigenstates, and evolve this state through time by approximating the mass eigenstates as plane waves. Then the neutrino state after a time t at a position $\mathbf{x}$ is

$$\left|\nu_L^{j'}(t,\mathbf{x})\right\rangle = U^{jk}\left|\nu_L^k\right\rangle e^{-i\phi_k}, \tag{2.60}$$

where k is an index that is summed over, and ϕ_k is given by

$$\phi_k = p^\mu x_\mu = E_k t - \mathbf{p}_k \cdot \mathbf{x}, \tag{2.61}$$

where E_k and $\mathbf{p}_k$ are the energy and three-momentum of the mass eigenstate ν_L^k, while p^μ and x^μ are its four-momentum and position four-vector. Using the inverse of the PMNS matrix, each of the mass eigenstates can be written as a linear combination of the three weak eigenstates, giving

$$\left|\nu_L^{j'}(t,\mathbf{x})\right\rangle = U^{jk}U^{lk*}\left|\nu_L^{l'}\right\rangle e^{-i\phi_k}, \tag{2.62}$$

where l is another index that is summed over. We can then find the probability of the neutrino being in one of the weak eigenstates after a time t as

$$\begin{aligned} P\left(\nu_L^{j'} \to \nu_L^{l'}\right) &= \left|\left\langle \nu_L^{l'} \middle| \nu_L^{j'}(t,\mathbf{x})\right\rangle\right|^2 = \left|U^{jk}U^{lk*}e^{-i\phi_k}\right|^2 \\ &= \left|U^{jk}U^{lk*}\right|^2 + 2\,\mathrm{Re}\Big\{U^{j1}U^{l1*}U^{j2*}U^{l2}e^{-i(\phi_1-\phi_2)}\Big\} \\ &+ 2\,\mathrm{Re}\Big\{U^{j1}U^{l1*}U^{j3*}U^{l3}e^{-i(\phi_1-\phi_3)}\Big\} + 2\,\mathrm{Re}\Big\{U^{j2}U^{l2*}U^{j3*}U^{l3}e^{-i(\phi_2-\phi_3)}\Big\}, \end{aligned} \tag{2.63}$$

where Re{} means the real part of whatever complex number is inside the brackets. Thus each of the three last terms will be some combination of sine and cosine functions of the phase differences $\phi_2 - \phi_1 = \Delta\phi_{21}$, $\phi_3 - \phi_1 = \Delta\phi_{31}$ and $\phi_3 - \phi_2 = \Delta\phi_{32}$. It can be shown[11] that when assuming that the momenta of the mass eigenstates are the same, and that the neutrinos are highly relativistic these phase differences can be written as

$$\Delta\phi_{ij} = \frac{m_i^2 - m_j^2}{2E} L, \tag{2.64}$$

where m_i is the mass of the mass eigenstate ν_L^i, E is the energy of the neutrino and L is the distance the neutrino has travelled since its production. This means that if there is no difference between the masses of the three neutrinos, and in particular if they have no mass at all, there will be no oscillations in the probability of observing the

[10]Although oscillations are observed in neutral K- and B-mesons, which both have two states with very similar masses.

[11]See for example section 13.4 and 13.5 of Thomson [4].

neutrino in any of the weak eigenstates. In fact, using the unitarity of the PMNS matrix, it is easy to see that in this case, $P\left(\nu_L^{j\prime} \to \nu_L^{j\prime}\right) = 1$, while $P\left(\nu_L^{j\prime} \to \nu_L^{l\prime}\right) = 0$ for $j \neq l$. Only if the mass differences are non-zero, and there is a mixing between the neutrino weak eigenstates and mass eigenstates (the PMNS matrix is not diagonal), will these probabilities change with the distance travelled. Then they will change periodically, such that the survival probability $P\left(\nu_L^{j\prime} \to \nu_L^{j\prime}\right)$ for example is 1 at time $t = 0$, then drops to a minimum at some distance, and then increases to 1 again. These are the oscillations that give neutrino oscillations their name. Now we have seen that only neutrinos with different, but nearly equal, masses will oscillate. But conversely, if oscillations were to be observed in experiments, the only logical explanation would be that neutrinos are not massless like they are in the SM.

2.5.2 Parameters of the PMNS matrix

As we have seen, non-zero neutrino masses imply that the neutrino states are mixed by the unitary PMNS matrix. In general, a 3 by 3 unitary matrix has nine real parameters. However, it can be shown that five of these parameters can be removed by transforming the lepton fields[12], without changing any of the physics. Therefore we can parameterise the PMNS matrix with four real parameters; three mixing angles and one complex phase. The mixing angles are commonly called θ_{12}, θ_{23} and θ_{13}, and the complex phase is δ. Then the PMNS matrix can be written in the form

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.65}$$

where $s_{ij} = \sin\theta_{ij}$ and $c_{ij} = \cos\theta_{ij}$. Just like for the CKM matrix, any complex elements of the PMNS matrix will lead to CP violation in the weak interaction of leptons. With the parametrisation of equation (2.65), the only possibility for complex elements comes from the factors containing the complex phase δ. Thus if $\delta = 0$ or $\delta = \pi$, CP is conserved in lepton interactions, and otherwise it parametrises the extent to which CP is violated. But as we will see in the following sections, the current measurements of δ lack the precision to determine whether or not CP is violated in the lepton sector.

2.5.3 Experimental evidence for neutrino oscillations

For an experiment to observe neutrino oscillations, it requires a strong source of neutrinos, a good understanding of the nature of the source (like its intensity and ratio of different flavours), and a large detector volume. The neutrino sources can either be constructed or natural. The most common ways to construct neutrino sources are in particle accelerators, where it is possible to make beams of neutrinos for example by colliding proton beams with a target, and nuclear reactors, which release electron neutrinos from beta decays. Since the 1970s these sources have been used to look for neutrino oscillations. However, early versions of these experiments were so-called short baseline experiments, where the measurement of the neutrinos was performed relatively close to the source, usually at

[12]This only applies if the neutrinos are Dirac fermions. If they are Majorana fermions, only three parameters can be removed, with three angles and three phases remaining, as is shown in section 14 of *Review of Particle Physics* [13].

distances of less than 1 km. This, combined with relatively low neutrino energies, meant that until the end of the 1990s such experiments where only sensitive to oscillations corresponding to squared mass differences of more than around 0.1 eV^2. At the end of the 90s these experiments had still not seen any conclusive evidence of neutrino oscillations [14] [15].

There are two natural sources of neutrinos that have also been used for oscillation experiments. These are the sun, producing what are called solar neutrinos, and cosmic ray interactions with the upper atmosphere of the earth, producing atmospheric neutrinos. Solar neutrinos are produced in very high numbers from different fusion and decay processes in the core of the sun. They are exclusively electron neutrinos, and have energies ranging from less than 0.1 MeV to around 15 MeV [16]. Atmospheric neutrinos are produced mainly in the decay of charged pions to muons and muon neutrinos. The muons also decay, resulting in another muon neutrino and an electron neutrino. The expected ratio of muon to electron neutrinos among atmospheric neutrinos is therefore around $2:1$ [17]. The energy of the atmospheric neutrinos depend on the energy of the incident cosmic ray, and they can range from around 1 MeV to as high as 1000 TeV.

Already in 1970, the Homestake Chlorine experiment started measuring the flux of electron neutrinos from the sun with energies greater than 0.814 MeV [17]. The measured flux in this and other similar experiments was consistently between 30% and 60% of what was expected from models of solar physics [16]. This discrepancy could be viewed as evidence for neutrino oscillations, since if some of the electron neutrinos oscillated to muon or tau neutrinos on their way from the sun to the detector, these would not be detected. However, they were not able to exclude the possibility that the discrepancy was caused by inaccuracies in the modelling of the sun's interior, so although promising, these results were not yet decisive.

The first decisive observation of neutrino oscillations came from atmospheric neutrino measurements, with a new detector type known as water Cherenkov detectors. Water Cherenkov detectors consist of a large volume of water and detect neutrinos indirectly from their elastic scattering off atomic electrons in the water or their charged-current interactions with nuclei. If the energy of the incoming neutrino is large enough, a charged lepton can be ejected from the water molecule, and if its speed is greater than the speed of light in water, it will emit photons at a fixed angle to its direction of motion, known as Cherenkov radiation. These photons are then detected by photo-multiplier tubes surrounding the water. The detectors are sensitive mainly to electron and muon neutrinos, since charged current interaction of tau neutrinos would require the production of a tau lepton, which has a much larger mass and therefore requires a very energetic neutrino. The charged lepton that is ejected will have the same flavour as the incoming neutrino, and electron and muon signals can be distinguished by the shape of the ring of light hitting the photo-multipliers. In this way, water Cherenkov detectors can detect electron and muon neutrinos separately. Crucially, they can also measure both the time of events, and the energy and direction of the incoming neutrino. The first experiment to use such a detector was the Kamiokande experiment in Japan. In 1996, this was replaced by the even larger Super-Kamiokande experiment. After less than two years of data taking, they already reported a significant lack of muon neutrinos compared to electron neutrinos among the neutrinos coming from the downward direction (from the direction of the centre of the earth) [2]. They concluded that this was evidence of oscillation of atmospheric muon neutrinos produced on the opposite side of the earth to either tau neutrinos or some other sterile neutrino during their passage through the planet. The probability of the observed data in the case of no oscillations was estimated to be less than 0.001%, which corresponds to a significance of 4.3σ.

This was a large breakthrough, but it still only confirmed oscillations of muon to tau neutrinos, so oscillations of electron to muon neutrinos and electron to tau neutrinos remained to be investigated. In addition, new experiments were needed to increase the precision of measurements of the PMNS parameters and square mass differences, which are related to the amplitude and frequency of oscillations respectively. The next step forward came from a new water Cherenkov detector in the SNO experiment. In this experiment, the detector is filled with heavy water instead of regular water, which makes interactions of neutrinos with nuclei possible even for the lower energy solar neutrinos, due to the low binding energy of the deuteron nucleus. Only electron neutrinos can interact via a charged current, because the production of muons or taus are not kinematically allowed. The neutral current interaction however, where the neutrino is simply scattered while the deuteron is broken into its constituent proton and neutron, is possible with all three neutrino flavours. This process is detected via the subsequent capture of the neutron by another deuterium atom, which releases a gamma ray, which again produces relativistic electrons that can give a Cherenkov signal. In this way, the SNO experiment could measure the total solar neutrino flux and the solar electron neutrino flux separately. This eliminates the uncertainty from modelling of the fusion reactions in the sun. After only a few years, in 2002, the experiment reported that the total neutrino flux was consistent with solar models, while the electron neutrino flux represented only around 35% of the total. The mu and tau neutrino flux combined was measured to be greater than zero with 5.3σ significance [18]. Thus, the oscillation of electron neutrinos to muon and/or tau neutrinos was as good as confirmed, and the solar neutrino problem was finally solved. Less than a year later, the reactor experiment KamLAND also reported observation of the disappearance of antielectron neutrinos from nuclear reactors at distances of the order of 100 km, providing even more evidence for oscillation of electron neutrinos [19].

2.5.4 Current state of neutrino oscillation parameter measurements

Since the initial discovery of neutrino oscillations, long baseline accelerator experiments such as K2K [20] and T2K [21] and short baseline experiments like Daya Bay [22] (which provided the first measurement of the θ_{13} parameter) have measured oscillations with greater precision. No clear contradiction to the three neutrino flavour mixing model with a unitary PMNS matrix have been found, and the experiments have gradually constrained the possible values of the PMNS mixing angles and the neutrino squared mass differences. The data from these experiments show a clear hierarchy in the mass differences, namely

$$\Delta m_{21}^2 \ll \left|\Delta m_{31}^2\right| \approx \left|\Delta m_{32}^2\right|, \tag{2.66}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$, and m_i as before means the mass of the eigenstate ν_L^i. However, the exact ordering of the neutrino masses is still unknown. The ordering is difficult to determine experimentally, because oscillation experiments do not measure the values of the masses themselves, but only the magnitude of the squared mass differences. Currently there are two different orderings that are non-equivalent, and have not been excluded by any experiments. These are called the normal ordering (NO)

$$m_1 < m_2 < m_3, \tag{2.67}$$

Table 2.2: The current values for the parameters of the PMNS matrix and two of the independent squared mass differences of the neutrinos, as given by the NuFIT collaboration [23, 24]. The values that fit data best assuming both normal mass ordering (NO) and inverted mass ordering (IO) are shown. The parameter Δm^2_{3l} represents the squared mass difference Δm^2_{31} in the case of NO and Δm^2_{32} in the case of IO.

	NO	IO
s^2_{12}	$0.303^{+0.012}_{-0.011}$	$0.303^{+0.012}_{-0.011}$
s^2_{13}	$\left(2.203^{+0.056}_{-0.059}\right) \times 10^{-2}$	$\left(2.219^{+0.060}_{-0.057}\right) \times 10^{-2}$
s^2_{23}	$0.572^{+0.018}_{-0.023}$	$0.578^{+0.016}_{-0.021}$
δ	$1.09^{+0.23}_{-0.14}\pi$ rad	$1.59^{+0.15}_{-0.18}\pi$ rad
Δm^2_{21}	$\left(7.41^{+0.21}_{-0.20}\right) \times 10^{-5}\text{eV}^2$	$\left(7.41^{+0.21}_{-0.20}\right) \times 10^{-5}\text{eV}^2$
Δm^2_{3l}	$\left(2.507^{+0.026}_{-0.027}\right) \times 10^{-3}\text{eV}^2$	$\left(-2.570^{+0.025}_{-0.028}\right) \times 10^{-3}\text{eV}^2$

and the inverted ordering (IO)

$$m_3 < m_1 < m_2. \tag{2.68}$$

For some of the PMNS mixing angles and squared mass differences, the values that fit experimental data best change slightly depending on the mass ordering. The current best values of these parameters are given by the NuFit collaboration [23, 24] and shown in table 2.2. The parameter Δm^2_{3l} represents the squared mass difference Δm^2_{31} in the case of NO and Δm^2_{32} in the case of IO. The squared mass difference that is not shown in the table can be calculated from the two that are given, since in both orderings $\Delta m^2_{31} = \Delta m^2_{32} + \Delta m^2_{21}$. We also notice that the best value for the complex phase δ is some distance away from both 0 and π in both orderings, and in particular in the inverted ordering. However, the uncertainty on this value is still so large that $\delta = \pi$, implying CP conservation, can currently not be excluded at $1 - 2\sigma$ confidence level [13]. There are however several planned experiments that are expected to be able to exclude CP conservation at $3 - 5\sigma$ for most true values of δ within the next 10-15 years. Experiments that are currently under construction include DUNE [25] and Hyper-Kamiokande [26], while other experiments are being planned, such as the ESSνSB [27]. The latter experiment, while still in a preparatory phase, would allow precise measurement of δ, assuming that a non-zero value of δ is observed by DUNE and/or Hyper-Kamiokande. This would reveal the extent of CP violation in the lepton sector, which could be important to understanding the origin of the observed asymmetry between matter and anti-matter in the Universe. These experiments are also expected to be able to determine the correct mass ordering.

2.5.5 Limits on neutrino masses

Since all neutrino oscillation experiments to this day seem to agree with the three flavour oscillation model, with parameters as given in the previous section, the only reasonable conclusion is that at least two of the neutrinos have non-zero masses (which are needed to get three distinct squared mass differences). The size of the measured squared mass differences also give a lower bound on these masses. In the normal mass ordering, the lower bound on the sum of the three masses is $\sum m_i > 5.87 \times 10^{-2}$ eV, while in the inverted ordering it is $\sum m_i > 9.92 \times 10^{-2}$ eV [28]. The lower bound is almost twice as large in the inverted ordering, since then both m_1 and m_2 must be relatively far above m_3, while in the normal ordering only m_3 must be significantly larger than the other two. Recently, the still on-going KATRIN experiment published their results of the measurement of the energy spectrum of tritium β-decay, which gave the strongest upper bound to date on the neutrino masses from direct kinematic experiments. They reported an upper bound on the weighted sum of squared masses defined by $m_\nu^2 = |U_{1i}|^2 m_i^2$ of $m_\nu < 0.8$ eV at 90% confidence level [29]. Using the unitarity of U, we find that $m_\nu \geq \min(m_i)$, so that independent of the PMNS parameter values, the limit from KATRIN is equivalent to $\min(m_i) < 0.8$ eV. This upper limit is much larger than the measured mass differences from oscillation experiments, so we can conclude that the upper limit on the sum of the three masses from KATRIN regardless of the mass ordering is $\sum m_i < 3 \times 0.8\,\mathrm{eV} = 2.4$ eV. However, the strictest upper limits on the sum of neutrino masses come from cosmology, where the masses of neutrinos affect the total amount of mass in the universe, which again affects things like the expansion rate of the universe and the Cosmic Microwave Background radiation[13]. An issue with cosmological limits is that they depend on the model of the evolution of the universe. If there for instance exist other particles with similar properties as neutrinos, such as sterile neutrinos or light thermal axions, limits based on the standard ΛCDM cosmological model may no longer apply. If the ΛCDM model is correct, one of the most recent limits from a study of large structure of the universe is $\sum m_i < 0.111$ eV at 95% confidence level [30]. However, if we take into account the uncertainty in the cosmological model by introducing more free parameters to the model, the bound loosens. By adding six new parameters to ΛCDM, an upper bound of $\sum m_i < 0.515$ eV was found in a study of Planck data [31].

The bounds on neutrino masses are improving quickly as many different experiments are measuring them with different approaches. It can be expected that within a few years the lower and upper bounds on the sum of neutrino masses will come so close that they allow the determination of the correct mass ordering, and the first precise measurements of each individual mass.

2.6 The problem

So far in this chapter we have looked at how fermions get their masses in the Standard Model through the BEH mechanism. We have also seen that there is clear evidence for at least two massive neutrinos, with masses of the order of 1 eV or smaller. All the while, we have not seemed to encounter any significant problem. The neutrinos are assumed to be massless in the SM, which does seem like an issue, since it is now in contradiction with several observations. But there is an obvious solution to this; simply add the right-handed neutrinos ν_R^i to the SM (which will have no other effects on the model due to their lack of interactions), and let the neutrinos have Yukawa couplings to the BEH field like all the other fermions. While such a model could be made to fit with all current experimental data, it is viewed by many to be an unsatisfactory explanation for the masses of the neutrinos, and it is commonly suggested that they get their mass through some other mechanism. One reason for thinking this comes down to the small scale of the neutrino masses. The lightest of the other fermions is the electron, with a

[13] See for example section 26 of the PDG review [13].

mass of around 511 keV [13]. As we saw in section 2.5.5, even the most conservative upper limits on the neutrino masses would put the largest neutrino mass almost six orders of magnitude below the electron mass. If neutrinos also get their mass through the BEH mechanism, there would be no reason to expect their masses, or equivalently their Yukawa couplings to the Higgs boson, to not be close to those of the charged leptons in size. So even if the neutrinos do couple to the Higgs boson, there would still be an open question connected to the smallness of their couplings. Additionally, as we will see in the next chapter, the special properties of neutrinos opens the possibility of another mass generating mechanism that is not available for the other fermions in the SM. This gives rise to a class of theories which could provide the neutrinos with small masses, and which in many ways seem to be more natural than the BEH mechanism alone.

Chapter 3

Neutrino mass models

In this chapter we introduce the three versions of the seesaw mechanism that we will study extensively in the second part of this thesis. We will show in particular how each version of the seesaw mechanism can give small masses to the SM neutrinos, and how the heavy neutrinos in these models could interact with SM particles. But first, we need to cover some general properties of Majorana fermions, which are central to the seesaw mechanism.

3.1 Majorana fermions

One thing that separates neutrinos from the other fermions is that it is not known what type of fermion they are. While quarks and charged leptons are known to be Dirac fermions, it is unknown whether neutrinos are Dirac or Majorana fermions [32]. The conceptual difference between Dirac and Majorana fermions is that Dirac fermions have distinct antiparticles, while Majorana fermions can be said to be their own antiparticles. To be more precise, we define the particle-antiparticle conjugate operator $\hat{C}$, which acts on a fermion spinor ψ as

$$\hat{C}(\psi) \equiv \psi^c = C\overline{\psi}^T, \tag{3.1}$$

where T is the matrix transpose, and C is the charge conjugation matrix satisfying the conditions [33]

$$C^{-1}\gamma_\mu C = -\gamma_\mu^T, \quad C^T = -C, \quad C^{-1} = C^\dagger. \tag{3.2}$$

Then the condition for a fermion spinor ψ_M to describe a Majorana particle is (up to an unphysical phase factor, which we set to 1 for simplicity)

$$\psi_M = \psi_M^c. \tag{3.3}$$

This condition limits the properties of any Majorana fermion. In particular, any global or local $U(1)$ gauge transformation of a Majorana field

$$\psi_M \to e^{iQ\theta}\psi_M \tag{3.4}$$

with $Q \neq 0$ would break the Majorana condition, since the conjugate field transforms as

$$\psi_M^c \to e^{-iQ\theta}\psi_M^c. \tag{3.5}$$

Therefore a Majorana fermion can not have any non-zero $U(1)$ charges, including for example electric charge and additive lepton number. The neutrinos are the only fermions in the SM that have no $U(1)$ charges[1], which is why they are also the only fermion that could possibly be of Majorana type.

As we saw in section 2.2, a fermion spinor can be decomposed into a left-handed and a right-handed chiral spinor using the projection operators P_L and P_R (see equation (2.13)). The chiral spinors still have four components, but only two degrees of freedom each. For Dirac fermions ψ_R is completely independent of ψ_L, which means that the complete fermion spinor has four independent degrees of freedom. For a Majorana fermion on the other hand, the Majorana condition gives

$$\psi_M = \psi_L + \psi_R = \psi_M^c = (\psi_L + \psi_R)^c = (\psi_L)^c + (\psi_R)^c. \tag{3.6}$$

We can compute $(\psi_L)^c$ using the definition of the particle-antiparticle conjugate

$$\begin{aligned}
(\psi_L)^c &= C\overline{\psi}_L^T = C\left[(P_L\psi_M)^\dagger \gamma^0\right]^T = C\left(\gamma^0\right)^T \left(\psi_M^\dagger P_L^\dagger\right)^T \\
&= C\left(\gamma^0\right)^T P_L^T \left(\psi_M^\dagger\right)^T = C\left(P_L\gamma^0\right)^T \left(\psi_M^\dagger\right)^T = C\left(\gamma^0 P_R\right)^T \left(\psi_M^\dagger\right)^T \\
&= CP_R^T \left(\gamma^0\right)^T \left(\psi_M^\dagger\right)^T = \frac{1}{2}\left(C + Ci\left[\gamma^0\gamma^1\gamma^2\gamma^3\right]^T\right)\overline{\psi}_M^T \\
&= \frac{1}{2}\left(C + iC\left(\gamma^3\right)^T\left(\gamma^2\right)^T\left(\gamma^1\right)^T\left(\gamma^0\right)^T\right)\overline{\psi}_M^T \\
&= \frac{1}{2}\left(C + i\left(-1\right)^4\gamma^3\gamma^2\gamma^1\gamma^0 C\right)\overline{\psi}_M^T = \frac{1}{2}\left(C + i\left(-1\right)^{10}\gamma^0\gamma^1\gamma^2\gamma^3 C\right)\overline{\psi}_M^T \\
&= P_R C\overline{\psi}_M^T = P_R\psi_M^c.
\end{aligned} \tag{3.7}$$

This shows that the particle-antiparticle conjugate of a left-handed chiral spinor is a right-handed chiral spinor, and a similar computation shows that the conjugate of a right-handed spinor is left-handed. For the Majorana condition (3.6) to be satisfied, this means that we must have

$$(\psi_L)^c = \psi_R, \tag{3.8}$$

which automatically gives

$$(\psi_R)^c = ((\psi_L)^c)^c = \psi_L, \tag{3.9}$$

and thus the Majorana condition is satisfied. But then we can write the complete Majorana spinor ψ_M in terms of only one of the chiral fields

$$\psi_M = \psi_L + (\psi_L)^c = \psi_R + (\psi_R)^c. \tag{3.10}$$

By defining the chiral decomposition of ψ_M^c to be

$$\psi_M^c = P_L\psi_M^c + P_R\psi_M^c = \psi_L^c + \psi_R^c, \tag{3.11}$$

we see from equation (3.7) (and the equivalent one for the opposite chirality) that $(\psi_L)^c = \psi_R^c$ and $(\psi_R)^c = \psi_L^c$. Then the Majorana field can also be written as

$$\psi_M = \psi_L + \psi_R^c = \psi_R + \psi_L^c. \tag{3.12}$$

[1] Neutrinos can be said to possess a non-zero lepton number, but the global lepton number conserving symmetry of the SM is accidental. If neutrinos are Majorana fermions, the consequence would be that this symmetry breaks down, and the concept of a lepton number no longer makes sense.

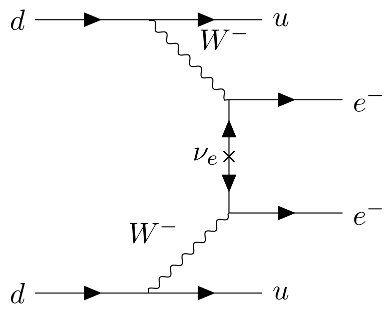

Figure 3.1: The Feynman diagram of the neutrinoless double beta decay process, mediated by a Majorana type electron neutrino ν_e.

Equation (3.10) shows that the entire Majorana spinor ψ_M is specified by a single chiral spinor ψ_L or ψ_R, meaning that a Majorana spinor only has two degrees of freedom. This is another distinction from Dirac fermions. The four degrees of freedom of a Dirac fermion can be interpreted as the two spin states $s = \pm 1/2$ of both the particle and antiparticle, while for a Majorana fermion the particle and antiparticle are the same, so it is logical that the field has a single pair of degrees of freedom corresponding to the spin states of this particle.

Majorana fermions also have special propagators that do not exist for Dirac fermions. This comes from the fact that the two-point correlation function

$$\langle 0| \, T \left[\psi_\alpha(x)\psi_\beta(y)\right] |0\rangle \tag{3.13}$$

is non-zero for Majorana fermions, which means that they have two additional propagators, corresponding to a Majorana particle turning into its antiparticle or vice versa. For the case of a Majorana neutrino, these propagators allow the famous neutrinoless double beta decay process, which is shown in figure 3.1[2]. The cross in the middle of the electron neutrino propagator indicate that it changes from particle to antiparticle, as is also indicated by the two outward pointing arrows. The experimental search for this process is one of the most promising ways to determine whether the SM neutrinos are of Majorana or Dirac type. There are other Beyond the Standard Model (BSM) processes that could give rise to neutrinoless double beta decay. However, it can be shown that regardless of the dominant process, the existence of this process implies that neutrinos are Majorana fermions [34]. And if it is assumed that the mechanism of figure 3.1 is dominant, the rate of this process will be [35]

$$\Gamma \propto \left(\frac{m_{ee}}{m_W}\right)^2, \tag{3.14}$$

where m_W is the mass of the W boson, and m_{ee} is the effective Majorana mass of the electron neutrino, defined by

$$m_{ee} = \left|\sum U_{1i}^2 m_i\right|. \tag{3.15}$$

One of the leading experiments in search of neutrinoless double beta decay is the GERDA experiment, which has recently set an upper limit of the order of 0.1 eV on m_{ee} [32]. However, from (3.15) we see that if the PMNS matrix U is not real, some terms in the

[2]This diagram, and all other Feynman diagrams shown in this thesis, are meant to be read with positive time direction from left to right.

effective mass sum can be negative or complex. This means that if the CP phase (in fact any of the three CP phases in the case of light Majorana neutrinos) of the PMNS matrix is non-zero, some of the terms can cancel each other out, which means that the upper limit on m_{ee} does not directly translate to upper limits on m_i without knowledge of the extent of CP violation in the PMNS matrix.

3.1.1 Majorana mass terms

The mass term in a Majorana fermion Lagrangian takes a slightly different form than for a Dirac fermion. The mass term for a Majorana field ψ_M can be written as

$$\begin{aligned}
\mathcal{L}_M &= -\frac{1}{2}M\overline{\psi}_M\psi_M = -\frac{1}{2}M\left(\overline{\psi}_L + \overline{(\psi_L)^c}\right)\left(\psi_L + (\psi_L)^c\right) \\
&= -\frac{1}{2}M\left(\overline{\psi}_L C\overline{\psi}_L^T + \overline{C\overline{\psi}_L^T}\psi_L + \overline{C\overline{\psi}_L^T}C\overline{\psi}_L^T\right) \\
&= -\frac{1}{2}M\left(\overline{\psi}_L C\overline{\psi}_L^T + \psi_L^T\gamma^0 C^\dagger\gamma^0\psi_L + \psi_L^T\gamma^0 C^\dagger\gamma^0 C\gamma^0\psi_L^*\right) \\
&= -\frac{1}{2}M\left(\overline{\psi}_L C\overline{\psi}_L^T - \psi_L^T C^\dagger\psi_L - \psi_L^T\gamma^0\psi_L^*\right) \\
&= \frac{1}{2}M\left(\psi_L^T C^\dagger\psi_L - \overline{\psi}_L C\overline{\psi}_L^T - \left(\overline{\psi}_L\psi_L\right)^*\right) \\
&= \frac{1}{2}M\left(\psi_L^T C^{-1}\psi_L - \overline{\psi}_L C\overline{\psi}_L^T\right),
\end{aligned} \tag{3.16}$$

where we introduced a conventional factor of 1/2, M is some complex constant, and we used the following properties of C:

$$\begin{aligned}
C^\dagger\gamma^0 &= -\gamma^0 C^\dagger \\
C^\dagger &= C^{-1}.
\end{aligned} \tag{3.17}$$

Alternatively, we can write the Majorana mass term without explicitly writing out the particle-antiparticle conjugation as

$$\mathcal{L}_M = -\frac{1}{2}M\overline{(\psi_L)^c}\psi_L + h.c. = -\frac{1}{2}M\overline{\psi^c}_R\left(\psi_R^c\right)^c + h.c. \tag{3.18}$$

Equation (3.18) demonstrates a very important property of the Majorana mass term. It can be written in terms of only one of the chiral spinors of the fermion! This is in contrast to the Dirac fermion mass term, which we saw in equation (2.16) must contain both the left-handed and right-handed spinors, which is the reason why the term breaks the SM gauge symmetry. The Majorana mass term, however, avoids this problem. If we think of the Majorana fermion as being the SM neutrino, we know that the right-handed field ψ_R has no charges under the SM gauge group. It is therefore trivial that the Majorana mass term for the neutrino expressed only in terms of ψ_R is gauge invariant, and thus allowed in the SM Lagrangian. This realisation is the foundation for most theories of neutrino mass generation beyond spontaneous symmetry breaking. In the following section, we look at a general example of three generations of neutrinos with Majorana mass terms, which will be highly relevant when moving to the so-called seesaw mechanism.

3.1.2 Dirac and Majorana mass terms

We will now look at a theory with three Majorana fermions (neutrinos) with left-handed chiral fields that in the weak eigenstate basis are called ν_L^l, where $l = e, \mu, \tau$ is a flavour index. These can be thought of as the SM neutrinos. Then we also introduce three separate Majorana fermions with right-handed fields that are called $N_R'^i$, $i = 1, \ldots, 3$ when they are not in the mass eigenstate basis. These correspond to the right-handed neutrinos that are not included in the SM, and have not yet been observed. It is important to note that the fields $N_R'^i$ are now not the same as the right-handed counterparts of ν_L^l, which are as before simply the particle-antiparticle conjugates of ν_L^l. As we have seen, it is possible to write down Majorana mass terms for these neutrino fields in the SM. In fact, it is possible to give both ν_L^l and $N_R'^i$ Majorana masses. Additionally, by combining the two types of fields, we can create what we will call "Dirac mass terms", which are on the same form as the Dirac fermion mass terms of the SM. In such a theory, the most general Lagrangian for the masses of the neutrinos is

$$\mathcal{L}_{\text{mass}}^{M+D} = - m_D^{il} \overline{N}_R'^i \nu_L^l - \frac{1}{2} M_L^{ll'} \overline{\left(\nu_L^l\right)^c} \nu_L^{l'} - \frac{1}{2} M_R^{ij} \overline{N}_R'^i \left(N_R'^j\right)^c + h.c. \tag{3.19}$$

Here, in each term there is an implicit sum over each repeated index, m_D is a complex 3 by 3 matrix, and M_L and M_R are complex and symmetric[3] 3 by 3 matrices. To find the transformation of the six Majorana fields to the mass eigenstate basis, we define the vectors

$$\begin{aligned} \nu_L &\equiv \begin{pmatrix} \nu_L^e \\ \nu_L^\mu \\ \nu_L^\tau \end{pmatrix}, \quad N_R' \equiv \begin{pmatrix} N_R'^1 \\ N_R'^2 \\ N_R'^3 \end{pmatrix}, \\ n_L &\equiv \begin{pmatrix} \nu_L \\ (N_R')^c \end{pmatrix} = \begin{pmatrix} \nu_L \\ N_L'^c \end{pmatrix}. \end{aligned} \tag{3.20}$$

Then it follows that

$$(n_L)^c = \begin{pmatrix} (\nu_L)^c \\ N_R' \end{pmatrix}, \tag{3.21}$$

and we see that we can write the full Majorana mass Lagrangian of equation (3.19) in matrix form as

$$\mathcal{L}_{\text{mass}}^{M+D} = -\frac{1}{2} \overline{(n_L)^c} \mathcal{M} n_L + h.c. \tag{3.22}$$

Here we introduce the six by six mass matrix $\mathcal{M}$ given in block form by

$$\mathcal{M} = \begin{pmatrix} M_L & m_D^T \\ m_D & M_R \end{pmatrix}, \tag{3.23}$$

from which we can see that $\mathcal{M}$ is complex and symmetric. It is easy to see that equation (3.22) gives the same Majorana mass terms as equation (3.19), but the Dirac mass term is not as straight forward, so we will show that it is in fact also the same in both equations.

[3]See [36] for a proof that the Majorana mass matrices must be symmetric.

Performing the matrix multiplication in equation (3.22), the terms involving m_D are

$$
\begin{aligned}
\mathcal{M}_{\text{mass}}^{M+D} &\sim -\frac{1}{2}\overline{(\nu_L)^c} m_D^T \left(N_R'\right)^c - \frac{1}{2}\overline{N'}_R m_D \nu_L + h.c. \\
&= \frac{1}{2}\nu_L^T C^{-1} m_D^T C \overline{N'}_R^T - \frac{1}{2}\overline{N'}_R m_D \nu_L + h.c. \\
&= \frac{1}{2}\nu_L^T m_D^T \overline{N'}_R^T - \frac{1}{2}\overline{N'}_R m_D \nu_L + h.c. \\
&= -\frac{1}{2}\left(\overline{N'}_R m_D \nu_L\right)^T - \frac{1}{2}\overline{N'}_R m_D \nu_L + h.c. \\
&= -\frac{1}{2}\overline{N'}_R m_D \nu_L - \frac{1}{2}\overline{N'}_R m_D \nu_L + h.c. = -\overline{N'}_R m_D \nu_L + h.c.
\end{aligned}
\tag{3.24}
$$

Here we used the properties of C to rewrite $\overline{(\nu_L)^c}$, and that when interchanging two fermion fields a negative sign is introduced, and that the transpose of a number is equal to the number itself. This shows that equations (3.19) and (3.22) are equivalent. Since $\mathcal{M}$ is symmetric, it can be diagonalised with a single unitary matrix [36], meaning that

$$\mathcal{M} = W^* \mathcal{M}_d W^\dagger, \tag{3.25}$$

where W is some six by six unitary matrix, and $\mathcal{M}_d$ is diagonal with real eigenvalues. Equation (3.22) can also be written on the form

$$\mathcal{L}_{\text{mass}}^{M+D} = \frac{1}{2} n_L^T C^{-1} \mathcal{M} n_L + h.c. \tag{3.26}$$

Then we see that the relation between n_L and the left-handed mass eigenstate fields that we will call χ_L is

$$n_L = W \chi_L, \tag{3.27}$$

and the Lagrangian mass terms in the mass eigenstate basis are

$$
\begin{aligned}
\mathcal{L}_{\text{mass}}^{M+D} &= \frac{1}{2}\chi_L^T C^{-1} \mathcal{M}_d \chi_L + h.c. = -\frac{1}{2}\overline{(\chi_L)^c} \mathcal{M}_d \chi_L + h.c. \\
&= \sum_{i=1}^{2n} -\frac{1}{2}\overline{\chi}^i M_d^{ii} \chi^i + h.c.
\end{aligned}
\tag{3.28}
$$

Here we have defined

$$\chi = \chi_L + (\chi_L)^c, \tag{3.29}$$

and we see that the masses of the six mass eigenstates are exactly the diagonal entries of $\mathcal{M}_d$. Equation (3.29) also shows that as long as the Majorana mass matrices M_L and M_R are non-zero, the mass eigenstates will satisfy the Majorana condition, and are therefore said to be of Majorana type. This can also be seen from the fact that as long as $\mathcal{M}_d$ has no degenerate eigenvalues, there are six distinct mass eigenstates, which is the same number as there are chiral fields ν and N in total. Thus each chiral field with only two degrees of freedom gets its own mass eigenstate, which must therefore be of Majorana type.

Finally, we want to see how the mixing between the two bases n_L and χ_L can give interactions to the otherwise sterile $N_R'^i$. To do this, we decompose the mixing matrix W into four 3 by 3 matrices, as follows

$$W = \begin{pmatrix} U' & V' \\ X & Y \end{pmatrix}. \tag{3.30}$$

The unitarity of W gives the useful relation

$$U'^{\dagger}U' + V'^{\dagger}V' = \mathbb{I}. \tag{3.31}$$

We can also separate the vector of left-handed mass eigenstates into the part containing the mass eigenstates of ν_L^l, which we will call ν_i, and the part containing the mass eigenstates of $N_L'^{i\,c}$, which we will call N_k. Then equation (3.27) gives

$$\nu_L^l = U'_{li}\nu_i + V'_{lk}N_k, \tag{3.32}$$

where there is a sum over i and k from 1 to 3. Now we assume that the left-handed fields ν_L^l have weak interactions with n flavours of charged leptons through the $W^{\pm}$ bosons and with themselves through the Z boson like the three SM neutrinos do. We also assume that the Dirac mass term arises from the same mechanism of spontaneous symmetry breaking as the fermion masses in the SM. This means that they originate from Yukawa interactions of the left-handed and right-handed neutrino fields with the Higgs boson. Then it can be shown[4] that the physical mass eigenstates of the right-handed neutrino fields enter the weak interaction and Yukawa interaction Lagrangian terms in the following way

$$\begin{aligned}\mathcal{L}_{\text{Int}} = &\frac{g}{\sqrt{2}}\sum_{k=1}^{3}\sum_{l=e}^{\tau} W_{\mu}^{-}\bar{l}_L\gamma^{\mu}V_{lk}N_k + \frac{g}{2\cos\theta_W}\sum_{k=1}^{3}\sum_{l=e}^{\tau} Z_{\mu}\bar{\nu}_L^l\gamma^{\mu}V_{lk}N_k \\ &+ \frac{gm_N^k}{2m_W}\sum_{k=1}^{3}\sum_{l=e}^{\tau} h\bar{\nu}_L^l V_{lk}N_k + h.c.\end{aligned} \tag{3.33}$$

Here g is the weak coupling constant, l_L is the left-handed field of the SM charged lepton of flavour l, $W_{\mu}^{\pm}$, Z_{μ} and h are the fields of the SM W, Z and Higgs bosons, m_N^k is the mass of the right-handed neutrino N_k and V_{lk} are elements of the mixing matrix V that is related to the matrix V' in equation (3.30) by a simple unitary rotation. This means in particular that the unitarity condition in equation (3.31) still holds with V' replaced by V, giving

$$U^{\dagger}U + V^{\dagger}V = \mathbb{I}, \tag{3.34}$$

where U is the mixing matrix that enters the weak interaction between the left-handed neutrinos and the charged leptons, and in the interactions of the left-handed neutrinos with themselves, similar to the PMNS matrix in the SM. The important point shown by the Lagrangian in equation (3.33) is that the physical particles of the fields $N_R'^{i}$ participate in the weak interaction through the single mixing matrix V, even if the weak eigenstates do not have any interaction terms in the Lagrangian.

3.2 Type 1 seesaw

The seesaw mechanism in general is a group of theories that attempt to explain the small masses of the three known neutrinos, usually by introducing a very heavy mass scale to the theory. The seesaw mechanism can naturally appear within the frameworks of Grand Unified Theories (GUTs), and it often makes use of the possibility of Majorana mass terms for neutrinos, as described in the previous section. What is now called the type

[4]See appendix A of *The Search for Heavy Majorana Neutrinos* [35].

1 seesaw mechanism, or just type 1 seesaw, is arguably the simplest of these theories, and was developed already in the 1970s, for example in the paper *$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?* [37] from 1977. To get small masses for the SM neutrinos ν_L^i through the type 1 seesaw, we simply need to add three right-handed neutrino fields N_R^i to the SM. The right-handed neutrino fields get Majorana mass terms and Yukawa couplings to the Higgs field and SM neutrinos, which results in a Dirac mass term. This therefore corresponds to the situation described in section 3.1.2, with $M_L = 0$. The Lagrangian terms for the neutrino masses in this case are

$$\mathcal{L}_M^{\text{Type 1}} = -\, m_D^{ij}\overline{N}_R^{i\prime}\nu_L^{j\prime} - \frac{1}{2}M_R^{ij}\overline{N}_R^{i\prime}\left(N_R^{j\prime}\right)^c + h.c. \tag{3.35}$$

Here we marked the neutrino fields with primes to indicate that they are the fields in the weak eigenstate basis, meaning the matrices m_D and M_R are in general not diagonal. The combined mass matrix for the six neutrino fields $\nu_L^{i\prime}$ and $N_R^{i\prime}$ is a six by six matrix $\mathcal{M}$ given by

$$\mathcal{M} = \begin{pmatrix} 0 & m_D^T \\ m_D & M_R \end{pmatrix}. \tag{3.36}$$

To get the desired small masses for ν_L^i, we need to assume that the Dirac masses are much smaller than the Majorana masses of the right-handed neutrinos: $m_D \ll M_R$. To show how this leads to small masses for ν_L^i after diagonalisation is not trivial for the case of three generations of neutrinos. We therefore begin by considering the simpler case of only one generation of active neutrino ν_L, and one generation of sterile neutrino N_R^i. In this case m_D and M_R are just numbers, so $\mathcal{M}$ is a 2x2 matrix. Assuming that the components of $\mathcal{M}$ are also real, we can diagonalise it with a simple rotation matrix

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}. \tag{3.37}$$

The diagonalisation of $\mathcal{M}$ from equation (3.25) for the case of a real valued U gives

$$\mathcal{M}_d = U^T\mathcal{M}U. \tag{3.38}$$

Performing this matrix multiplication, we can find the four components of $\mathcal{M}_d$

$$\begin{aligned} \mathcal{M}_d^{11} &= -m_D\sin 2\theta + M_R\sin^2\theta \\ \mathcal{M}_d^{12} = \mathcal{M}_d^{21} &= -\frac{1}{2}M_R\sin 2\theta + m_D\cos 2\theta \\ \mathcal{M}_d^{22} &= M_R\cos^2\theta + m_D\sin 2\theta. \end{aligned} \tag{3.39}$$

Requiring the off-diagonal elements to be zero, we get

$$\begin{aligned} -\frac{1}{2}M_R\sin 2\theta + m_D\cos 2\theta = 0 \\ M_R\sin 2\theta =& 2m_D\cos 2\theta \\ \tan 2\theta =& \frac{2m_D}{M_R}. \end{aligned} \tag{3.40}$$

Using the assumption $m_D \ll M_R$, we get

$$\theta \approx \frac{m_D}{M_R}, \tag{3.41}$$

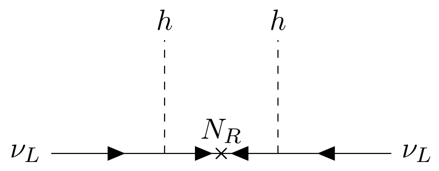

Figure 3.2: The Yukawa interactions of the left- and right-handed neutrinos that give the left-handed neutrinos small masses in the type 1 seesaw mechanism after spontaneous symmetry breaking of the BEH field.

since $\tan(x) \approx x$ for small x. Inserting this back into the expression for the diagonal elements of $\mathcal{M}_d$ and using the smallness of θ to only keep terms that are first order in θ, we get the mass eigenvalues

$$\begin{aligned} \mathcal{M}_d^{11} &\approx -2\frac{m_D^2}{M_R} + \frac{m_D^2}{M_R} = -\frac{m_D^2}{M_R} \\ \mathcal{M}_d^{22} &\approx 2\frac{m_D^2}{M_R} + M_R - \frac{m_D^2}{M_R} \approx M_R. \end{aligned} \tag{3.42}$$

Now we see that the mass of the active neutrino ν_L becomes very small; of the order of m_D^2/M_R. The heavy sterile neutrino retains its large mass, with corrections of the same order as the active neutrino mass.

Although the case of three generations of ν_L^i and N_R^i is not as straight forward, it can be shown that the exact same suppression occurs for the masses of ν_L^i. The three left-handed neutrinos get a mass matrix m_ν that is given by [33]

$$m_\nu \approx -m_D^T M_R^{-1} m_D, \tag{3.43}$$

which will have eigenvalues that are suppressed by the size of the matrix elements of M_R compared to the matrix elements of m_D^2. So this will again automatically give very small neutrino masses if $m_D \ll M_R$ (taking now m_D and M_R to be the scale of matrix elements rather than the whole matrices). Since we assume that the Dirac mass terms arise from the normal spontaneous symmetry breaking mechanism of the SM via Yukawa couplings to the Higgs boson, we can illustrate the neutrino mass matrix in equation 3.43 diagrammatically. The lowest order diagram giving the light neutrino masses in type 1 seesaw is shown in figure 3.2.

An important question to ask is how large the scale of M_R must be in order to give active neutrino masses that agree with the current experimental bounds. To get active neutrinos with masses of at most 1 eV, assuming that the largest eigenvalue of m_D is of the order of the electroweak scale of 100 GeV, we get

$$M_R = \frac{m_D^2}{m_\nu} \approx \frac{(100 \times 10^9)^2}{1}\,\text{eV} = 10^{22}\,\text{eV} = 10^{13}\,\text{GeV}. \tag{3.44}$$

This is close to the expected scale of the proposed grand unification of electroweak and strong interactions, which explains why seesaw mechanisms are often proposed within a grand unified theory (GUT) framework [38]. The discovery of the neutrino masses therefore made the idea of grand unification even more appealing from a theoretical viewpoint. Unfortunately, the Majorana mass scale $M_R \sim 10^{13}$ GeV required in this case is far higher than what we are able to probe directly at particle accelerators. This scenario therefore makes it very difficult to get any predictions from type 1 seesaw that can be tested in experiments. If the eigenvalues of m_D are instead closer to the electron mass of ~ 1 MeV, the Majorana mass scale required to give 1 eV neutrino masses reduces to $M_R \sim 1$ TeV. This is within the reach of current collider experiments, such as the LHC.

Regardless of the mass scale required, the type 1 seesaw gives an attractively simple explanation for the small neutrino masses. But there are also other versions of the seesaw mechanism that give the same natural suppression of the active neutrino masses, while also predicting new physics, including new neutrinos, that could appear at lower energy scales.

3.3 Inverse seesaw

The "inverse seesaw" mechanism is a promising alternative to the basic seesaw mechanism described in the previous section. It was first proposed within the frameworks of superstring theories [39], but this framework is not required for the mechanism to occur. One thing that most realisations of the inverse seesaw mechanism have in common is that they introduce an explicit lepton number symmetry to the theory, usually in terms of a $U(1)_{B-L}$ gauge symmetry [40]. Here B is the baryon number and L is the lepton number. This symmetry is subsequently broken in order to give masses to the neutrinos.

Here we will look at a simple realisation of the inverse seesaw mechanism, which is presented in the paper *Heavy Neutrinos with Dynamic Jet Vetoes* [41]. In this theory, the SM is extended by two sets of new fermion singlets, each coming in three generations. One of these sets is equivalent to the right-handed neutrinos N_R^i we have introduced before, but the other triplet X_R^i is different in that it has the opposite lepton number, $L = -1$. These opposite lepton numbers are what give this model an explicit almost-conserved lepton number. In this theory, both the N_R^i and X_R^i get Majorana mass terms. There is also a Dirac-like mass term with combinations of N_R^i and X_R^i fields, in addition to the Dirac mass term of ν_L and N_R that we remember from the type 1 seesaw. In total, the neutrino mass Lagrangian can be written as

$$\begin{aligned}\mathcal{L}_{\text{mass}} = &-\overline{\nu}_L m_D N_R - \overline{N}_L^c m_R X_R \\ &- \frac{1}{2}\overline{X}_L^c \mu_X X_R - \frac{1}{2}\overline{N}_L^c \mu_R N_R + h.c.\end{aligned} \tag{3.45}$$

Here the Majorana mass matrices μ_X and μ_R can be interpreted as the masses of the new particles N_R^i and X_R^i, but also as coupling constants for the lepton number violating Majorana self-interactions. Therefore the elements of these matrices must be small compared to the electroweak scale, since no such lepton number violation has been observed. In fact, according to the paper [41], the sterile neutrino masses from the matrix μ_R only contribute to light neutrino masses at sub-leading level, and they can therefore be ignored without any significant impact on the mass generation or on experimental signature. Setting $\mu_R = 0$ and defining the basis

$$n_R^T \equiv \left(\nu_R^c,\, N_R,\, X_R\right) \tag{3.46}$$

the mass terms can be written in matrix form as

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} n_R^T \mathcal{M}_{\text{ISS}} n_R + h.c. \tag{3.47}$$

Where $\mathcal{M}_{\text{ISS}}$ is a 9x9 mass matrix of the form

$$\mathcal{M}_{\text{ISS}} = \begin{pmatrix} 0 & m_D & 0 \\ m_D^T & 0 & m_R \\ 0 & m_R^T & \mu_X \end{pmatrix}. \tag{3.48}$$

This is another symmetric mass matrix, and it can be diagonalised by a single unitary matrix. This gives the following 3x3 mass matrix for the SM neutrinos [41]

$$m_\nu^{\text{light}} \approx m_D (m_R^T)^{-1} \mu_X m_R^{-1} m_D^T, \tag{3.49}$$

which is reminiscent of the equivalent matrix in type 1 seesaw, however it is now the smallness of μ_X that causes the neutrino masses to be small, instead of the largeness of m_R. It can also be shown that in the $\mu_X \to 0$ limit, the heavy neutrino mass eigenstates of N_R^i and X_R^i become pairwise degenerate. Since the energy scales at particle accelerators are much larger than the light neutrino masses, and therefore also μ_X, $\mu_X = 0$ is an excellent approximation for the purposes of such experiments. Since the Majorana fields N_R^i and X_R^i have the same values for all charges except lepton number, they are essentially indistinguishable when their masses become degenerate. This means that we can think of each pair of N_R^i and X_R^i as a single particle. Since each of the Majorana fields have two degrees of freedom each, this new particle gets four degrees of freedom, just like a Dirac fermion. The resulting states are therefore referred to as pseudo-Dirac particles (neutrinos). Calling these pseudo-Dirac states N'_m for $m = 1, 2, 3$, we can then reduce the mass matrix to a 6x6 matrix for the masses of the three light neutrinos and the three heavy pseudo-Dirac neutrinos. This situation is exactly the same as the one we studied in section 3.1.2, so we know that the mass eigenstates of the light and heavy neutrinos can be written as a unitary rotation of their weak eigenstates as in equation (3.27). The couplings of the pseudo-Dirac neutrinos to the SM leptons through the weak interaction and the Yukawa interaction also have the exact same form as the Lagrangian in equation (3.33). The pseudo-Dirac neutrinos in this version of the inverse seesaw mechanism therefore behave very similarly to the Majorana neutrinos of the type 1 seesaw. But there is one key difference that makes it possible to distinguish the two cases in experiments. In the limit $\mu_X, \mu_R \to 0$, lepton number is conserved in the inverse seesaw theory. This means that the pseudo-Dirac neutrinos also must conserve lepton number in their interactions, as opposed to Majorana neutrinos which can change lepton number by two units (for example in the neutrinoless double beta decay process shown in figure 3.1). A N_m particle can therefore only decay to a boson and a lepton, while a N_m antiparticle can only decay to a boson and an antilepton. Lepton flavour must not be conserved, as we see from equation (3.33) if the mixing matrix V is not diagonal.

Just like in the type 1 seesaw, we can ask ourselves how small the parameter μ_X must be in order to give SM neutrino masses within current constraints. If we want the pseudo-Dirac neutrinos to be observable at collider experiments, their masses, given by m_R, must be at most a few TeV. If we again assume $m_D \sim 100$ GeV, we get an estimated upper limit of

$$m_\nu^{\text{light}} < 1\,\text{eV} \implies \mu_X < \frac{(1\,\text{eV})\left(10^{12}\,\text{eV}\right)^2}{\left(100 \times 10^9\,\text{eV}\right)^2} \approx 100\,\text{eV}. \tag{3.50}$$

Now it is tempting to ask whether the inverse seesaw mechanism actually can solve the neutrino mass problem, or if it simply transfers the problem of requiring a very small Yukawa coupling for neutrinos to instead requiring a very small parameter μ_X. There is however an important difference between the two. In the latter case, the small parameter μ_X represents a small breaking of the lepton number symmetry. The smallness of μ_X therefore meets 't Hooft's criteria for "naturalness" [42], which states that a very small parameter can be natural if that parameter being equal to zero would increase the symmetry of the theory. However, it should be acknowledged that this principle is not universally accepted [43], and to some the inverse seesaw mechanism might therefore seem like an unsatisfactory solution to the neutrino mass problem.

3.4 Constraints in a general seesaw scenario

In the paper by Pascoli et al. [41], they also present the experimental constraints on the non-unitarity of the light neutrino mixing matrix U in a general seesaw scenario. Due to the unitarity condition in equation (3.34), this in turn constrains the heavy-light mixing matrix elements that we called V_{lk} in section 3.1.2. Assuming that the mixing between the active and sterile neutrinos is dominated by the lightest of the sterile neutrinos, which we call N_1, the constraints at 95% confidence level on the parameters coupling N_1 to each flavour of leptons are

$$|V_{e1}| < 0.050, \quad |V_{\mu 1}| < 0.021, \quad |V_{\tau 1}| < 0.075 \tag{3.51}$$

3.5 The left-right symmetric model

Another theoretical framework that is commonly used to explain the small neutrino masses is the Left-Right Symmetric Model (LRSM). The LRSM was originally suggested as an extension of the Standard Model in an attempt to explain the parity violation of the weak interaction. In the LRSM, parity violation arises by introducing an exact parity symmetry at high energies that is spontaneously broken to give the maximally parity violating weak interactions of the SM. The model was first introduced in the 1970's by Pati, Salam, Mohapatra and Senjanović [44, 45, 46]. In the LRSM, the electroweak gauge group of the SM is extended to become

$$SU(2)_L \otimes SU(2)_R \otimes U(1)_{B-L}. \tag{3.52}$$

The right-handed fields of fermions then form doublets under the $SU(2)_R$ symmetry, just like the left-handed fields for $SU(2)_L$. This requires the inclusion of right-handed neutrino fields N_R^i in these models, in order to make the lepton $SU(2)_R$ doublet

$$l_R = \begin{pmatrix} N_R^i \\ e_R^i \end{pmatrix}. \tag{3.53}$$

The new $SU(2)_R$ group also introduces three new gauge bosons, which we call $W_R^\pm$ and Z_R. The observed suppression of right-handed charged current weak interactions requires the masses of these new gauge bosons to be much larger than their SM counterparts $W^\pm$ and Z. In order to break the left-right symmetry and give mass to the new gauge bosons, the Higgs sector of the SM must also be extended. In the minimal LRSM, the Higgs sector consists of a pair of $SU(2)_L$ and $SU(2)_R$ triplet scalar fields called Δ_L and Δ_R, and a bi-doublet scalar field Φ [47]. Then the breaking of the left-right symmetry and the $B - L$ symmetry is caused by a non-zero vev of the Δ_R field, which we call v_R. Since we observe maximally violated left-right symmetry at the electroweak scale, the scale of v_R must be much larger than the electroweak scale of ~ 100 GeV. The vev v_L of Δ_L, on the other hand, will affect the mass of the SM W boson, and is therefore constrained by precision measurements to be much smaller than the EW scale [48]. Finally, the bi-doublet acquires two vevs k_1 and k_2 that play the role of the SM Higgs vev, namely breaking the electroweak symmetry and giving mass to the W and Z bosons and the SM fermions. Both the SM neutrinos ν_L^i and the right handed neutrinos N_R^i can get Majorana mass terms from their Yukawa couplings to the $\Delta_{L,R}$ triplets, which after spontaneous symmetry breaking become

$$\mathcal{L}_{\text{mass}}^M = -\frac{1}{\sqrt{2}} v_L \lambda_L^{ij} \overline{\nu_L^i} \nu_R^{cj} - \frac{1}{\sqrt{2}} v_R \lambda_R^{ij} \overline{N_R^i} N_L^{cj} + h.c. = -\frac{1}{2} M_L \overline{\nu_L} \nu_R^c - \frac{1}{2} M_R \overline{N_R} N_L^c + h.c. \tag{3.54}$$

Here $\lambda_{L,R}^{ij}$ are the Yukawa couplings of the $\Delta_{L,R}$ fields to the neutrinos, and in the second equality we write the mass terms in matrix form by defining the mass matrices $M_{L,R} = \sqrt{2} v_{L,R} \lambda_{L,R}$. These Majorana mass terms mean that the SM neutrinos and the right-handed neutrinos must be of Majorana type in the LRSM, as we showed in section 3.1. A Dirac mass term for the neutrinos also arises naturally from their Yukawa coupling to Φ. In matrix form, this looks like

$$\mathcal{L}_{\text{mass}}^D = -m_D \overline{\nu_L} N_R + h.c. \tag{3.55}$$

Here m_D is a linear combination of the vevs of Φ and its Yukawa couplings [47]. Combining the mass terms from equations (3.54) and (3.55), we see that we have the exact situation described in section (3.1.2). We therefore know that the six by six mass matrix for the neutrinos will have the form

$$\mathcal{M} = \begin{pmatrix} M_L & m_D^T \\ m_D & M_R \end{pmatrix}. \tag{3.56}$$

Now, in the natural limit $v_L \to 0$, we have $M_L \sim 0$, and the neutrino mass matrix reduces to the one of the type 1 seesaw (see equation (3.36)):

$$\mathcal{M} = \begin{pmatrix} 0 & m_D^T \\ m_D & M_R \end{pmatrix}. \tag{3.57}$$

We therefore immediately know that after diagonalisation, the SM neutrino mass matrix will be

$$m_\nu \approx -m_D^T M_R^{-1} m_D, \tag{3.58}$$

while the mass matrix for the right-handed neutrinos is

$$m_N \approx M_R. \tag{3.59}$$

Since M_R is proportional to the large vev v_R, it is naturally much larger than m_D in the LRSM, giving a type 1 seesaw-like suppression of the SM neutrino masses. Since both the right-handed neutrino masses and the W_R mass are proportional to v_R, it is also natural to assume that $m_N \sim M_{W_R}$, where M_{W_R} is the mass of W_R.

Finally, due to the mixing between the weak eigenstates ν_L and N_R and their mass eigenstates caused by the non-diagonal mass matrix $\mathcal{M}$, the mass eigenstates ν_m and N_m enter the right-handed weak interactions in the following way [49]

$$\mathcal{L}_{\text{Int}}^R = -\frac{g_R}{\sqrt{2}} \sum_{l=e}^{\tau} W_{R\mu}^+ \left[\sum_{m=1}^{3} \overline{\nu_m^c} X_{lm} + \sum_{n=1}^{3} \overline{N}_n Y_{ln} \right] \gamma^\mu l_R^- + h.c. \tag{3.60}$$

Here g_R is the right-handed weak coupling constant, which must be equal to the SM weak coupling constant g at tree level due to the left-right symmetry [50], and X and Y are mixing matrices expressing the mixing of the light and heavy neutrino mass eigenstates respectively with the right-handed weak eigenstates N_R. In the minimal LRSM, these mixing matrices naturally scale as $X \sim m_\nu/m_N \ll \mathbb{I}$ and $Y \sim \mathbb{I}$ [49]. The W_R boson also couples to the right-handed fields of the SM quarks in the following way

$$\mathcal{L}_{\text{Int}}^{R} = -\frac{g_R}{\sqrt{2}} \sum_{i=1}^{3} \sum_{j=1}^{3} W_{R\,\mu}^{+} \overline{u}_R^{i} V_{\text{CKM}}^{\prime ij} \gamma^\mu d_R^{j} + h.c. \tag{3.61}$$

Here V'_{CKM} is the right-handed version of the CKM matrix, which is expected to obey $V'_{\text{CKM}} \sim V_{\text{CKM}} \sim \mathbb{I}$ [47], where V_{CKM} is the SM CKM matrix, which is known to be nearly diagonal.

In this chapter we have presented a few simple realisations of the seesaw mechanism, which all result in small masses of the SM neutrinos by introducing new, heavy neutrinos. In particular, we have seen how the heavy neutrinos in each of the models could interact with the SM particles. These interactions make it possible to either discover or falsify the theories, but in order to observe the interactions of these heavy particles a very special kind of experiment is needed. In the next chapter, we will introduce one of the few experiments in the world that could be sensitive to the interactions of the heavy neutrinos in the seesaw models.

Chapter 4

LHC and the ATLAS detector

The LHC is the world's largest and highest energy particle accelerator. It is therefore among the experiments that are best suited to search for new particles and phenomena, and in particular new heavy particles such as the heavy neutrinos that are the focus of this thesis. In our analysis in the following part, we will make extensive use of simulations of events at the LHC, with particular focus on the ATLAS detector, and we will therefore summarise the basic workings of the experiment and useful concepts connected to it.

The LHC consists of a circular beampipe, where two beams of protons are accelerated to near the speed of light in opposite directions. These beams are then made to collide within the regions of the detectors, among which is the ATLAS detector. The energy of the proton beams is usually defined in the reference frame of the detector, and is then called the centre of mass energy, with the symbol $\sqrt{s}$. At the current Run 3 of the LHC, the collider is operating at $\sqrt{s} = 13.6$ TeV. Since each of the two colliding proton beams have the same speed, they have an energy of $\sqrt{s}/2$ each in the detector frame. But since protons are composite particles, the total energy in each individual proton-proton interaction is not equal to the centre of mass energy. At such high energies, it is the constituents of the protons, the quarks and gluons, collectively referred to as "partons", that interact with each other. The true energy of the interaction depends on the fraction of the protons' energy that are carried by the interacting partons at the moment of the interaction. Since the protons are highly relativistic, their momentum and energy are very close to equal, and therefore the energy carried by a parton is usually given in terms of the ratio of the parton's momentum to the proton's total momentum, which is called x. Then the invariant mass M of the interaction between a parton with momentum fraction x_1 from one beam and a parton with momentum fraction x_2 from the other beam is $M = \sqrt{x_1 x_2 s} < \sqrt{s}$. This invariant mass is the highest mass of any particle that can possibly be created in that specific interaction.

There are many different interactions that could occur in the collision, and the probability of any given interaction to occur is given by the cross section σ. It is defined as the ratio of the number of times that specific interaction occurs per unit time divided by the flux of particles in the beams. Since flux has units of particles per area per time, the cross section has units of area. At the LHC, cross sections are typically measured in units of barn, with symbol b, defined as 1 b $= 10^{-28}$ m^2. The proton flux is commonly referred to as the luminosity $\mathcal{L}(t)$. Multiplying the luminosity by the cross section of a process we get the expected number of occurrences of that process per unit time. If we instead multiply the cross section by the integral of the luminosity over time, we get

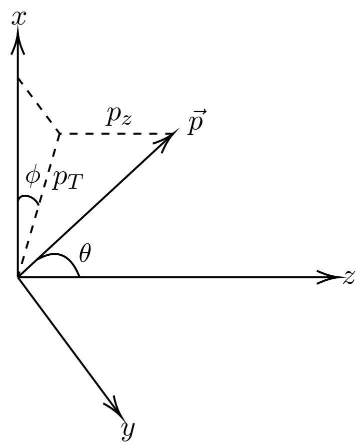

Figure 4.1: This figure illustrates the coordinate system we use to describe the kinematics of particles in a collider event. Here, a particle is being created at the primary vertex in the origin, and its momentum is shown as the vector $\vec{p}$. The z-axis is pointing in the beam direction, and the angle between $\vec{p}$ and the z-axis is the polar angle θ. The xy-plane is then what we call the transverse plane, and the projection of $\vec{p}$ into the transverse plane is the transverse momentum p_T. The angle between the transverse momentum vector and the x-axis is the azimuthal angle ϕ. The z-component of the momentum, p_z, is also labelled in the figure.

the expected number of occurrences over that period of time. This gives the important relation

$$N = \sigma \int_{t_0}^{t_1} \mathcal{L}(t)\, \mathrm{d}t = \sigma L, \tag{4.1}$$

where N is the expected number of a certain kind of event between the times t_0 and t_1, σ is the cross section of that event, and we defined the integrated luminosity L to be the integral of the luminosity over time. Since N is dimensionless, L has the inverse units of σ, and is therefore commonly measured in inverse barn, b^{-1}.

4.1 LHC kinematics and kinematic variables

All of the physical information about a collider event can be given in terms of kinematic variables. These variables can be loosely categorised in two groups; primary variables that are measured directly in the experiment, and secondary variables which are calculated from the primary ones. To define the primary kinematic variables of LHC events, we first need to specify the coordinate system used to describe the motion of particles at the LHC. This is shown in figure 4.1. The z-axis of the coordinate system is pointing in the direction of the beampipe, and the xy-plane is called the transverse plane. The angle between the z-axis and a particle's momentum vector is the polar angle called θ. The projection of the particle's momentum into the transverse plane is the transverse momentum vector $\vec{p}_T$, whose magnitude is the transverse momentum p_T. The angle between the transverse momentum vector and the x-axis is the azimuthal angle called ϕ. The transverse momentum p_T is one of the most important kinematic variables in particle colliders. One of the reasons for this is that since the incoming beam particles have negligible momentum in the transverse direction and momentum is conserved, the vector sum of the transverse momenta of all particles coming out of an event must be zero. Particles with large p_T are often called "hard", while particles with low p_T are called "soft". We will also refer to the particle with the largest p_T in an event as "leading", for example is the "leading lepton" the lepton with the largest p_T. The angles θ and ϕ are also important variables. However, θ is usually replaced by a similar variable called the pseudorapidity η, which is defined as

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \tag{4.2}$$

This means that as θ goes to zero, η goes to infinity, and when $\theta = \pi/2$, meaning that the momentum of the particle is purely in the transverse direction, $\eta = 0$. In the opposite direction, when θ goes to π, η goes to negative infinity. The pseudorapidity is symmetric around the transverse plane, which is a useful property. The azimuthal angle ϕ of a particle is usually not interesting in itself, because the distribution of ϕ is almost always uniform, due to the rotational symmetry of the event. Much more interesting is the difference in ϕ of different particles in an event. This difference is called $\Delta\phi$, and if the azimuthal angles of the two particles are called ϕ_1 and ϕ_2, we compute the $\Delta\phi$ as

$$\Delta\phi = \min\left[|\phi_1 - \phi_2|, 2\pi - |\phi_1 - \phi_2|\right], \tag{4.3}$$

such that $\Delta\phi \in [0, \pi]$.

Now we move to the secondary kinematic variables, which are not measured directly, but derived from the measured quantities energy, momentum, charge, p_T, η and ϕ. One of the most important of these, which we have already mentioned, is the invariant mass M of a system of particles. It can be defined as the total mass of that system of particles in the reference frame where the vector sum of their momenta is zero. In other reference frames, the relativistic mass of the particles will be different, but the invariant mass is a Lorentz invariant quantity, and is therefore the same in all reference frames. It can be computed from the energies E_i and the three-momenta $\vec{p}_i$ of a set of n particles as

$$M^2 = \left(\sum_{i=1}^{n} E_i\right)^2 + \left(\sum_{i=1}^{n} \vec{p}_i\right)^2. \tag{4.4}$$

A property of the invariant mass that is very useful in particle physics is that if a particle with mass m decays into some new set of particles, the invariant mass of that set of particles will be equal to m in any reference frame. This can easily be seen using energy conservation in the rest frame of the decaying particle.

Two other important variables are the transverse energy E_T and the missing transverse energy E_T^{miss}. The transverse energy of a particle is the energy it would have if it had no momentum component is the z direction, meaning that

$$E_T^2 = m^2 + \vec{p}_T^{\,2}, \tag{4.5}$$

where m is the mass of the particle. The missing transverse energy is used to get information about particles that otherwise can not be measured by the detector, such as neutrinos. Due to conservation of energy and momentum, we can infer some things about these "missing particles" from the measurements of all the other particles in the event. We first define the missing transverse momentum vector $\vec{p}_T^{\text{miss}}$ as the vector that when summed with the transverse momentum vectors of all other particles in an event would give zero. This is therefore equal to the transverse momentum vector that a missing particle, such as a neutrino, must have had in order to give the expected net 0 transverse momentum in the event. This definition of the missing transverse momentum vector gives

$$\vec{p}_T^{\text{miss}} = -\sum_i \vec{p}_{T,i}, \tag{4.6}$$

where $\vec{p}_{T,i}$ is the transverse momentum vector of particle number i, and the sum is over all particles from the event that reach the detector[1]. The missing transverse energy is then defined as the magnitude of the missing transverse momentum vector

$$E_T^{\text{miss}} = \left|\vec{p}_T^{\text{miss}}\right|. \tag{4.7}$$

This means that the missing transverse energy is only truly the transverse energy of the missing particle if the missing particle has a negligible mass compared to its momentum. But since the main source of missing transverse energy at the LHC are neutrinos, this is a very valid assumption.

Next, we want to define a variable known as transverse mass, m_T. This is calculated in a very similar way to the invariant mass, but with the energies and momenta replaced by the transverse energies and transverse momenta. For a two-body system, this gives the transverse mass

$$m_T^2 = \left(E_{T,1} + E_{T,2}\right)^2 - \left(\vec{p}_{T,1} + \vec{p}_{T,2}\right)^2, \tag{4.8}$$

where $\vec{p}_{T,i}$ and $E_{T,i}$ are the transverse momentum and transverse energy of particle i. Multiplying out the parentheses in equation (4.8) and using equation (4.5) to replace the terms with squared transverse energy, we get

$$m_T^2 = m_1^2 + m_2^2 + 2\left(E_{T,1}E_{T,2} - \vec{p}_{T,1}\vec{p}_{T,2}\right). \tag{4.9}$$

Now, if the two particles are massless or their masses are negligible, we have

$$|\vec{p}_{T,i}| = E_{T,i}, \tag{4.10}$$

and the expression for the transverse mass reduces to

$$m_T^2 = 2E_{T,1}E_{T,2}\left(1 - \cos\left(\Delta\phi\right)\right), \tag{4.11}$$

where $\Delta\phi$ is the angle between the transverse momenta of the two particles, which by definition of the transverse momentum must be an angle in the transverse plane. In many cases, one of the two particles we want to calculate the transverse mass of is not a detected particle, but the missing momentum vector. In that case, if the mass of the other particle can also be neglected, the transverse mass is computed as

$$m_T = \sqrt{2p_T E_T^{\text{miss}}\left(1 - \cos\left(\Delta\phi\right)\right)}, \tag{4.12}$$

where p_T is the transverse momentum of the measurable particle and $\Delta\phi$ is the angle between that particle and the missing momentum vector in the transverse plane. It is

[1]The sum over transverse momentum vectors includes tracks with very low p_T that are not associated with a reconstructed particle. These are called soft terms.

interesting to compare the transverse mass to the invariant mass of the two particles. Equation (4.4) for the case of two particles gives

$$M^2 = (E_1 + E_2)^2 - (\vec{p}_1 + \vec{p}_2)^2 = m_1^2 + m_2^2 + 2\,(E_1E_2 - \vec{p}_1 \cdot \vec{p}_2)\,. \tag{4.13}$$

If we again use the assumption that the masses of the two particles are negligible, we get

$$M^2 = 2E_1E_2\,(1 - \cos(\Delta\alpha))\,, \tag{4.14}$$

where $\Delta\alpha$ is the angle between the momenta of the two particles, which is in this case not restricted to the transverse plane. Comparing equations (4.14) and (4.11), we see that they have the same form, but that the general E_1, E_2 and $\Delta\alpha$ in equation (4.14) are replaced by their transverse equivalents in equation (4.11). In fact, it is easy to see that the invariant mass is the largest possible value of the transverse mass, the two being equal in the special case $E_1 = E_{T,1}$, $E_2 = E_{T,2}$ and $\Delta\alpha = \Delta\phi$, which occurs when both particles move only in the transverse plane. To see this, we first notice that $E_T \leq E$, since

$$E_T^2 = m^2 + p_T^2 = E^2 - p_z^2, \tag{4.15}$$

and p_z^2 is of course a positive number. When it comes to the angles $\Delta\phi$ and $\Delta\alpha$, they are by definition between 0 and π radians. And we must have $\Delta\phi \leq \Delta\alpha$, since the separation in 3D space must be at least as large as the separation in any given plane. Because the function $1 - \cos x$ is increasing in the interval $x \in [0, \pi]$, this implies that $1 - \cos\Delta\phi \leq 1 - \cos\Delta\alpha$. This shows that $m_T \leq M$, and that they are only equal when all three variables are equal. This makes the transverse mass a useful variable for approximating the invariant mass of two particles when one of them can not be measured directly, like a neutrino. The definition of the transverse mass can also be extended to n particles, in which case

$$m_T^2 = (E_{T,1} + \cdots + E_{T,n})^2 - (\vec{p}_{T,1} + \cdots + \vec{p}_{T,n})^2\,, \tag{4.16}$$

which has an upper bound of the invariant mass of the same n particles.

The hadronic activity h_T is defined as the scalar sum of the transverse momenta of all the charged leptons and jets[2] in an event

$$h_T = \sum_l p_{T,l} + \sum_j p_{T,j}, \tag{4.17}$$

where the two sums are over all the leptons and all the jets respectively.

Finally, we define the angular separation ΔR between two particles as

$$\Delta R^2 = \Delta\phi^2 + \Delta\eta^2, \tag{4.18}$$

where the $\Delta\phi$ between the two particles is as given in equation (4.3), and $\Delta\eta$ is simply given by $\Delta\eta = |\eta_1 - \eta_2|$, where η_1 and η_2 are the pseudorapidities of each of the particles.

[2]"Jet" is a term used to describe the cone of a large number of hadrons that result from the production/radiation of a quark or gluon in the proton-proton interaction, due to the special confinement property of QCD. In order to determine the energy and momentum of the initial state quark or gluon, one therefore has to measure the total energy and momentum of all the hadrons in the jet.

4.2 The ATLAS detector

The ATLAS detector is designed to detect the particles emerging from the proton-proton collisions at the LHC, including electrons, muons[3], photons and hadrons, and measure their energy and momenta. The detector has a cylindrical shape around the beampipe, covering all values of ϕ. But since the beampipe itself is in the way, it is impossible for the detector to cover all values of η. The ATLAS detector consists of several subdetectors: The inner detector, the electromagnetic calorimeter, the hadronic calorimeter and the muon spectrometer. The inner detector is, as the name suggests, the part of the detector that is closest to the beampipe. It consists of a pixel detector and tracking systems that give information about electrically charged particles' momenta and the sign of their charge by bending their trajectories in a strong magnetic field. It covers the range $|\eta| < 2.5$. The next layer is the electromagnetic calorimeter (ECAL), in which all electrically charged particles leave signals. Photons and electrons are stopped entirely in the ECAL, which allows their energy to be measured from the amplitude of the signal. Muons and electrically charged hadrons usually lose only a small fraction of their energy in the ECAL. The ECAL covers the range $|\eta| < 3.2$. Outside of the ECAL, there is the hadronic calorimeter (HCAL). This is designed to stop all hadrons, and thereby measure their energy, like the ECAL does for photons and electrons. Different sections of the HCAL combine to cover the range $|\eta| < 4.9$. The final detector layer of ATLAS is the muon spectrometer. This is a tracking detector with multiple layers in a strong magnetic field, that allows the measurement of the momentum and charge of the muons, which tend to pass through the other parts of the detector. The muon spectrometer covers the range $|\eta| < 2.7$.

4.2.1 Particle reconstruction and identification

When a particle moves through the ATLAS detector it can leave several signals in different parts of the detector. In order to interpret the data, we need a way to translate the raw data from the detector back to information about the particle. Ideally, we want to know the particle type, the energy and momentum, and also the point of origin of all the particles going through the detector. The first step from detector output to particle information is called ”reconstruction”. Reconstruction involves a set of algorithms that attempt to infer the path of a particle through the detector. For all particles, the calorimeter information can be used to infer the direction of motion of the particle. For charged particles, the calorimeter information is matched with a track from the tracking detector, and for muons the track from the inner detector is additionally extrapolated to a track in the muon spectrometer. By following the reconstructed path of the particle backwards, it is also possible to reconstruct the point of origin of the particle. Finally, the type of particle can also be reconstructed, for example based on which parts of the detector produce a signal. So the output of reconstruction are reconstructed objects, e.g. a reconstructed electron. However, reconstruction is not 100% accurate, and in some analyses we are not interested in all reconstructed objects of a certain type. Therefore a second step, called particle identification, is performed after reconstruction. Particle identification works differently, and presents different challenges, based on what type of particle is to be identified. For our analysis in the following part, we will mainly be

[3] Tau leptons have such short lifetimes that they typically decay before reaching the active part of the ATLAS detector, and they therefore have to be detected indirectly from their decay products [51].

concerned with the identification of electrons and muons from a signal process. These electrons and muons originate from the main proton-proton interaction (or from very close to it, for example from W or Z boson decay), and are called "prompt" leptons. These prompt leptons are only a subset of the total reconstructed lepton objects. Other reconstructed leptons can either originate from true leptons from "non-prompt" sources such as leptonic decays of hadrons or photons converting to e^+e^- pairs, or from hadronic jets or photons that are wrongly reconstructed as leptons. These wrongly reconstructed leptons are often called "fake leptons". To reduce the number of non-prompt and fake leptons in our analyses, there are different requirements put on reconstructed leptons in the identification. The requirements also come in different levels of strictness, which define so-called identification working points. For electrons, there are four commonly used identification working points, which are in increasing order of strictness called "VeryLoose", "Loose", "Medium" and "Tight". The exact requirements that must be fulfilled by a reconstructed electron in order to be identified under each of these working points are given by the ATLAS collaboration [52]. Using working points with greater strictness means that a lower fraction of reconstructed electrons will be identified as electrons. This can help to reduce background from non-prompt electrons. Similarly, there are three identification working points for muons that in increasing order of strictness are called "Loose", "Medium" and "Tight", and there are also special working points called "Low-p_T" and "High-p_T" that are specialised for identification of muons with low and high transverse momenta respectively [53].

To reduce the background from fake and non-prompt leptons even further, we can use the fact that prompt leptons tend to be "isolated", meaning that there is little activity within an area $\Delta\phi \times \Delta\eta$ of the lepton, both in the inner detector and the calorimeters. Non-prompt leptons originating from jets, on the other hand, will be associated with a lot of activity in the HCAL in the same direction as the leptons' trajectories. There are two types of isolation; calorimeter-based and track-based. The calorimeter-based isolation estimates the total transverse energy deposited from all other particles in both calorimeters within a cone of some size ΔR centered at the path of the lepton candidate. Track-based isolation uses the scalar sum of the transverse momenta of tracks originating from the primary vertex within some cone of size ΔR of the lepton candidate. The size of the cone can also be varied depending on the transverse momentum of the lepton candidate, to optimize results for high p_T. Requirements on calorimeter- and track-based isolation variables are also given in terms of a set of working points that are defined by the ATLAS collaboration [52, 53].

Finally, one factor that complicates the reconstruction of an event of interest is the fact that each time the two proton beams intersect, there are multiple interactions happening between the protons. The average number of interactions per beam crossing is referred to as "pile-up". The pile-up can affect things like the number of fake leptons and energy measurements of jets.

We have now covered the main motivations for and the theoretical background required to understand our analysis, which will be presented in the remaining part of the thesis.

Part II

Analysis and Results

Chapter 5

Signal and background processes

In this chapter we present the specific signal processes that we will study in our following analysis and some of the Standard Model processes that are similar to these. We begin by introducing the software we used to perform simulations of the signal processes.

5.1 Monte Carlo simulations using MadGraph

Monte Carlo (MC) simulations in general are computer simulations that use sampling of (pseudo-) random numbers to obtain a numerical result. This means that MC simulation of a single process can give many different results when it is repeated multiple times, even when the initial conditions are the exact same each time. In particle physics meanwhile, the parameters we can obtain from a theory, like the (differential) cross sections and branching ratios of different processes, only represent the probability of an event to occur in a certain way. Particle physics events therefore have the same property of producing a variety of results from identical initial conditions. This is why MC simulations are well suited to simulate particle physics events. More specifically, in this thesis we will use numerical algorithms to create MC simulations of events at the LHC. Such algorithms are called MC event generators. The input of a particle physics MC event generator is a set of initial state particles and their energy and momentum. Then, the event generator can compute the probability distributions of all the allowed final states of the interaction, based on some theoretical framework (for example the SM, or an extension of it). Using some method of random number generation, the event generator picks one of these final states for each event. The output of the event generator is a list of the physics objects in the final state such as charged leptons and jets, with precisely defined four-momenta, charges and other properties.

In this thesis we use `MadGraph5_aMC@NLO` [54] version 3.3.1 for MC simulations of proton-proton collider events. Unless stated otherwise, all processes are simulated at next-to leading order (NLO) in QCD. We also need to simulate parton showering of colour charged objects in the final state. For this we use `PYTHIA8` [55] version 8.308. We run `MadGraph` through ATLAS software, using `Athena`[1] release 22.

When performing MC event generation at higher orders of perturbation, not all events will represent an equal share of the total cross section of the simulated process. To account for this, each event is given a weight, and in particular, some events are given negative weights [56, 57]. These negatively weighted events represent corrections from

[1] `Athena` is an ATLAS software package, and can be found at `https://gitlab.cern.ch/atlas/athena`.

the higher order calculations. The physical integrated luminosity of the simulation is then[2]

$$L = \sum_{i=1}^{n} w_i \frac{1}{\sigma}, \tag{5.1}$$

where n is the number of simulated events, w_i is the weight given to event number i and σ is the cross section of the simulated process. This replaces the naïve expression

$$L = \frac{n}{\sigma}, \tag{5.2}$$

which is only correct if all events have the same weight. The presence of negative weights can reduce the statistics of a MC simulation, and this is one of the main challenges with MC event generation at beyond leading order.

5.2 Signal models and processes

In this section we describe the specific heavy neutrino models we will study, and the processes within the frameworks of these models that we will focus on. For all the models we look at, the signal processes will be ones that can produce a heavy neutrino at a proton-proton collider.

5.2.1 Majorana and pseudo-Dirac model

The first type of signal model we will study is a simple extension of the Standard Model by three heavy neutrinos of either Majorana or pseudo-Dirac type, one of which having a mass in the 100 GeV to TeV range. We call these the Majorana model and pseudo-Dirac model respectively. These models are examples of the theory discussed in section 3.1.2, and can be thought of as minimal realisations of the type 1 seesaw and the inverse seesaw mechanisms. For the simulations of processes within these models with `MadGraph`, we use the two sets of publicly available model files developed by Alva et al, Degrande et al. and Pascoli et al [58, 59, 41]. These `MadGraph` models are based on the exact interaction Lagrangian that we presented in equation (3.33), the only difference between the two sets of model files being that the heavy neutrinos are either pseudo-Dirac or Majorana type, meaning that their interactions either conserve lepton number, or have the possibility of changing it by two units.

To find the possible production processes for these heavy Majorana or pseudo-Dirac neutrinos at the LHC, we need to look at the interaction Lagrangian of equation (3.33). It is easy to see that these interaction terms give three possible first order production processes from initial states consisting of quarks or gluons. These are all so-called s-channel processes, where one quark from each proton interact to give an intermediate boson, which then decays to a heavy neutrino. The three different possibilities are then characterised by which boson makes up the intermediate state. When the intermediate boson is a Z or Higgs, their decay will result in a final state SM neutrino, along with the heavy neutrino. The Feynman diagrams of these two processes are shown in figure 5.1. The incoming quarks are both labeled q to show that they are a quark-antiquark pair of the same flavour. The heavy neutrino labelled N can be of either Majorana or Dirac type.

[2] Assuming that the event weights are scaled such that the weight of an event that occurs with exactly the expected probability is $w_i = 1$.

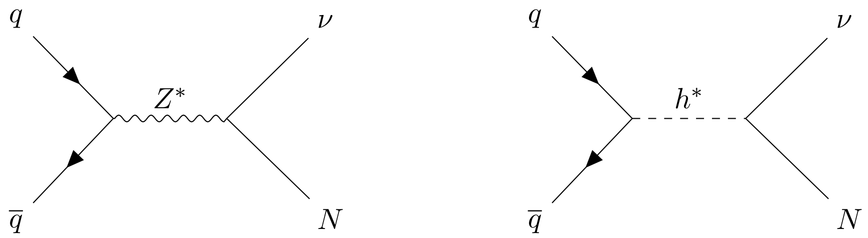

(a) Production through virtual Z^* (b) Production through virtual Higgs

Figure 5.1: Two of the first order production processes for heavy Majorana or Dirac neutrinos at proton-proton colliders, through a virtual Z^* boson (a) or a virtual Higgs boson h^* (b). In both of these processes, the heavy neutrino N is produced along with a Standard Model neutrino ν.

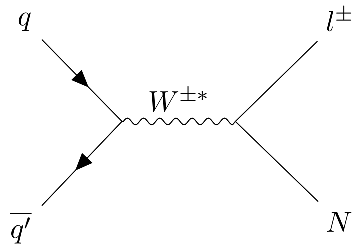

Figure 5.2: One of the three first order production processes for heavy Majorana or Dirac neutrinos at proton-proton colliders, in which the heavy neutrino is produced along with a charged lepton.

The Z and Higgs bosons are marked with asterisks to indicate that they must be off-shell when the mass of the heavy neutrino is larger than the Z mass or Higgs mass respectively. These two processes are not very interesting from an experimental perspective, since the SM neutrino can not be measured directly, so much of the information about the event is lost. Additionally, the small coupling of the constituent quarks of the proton to the Higgs boson means that the cross section of the second process is very small (although the coupling between the Higgs and N can be relatively large for large neutrino masses). The final first order production process is through a virtual W^*-boson, and results in a charged lepton along with the heavy neutrino. The Feynman diagram is shown in figure 5.2. Here the initial state quarks are labelled q and q' to show that they are of different flavour (one is down type and the other is up type). This will be the signal process that we focus on for both the Majorana and pseudo-Dirac models, since the charged lepton is relatively easy to detect, at least if it is an electron or muon, and it has a cross section comparable to the first process in figure 5.1.

Decays

For most sets of neutrino masses and mixing parameters we will look at, the heavy neutrinos will have very short lifetimes. This means that there is more to the signal process than simply the production of the neutrino, since the neutrino will also decay before reaching the detector, and its decay products will be measured instead.

The possible two-body decays of the neutrinos in the Majorana and pseudo-Dirac

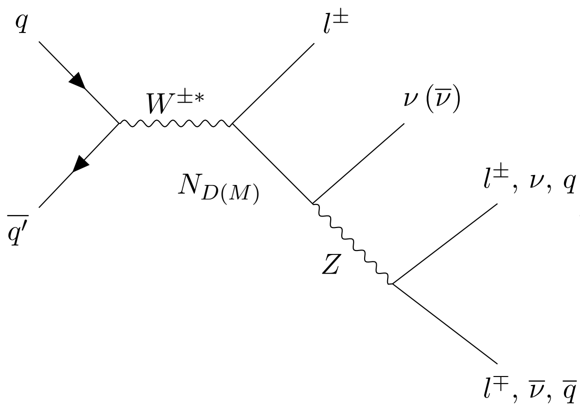

(a) Decay via Z

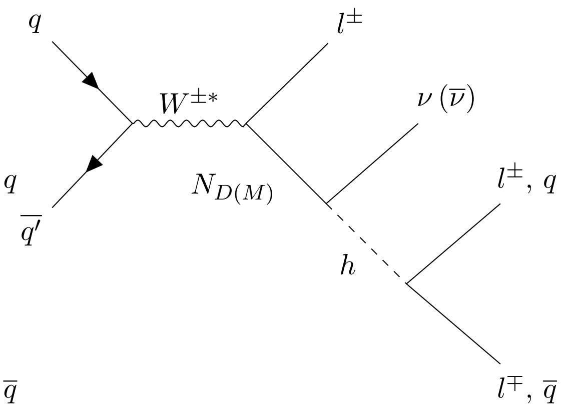

(b) Decay via Higgs

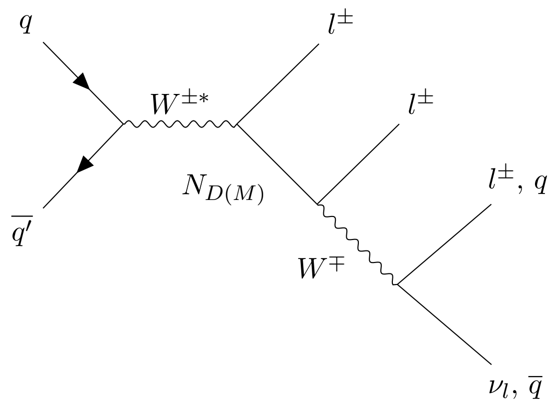

(c) Decay via W

Figure 5.3: All the possible final states resulting from the first order production of a Majorana or pseudo-Dirac heavy neutrino through a virtual W^* boson at a proton-proton collider, and the subsequent first order decay of the neutrino.

models are determined by the same interaction terms in the Lagrangian as the first order productions. So there are again three possibilities, assuming that $m_N > m_h$, where m_N is the mass of the heavy neutrino and m_h is the mass of the Higgs boson. These three are the decay to a W, Z or Higgs boson, along with a lepton. Each of these decays will lead to very different kinematics in the final states, and therefore require different strategies when searching for these processes in an experiment. In fact, since both the W, Z and Higgs bosons are also highly unstable, these will decay to different sets of particles, giving even more possible final states. Feynman diagrams for all these final states (only including two-body decays of the neutrino and bosons) are shown in figure 5.3. All of these processes are well suited to search for the heavy neutrino, except for the case in figure 5.3a when the Z boson decays to two neutrinos. This process would give a signature of only a single charged lepton and missing energy, and would therefore give almost no information about the presence of the heavy neutrino. For the other decay modes of the Z, the signature would be more interesting. The neutrino from the decay of N would result in a relatively large E_T^{miss} in such an event. In the case when the Z decays to two charged leptons, the final state would contain a total of three charged leptons, two of which would have an invariant mass close to the Z mass. The decay of

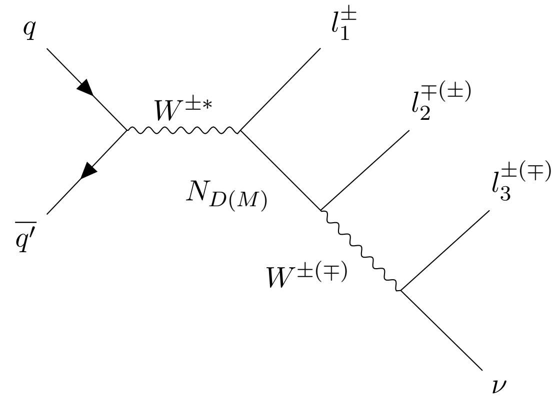

Figure 5.4: The lowest order Feynman diagram of the signal process we will study for the Majorana and pseudo-Dirac models. For the pseudo-Dirac model, the allowed particle charge combinations are shown with the signs that are not in parentheses. For the Majorana model, both the signs inside and outside of the parentheses are possible, but the choice must be consistent in a single event.

the Z to quarks would result in one charged lepton and two jets with invariant mass close to the Z mass. The neutrino decay through the Higgs gives very similar kinematics, only replacing the mass of the Z with the mass of the Higgs. However, the three lepton final state will not be observable in this case, due to the tiny coupling of the Higgs to electrons and muons. The most prominent first order decays of the Higgs would result in pair of b-jets with invariant mass close to m_h or a pair of tau leptons, which subsequently decay. Additionally, the Higgs has relatively large branching ratios for three-body decays and decays through loops, for example $h \to W^*W$ and $h \to gg$, which would give even more possible final states.

Although many of these processes present promising opportunities to search for heavy neutrinos, we can not cover them all in this thesis. We therefore choose to study only the decay of the heavy neutrino through a W boson. As we see from figure 5.3c, this decay mode gives two possible final states. The first possibility consists of three charged leptons and a neutrino, while in the other case it is two charged leptons and two quarks. Both of these decay modes have their drawbacks when it comes to experimental detection. For the leptonic decay of the W boson, the main problem is again that the neutrino can not be detected, and therefore the total energy and momentum of the process is unknown. This makes it more difficult to for example reconstruct the mass of the heavy neutrino from signal events. For the hadronic decay, the total energy of the event can be determined, but the quarks will create jets of hadrons that are not measured as precisely as charged leptons. Additionally, the Standard Model background is much larger for events with two charged leptons and two jets in the final state than for three charged leptons, which makes the signal more difficult to discover. We will therefore narrow our scope even further, and only look at the leptonic decays of the W for the Majorana and pseudo-Dirac models. It would be interesting to perform a similar analysis for the hadronic decay as well.

With the leptonic decay of the W, the full signal processes for the Majorana and pseudo-Dirac models become as shown in figure 5.4. A key difference between the Majorana and pseudo-Dirac models is the allowed sign combinations of the charges of

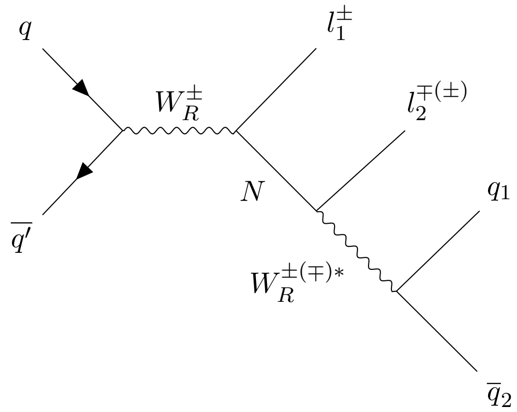

Figure 5.5: The Feynman diagram for the Keung-Senjanović process of production of a right-handed W_R boson and a heavy Majorana neutrino N at a proton-proton collider.

leptons l_1 and l_2. In the pseudo-Dirac case, they must have opposite signs, while in the Majorana case they can have either opposite or the same sign, as indicated by the parentheses in figure 5.4.

5.2.2 LRSM model

We will also study a simple realisation of the Left-Right Symmetric Model (LRSM), based on the minimal LRSM described in section 3.5. We will focus on models where the mass of the W_R boson is $M_{W_R} \gtrsim 5$ TeV, and the natural case of a heavy Majorana neutrino that has a mass m_N that is of the same order as M_{W_R}. We will also consider only models where $m_N < M_{W_R}$. This gives the possibility for the so-called Keung-Senjanović process [60] for producing an on-shell W_R boson at the LHC, with a subsequent decay through the heavy Majorana neutrino, which is shown in figure 5.5. This process has been extensively studied, for example in a recent study by the ATLAS collaboration that set the most stringent lower limits on the W_R and N masses to date [61]. In this thesis, we will put more emphasis on studying the kinematics of the Keung-Senjanović process for different masses using MC simulations. We will then use what we learn to attempt to improve the search method for this process, and potentially show how properties of the W_R and N particles can be deduced if this signal were to be discovered.

For the MC simulations with `MadGraph`, we will use the publicly available effective LRSM model developed by Mattelaer et al [49]. As is mentioned in their paper, this effective model does not include the new scalar fields of the LRSM. However, since these fields can mediate flavour changing neutral currents, their masses must be larger than what is accessible at the LHC [62]. The exclusion of these fields from the model will therefore have little effect at LHC energies. In the `MadGraph` model, g_R is set equal to the SM weak coupling constant, but this can easily be modified by simply scaling the cross sections. Finally, we use the simplifications $V'_{\text{CKM}} = \mathbb{I}$, and assume the mixing between the heavy and light neutrinos to be $X = 0$ and that the heavy neutrino mass eigenstates are equal to their right-handed weak eigenstates; $Y = \mathbb{I}$. These assumptions are motivated by the natural scalings given in section 3.5. This means that in our simulations, the W_R does not couple to pairs of SM leptons, while it does couple to SM quarks with equal strength to each generation. And in the coupling of the W_R

to one of the heavy neutrinos N_m no cross generational couplings are allowed. For the Keung-Senjanović process this means that unless some of the neutrino masses are degenerate, the leptons l_1 and l_2 must be of the same flavour in this model.

5.3 Signal cross sections

The most important variable determining whether a signal can be discovered at a collider or not is its cross section. The cross section of our signal processes depend mainly on the masses of the heavy neutrinos, their mixing with active neutrinos, and of course on the energy of the collider. Since protons are composite particles, the cross sections also depend on their parton distribution functions (PDFs). A PDF gives the probability distribution for the momentum fraction x of the proton carried by one of the partons. This additional complexity due to the composite nature of protons makes it difficult to compute cross sections for proton-proton collider events analytically. We therefore use `MadGraph` to compute the cross sections numerically. This also allows us to include next-to leading order (NLO) QCD contributions to the cross section.

5.3.1 Majorana and pseudo-Dirac model

We first look at the signal process for the Majorana and pseudo-Dirac models through the process shown at leading order (LO)[3] in figure 5.4. It turns out that in the simple models we study, the cross section of this process is independent of the Dirac/Majorana nature of the neutrino, as long as the mass and mixing parameters are the same. The reason for this is that the interaction terms of the Lagrangian look the same in both models. In the Dirac case there are two different heavy neutrinos that can be created (neutrino or antineutrino), but the neutrino can only decay to a charged *lepton*, while the antineutrino can only decay to a charged antilepton, as indicated in figure 5.4 by the signs that are not in parenthesis. In the Majorana case, on the other hand, there is only one neutrino, but it couples to both kinds of charged lepton (and in the model we look at the two couplings are equal), so the total number of possible interactions is the same, giving the same cross section.

We now assume for simplicity that only one of the three heavy neutrinos has a mass that makes it suitable for detection, and that the other two have very different masses, such that any interference between the interactions can be neglected. We also assume that this heavy neutrino mixes equally with the three light neutrinos, meaning

$$|V_{e1}|^2 = |V_{\mu 1}|^2 = |V_{\tau 1}|^2, \tag{5.3}$$

where we have chosen the detectable neutrino to be mass eigenstate number $N = 1$, which we can do without loss of generality. It is then easy to see from the Lagrangian in equation (3.33) that the cross section of the production process in figure 5.2 will simply be proportional to $|V_{lk}|^2$.

For the full leptonic decay process shown in figure 5.4, one would expect that the cross section is proportional to $|V_{lk}|^4$, since the diagram contains two identical vertices coupling the heavy neutrino to the W boson. However, we need to take into account that the heavy neutrino is an unstable particle with some decay rate Γ_N. Because the

[3] This type of diagram is also called a ”tree level” diagram or a ”Born” diagram. Both of these terms mean that the diagram contains no loops.

heavy neutrino only interacts weakly, we can assume that $\Gamma_N \ll m_N$, and if m_N is much larger than the masses of the other particles in the process, all the conditions for using the so-called narrow width approximation (NWA) are satisfied [63]. The NWA corresponds to choosing the unstable intermediate particle (the heavy neutrino in our case) to be exactly on-shell. Then the total cross section for the production and decay of the heavy neutrino is given by the cross section for the production alone, multiplied by the branching ratio (BR) of the decay. The BR of the decay of a particle X with total decay rate Γ_X to some final state Y is defined as

$$BR(X \to Y) = \frac{\Gamma\,(X \to Y)}{\Gamma_X}. \tag{5.4}$$

As mentioned in section 5.2.1, we know all the possible decays of the heavy neutrino, and they all occur through its mixing with the active neutrinos. This means that the rate of each individual decay process must be proportional to $|V_{lk}|^2$. But then the total decay rate Γ_N, which is simply the sum of the rates of all possible decays, must also be proportional to $|V_{lk}|^2$. This means that the BRs for each of the possible decays of the heavy neutrino are independent of the mixing matrix elements. Therefore, as long as the NWA applies the leptonic decay cross section will also be proportional to $|V_{lk}|^2$, and it will be approximately equal to the cross section for the production of N through a W boson, multiplied by the factor

$$\frac{\sigma\,(p\,p \to N\,l \to l\,l\,l\,\nu)}{\sigma\,(p\,p \to N\,l)} \approx BR\,(N \to l\,l\,\nu) \approx BR\,(N \to W\,l) \times BR\,(W \to l\,\nu)\,, \tag{5.5}$$

where we used the NWA again for the W boson. If we look at the scaled cross section

$$\sigma' = \frac{\sigma}{|V_{lk}|^2}, \tag{5.6}$$

this will be independent of the size of the mixing parameters for both the production process and the full process with leptonic decay. We calculate this scaled cross section for a few different heavy neutrino masses and centre of mass energies using `MadGraph`. The calculations are done at NLO in QCD, meaning that in addition to the Feynman diagram shown in figure 5.2, there will be contributions from diagrams with one or two additional QCD vertices. Diagrams with one or two QCD vertices correspond to FKS-subtracted real emission and single loops with colour charged particles respectively, as explained in the paper on the release of `MadGraph5` [54]. We do not include NLO QED contributions, since these will be suppressed by at least a factor $\alpha_{\mathrm{EM}} \approx 1/137$. Examples of diagrams that are included in the computation of the cross sections are shown in figure 5.6. The resulting cross sections are shown in table 5.1. `MadGraph` also allows us to estimate the systematic uncertainty in the cross section as a result of scale variation and PDF variation. Scale variation refers to the change in the cross section due to the varying energy of each event, which affects the size of the coupling constants $\alpha_{\mathrm{S}}(\mu)$ and $\alpha_{\mathrm{EM}}(\mu)$. The PDF variation comes from the fact that we do not have a very good model of the QCD processes within protons, and therefore parton distribution functions have large uncertainties. In `MadGraph`, these systematic uncertainties are estimated using reweighting techniques, as is explained in the `MadGraph5` release paper [54]. We combined these two sources of uncertainty to get the errors listed in table 5.1, and we see that they are generally between 2% and 6% of the value of the cross section.

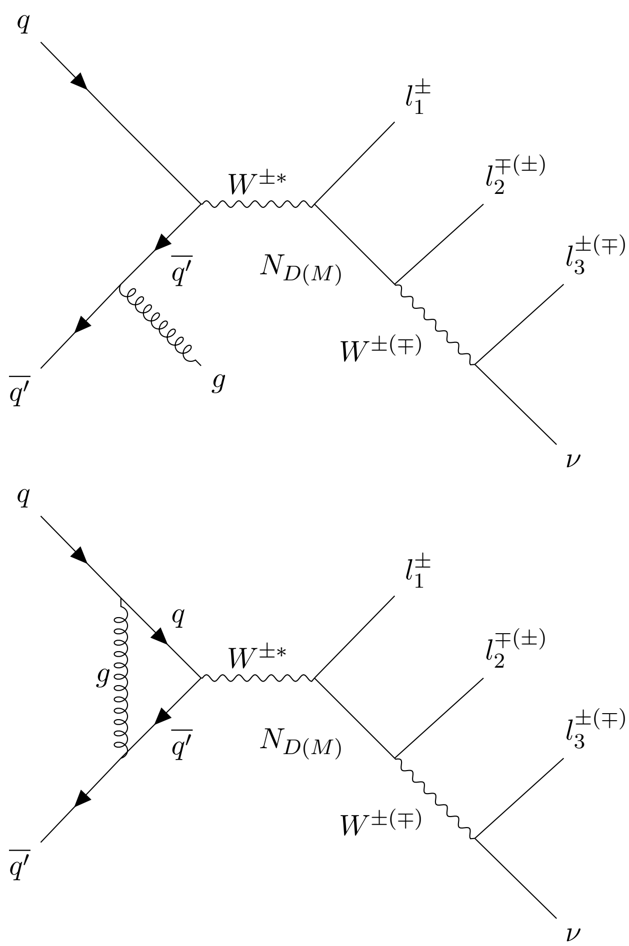

Figure 5.6: Two Feynman diagrams of the production of a Majorana or Dirac heavy neutrino with subsequent fully leptonic decay. Both contribute to the next-to leading order QCD correction, since they contain QCD vertices not found in the leading order diagram. In the top diagram, a gluon is radiated from one of the initial state quarks. In the bottom diagram, a gluon is exchanged between the two quarks, making a loop in the diagram and giving two QCD vertices.

Table 5.1: The scaled cross section σ', as defined in equation (5.6), in units of pb for the production and fully leptonic decay of a Majorana or pseudo-Dirac heavy neutrino for three different centre of mass energies $\sqrt{s}$ of the proton-proton collider and five different heavy neutrino masses m_N. The cross sections are computed to NLO in QCD using MadGraph, and it is assumed that the heavy neutrino mixes equally to the three SM neutrinos.

	$\sqrt{s}$ [TeV]		
	13.0	13.6	14.0
m_N [GeV]	σ' [pb]		
200	1.00 ± 0.05	1.075 ± 0.029	1.07 ± 0.04
400	$(7.23 \pm 0.27) \times 10^{-2}$	$(7.62 \pm 0.17) \times 10^{-2}$	$(8.16 \pm 0.22) \times 10^{-2}$
600	$(1.46 \pm 0.04) \times 10^{-2}$	$(1.55 \pm 0.04) \times 10^{-2}$	$(1.72 \pm 0.06) \times 10^{-2}$
900	$(2.36 \pm 0.11) \times 10^{-3}$	$(2.74 \pm 0.09) \times 10^{-3}$	$(2.86 \pm 0.18) \times 10^{-3}$
1200	$(6.38 \pm 0.34) \times 10^{-4}$	$(6.97 \pm 0.28) \times 10^{-4}$	$(7.7 \pm 0.4) \times 10^{-4}$

From the table we also see that as expected, the cross sections become smaller for larger heavy neutrino masses. The cross sections also become larger for larger centre of mass energies $\sqrt{s}$. The three values for $\sqrt{s}$ used here correspond to the centre of mass energies of the LHC during Run 2 (13 TeV), during the current Run 3 (13.6 TeV) and the future High Luminosity LHC (HL-LHC) runs (14.0 TeV). The increase in energy has a greater effect on the cross sections for the heavier neutrinos. For $m_N = 200$ GeV, we see that the cross section only increases by 7% from 13.0 TeV to 14.0 TeV, while for $m_N = 1200$ GeV, it increases by around 20%.

For comparison, we also computed the cross section for $m_N = 200$ GeV and $\sqrt{s} = 13.6$ TeV at LO, and got the result $\sigma' = 0.828 \pm 0.024$ pb. This represents a difference of 30% from the NLO result given in table 5.1. This can also be expressed in terms of the commonly used K-factor, which is defined to be the ratio between the NLO and LO cross sections of a given process. This particular process therefore has a K-factor of 1.30.

We can also use MadGraph to test the validity of the NWA for the heavy neutrinos, as it allows us to compute the decay rate of the neutrinos[4]. The result is with the same mixing parameters defined in (5.3) and $m_N = 200$ GeV

$$\Gamma_N^{200} = 24.7|V_{lk}|^2\text{GeV}, \tag{5.7}$$

while $m_N = 1200$ GeV gives

$$\Gamma_N^{1200} = 6.81 \times 10^3|V_{lk}|^2\text{GeV}. \tag{5.8}$$

So for the neutrino mass of 200 GeV, the condition $\Gamma_N \ll m_N$ for use of the NWA is satisfied for more or less all values of the mixing parameters (the unitarity condition of V from equation (3.34) gives $|V_{lk}|^2 \leq 1/3$). For the mass of 1200 GeV, the mixing

[4]It is of course relatively straight forward to compute the decay rate analytically as well, at least if only the two-body decays are taken into account.

parameters must be smaller than around 0.1 to have Γ_N/m_N of the order of 5% or less. But mixing parameters greater than 0.1 would be in contradiction with the limits given in equation (3.51), so we can conclude that the NWA is a good approximation in most reasonable realisations of the model.

5.3.2 LRSM model

The cross section for the Keung-Senjanović process shown in figure 5.5 depends mainly on the mass of the W_R and the Majorana neutrino N. For a given neutrino mass m_N, the cross section is independent of which neutrino mass eigenstate N_m, $m = 1, \ldots, 3$ is considered, since each mass eigenstate couples equally to a flavour of the leptons l_1 and l_2, and the mass differences between the lepton flavours is irrelevant at these energies. We therefore only need to compute the cross section for one mass eigenstate, and we choose N_1, which is the one that couples to the electron. We again use `MadGraph` to compute the cross section at NLO in QCD for the centre of mass energies of Run 2, Run 3 and the HL-LHC. As mentioned before, the model assumes equality of the right- and left-handed weak couplings; $g_R = g$. However, if we want to be general, we can again use the NWA to argue that the cross section of the process will be proportional to g_R^2. This case is slightly different, since we now want to use the NWA for the W_R, but the argument is the same as in the Majorana and pseudo-Dirac models. The W_R only decays through couplings that are proportional to g_R, and therefore its branching ratios are independent of g_R. If the NWA is valid, the only dependence on g_R therefore comes from the coupling of the W_R to the initial state quarks, which gives a g_R^2-dependence in the cross section of the entire Keung-Senjanović process. To see whether the NWA is valid, we again use `MadGraph` to compute the decay width of the W_R for a few mass values. For $M_{W_R} = 3000$, 5000 and 6000 GeV, the decay widths are approximately 110, 180 and 220 GeV respectively for $g_R = g$. These are all less than 5% of the respective masses, and so the NWA is a good approximation, and we will assume $\sigma \propto g_R^2$. We define the relative strength of the right-handed weak interaction using the parameter κ_R, defined such that

$$g_R = \kappa_R g. \tag{5.9}$$

Then we define the scaled cross section σ' as

$$\sigma' = \frac{\sigma}{\kappa_R^2}, \tag{5.10}$$

such that σ' is approximately independent of the size of g_R. The results for this scaled cross section computed with `MadGraph` for six pairs of mass parameters are given in table 5.2. We see that the cross sections are quite sensitive to both the W_R and neutrino masses, and decrease for larger masses, as expected. We also see that the increase of $\sqrt{s}$ affects the models with larger masses more. For $m_N = 5$ TeV and $M_{W_R} = 6$ TeV, the cross section is more than doubled from $\sqrt{s} = 13$ TeV to $\sqrt{s} = 14$ TeV, while for $m_N = 1.5$ TeV and $M_{W_R} = 3$ TeV, it only increases by around 40%. Rather alarmingly, we notice that the uncertainty in the cross sections becomes very large for the larger W_R masses. For $M_{W_R} = 3$ TeV, the uncertainties are of the order of $5 - 10\%$, as before. For $M_{W_R} = 5$ TeV, however, the uncertainties increase to around 30%, and for $M_{W_R} = 6$ TeV, they are over 50% in some cases. The cause of this large uncertainty is mainly the PDF variations, which become very large for these models. The reason for this is probably that when the mass of the W_R comes close to half of the centre of mass energy

Table 5.2: The scaled cross section σ', as defined in equation (5.10), in units of fb of the Keung-Senjanović process for six sets of W_R and heavy Majorana neutrino masses and three different centre of mass energies $\sqrt{s}$ of the proton-proton collider. The cross sections are computed to NLO in QCD using `MadGraph`, and it is assumed that the heavy neutrino only couples to one flavour of SM charged lepton and that the right-handed weak coupling is equal to the SM weak coupling.

		$\sqrt{s}$ [TeV]		
		13.0	13.6	14.0
m_N [TeV]	M_{W_R} [TeV]	σ' [fb]		
1.5	3	10.0 ± 0.8	12.2 ± 0.9	13.9 ± 0.8
2	3	5.7 ± 0.5	7.1 ± 0.6	7.9 ± 0.5
4	5	$(3.2 \pm 1.2) \times 10^{-2}$	$(5.1 \pm 1.5) \times 10^{-2}$	$(6.9 \pm 1.9) \times 10^{-2}$
4.5	5	$(9.3 \pm 3.5) \times 10^{-3}$	$(1.5 \pm 0.4) \times 10^{-2}$	$(2.0 \pm 0.6) \times 10^{-2}$
4	6	$(5 \pm 3) \times 10^{-3}$	$(9 \pm 4) \times 10^{-3}$	$(1.5 \pm 0.7) \times 10^{-2}$
5	6	$(1.7 \pm 1.2) \times 10^{-3}$	$(2.9 \pm 1.5) \times 10^{-3}$	$(4.6 \pm 1.9) \times 10^{-3}$

of the collider, the momentum fractions x of the two quarks must be very large. The probability for a parton to have a large x is very small, and therefore a small change in this probability (from a slightly different derivation of the PDF) will yield a very different cross section. We will need to keep this large uncertainty in mind for the rest of our analysis.

5.4 Background processes

In this section we look at some of the SM processes that give similar final states as the signal processes we are studying. The amount of SM background in an experiment is one of the limiting factors for discovery of a signal, and it is therefore important to know how much background there is, and some of its characteristics.

5.4.1 Three lepton final state

For the signal process of the Majorana and pseudo-Dirac models shown in figure 5.4, the background will consist of processes that can result in final states with three charged leptons and missing energy. There are relatively few SM processes giving three prompt leptons, and the largest contribution by far comes from a subset of the so-called ”diboson” processes. The diboson processes are ones where either a pair of W bosons, a pair of Z bosons or one W and one Z boson are produced. The final state of three prompt leptons and missing energy occurs in the case when one W and one Z both decay leptonically. Two such processes are shown in figure 5.7. The diboson three lepton background is characterised by one of the charged lepton pairs having an invariant mass close to the Z boson mass. This background can therefore be significantly reduced by disregarding events where any pair of charged leptons in the final state has an invariant mass that is

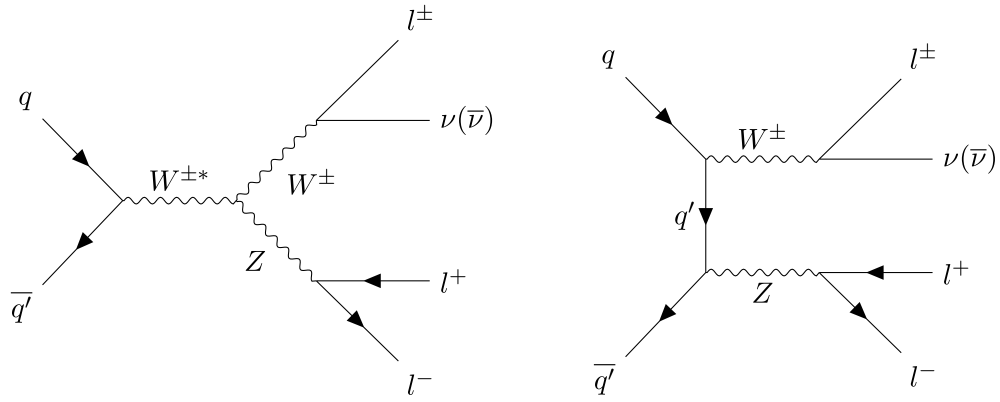

(a) s-channel $W\,Z$ (b) t-channel $W\,Z$

Figure 5.7: The two diboson processes that are the dominant SM background of three prompt leptons and missing energy.

within some small distance of the Z mass.

Another possible source of three prompt leptons is the similar $Z\,Z$ diboson process where both Z bosons decay to two leptons, but one of the four leptons escapes detection. This leads to the same signature of three leptons and missing energy, but can also be reduced by removing events where two leptons have invariant mass close to the Z mass. The remaining background for the three lepton final state mainly consists of processes with at least one non-prompt lepton. Examples of such backgrounds are the two classes of processes called "Z+jets" and "ttbar". In figure 5.8, we show examples of Z+jets processes that give two prompt leptons and a jet after hadronisation of the final state quark or gluon. The jet gives some small probability for a third lepton to be reconstructed in the event, and since these processes have relatively large cross sections they turn out to be a significant background for the three lepton final state. These processes are also characterised by the invariant mass of the prompt lepton pair being close to the Z mass, and can therefore be reduced by the same method as the diboson background. In addition to this, the Z+jets events do not contain real missing energy, and can therefore be reduced by requiring some minimum E_T^{miss}. In figure 5.9, we show one of several ttbar background processes that result in at least two charged leptons in the final state. In addition to the two charged leptons there are two b-quarks in the final state, which result in so-called b-jets after hadronisation. These b-jets can give additional leptons, either from leptonic decays of hadrons or from misidentifications. This background can not be reduced using invariant masses, since none of the charged lepton pairs come from the direct decay of a single particle. Instead, the ttbar background can be reduced by disregarding events with b-jets, or by increasing the requirements for lepton identification, for example putting strict requirements on the isolation of leptons. This reduces the probability of a lepton from one of the jets to pass the identification criteria.

5.4.2 2l+2j final state

After hadronisation, the Keung-Senjanović process of figure 5.5 will result in a final state of two charged leptons and two jets. This final state has significantly more SM background than the three lepton and missing energy final state. Some of the backgrounds for these

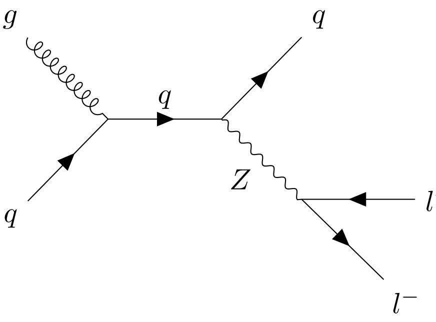

(a) s-channel Z+jets

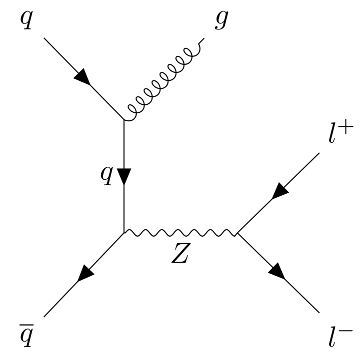

(b) t-channel Z+jets, gluon jet

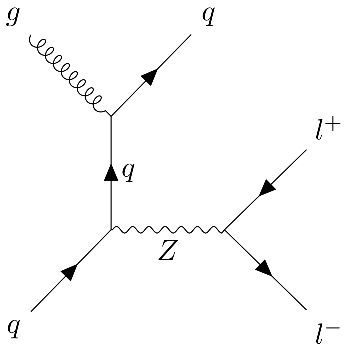

(c) t-channel Z+jets, quark jet

Figure 5.8: Three examples of processes making up the Z+jets background, which give two prompt leptons and a jet after hadronisation of the final state quark or gluon. The jet can potentially give a third non-prompt lepton.

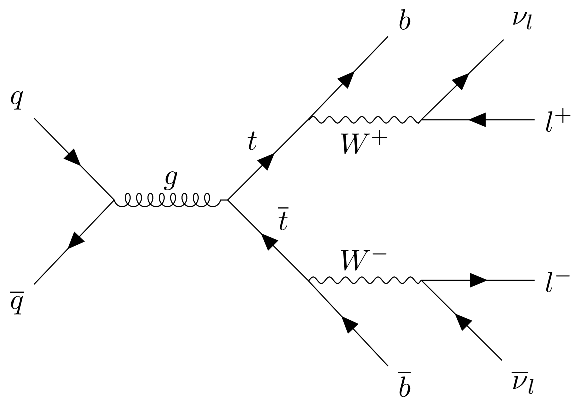

Figure 5.9: One of the processes making up the ttbar background, which gives two leptons and two b-jets after hadronisation. The b-jets can potentially result in a third lepton being reconstructed.

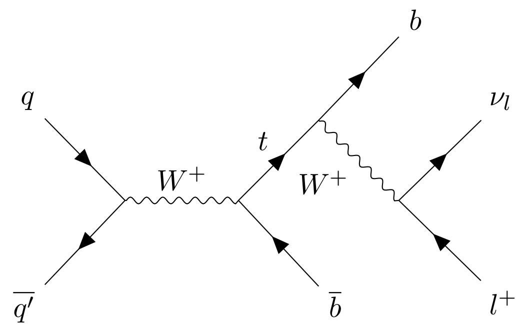

Figure 5.10: One of the processes making up the single top background, which gives at least one lepton and two b-jets after hadronisation. The b-jets can potentially result in a second lepton being reconstructed.

two final states are similar. For example, the 2l+2j final state also has background from the diboson processes, including the ones shown in figure 5.7, but with the W decaying to two quarks instead of a lepton and a neutrino. These two quarks will produce two jets after hadronisation, so the final state has the form 2l+2j. In addition to this, the production of two Z bosons, with one decaying hadronically and the other leptonically, will also give background for the 2l+2j final state. The diboson background is therefore larger for the 2l+2j final state, both due to the additional contribution from ZZ production and since the branching ratios of the W and Z bosons to quarks are significantly larger than their branching ratios to leptons. However, the diboson background can also in this case be greatly reduced by removing events where the lepton pair has an invariant mass close to the Z mass. The background from Z+jets and ttbar processes are even more important for the 2l+2j final state, since no additional leptons to the two shown in figure 5.8 and 5.9 are needed. It seems like the Z+jets processes of figure 5.8 can only give a single jet, but another jet can easily appear from higher order QCD corrections, for example. The Z+jets background can again be suppressed using the Z peak in the invariant mass, and the ttbar background by removing events with b-jets, if the signal process is not expected to give many b-jets.

The 2l+2j final state also has additional background types that are less relevant for the three lepton final state. These include the so-called "single top" and "W+jets" backgrounds. The W+jets background consists of processes like the ones shown in figure 5.8, but with the Z boson replaced by a W boson. If the W decays leptonically, there will be one prompt lepton in these final states, such that one additional non-prompt lepton and an additional jet will lead to a 2l+2j final state. In figure 5.10, we show one of the processes contributing to the single top background. We see that at least one lepton can come from the decay of the top quark, and if an additional non-prompt lepton is reconstructed, the final state can become 2l+2j. For the W+jets and single top backgrounds, the main method for reducing the background is to use strict lepton identification requirements, since neither of them give two prompt leptons.

Chapter 6

Truth analysis

In this chapter we will study the kinematics of the signal processes we described in the previous chapter using our Monte Carlo simulations. In the first section, we present the tools and methods used for this analysis. In the second section, we show the most relevant kinematic distributions for the Majorana model and discuss the characteristics of the distributions and what can be learned from them. In the third section we show the main difference between the Majorana and pseudo-Dirac models. In the fourth section we use what we learned in the previous two sections to develop a method for identifying the vertex of origin of the final state charged leptons of the signal process. Finally, in the fifth section we show the most relevant kinematic distributions for the LRSM signal process and discuss these.

6.1 Event selection and smearing using SimpleAnalysis

When studying either real experimental data from the LHC or MC simulated data, we often want to choose a subset of the total number of events in the data to study, based on some criteria. The process of choosing a subset of events is called event selection, and the criteria used are called selection criteria, or "cuts". In this thesis, we use the program `SimpleAnalysis` [64] to perform cuts on MC simulated events of both signal and background processes, and to create kinematic distributions from the selected events. The specific cuts may differ in each application, and depend for example on the final state of the process we are interested in and the parameters of the signal model we are studying. There are however a few cuts that will be applied in all our analyses. Since our end goal is to say something about the discoverability of certain signals using the ATLAS detector at the LHC, we put some ATLAS-specific requirements on the particle objects from MC simulations. In order to be detectable by ATLAS at all, we require any charged lepton object to have $p_T > 5$ GeV and any jet objects to have $p_T > 20$ GeV in order to be identified as true charged leptons and jets. We also require all charged leptons and jets to have $|\eta| < 2.5$, corresponding to the range of the inner pixel and tracking detectors of ATLAS, meaning that objects within this range can be reconstructed with higher confidence. The objects that do not fulfill these requirements are simply ignored, but the events containing them are not necessarily discarded. In addition to this, `SimpleAnalysis` lets us perform so-called overlap removal on the objects from the MC simulations. This mimics the procedure with the same name that is performed in real experiments due to the fact that the same physical object can be reconstructed as multiple objects. Overlap removal attempts to reduce over-counting of objects by removing reconstructed objects that are too close to other objects. In our analyses, we will require that any electron candidate must have a separation of $\Delta R > 0.01$ from any muon candidates. Jets within $\Delta R < 0.2$ of electrons or muons are removed, while for electrons and muons we define a minimal ΔR separation requirement from jets that changes with their p_T. An electron or muon with transverse momentum p_T is removed if it has a distance from a jet that satisfies

$$\Delta R < \min\left[0.4, 0.04 + \frac{10\,\text{GeV}}{p_T}\right]. \tag{6.1}$$

As mentioned in section 5.1, the output of the MC simulations are lists of physics objects with precisely defined properties. This is what we call ”truth” information, and it is very useful for determining the characteristic signature of the different signal models which will be the focus of the remainder of this chapter. However, if we want to perform an analysis that replicates the analysis of real experimental data, we need to modify the truth information before studying event kinematics. In a real experiment, there are uncertainties in each direct measurement of an observable, and no particle can be identified with 100% certainty, but instead need to be reconstructed and identified as described in section 4.2.1. In analyses using MC simulated data, this behaviour is commonly replicated by entering the MC output into a full simulation of the real world detector. This gives an accurate estimate of measurement uncertainties and reconstruction efficiencies, but it is also a very computationally expensive procedure. Instead of a full detector simulation, we will therefore use a feature in `SimpleAnalysis` called ”smearing”, which is supposed to approximate a detector simulation of the ATLAS detector by using parametrisations for object reconstruction and identification efficiencies as well as object resolutions. We will only be applying smearing on electrons and muons. One of the effects of smearing is that the momenta and energies of the truth leptons are changed from the exact truth value to Gaussian distributions with means at the truth value and standard deviations determined by a function of p_T and η, in an attempt to approximate the measurement resolutions of the ATLAS detector. The probability of finding fake leptons in a jet or an electron being wrongly identified as a photon, what is collectively referred to as ”fake rates”, are also estimated. Smearing also allows us to implement the isolation working points for electrons and muons, as described in section 4.2.1. Finally, the probability that a given true electron or muon is actually identified is determined using a function of the particle’s p_T and η, and the identification working point that is chosen. This means that if we use the Tight identification working point for electrons for example, a relatively large fraction of electron truth objects will be removed from the events. If the Loose working point is used instead, a smaller fraction will be removed. This can be used to regulate the amount of background for final states with leptons.

In the following sections, we use `SimpleAnalysis` to study the truth information from MC simulations of some selected signal models, to learn about their characteristics.

6.2 Majorana model

In this section, we will study the kinematics of the signal process shown in figure 5.4, for the case when the heavy neutrino is of Majorana type. We again begin by performing

Monte Carlo simulations of the process using `MadGraph`. For this analysis, we set the centre of mass energy of the proton-proton collider to $\sqrt{s} = 13.6$ TeV, corresponding to run 3 of the LHC. As before, we assume that only the mass eigenstate N_1 contributes to the process (the other two mass eigenstates having either much larger or much smaller masses), but now we assume that the mixing matrix V between the active and sterile neutrinos is close to diagonal. We let

$$|V_{e1}|^2 = 10|V_{\mu 1}|^2 = 100|V_{\tau 1}|^2, \tag{6.2}$$

meaning that the couplings of the heavy neutrino to the muon and tau are suppressed by factors of 10 and 100 respectively compared to the coupling to the electron (after squaring the matrix element). The size of $|V_{e1}|^2$ then determines the total cross section, but that is irrelevant for this part of the analysis, since we do not care about the effective luminosity of the MC simulations. We perform the simulation for four different N_1 masses; $m_N = 200, 400, 800, 1200$ GeV, and in each of these simulations we generate 10 000 events. These four mass values were chosen because they allow for production at the LHC, and they give a wide range of kinematics, due to their different sizes relative to the W boson mass. As before, the simulations are done to NLO in QCD, meaning that processes like the ones shown in figure 5.6 are taken into account. The `MadGraph` code also includes a three-lepton filter, which ensures that all 10 000 events have three charged leptons in the final state. To make histograms for the kinematic variables, we use `SimpleAnalysis`. In `SimpleAnalysis` we also perform a few additional cuts, requiring a charged lepton to have $p_T > 5$ GeV and $-2.5 < \eta < 2.5$ to be labelled as a signal lepton. Then only events with at least three signal leptons, and where the three signal leptons with the highest p_T, what we call the three leading leptons, are exactly the leptons l_1, l_2 and l_3 in figure 5.4, in arbitrary order. These requirements reduce the number of events only slightly, and we are left with between 7 523 ($m_N = 200$ GeV) and 8 329 ($m_N = 1200$ GeV) events. In the following plots, the histograms are scaled such that the sum of all the bin entries is equal to the respective number of events remaining after the cuts.

In figure 6.1 we show the number of signal leptons after the cuts. We see that for all four heavy neutrino masses, the overwhelming majority of events have exactly three signal leptons. For all m_N, the number of events with more than three signal leptons is of the order of 100, meaning at most a few percent of the total number of events. The largest heavy neutrino mass m_N is the only case in which there are some events with six signal leptons, but there are less than 10 of these events.

In figure 6.2, we show the distributions of the transverse momenta of the three final state leptons for each of the four mass values. The leptons are ordered by their vertex of origin, meaning that $p_{T,1}$, $p_{T,2}$ and $p_{T,3}$ are the transverse momenta of the leptons l_1, l_2 and l_3 from figure 5.4 respectively. The knowledge of which lepton comes from which vertex is something that we would not have in a real experiment, so we can only produce these distributions because the events were generated on a computer and we have access to the truth information. From the figure, we see that the distributions have relatively similar shapes for all four mass values. But when the mass of the heavy neutrino increases, we see that the distributions for all three leptons become wider (note the different scales on the x-axes), meaning that there is a higher probability for the leptons to have large transverse momenta . Interestingly, the distribution of $p_{T,2}$ is shifted slightly more to the right compared to the other two distributions for larger m_N. Both these effects are expected, since in order to create a heavier neutrino, the

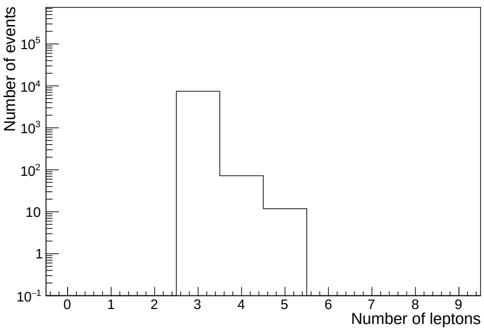

(a) $m_N = 200$ GeV

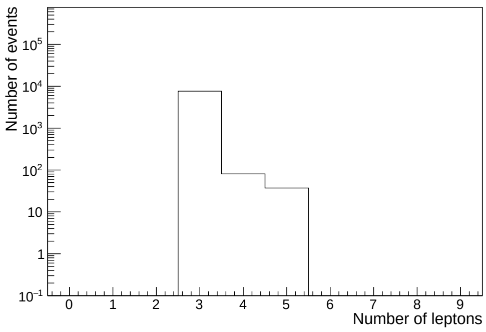

(b) $m_N = 400$ GeV

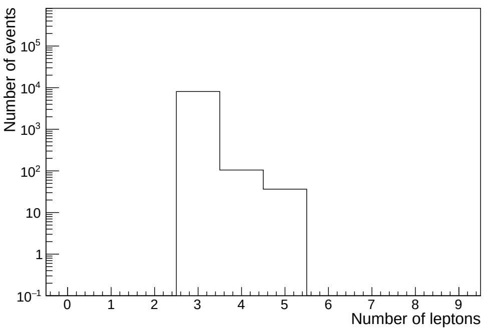

(c) $m_N = 800$ GeV

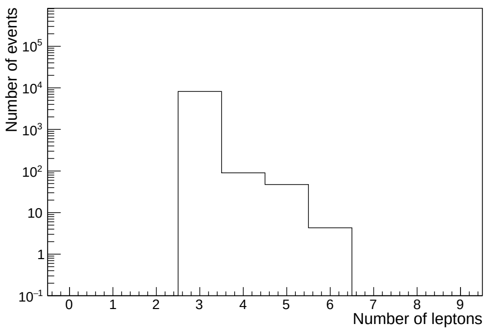

(d) $m_N = 1200$ GeV

Figure 6.1: The distributions of the number of signal leptons passing the cuts in each event for different heavy Majorana neutrino masses m_N.

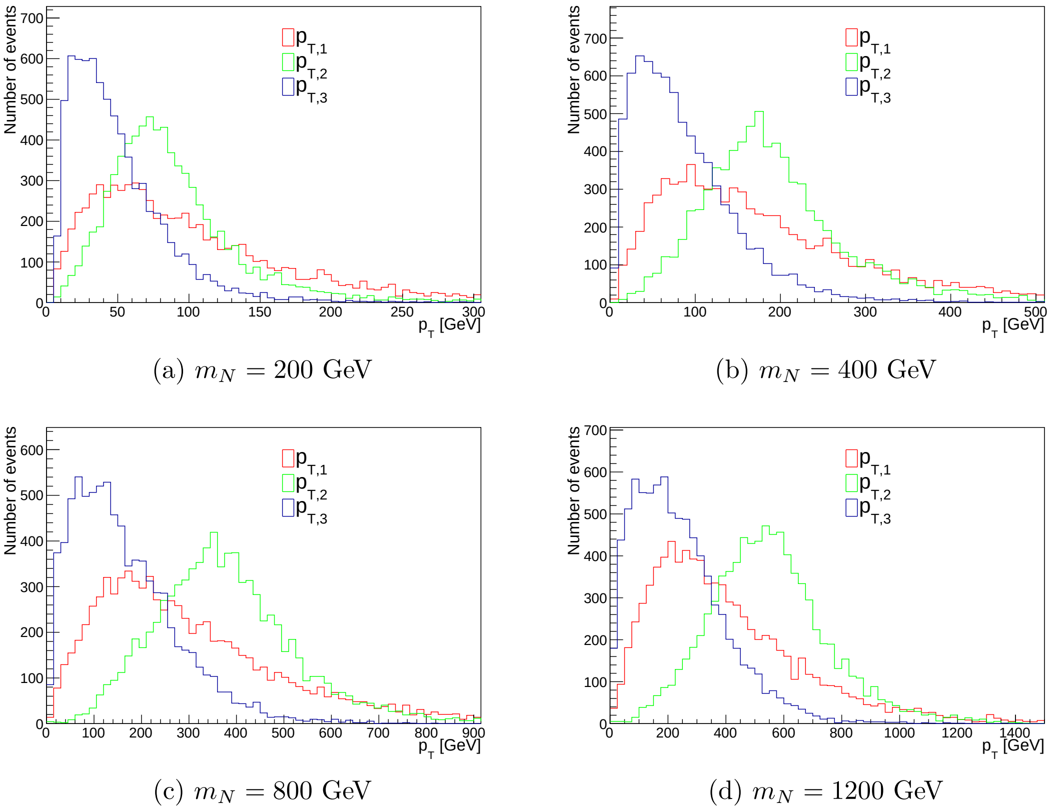

(a) $m_N = 200$ GeV

(b) $m_N = 400$ GeV

(c) $m_N = 800$ GeV

(d) $m_N = 1200$ GeV

Figure 6.2: The distributions of the transverse momenta of the three leading signal leptons in each event, ordered by their creation vertex in the Feynman diagram, for different heavy Majorana neutrino masses m_N.

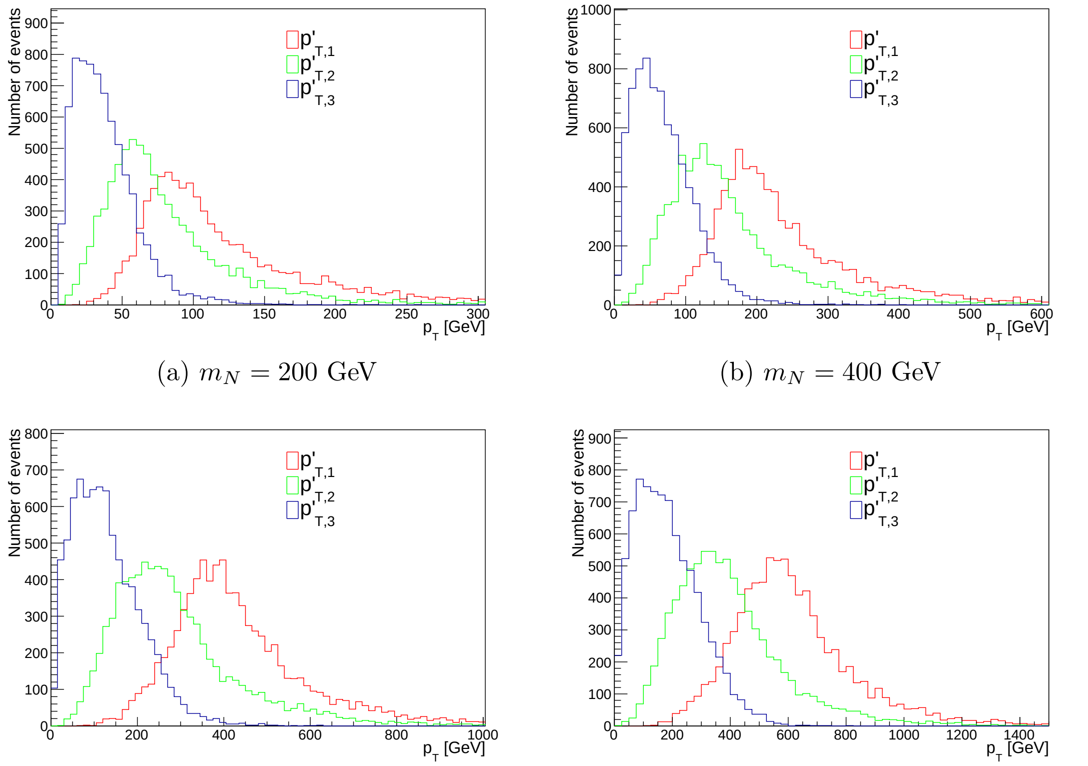

(a) $m_N = 200$ GeV

(b) $m_N = 400$ GeV

(c) $m_N = 800$ GeV

(d) $m_N = 1200$ GeV

Figure 6.3: The distributions of the transverse momenta of the three leading signal leptons in each event, ordered from highest to lowest p_T, for different heavy Majorana neutrino masses m_N.

process has to be very energetic (the initial state partons must have had large momentum fractions x), which gives all the particles in the event higher momenta. The large mass of the neutrino will give its two decay products particularly large momenta, since the mass of the charged leptons and the W boson are much smaller than for example 800 GeV. This explains why l_2 is the hardest lepton in most events for $m_N = 800$ GeV and $m_N = 1200$ GeV. Even when $m_N = 200$ GeV, the peak of the $p_{T,2}$ distribution is at a slightly larger transverse momentum value than the peak of the $p_{T,1}$ distribution, but $p_{T,1}$ has a relatively large tail at large values meaning that in a few events the first lepton can be very hard, and harder than l_2. For the larger masses, the tail of the $p_{T,1}$ and $p_{T,2}$ distributions have approximately the same size. For all mass values, l_3 is the lepton that has the lowest transverse momentum in most events.

In figure 6.3 we show similar distributions of the transverse momenta of the final state leptons, but in these figures the leptons are ordered by the size of their p_T in each event. These distributions are therefore possible to make with real data, since a detector can measure (with some error) the transverse momenta of charged leptons. We label these transverse momenta with primes to separate them from $p_{T,i}$, which are the transverse momenta of lepton l_i. We see that these distributions are quite similar to the distributions ordered by vertex of origin, with the replacements $p_{T,1} \to p'_{T,2}$,

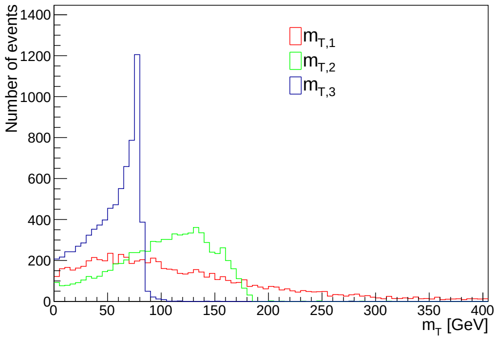

(a) $m_N = 200$ GeV

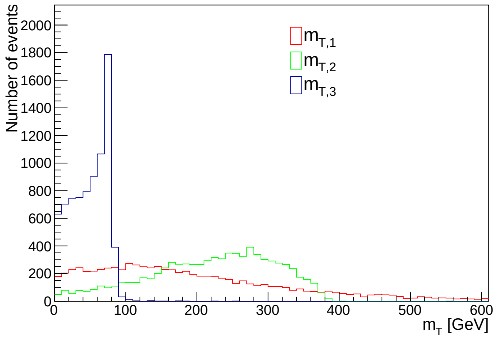

(b) $m_N = 400$ GeV

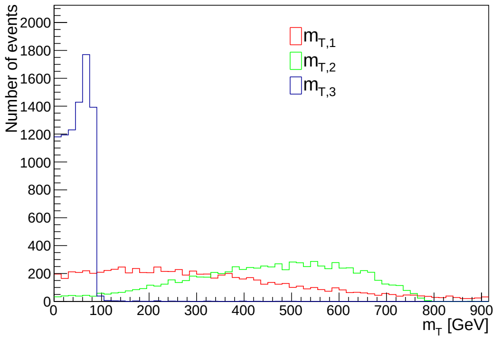

(c) $m_N = 800$ GeV

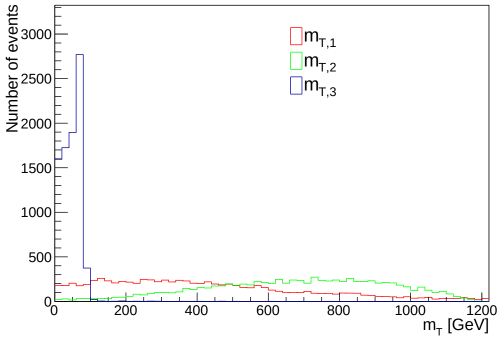

(d) $m_N = 1200$ GeV

Figure 6.4: The distributions of the transverse masses of the signal leptons l_i and the final state neutrino, which we call $m_{T,i}$, for different heavy Majorana neutrino masses m_N.

$p_{T,2} \to p'_{T,1}$ and $p_{T,3} \to p'_{T,3}$. This is not surprising, since the most common ordering of the transverse momenta has the largest p_T from vertex 2, followed by vertex 1 and vertex 3. We notice that the distributions of $p'_{T,i}$ are more clearly separated from each other than the distributions of $p_{T,i}$, which is also natural, since for example $p'_{T,1}$ in fact picks the largest p_T in every single event, so this will be shifted more to the right than the distribution of $p_{T,2}$. Similarly, $p'_{T,3}$ always picks the smallest p_T, and will therefore be shifted left compared to the distribution of $p_{T,3}$.

In figure 6.4 we show the transverse masses of each of the charged leptons l_i with the final state neutrino. The transverse mass of l_i and ν is named $m_{T,i}$. We see that the distributions of $m_{T,1}$ and $m_{T,2}$ are quite similar to the transverse momentum distributions of the respective charged leptons shown earlier. For the smaller neutrino masses $m_N = 200$ GeV and $m_N = 400$ GeV, the $m_{T,2}$ distribution has a peak at a higher value than $m_{T,1}$, but the latter has a relatively large tail at large values of m_T. For the two larger neutrino masses, both leptons get a wider distribution of m_T, but the $m_{T,2}$ distribution is significantly higher at large values. These similarities are expected, since the transverse mass is proportional to the transverse momentum of the charged lepton, as seen from equation (4.12). So when there is no particular correlation between the transverse energies of the charged lepton and neutrino and the angle between them, it is

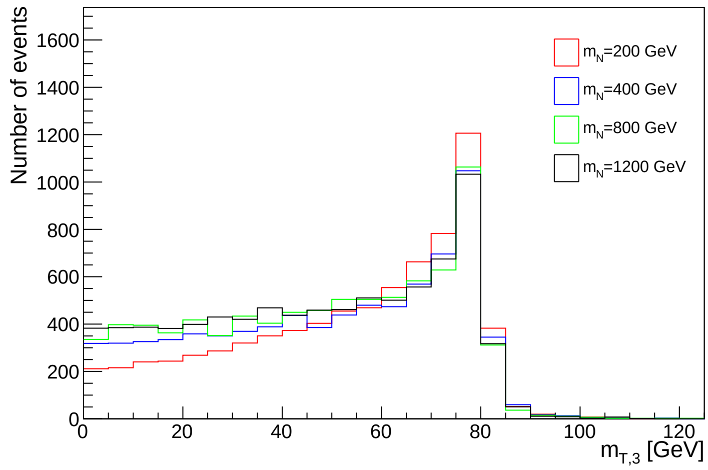

Figure 6.5: The distributions of the transverse mass of l_3 and the final state neutrino for each of the four heavy neutrino masses.

logical that the distribution of $m_{T,i}$ will have approximately the same shape as the $p_{T,i}$ distribution for the same charged lepton.

The main difference from the transverse momentum plots is seen in the distribution of $m_{T,3}$. We see that for all heavy neutrino masses, this distribution has a sharp peak at relatively low m_T, which looks different from the peak in the transverse momentum plots. The reason for this difference is that lepton 3 does in fact originate in a two-body decay, meaning that in all events where the mother W boson is on-shell, the transverse mass of the lepton is limited by the mass of the W boson. It therefore makes sense that the peaks in the $m_{T,3}$ distributions fall off at around the W mass of approximately 80 GeV. Interestingly, the peak in the $m_{T,3}$ distribution is slightly broader towards lower m_T and not as tall for the heavier neutrino masses. This is not easy to see from figure 6.4, so we make a new plot with only the $m_{T,3}$ distributions for each heavy neutrino mass. This plot is shown in figure 6.5. We see that with a bin width of $5GeV$, the tallest bin of the distribution for $m_N = 200$ GeV is around 20% taller than the tallest bin for $m_N = 1200$ GeV, while the distribution for $m_N = 1200$ GeV is around twice as high as $m_{T,3}$ goes to zero. That the distributions for larger m_N are higher towards low m_T values means that in these models, it is less likely that a large fraction of the momenta of l_3 and the final state neutrino are in the transverse direction.

Since the transverse mass distribution of l_3 is quite easily distinguishable from the other two leptons, this is a promising variable for identifying l_3 when we do not have the information about the true vertex of origin.

In figure 6.6, we show the distribution of the $\Delta\phi$ between each of the three charged leptons and the missing transverse energy in each event. We see that the shapes of the distributions change significantly with the heavy neutrino mass. We call the $\Delta\phi$ of l_i and the missing transverse momentum vector $\Delta\phi_i$. For $m_N = 200$ GeV, $\Delta\phi_1$ and $\Delta\phi_2$ have very similar distributions, which increase almost linearly towards large values of $\Delta\phi$. Meanwhile, $\Delta\phi_3$ has a more symmetrical distribution, but it is slightly skewed towards lower values. When m_N increases, we see some clear trends in how the distributions

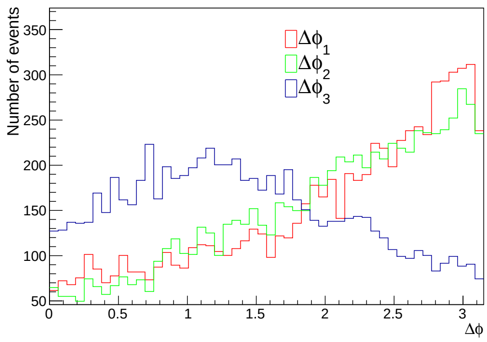

(a) $m_N = 200$ GeV

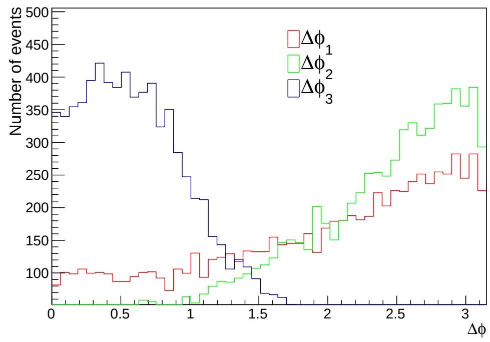

(b) $m_N = 400$ GeV

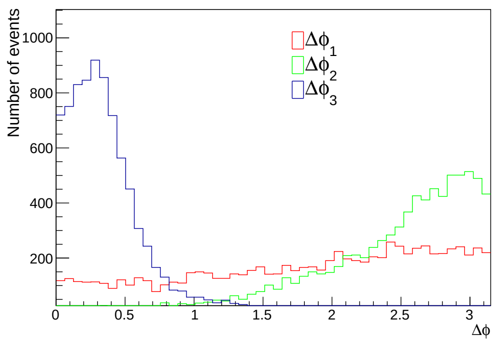

(c) $m_N = 800$ GeV

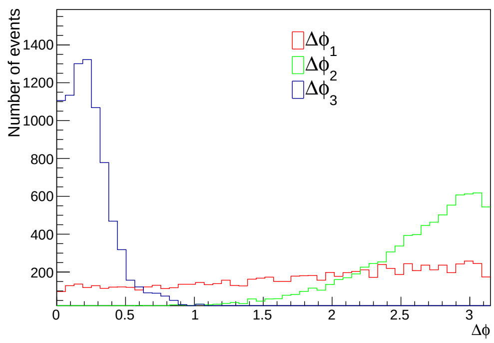

(d) $m_N = 1200$ GeV

Figure 6.6: The distributions of $\Delta\phi_i$, defined as the $\Delta\phi$ of the signal lepton l_i and the missing transverse momentum vector, for different heavy Majorana neutrino masses m_N.

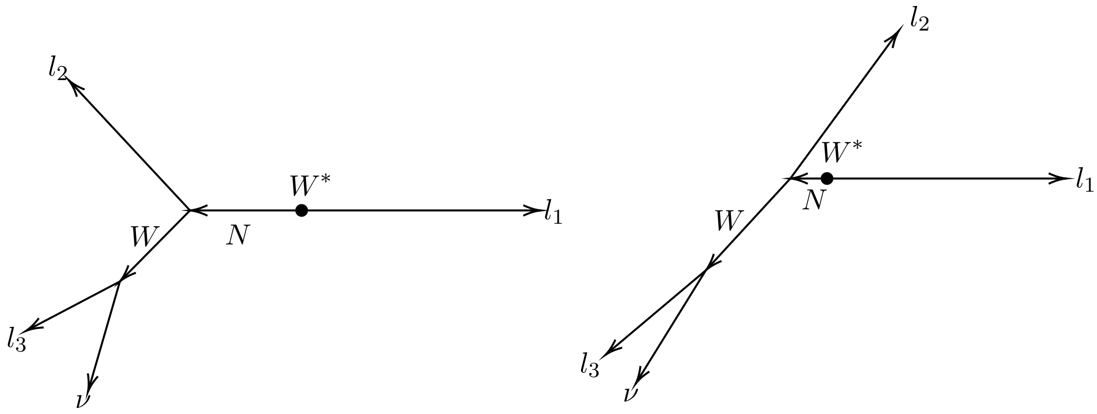

(a) Small m_N (b) Large m_N

Figure 6.7: Two examples of the event geometry in the production and leptonic decay of a heavy Majorana neutrino. Both diagrams are shown in the transverse plane of the collider reference frame. The direction and length of arrows represent the velocities of the particles projected onto the transverse plane. In the left diagram, the heavy neutrino N has a mass that is not much larger than the mass of the W, while in the diagram to the right, $m_N \gg m_W$.

change. The distribution of $\Delta\phi_3$ shifts more and more towards zero, and for $m_N = 1200$ GeV, almost all events have $\Delta\phi_3 < 0.5$. The distribution of $\Delta\phi_2$, on the other hand, shifts towards the maximum of $\Delta\phi = \pi$, and most events have $\Delta\phi_2 > 2$ for $m_N = 1200$ GeV. The distribution of $\Delta\phi_1$ appears to become flatter for higher neutrino masses, although it is still slightly skewed towards large values for $m_N = 1200$ GeV. We can explain all of these effects by looking at the geometry of the events. In figure 6.7 we show representative examples of the event geometry for the case of relative small m_N (say $m_W < m_N < 300$ GeV) and for the case of large m_N ($m_N > 1000$ GeV). In the figure, the arrows show the direction of motion of the different particles, and the lengths of the arrows indicate the speed of the particle in the collider frame. The plane of view in the figure is the transverse plane, so it is more accurate to say that the arrows are the projection of the particle velocities in the transverse plane. This means that the angle between two arrows in the figure correspond to the difference in ϕ of the two particles, and when the two particles are l_i and ν it is exactly $\Delta\phi_i$. Since the partons of the protons have negligible momentum in the transverse direction, and momentum is conserved, the virtual W^* boson will be approximately stationary in the transverse plane, as is indicated by the points in figure 6.7. Then, to conserve momentum, the heavy neutrino and l_1 must be back-to-back in the transverse plane and have the same transverse momentum. For the same W^* invariant mass and polar angle of the neutrino in the W^* frame, the neutrino (N) with lower mass will get a higher transverse velocity, both because the transverse momentum becomes larger[1] and because the velocity for a given momentum is larger for a particle with lower mass. Although the lighter heavy neutrino might not get the larger transverse velocity in some single pair of events where the W^* invariant

[1]It is easy to show that in the W^* frame the transverse momentum of the heavy neutrino when neglecting the mass of l_1 is $p_T = \frac{W^2 - m_N^2}{2W} \sin\theta$, where W is the invariant mass of the off-shell W^* boson and θ is the polar angle of the neutrino in this frame. Since the boost to the collider frame is perpendicular to the transverse plane, this is also the transverse momentum in the collider frame.

mass and neutrino polar angle are different, it will be true when averaging over many events. To indicate this, the neutrino in the left diagram has a longer arrow than on the right. A consequence of this is that the reference frame of the neutrino is in general more boosted in the transverse direction for lower masses. We know that the more boosted a particle is, the more "forward" its daughters become, meaning that they tend to move in the same direction as the mother with a small angle between them in the collider frame. For $m_N \approx 200$ GeV, the daughter W boson will also get a relatively high velocity, and therefore l_3 and ν will have a relatively small angle between them. This now explains all the features of the plot in figure 6.6a. The $\Delta\phi$ of l_1 tends to be large because l_1 is back-to-back with the heavy neutrino in the transverse plane, and for $m_N \approx 200$ GeV, the $\Delta\phi$ between ν and N tends to be smaller than 90 degrees. Because the on-shell W gets a relatively large velocity, the final state neutrino (or missing energy) will tend to move in the same direction as the W, and the angle between l_3 and ν tends to be small. This explains why the distribution of $\Delta\phi_3$ is skewed towards low values. Finally, since the heavy neutrino's frame is slightly boosted compared to the collider frame, the angle between l_2 and the W boson (and therefore also the SM neutrino) is smaller than 180 degrees, which is why $\Delta\phi_2$ is not always close to π, but has a relatively large tail towards lower values.

Now, in figure 6.7b, we see what happens when the heavy neutrino mass is so large that the heavy neutrino will almost always get a very low transverse velocity. Then the collider frame and the heavy neutrino frame are approximately the same. Therefore, l_2 and W will be almost exactly back-to-back in the transverse plane, and their ϕ (not $\Delta\phi$) will be approximately uniformly distributed. For large m_N, the W boson is more boosted, so both the SM neutrino and l_3 will move in the same direction as the W boson. This explains why $\Delta\phi_3$ tends towards 0 as m_N increases, while $\Delta\phi_2$ tends towards π. And when the centre of mass frame of the heavy neutrino is approximately equal to the collider frame, the l_2 and W can together be rotated in an arbitrary direction, which is why the distribution of $\Delta\phi_1$ becomes more uniform.

In figure 6.8, we show the distributions of the invariant masses of the three pairs of signal leptons in each event. We label the invariant mass of l_1 and l_2 as $m_{ll,1}$, the invariant mass of l_1 and l_3 as $m_{ll,2}$ and the invariant mass of l_2 and l_3 as $m_{ll,3}$. We see that the three distributions get wider and their centres move to larger values when the heavy neutrino mass is increased. This is expected, since all three leptons are more energetic on average when m_N is larger. For all m_N values, the distribution of $m_{ll,1}$ has the largest tail towards high values, which makes sense because both l_1 and l_2 tend to be very energetic. The distribution of $m_{ll,2}$ has almost the same shape as the one of $m_{ll,1}$ for $m_N = 200$ GeV, but when m_N increases, the $m_{ll,2}$ distribution shifts to the left compared to the other two. This is connected to what we saw in figure 6.2; that for the lowest heavy neutrino masses, l_1 is approximately as energetic as l_2, while for larger m_N, l_2 becomes the most energetic of the leptons. Since $m_{ll,2}$ is the only invariant mass not involving l_2, it makes sense that it becomes smaller compared to the others. Finally, the distribution of $m_{ll,3}$ has a very distinct shape in all four plots. It has a higher peak than the other two distributions, and a sharp edge towards larger invariant masses. This can be explained by the fact that the invariant mass of l_2 and l_3 is almost the same as the invariant mass of the heavy neutrino, since it only misses the contribution from the final state neutrino. Therefore, since the heavy neutrino is on-shell, $m_{ll,3}$ must be smaller than m_N, and the difference between these two values is due to the energy carried by the neutrino. This

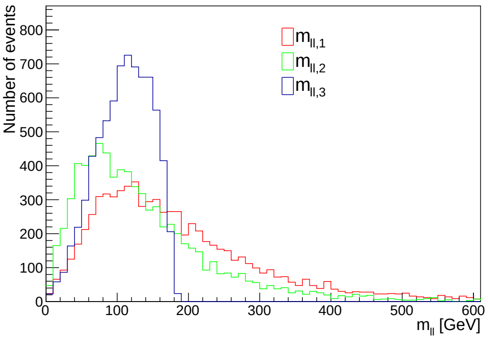

(a) $m_N = 200$ GeV

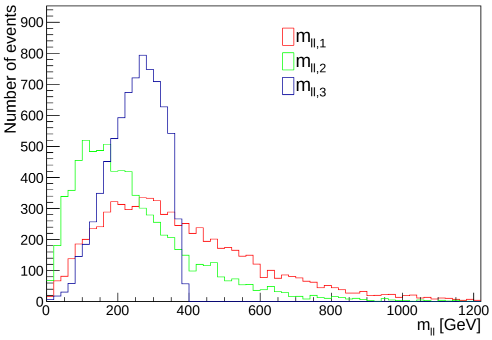

(b) $m_N = 400$ GeV

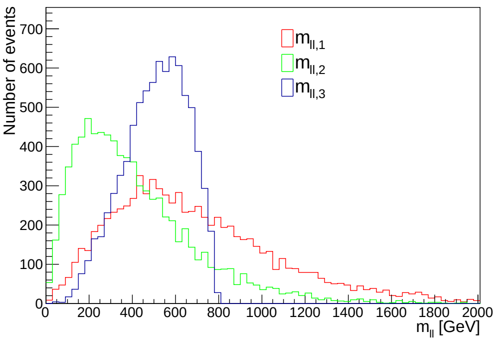

(c) $m_N = 800$ GeV

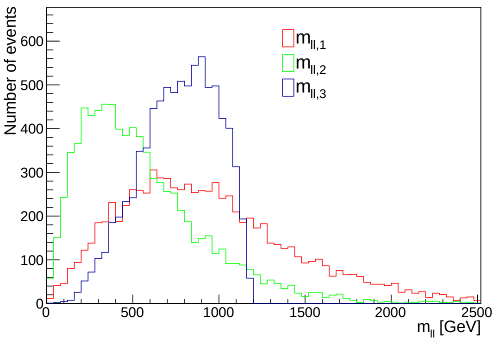

(d) $m_N = 1200$ GeV

Figure 6.8: The distributions of the invariant mass of the three pairs of signal leptons in each event for different heavy Majorana neutrino masses m_N. Here, $m_{ll,1}$ is the invariant mass of leptons l_1 and l_2, $m_{ll,2}$ is the invariant mass of l_1 and l_3 and $m_{ll,3}$ is the invariant mass of l_2 and l_3.

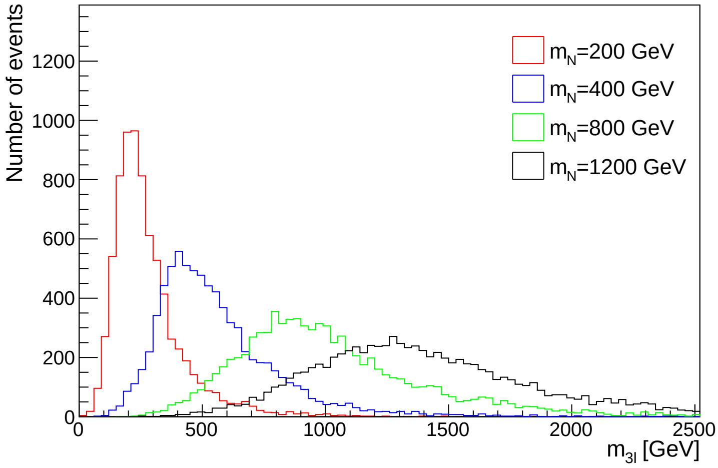

Figure 6.9: The distributions of the invariant mass of the three leading signal leptons in each event for different heavy Majorana neutrino masses m_N.

effect is reminiscent of the effect that lead to the first evidence of the existence of the SM electron neutrino from the energy spectrum of beta-decay electrons, which has a similar edge due to the energy carried by the neutrino. We see that the edges of the $m_{ll,3}$ distributions line up very well with m_N in each plot, which fits with this explanation.

In figure 6.9 we show the distributions of the invariant mass of the combination of all three charged leptons, which we call m_{3l}, in each event. Since this does not include the neutrino, it is a variable that can be measured experimentally. On the other hand, not taking the neutrino into account means that this three-lepton invariant mass does not have direct physical meaning. It does however give some indication of how much mass and/or energy is present in an event. In fact, we can clearly see this from the figure. For each heavy neutrino mass, the peak of the invariant mass distribution is quite close to m_N. Therefore the invariant mass of the three leptons can be used to get information about the mass of the heavy neutrino. We notice that as well as being shifted towards larger invariant masses, the distributions also become broader and more symmetrical for larger heavy neutrino masses.

Figure 6.10 shows the invariant mass of only lepton l_2 and l_3, and the final state neutrino. We call this invariant mass $m_{ll\nu}$. This variable is impossible to calculate from experimental data, since it requires knowledge of the full momentum of the neutrino, while only the transverse component can be measured. If this variable could be calculated it would be very useful, because it has a clear physical significance, in contrast to the invariant mass of the three charged leptons. The reason for this significance is that the three particles included in this invariant mass are exactly the three decay products of the heavy neutrino in the final state. This means that their combined invariant mass should equal the mass of the neutrino, giving very clear resonances. This is exactly what we see in figure 6.10, where there are very clear peaks right at the respective m_N values. These distributions of $m_{ll\nu}$ could also be used to measure the decay width of the heavy neutrino, from the width of the peaks.

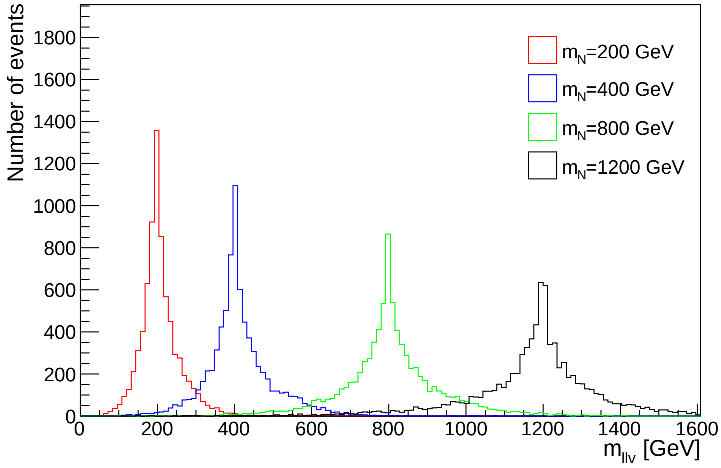

Figure 6.10: The distributions of the invariant mass of the signal leptons l_2 and l_3, and the final state neutrino ν, for different heavy Majorana neutrino masses m_N.

A variable that is similar to $m_{ll\nu}$, but can be calculated from experimental measurements, is the transverse mass of the same set of particles, $m_{T,23}$. The distributions of $m_{T,23}$ are shown in figure 6.11. As noted in section 4.1, the transverse mass of three particles will have an upper limit at the invariant mass of the three particles. We therefore expect these distributions to have a wide peak at lower invariant masses, and a sharp edge at the respective heavy neutrino masses, which are equal to the invariant mass of the three particles. This is more or less what we see in figure 6.11. For all four heavy neutrino masses, the distributions have a peak at around half of m_N, and fall off to zero at very close to m_N, as expected. It appears that the peak of the distribution moves slightly towards lower values relative to m_N and that the downward slope towards m_N becomes relatively less steep for larger heavy neutrino masses. These plots show that $m_{T,23}$ can be used to estimate the mass of the heavy neutrino, by looking at the point where the distribution drops towards zero, but this might be less accurate for larger m_N, when the edge of the peak seems to become less sharp. The accuracy with which m_N can be determined from these distribution in a real experiment is also limited by the momentum and energy resolutions of the detector, since worse resolution would make the edges even less sharp.

In figure 6.12, we show the distributions of the missing transverse energy in each event. This should be equal to the transverse energy/momentum of the final state neutrino ν, as long as there are no other significant contributions to missing energy. We see that the missing energy tends to become larger for larger m_N, which is again a result of the events becoming more energetic on average.

Although the tree-level signal process shown in figure 5.4 has only leptons in the final state, our simulated events can also contain quarks and/or gluons in the final state, which result in hadronic jets after hadronisation. The origin of these quarks and gluons can be either NLO radiation processes like the one shown in the top diagram of figure 5.6, or

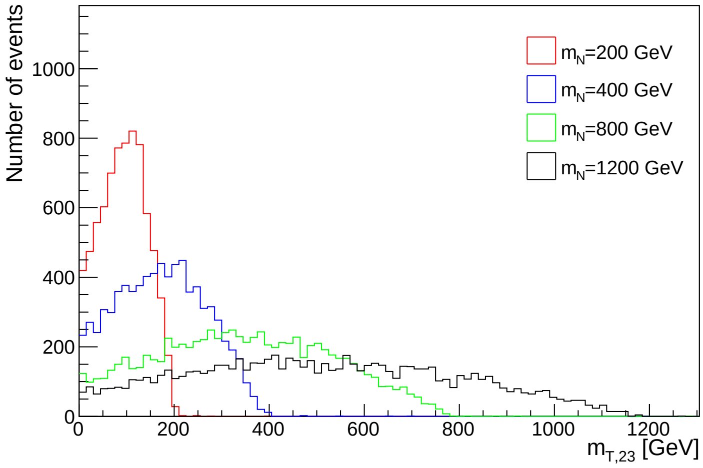

Figure 6.11: The distributions of the transverse mass of the signal leptons l_2 and l_3, and the final state neutrino ν, for different heavy Majorana neutrino masses m_N.

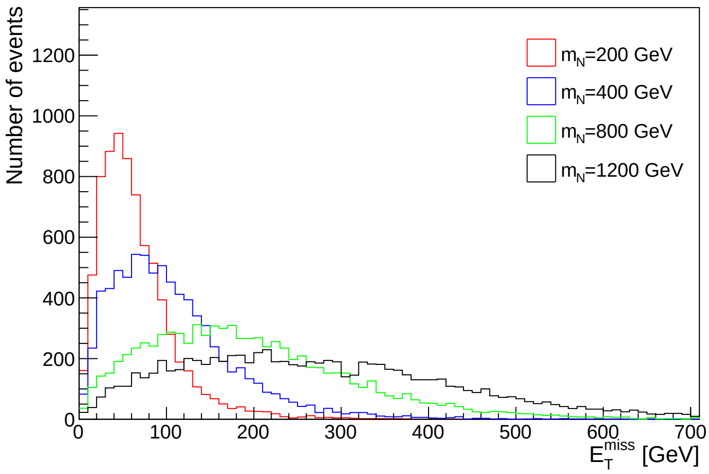

Figure 6.12: The distributions of the missing transverse energy E_T^{miss} in each event for different heavy Majorana neutrino masses m_N.

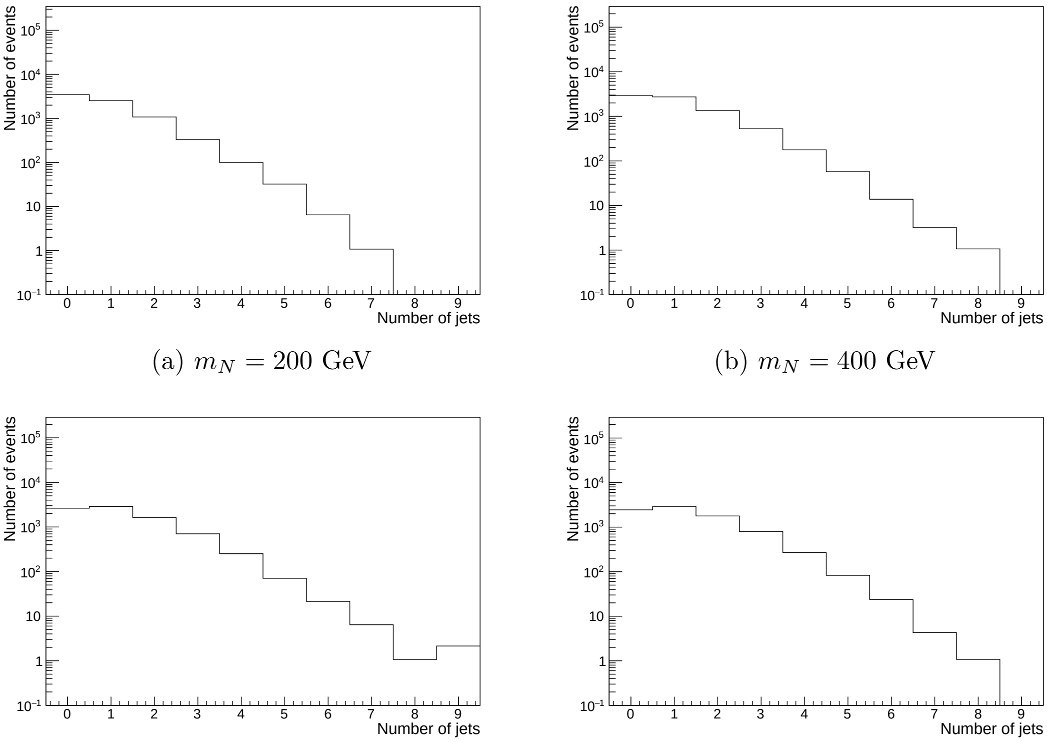

(a) $m_N = 200$ GeV

(b) $m_N = 400$ GeV

(c) $m_N = 800$ GeV

(d) $m_N = 1200$ GeV

Figure 6.13: The distributions of the number of signal jets passing the cuts in each event for different heavy Majorana neutrino masses m_N.

the hadronic decay of a signal tau lepton. To label a jet as a signal jet, we require that it passes the cuts $p_T > 20$ GeV and $|\eta| < 2.5$. The distributions of the number of such signal jets in each event are shown in figure 6.13. We see that unlike the distributions of number of signal leptons, the number of signal jets tends to increase with the heavy neutrino mass. For $m_N = 200$ GeV, almost half the events have no jets, and only a little over 1 000 events have more than one jet. For $m_N = 400$ GeV, there are less events with no jets, and in this case there is approximately the same number of events with zero jets as events with one jet. For the two largest m_N, the most common number of jets in an event is one, and more than a third of the events have more than one jet. The reason why there tend to be more jets for higher m_N is probably also due to the higher energy in the events, which might increase the chance of QCD radiation, and cause more than one jet to be formed from each radiated quark or gluon.

In figure 6.14, we show the distributions of the number of signal b-jets in each event. B-jets are jets that originate from a b quark, which can be identified in experiments using the properties of B hadrons [65]. Since we have the truth information from the MC simulation, we know exactly which jets are b-jets, and so signal b-jets are simply the b-jets that pass the same cuts as the regular jets. We see that there are very few b-jets in all four simulations. The number of events with b-jets increases slightly for the larger heavy neutrino masses, but it does not get larger than around 100. For $m_N = 1200$ GeV, there is around one event that has three b-jets, which is the largest number of b-jets in

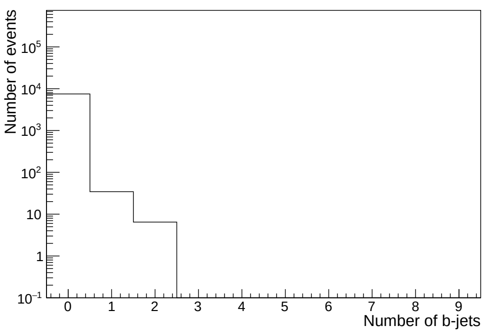

(a) $m_N = 200$ GeV

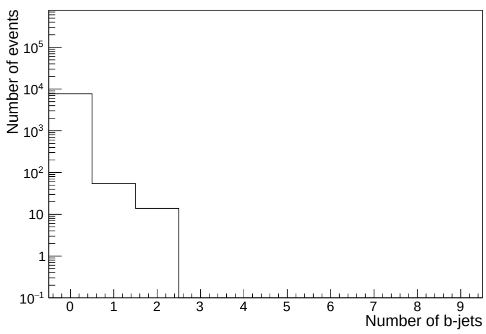

(b) $m_N = 400$ GeV

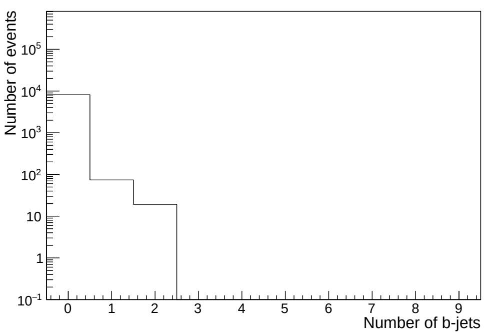

(c) $m_N = 800$ GeV

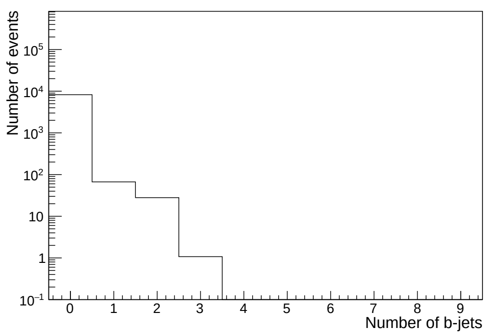

(d) $m_N = 1200$ GeV

Figure 6.14: The distributions of the number of signal b-jets passing the cuts in each event for different heavy Majorana neutrino masses m_N.

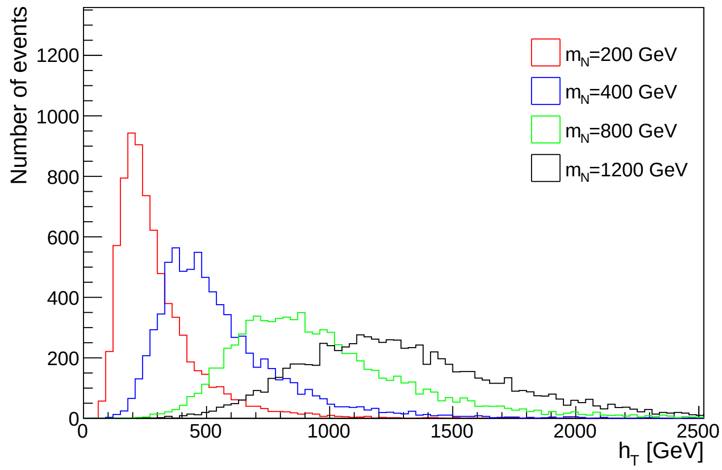

Figure 6.15: The distributions of the hadronic activity h_T of all leptons and jets that passed the cuts in each event for different heavy Majorana neutrino masses m_N.

any of the events.

In figure 6.15, we show the distributions of the hadronic activity h_T, given by the scalar sum of the transverse momenta of all the signal leptons and signal jets in each event. We see that these distributions are very similar to the distributions of the invariant mass of the three leading signal leptons shown in figure 6.9. The h_T distributions also become wider and move to larger values for higher m_N, and the distributions become more symmetric. The peaks of the distributions are also located at approximately the same values as in the m_{3l} distributions, which is around the heavy neutrino mass. This might not be that surprising, since we expect the leptons to carry most of the transverse momentum in each event, and since the leptons have negligible masses the scalar sum of their p_T is equal to the sum of their (transverse) energies. If the vector sum of the three leptons' transverse momenta is zero, we see from equation 4.16 that the sum of their transverse energy is equal to the transverse mass of the three leptons. However, we know that the vector sum of the three lepton transverse momenta is not zero, but approximately equal to the missing transverse momentum, due to conservation of transverse momentum. Therefore the transverse mass of the three leptons is smaller than the scalar sum of their p_T by approximately this contribution from E_T^{miss}. It is therefore natural that the scalar sum of p_T is approximately equal to the invariant mass, which we know is also slightly larger than the transverse mass.

In the model we use, the heavy Majorana neutrino couples equally to both signs of any of the charged leptons. This means that we expect approximately the same number of events where l_1 and l_2 have the same sign (SS) as events where they have opposite sign (OS). If lepton number was conserved, l_1 and l_2 would be forced to be OS, so the possibility of this pair being SS is a unique signature of a Majorana type neutrino. To verify that the heavy Majorana neutrino breaks lepton number conservation in approximately half of the events, we plot the product of the charges of l_1 and l_2 in units of the elementary charge. This product being equal to -1 is equivalent to l_1 and l_2 being OS, while when

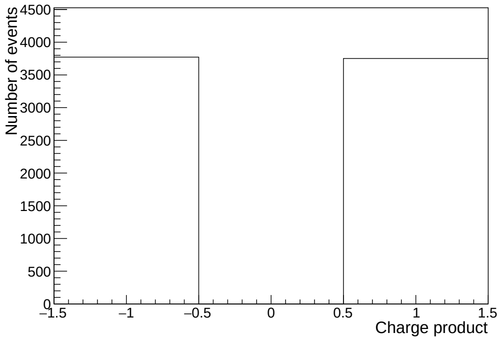

(a) $m_N = 200$ GeV

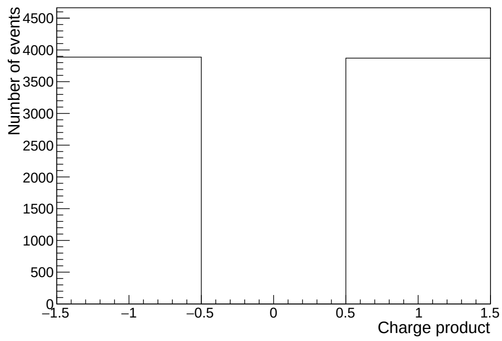

(b) $m_N = 400$ GeV

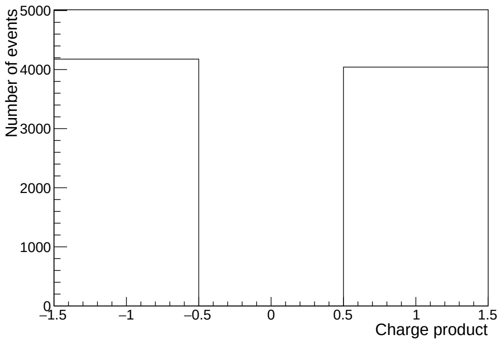

(c) $m_N = 800$ GeV

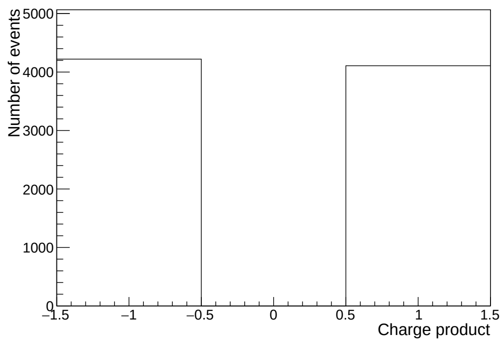

(d) $m_N = 1200$ GeV

Figure 6.16: The distributions of the product of the charges of l_1 and l_2 in each event for different heavy Majorana neutrino masses m_N. This shows whether the leptons have the same charge $(+1)$ or the opposite charge (-1).

the product is $+1$ the lepton pair is SS. The plots of the charge products are shown in figure 6.16. In the figure we see that the number of events with OS and SS l_1 and l_2 is approximately equal for all the heavy neutrino masses, as expected. However, there seems to be a slight excess of events with OS leptons, in particular for the masses $m_N = 800$ GeV and $m_N = 1200$ GeV. To find out whether the excess is significant, we can compute the probability of seeing this excess if each event does have an exact 50% probability of both SS and OS, and that the choice is made independently in each event. If we assume this to be the case, the probability of seeing x events with OS will follow a binomial distribution with success rate $p = 1/2$:

$$P\left(X = x\right) = \binom{n}{x} \left(\frac{1}{2}\right)^n, \tag{6.3}$$

where n is the total number of events. Although we can not know for certain without knowing the exact inner workings of `MadGraph`, the choice of sign for the leptons is most likely made using its pseudo-random number generator, and in a good pseudo-random number generator there is very little correlation between consecutive values. Therefore the use of the binomial distribution should be justified. However, for each simulation of 10 000 events we use the same seed for the random number generator in `MadGraph`, and we can therefore expect some correlations between the four simulations. We therefore

Table 6.1: The number of events x where the leptons l_1 and l_2 have opposite sign out of the n events that were simulated and passed cuts for each heavy neutrino mass value. The last column gives the probability of seeing this many opposite sign events or more, assuming that x follows the binomial distribution (6.3).

m_N [GeV]	n	x	$P(X \geq x)$
200	7 523	3 772	0.41
400	7 757	3 886	0.44
800	8 218	4 176	0.07
1200	8 329	4 221	0.11

look at each of the four simulations separately. To find out whether the deviations from equal numbers of OS and SS events for each simulation are significant we need to sum the distribution (6.3) from the observed $X = x$ numbers of OS events to $X = n$, to find the probabilities of getting that many OS events or more by random chance. The values of n, x and $P(X \geq x)$ for each simulation are given in table 6.1. The values of x have to be rounded to the nearest whole number, because the MC event weights cause the number of entries in the histogram bins to not be whole numbers. In table 6.1, we see that the excesses in OS events for $m_N = 200$ GeV and $m_N = 400$ GeV are not significant, as there is more than 40% chance of getting this many OS events or more by chance. The excesses for $m_N = 800$ GeV and $m_N = 1200$ GeV are more significant, but still not to the extent that it could not happen by chance. It is interesting that all four simulations have an excess of OS events. The reason for this might be the previously mentioned fact that we used the same seed for the MC simulations, and that this introduced some correlation in the choice of signs between the four simulations. Another interesting thing to note is that when only counting unweighted events, the excess becomes less significant. For $m_N = 800$ GeV, 4 164 out of the 8 218 events have OS leptons when not applying the event weights. This gives a probability $P(X \geq 4\ 164) \approx 0.11$. For $m_N = 1200$ GeV, 4 213 out of the 8 329 events have OS lepton, which also increases the probability compared to the weighted case. Now we get $P(X \geq 4\ 213) \approx 0.15$. This means that in both of these simulations, the OS events for some reason get have higher average event weights, which contributes to the excess seen in figure 6.16 and table 6.1. It is reasonable to conclude that the choice of OS or SS leptons is in fact 50/50 as we expect from the model, and that the MC simulations work correctly, within expected statistical fluctuations.

6.3 Pseudo-Dirac model

Since the pseudo-Dirac model that we use has the exact same Lagrangian as the Majorana model, the kinematics of the signal process will be the same for the case of a pseudo-Dirac heavy neutrino with the same masses and mixing parameters. We will therefore not look at all the kinematic distributions for the pseudo-Dirac model. However, we know that one difference between the two models is that the pseudo-Dirac heavy neutrino should conserve lepton number. This means that the charges of l_1 and l_2 in figure 5.4 must have opposite signs. The only distribution we will show for the pseudo-Dirac model is therefore the distribution of the product of the charges of l_1 and l_2, to verify that they in

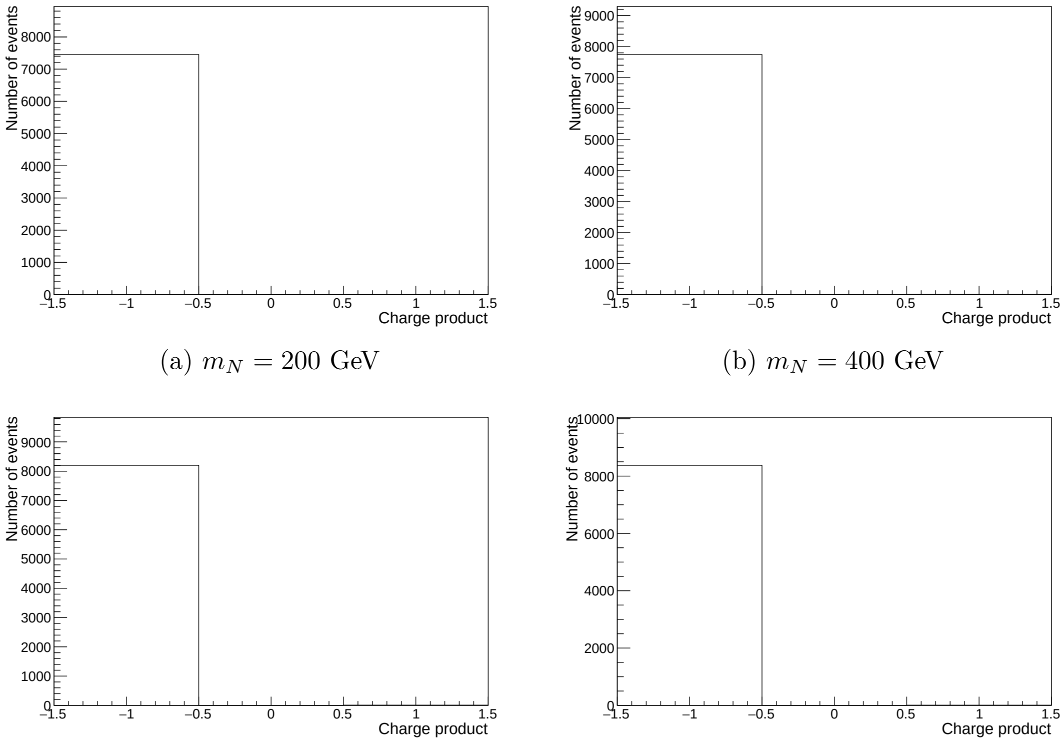

(a) $m_N = 200$ GeV (b) $m_N = 400$ GeV (c) $m_N = 800$ GeV (d) $m_N = 1200$ GeV

Figure 6.17: The distributions of the product of the charges of l_1 and l_2 in each event for different heavy pseudo-Dirac neutrino masses m_N. This shows whether the leptons have the same charge (+1) or the opposite charge (−1).

fact have the opposite sign. The distributions for $m_N = 200$, 400, 800 and 1200 GeV are shown in figure 6.17. We see that as expected, all the events have a product of charges equal to −1, which is equivalent to l_1 and l_2 having opposite charge.

6.4 Identifying the lepton vertices of origin

In this section, we want to use what we learned in section 6.2 to develop a method to identify the vertex of origin for the three leptons in a signal event of the Majorana or Dirac model, using only variables that are measurable in collider experiments. We use the notation of figure 5.4 to refer to the correct identification for the leptons; l_1 is the lepton from the production vertex of the heavy neutrino, l_2 is the lepton from the decay vertex, and l_3 is the lepton from the decay of the W boson. An efficient method for identifying l_1, l_2 and l_3 in experiments can be useful in several ways. First of all, it can increase the sensitivity of the experiment to the pseudo-Dirac model, since it would allow us to perform a cut requiring that the charge of the leptons identified as l_1 and l_2 must be opposite. A similar idea could also be used to identify the Majorana or pseudo-Dirac nature of the heavy neutrino, by counting both the number of events where l_1 and l_2 have the opposite charge and where they have the same charge. If the heavy neutrino is of pseudo-Dirac type we expect an excess of the first type of events, but the number of same charge events should match the expectation from the background. If it is of Majorana type, on the other hand, we expect to see an excess of both kinds of events compared to the background expectation.

An important thing to note is that for both of these applications we only need to know the product of the charges of l_1 and l_2. We therefore do not need to identify both of them precisely, since it will suffice to know which pair of leptons is the pair l_1 and l_2. One way to identify the correct pair of l_1 and l_2 could therefore be to instead identify l_3 in each event, and assume that the other two leptons with the highest p_T are l_1 and l_2. Another idea could be to try to identify the pair of leptons l_1 and l_2 out of the three leading leptons. With all this in mind, there are a few kinematic variables that appear to be good candidates for lepton identification. There is the transverse momentum, which as we saw in figure 6.2 tends to be largest for l_2 and l_1, especially for larger neutrino masses. Another promising variable is the transverse mass, since as we saw in figure 6.4, the transverse mass of l_3 has a sharp edge at the W mass, while $m_{T,1}$ and $m_{T,2}$ tend to be much larger. There is also the $\Delta\phi$ between each lepton and the missing transverse momentum vector, which tends to be smallest for l_3, particularly for larger m_N. We also saw in figure 6.8 that the pair l_1 and l_2 has very large m_{ll} values in some events, so we might be able to use the maximum m_{ll} of all three pairs of leptons to identify the pair we are looking for. Finally, we could try to use some combination of several of these variables. In particular, since both p_T, m_T and $\Delta\phi$ tend to be smallest for l_3, it is tempting to try to order the leptons by something like $m_T \times \Delta\phi$, with the goal of creating a variable that has an even higher chance of being smallest for l_3. To test all of these options, we use the same MC samples of Majorana models that we studied in section 6.2, and compare each method to the truth information. For this analysis, we include all events with at least three leptons that fulfill $p_T > 5$ GeV and $|\eta| < 2.5$ (unlike in section 6.2 we do not require that the three leading leptons are l_1, l_2 and l_3, since this cut can not be done in experiments). The results of the analysis in terms of the success rate of each identification method are given in table 6.2. For all variables except m_{ll}, the identification method we used was to order the three leading leptons in each event by the given variable, and identify l_3 as the lepton with the smallest value of the variable. The leptons l_1 and l_2 are then assumed to be the other two of the leading leptons, and the identification is called a success regardless of the order of l_1 and l_2. For m_{ll}, the pair among the three leading leptons with the largest value of m_{ll} was identified as the l_1 and l_2 pair.

We see that out of the three variables p_T, m_T and $\Delta\phi$, it is m_T that is best for identification. It is approximately as efficient as p_T and more efficient than $\Delta\phi$ for small m_N, and much more efficient than p_T and slightly more efficient than $\Delta\phi$ for large m_N. The m_{ll} method performs very poorly, as it does not even have a success rate of 1/3, meaning that the other two lepton pairs are more likely to have the largest m_{ll} than l_1 and l_2 are. But the variable that performed best for every model out of the ones we tested is the product $p_T \times \Delta\phi$. That this product has a higher success rate than both p_T and $\Delta\phi$ on their own must mean that there is some correlation between the leptons' values of p_T and $\Delta\phi$. For example, it could be that in events where l_3 has a relatively large p_T it tends to have particularly small $\Delta\phi$, such that the product of the two is smaller than the same product for both l_1 and l_2. To test this, we can plot the two dimensional distributions of the p_T and $\Delta\phi$ of l_3 in the Majorana model. These are shown in figure 6.18. We see that in fact there seems to be a small negative correlation between $p_{T,3}$ and $\Delta\phi_3$, indicated by the negative slope of the yellow regions in some

Table 6.2: The success rates of our identification methods for the lepton vertices of origin for Majorana models with different m_N. For all the variables except m_{ll}, the lepton with the lowest value of the given variable is assumed to be l_3, and the other two are assumed to be l_1 and l_2. For m_{ll}, the lepton pair with the largest m_{ll} is assumed to be l_1 and l_2. In all cases, the identification is counted as a success as long as l_3 is identified correctly, and the pair l_1 and l_2 are identified correctly, without needing to distinguish between the two.

	$m_N = 200$ GeV	$m_N = 400$ GeV	$m_N = 800$ GeV	$m_N = 1200$ GeV
Variable	Success rates			
p_T	0.65	0.70	0.69	0.66
m_T	0.65	0.82	0.90	0.93
$\Delta\phi$	0.52	0.75	0.86	0.90
m_{ll}	0.26	0.24	0.25	0.27
$p_T \times \Delta\phi$	0.67	0.84	0.91	0.93
$m_T \times \Delta\phi$	0.59	0.80	0.89	0.92

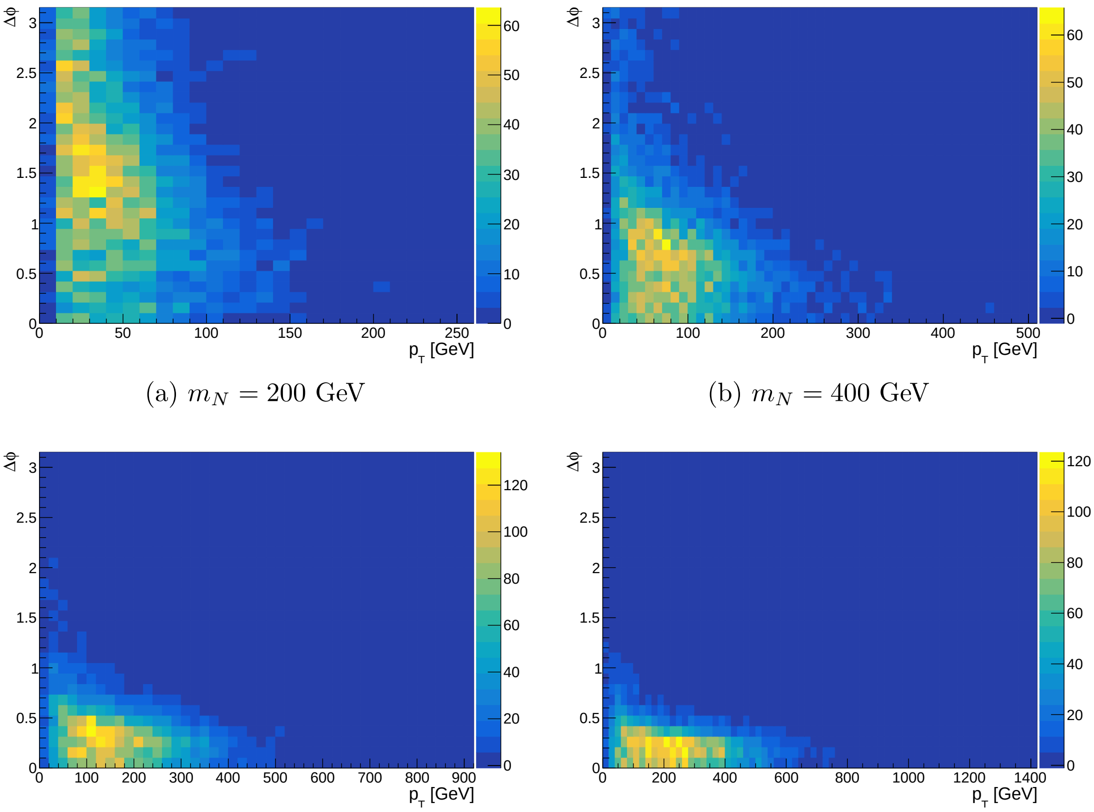

(a) $m_N = 200$ GeV

(b) $m_N = 400$ GeV

(c) $m_N = 800$ GeV

(d) $m_N = 1200$ GeV

Figure 6.18: The distributions of the transverse momentum versus the $\Delta\phi$ of l_3 for different heavy Majorana neutrino masses m_N.

Table 6.3: The mass values of the heavy Majorana neutrino and the right-handed W boson that were used in our simulations of the Keung-Senjanović process.

Model number	m_N [TeV]	M_{W_R} [TeV]
1	1.5	3
2	2	3
3	4	5
4	4.5	5
5	4	6
6	5	6

of the plots in figure 6.18. The negative correlation does not seem to be as strong for $m_N = 200$ GeV, which might explain why the $p_T \times \Delta\phi$ method performs the worst for this model.

The same correlation effect does not seem to apply to the variable $m_T \times \Delta\phi$, which performs worse than m_T alone in all models. This is not surprising, since as we saw in section 4.1, m_T is an increasing function of $\Delta\phi$, meaning that we expect a strong positive correlation between the two. Therefore events in which l_3 has unusually large m_T it will tend to also have large $\Delta\phi$, and so the product of the two can more easily be larger than the same product for the other two leptons.

We also tried a few other combinations of these variables such as $(m_T + p_T) \times \Delta\phi$ and $m_T \times p_T \times \Delta\phi$, but none of these gave better results than $p_T \times \Delta\phi$. We therefore conclude that the best method for identifying the pair of leptons l_1 and l_2 is to choose the two leptons out of the three with largest p_T that have the largest value of $p_T \times \Delta\phi$. This method has a success rate of around 2/3 for the smallest neutrinos mass we tested, $m_N = 200$ GeV, and the success rate improves drastically for larger masses. For masses of 800 GeV or more, the success rate is above 90%.

6.5 Left-right symmetric model

Finally, we will study the kinematics of the Keung-Senjanović process in the framework of the LRSM. We again let $\sqrt{s} = 13.6$ TeV, and perform MC simulations at NLO in QCD. We use the assumption that the right-handed neutrino mixing matrix is diagonal, and assume that N_1 is much lighter than the other right-handed neutrinos, such that it is the only one that contributes to the process. This right-handed neutrino then couples only to the electron through the right-handed charged current interaction, so we will only have electrons as the two signal leptons l_1 and l_2. We perform six different simulations with different W_R and heavy Majorana neutrino masses, and generate 10 000 events for each of these. The mass values used are shown in table 6.3, along with the model number that we will use to refer to them by. Models 1 and 2 are excluded under the assumption that the right-handed weak coupling of the W_R boson is the same as the SM weak coupling from a study using data from the ATLAS detector from Run 2 of the LHC [61]. We still include these models, because we hope to put upper limits on the right-handed weak coupling in these models. The four other models are just outside of the current excluded region of masses, and are therefore of special interest. In all six cases, the heavy neutrino mass is smaller than M_{W_R}, meaning that the first W_R boson in the Feynman diagram will be on-shell. At the same time, m_N is never smaller than $M_{W_R}/2$, meaning that we do not enter the regime where the jets from the W_R decay will be collimated. In `SimpleAnalysis`, we implement cuts on leptons of $p_T > 5$ GeV and $|\eta| < 2.5$, and cuts on jets of $p_T > 20$ GeV and $|\eta| < 2.5$. Then only events with at least two remaining signal leptons and two signal jets are kept, and the two leading leptons and two leading jets are also required to be the ones from the tree-level signal process. After these cuts, the number of remaining events is between 8 871 (model 4) and 9 182 (model 2). We begin the analysis by plotting the number of signal leptons, jets and b-jets in each event. These distributions are shown in figures 6.19–6.21. We see that just like in the models without W_R, the number of signal leptons is not affected much by changing the new particles' masses. In all models, most events have two signal leptons, while between 1 000 and 2 000 events have three signal leptons, and less than 1 000 events have four leptons. The maximum number of signal leptons is eight, which occurs in a few events in models 3, 4 and 5. From the other two figures we see that the number of signal jets and b-jets also do not change that much between these simulations. For each model, the distribution of the number of signal jets has a peak at three jets, with a few thousand events having this number of signal jets. The distributions then decrease gradually as the number of signal jets increases. For all models, there are tens of events that have ten signal jets or more. The maximum number of jets in one event is 15, which occurs a few times in model 6. From figure 6.21, we see that a very small fraction of the jets are b-jets. In all models, over half the events have no signal b-jets, while between 1 000 and 2 000 events have one or two b-jets. The maximum number of b-jets in an event is five, occurring in models 3 and 4.

In figure 6.22, we show the transverse momenta of l_1 and l_2, but ordered by their p_T. We call these transverse momenta $p'_{T,li}$, to distinguish them from the p_T of l_i, which we call $p_{T,li}$. These distributions have some very interesting features, and to understand them we need to think about the kinematics of this specific process. In the case when both the first W_R boson and the heavy neutrino are on-shell, the transverse momentum of l_1 has a maximum allowed value that depends only on the masses of W_R and N. To see this, we find the four-momenta of the W_R boson and its decay products in the W_R centre of mass frame. The four-momentum of the W_R boson in its centre of mass frame is trivial:

$$p_W^\mu = (M_{W_R}, 0, 0, 0)\,. \tag{6.4}$$

To get the maximum possible transverse momentum for l_1, we want it to move only in the transverse direction. We choose its momentum to point along the x-axis. Then the four-momenta of l_1 and N in the W_R frame can be written as

$$\begin{aligned} p_{l1}^\mu &= (p_{l1}, p_{l1}, 0, 0) \\ p_N^\mu &= \left(\sqrt{m_N^2 + p_{l1}^2}, -p_{l1}, 0, 0\right), \end{aligned} \tag{6.5}$$

where we used that the mass of l_1 is negligible, and we call the magnitude of its momentum p_{l1}. Conservation of energy gives

$$M_{W_R} = \sqrt{m_N^2 + p_{l1}^2} + p_{l1}, \tag{6.6}$$

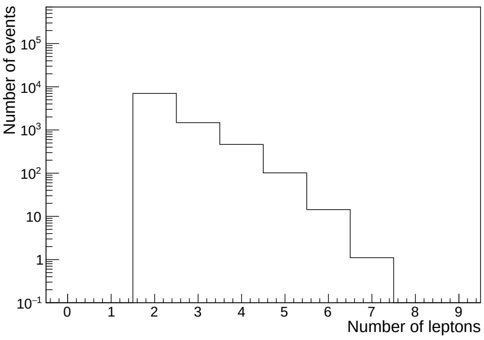

(a) Model 1

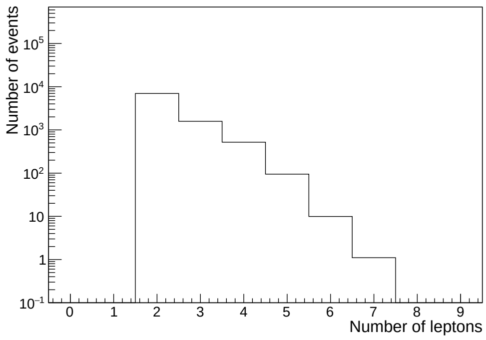

(b) Model 2

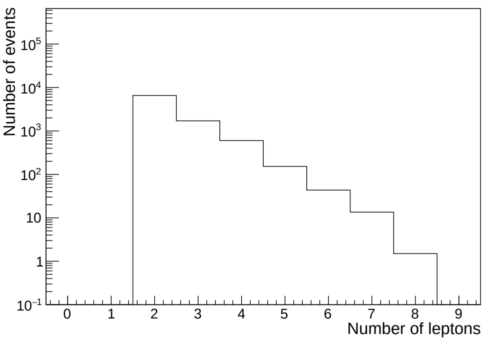

(c) Model 3

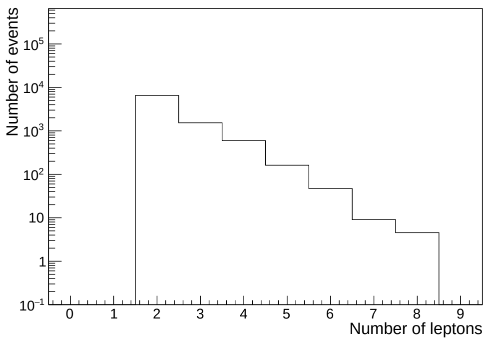

(d) Model 4

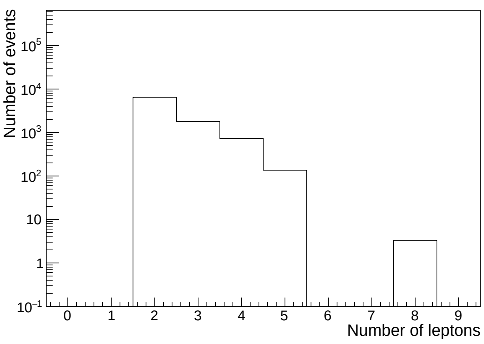

(e) Model 5

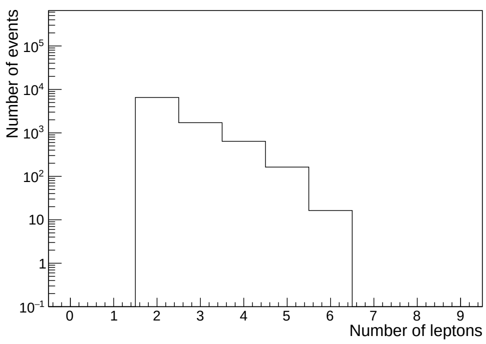

(f) Model 6

Figure 6.19: The distributions of the number of signal leptons in each event for the different simulated LRSM models.

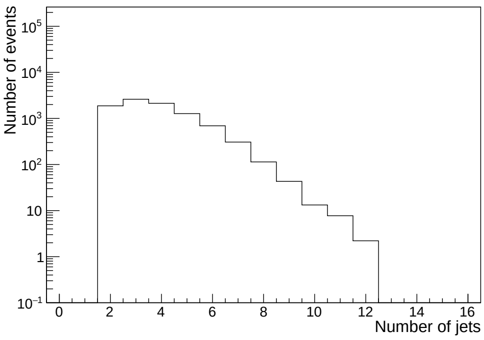

(a) Model 1

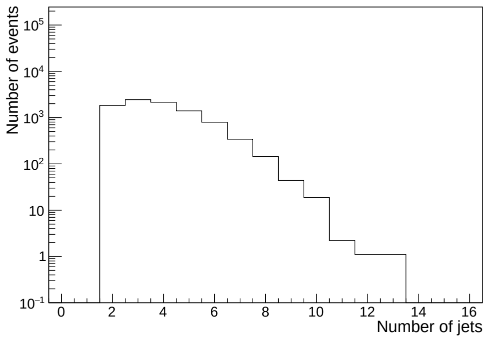

(b) Model 2

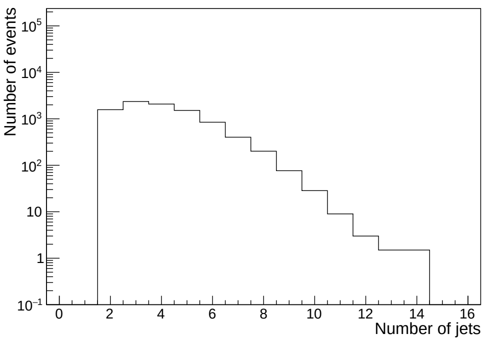

(c) Model 3

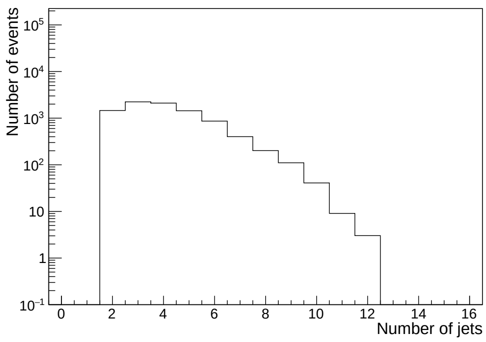

(d) Model 4

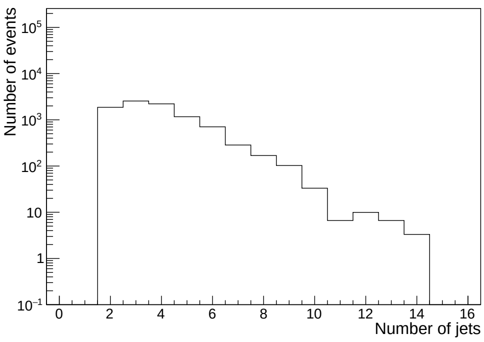

(e) Model 5

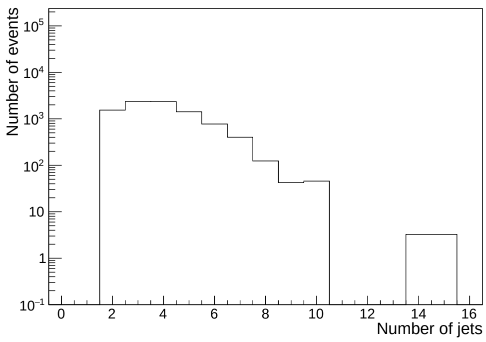

(f) Model 6

Figure 6.20: The distributions of the number of signal jets in each event for the different simulated LRSM models.

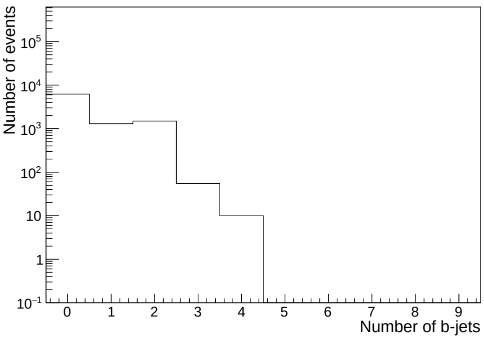

(a) Model 1

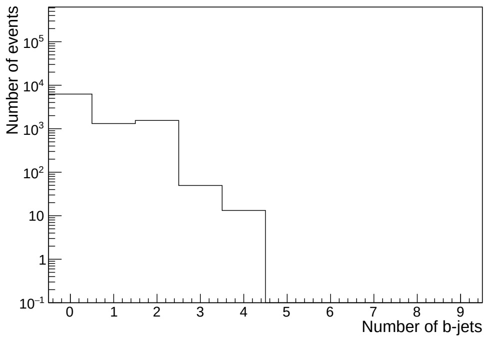

(b) Model 2

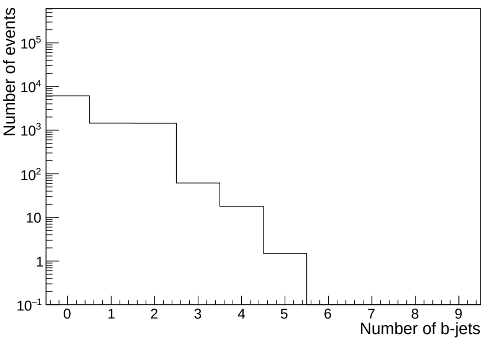

(c) Model 3

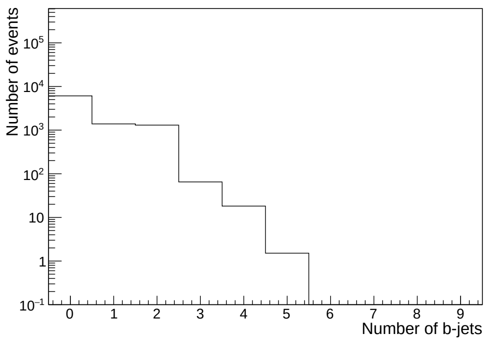

(d) Model 4

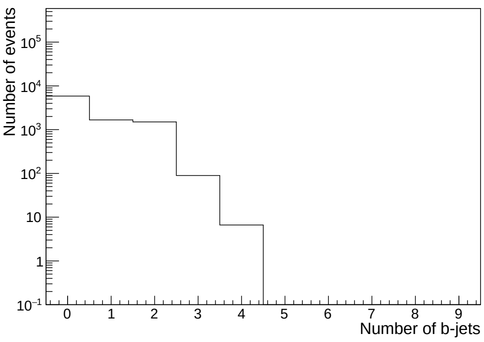

(e) Model 5

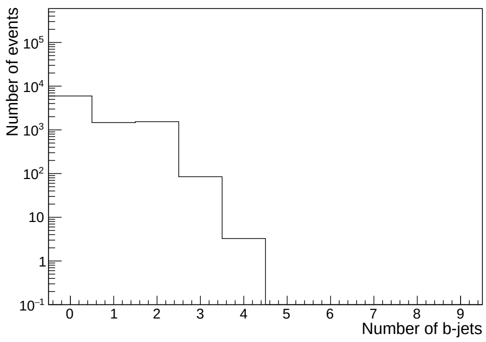

(f) Model 6

Figure 6.21: The distributions of the number of signal b-jets in each event for the different simulated LRSM models.

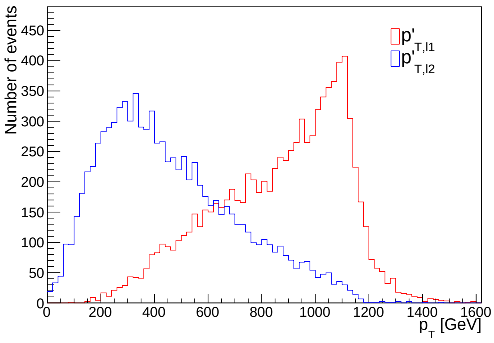

(a) Model 1

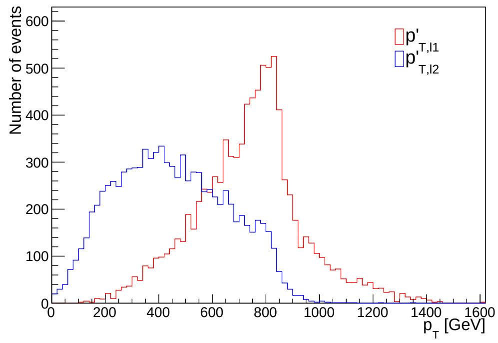

(b) Model 2

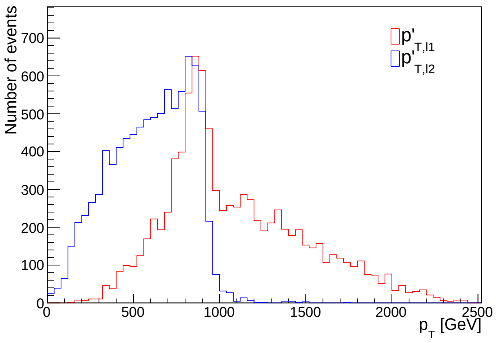

(c) Model 3

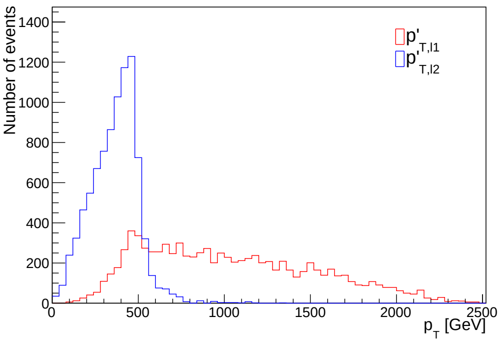

(d) Model 4

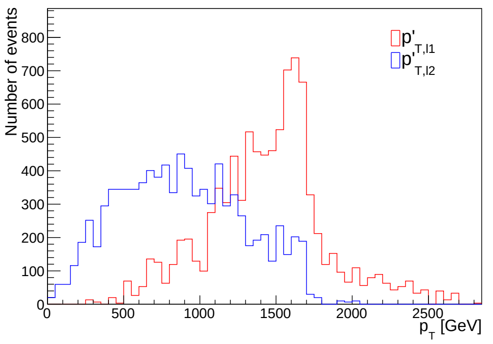

(e) Model 5

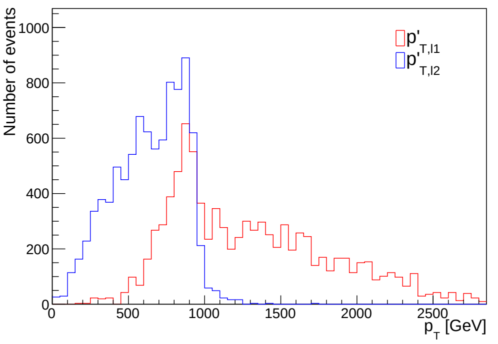

(f) Model 6

Figure 6.22: The distributions of the transverse momenta of the two leading signal leptons in each event, ordered by p_T, for the different simulated LRSM models.

Table 6.4: The maximum transverse momenta of the lepton from the production vertex of the heavy neutrino (the one labelled l_1 in the Feynman diagram of figure 5.5) that are kinematically allowed when both the W_R boson and the heavy neutrino are exactly on-shell. The maximum depends on the masses of the W_R and N, and therefore vary between models.

Model number	$p_{T,l1}^{\max}$ [GeV]
1	1 125
2	833
3	900
4	475
5	1 667
6	917

which can be solved for p_{l1} to give

$$p_{l1} = \frac{M_{W_R}^2 - m_N^2}{2M_{W_R}}. \tag{6.7}$$

As we argued in section 6.2, the centre of mass frame of the W_R boson is related to the collider frame by a boost in only the z direction, due to the negligible transverse momenta of the partons. Therefore equation 6.7 gives the maximum transverse momentum of l_1 in the collider frame as well. If we call the mass difference between the W_R boson and the heavy neutrino Δm, we have the relation (assuming $m_N < M_{W_R}$, which is the case in all our simulations)

$$m_N = M_{W_R} - \Delta m. \tag{6.8}$$

Then equation (6.7) can be written as

$$p_{T,l1}^{\max} = \frac{(2M_{W_R} - \Delta m)\,\Delta m}{2M_{W_R}} = \Delta m \left(1 - \frac{\Delta m}{2M_{W_R}}\right). \tag{6.9}$$

From this expression, we see that if $\Delta m \ll 2M_{W_R}$, then

$$p_{T,l1}^{\max} \approx \Delta m. \tag{6.10}$$

This shows that when the difference between the masses of the W_R boson and the heavy neutrino is small compared to the W_R mass, the maximum possible transverse momentum of l_1 will be approximately equal to the mass difference. Using equation (6.9), which holds for any mass difference, we can compute $p_{T,l1}^{\max}$ for each of the six models. The results are shown in table 6.4. Comparing the values in the table with the plots in figure 6.22, we see that the values explain many features of the p_T distributions. A consequence of the upper limit on the transverse momentum of l_1 is that $p'_{T,l2}$ can never be larger than $p_{T,l1}^{\max}$, since then both leptons must have p_T above the limit. We see that in all the plots, the distribution of $p'_{T,l2}$ does in fact stay almost exclusively beneath the respective upper limit. There are some events where the distribution is slightly above the limit, and this is probably caused by the fact that the particles are never perfectly on-shell, but their invariant mass follows a distribution with finite width. We also see that in each plot one of the two distributions or both have a clear edge around the value of $p_{T,l1}^{\max}$. This shows that in many of the events, the transverse momentum of l_1 is very close to the upper limit. For models 1 and 2, it is $p'_{T,l1}$ that has the clearest edge around $p_T = 1125$ GeV and $p_T = 833$ GeV respectively, meaning that in these models l_1 tends to be the lepton with the highest p_T, even though the transverse momentum of l_2 has no upper limit. In model 3, both distributions have peaks with edges at around 900 GeV, so here l_1 can have either larger or smaller p_T than l_2, with approximately the same probability. From the relatively large tail of the $p'_{T,l1}$ distribution, we see that in a significant number of events the p_T of l_2 is much larger than that of l_1. In model 4, the W_R mass is the same as in model 3, meaning that the energy of the events are approximately the same, but here the upper limit on $p_{T,l1}$ is only 475 GeV, as seen in table 6.4. We see that this causes l_1 to almost always have the smallest p_T, which gives a very clear edge in the $P'_{T,l2}$ distribution around 475 GeV. The transverse momentum of l_2 on the other hand can get larger than 2 TeV in some events. Model 5 has the largest $p_{T,l1}^{\max}$, and we see that the p_T distributions are very similar to the ones in model 2, only moved towards larger values. Due to the large W_R mass in model 5, the events are more energetic which gives larger p_T on average, but since the upper limit on $p_{T,l1}$ is so large, l_1 still tends to be the leading lepton in the events. The distributions of model 6 are relatively similar to those in model 3, which also makes sense since $p_{T,l1}^{\max}$ is slightly larger in model 6 than model 3, but M_{W_R} is also larger. So in model 6, we see that l_1 again can have either larger or smaller p_T than l_2, but it seems slightly more likely to be smaller in this case. From all the plots in figure 6.22, we see a promising way of determining a relationship between the masses of W_R and N in an experiment. If the distributions of $p'_{T,li}$ show a clear edge, the approximate p_T value of this edge can be inserted into equation (6.7) to get a value for the squared mass difference divided by the W_R mass. In particular, if either M_{W_R} or m_N is already known, this method can be used to determine the other mass.

Finally, we note that in figure 6.22 and many of the following figures it appears that the simulations of models 5 and 6 have much lower statistics than the other four, as their distributions fluctuate a lot. This is again connected to the difficulties of simulating events with very high masses, as we also saw when computing the cross sections of models 5 and 6. In particular, we note that for models 5 and 6, as much as 30% of events get negative weights from the MC simulation, while model 1, for example, has less than 5% negative weights. As mentioned in section 5.1, a large fraction of negative weights causes a decrease in statistics. To get similar results for models 5 and 6 as for the other four would therefore require simulating more events.

In figure 6.23, we show the transverse momenta of the two leading jets in each event ordered by their p_T. Neither of the jets have any upper limit on their transverse momentum, like l_1 does, and we see that the distributions are therefore not as sensitive to changes in the heavy neutrino mass. It is the mass of the W_R that is most directly linked to the energy of the events, since it is the heaviest particle that is being produced. We therefore see that both $p'_{T,ji}$ distributions are shifted towards larger values when the W_R mass increases. But within the pairs of models that have the same M_{W_R}, the distributions have very similar shapes, while extending only slightly more towards larger values in the models where m_N is larger. For example, comparing models 1 and 2, we see that in model 1, the second jet never has a p_T greater than around 750 GeV, while in model 2 it extends to around 900 GeV. Interestingly, in all the models, the first bin

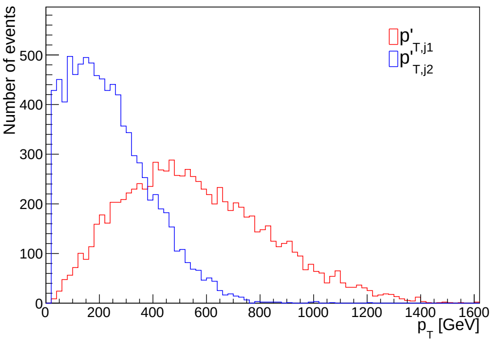

(a) Model 1

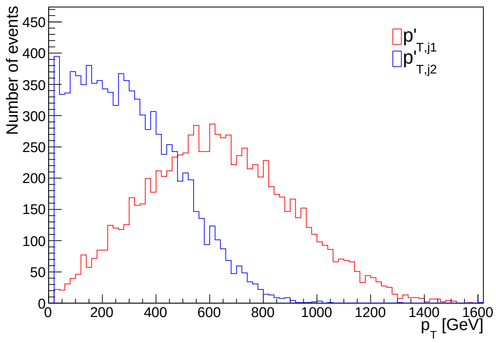

(b) Model 2

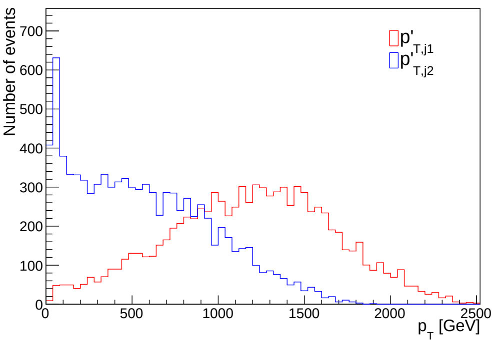

(c) Model 3

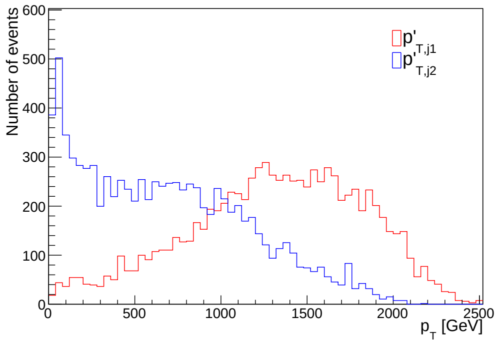

(d) Model 4

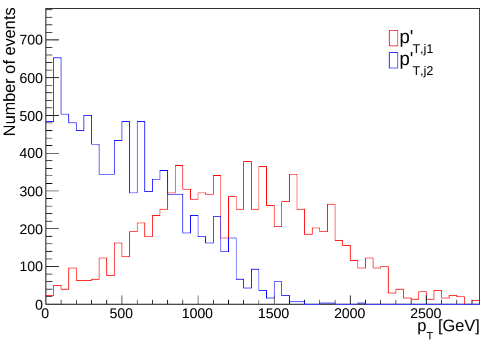

(e) Model 5

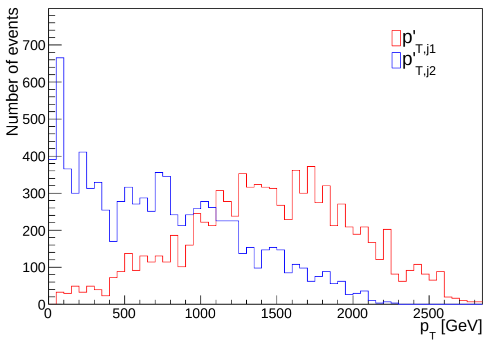

(f) Model 6

Figure 6.23: The distributions of the transverse momenta of the two leading signal jets in each event, ordered by p_T, for the different simulated LRSM models.

of the $p'_{T,l2}$ distribution is very low, which is probably due to the p_T cut of 20 GeV on jets. Then the second bin is the highest or one of the highest in the entire distribution, meaning that the second jet has relatively low transverse momentum in many events.

In figure 6.24, we show the $\Delta\phi$ between l_1 and l_2, and between j_1 and j_2. We call these $\Delta\phi_{ll}$ and $\Delta\phi_{jj}$ respectively. Comparing the distributions in model 1, where the relative difference between the masses of W_R and N is the largest, to those in model 4, where the relative difference is smallest, we see something interesting. It seems like when the mass difference is large, $\Delta\phi_{jj}$ is approximately uniformly distributed, while $\Delta\phi_{ll}$ is skewed towards large values, while when the mass difference is small, $\Delta\phi_{jj}$ is skewed towards large values and $\Delta\phi_{ll}$ is approximately uniformly distributed. This can be explained using similar geometric arguments as we used in section 6.2 for the process with only SM W bosons and leptonic decay. When the mass difference between W_R and N is large (and in particular large compared to M_{W_R}), we see from equation (6.7) that the heavy neutrino will get a sizeable momentum compared to its mass. Therefore the heavy neutrino is boosted in the transverse direction, and its decay products will tend to move in the same direction as it, which explains why $\Delta\phi_{ll}$ tends to be large in this case. Now, the W_R^* boson coming from the heavy neutrino decay can in principle have any invariant mass that is less than m_N, since it is off-shell. We also know that a large invariant mass means that the W_R^* gets a small velocity, while a small invariant mass gives a large velocity. The W_R^* therefore gets an arbitrary amount of boost in the transverse direction, which explains why the two jets have an almost uniform distribution of $\Delta\phi$. In the case when the mass difference between W_R and N is small compared to M_{W_R}, we know that N gets a maximum transverse momentum that is approximately equal to the mass difference, and therefore much smaller than m_N. The transverse velocity of N will therefore be very small, and l_2 and W_R^* will be back-to-back in the transverse plane, with an arbitrary orientation. Therefore, $\Delta\phi_{ll}$ becomes approximately uniformly distributed. The invariant mass of the off-shell W_R^* is still limited by m_N, but this is now very close to the true mass of the W_R. The fact that the $\Delta\phi_{jj}$ distribution becomes skewed towards larger values indicates that the W_R^* tends to get a small transverse velocity, such that j_1 and j_2 are almost back-to-back. Equivalently, this means that the W_R^* tends to get a large invariant mass compared to its momentum when the mass difference between the neutrino and W_R is small. We do not have a good explanation for why this is the case. In models 2, 3, 5 and 6 we see some intermediate situations, where both distributions are relatively close to uniform, but both are slightly skewed towards large values.

In figure 6.25, we show the distributions of the invariant masses of both the lepton pair l_1 and l_2, and the jet pair j_1 and j_2, which we call m_{ll} and m_{jj} respectively. Here, m_{jj} is the easiest to interpret physically, since it is equal to the invariant mass of the W_R^* boson. We therefore know that it has an upper limit of m_N, when the heavy neutrino is on-shell. We see that m_N coincides very closely with the end of the m_{jj} distribution in all six plots, so this does seem to be the case. We see that in model 4, when m_N is closest to M_{W_R}, the distribution of m_{jj} is the most uniform. This might help to explain why $\Delta\phi_{jj}$ tends to take larger values in this model, since there are more events where the invariant mass of the W_R^* is close to m_N than in model 1 for example. These events must have a W_R^* with very low velocity, and therefore give large $\Delta\phi_{jj}$. The distributions of m_{ll} are more difficult to interpret, but we do know that they have an upper limit of M_{W_R}, since the invariant mass of the two leptons and two jets combined is equal to that of W_R. We

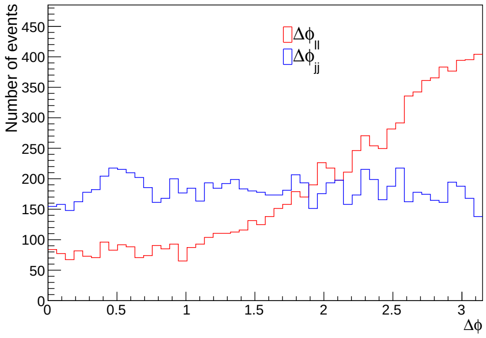

(a) Model 1

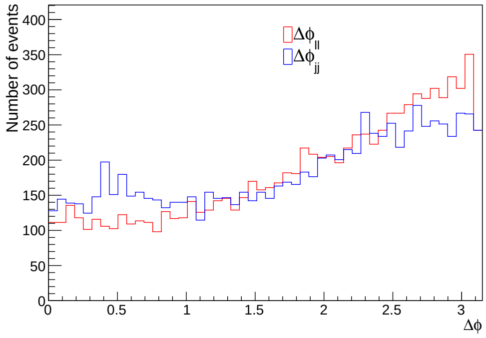

(b) Model 2

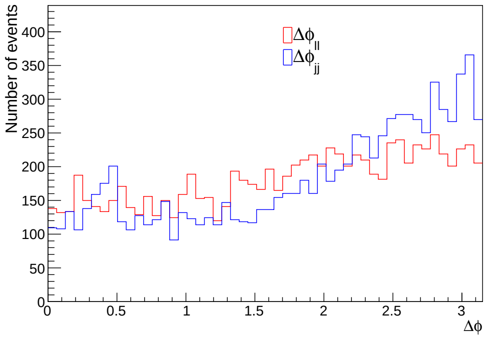

(c) Model 3

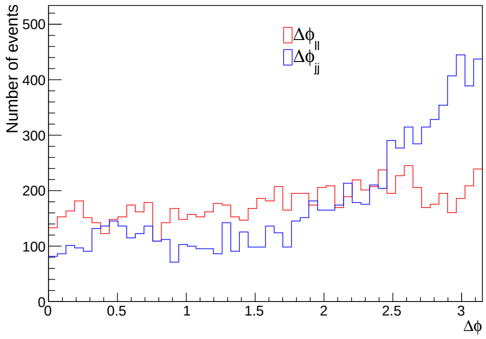

(d) Model 4

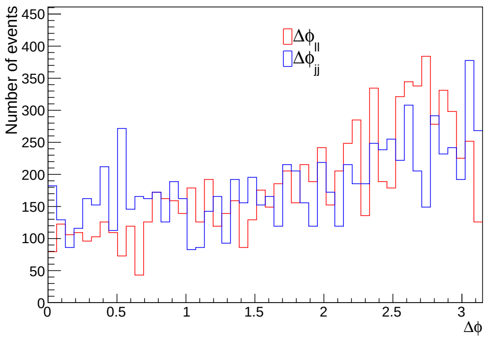

(e) Model 5

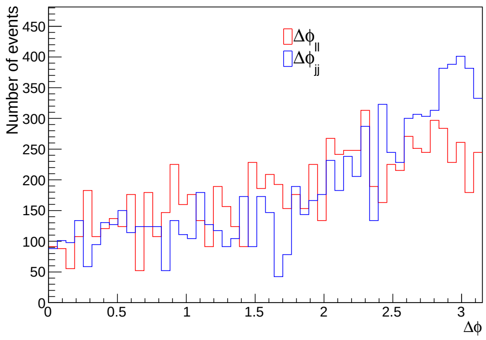

(f) Model 6

Figure 6.24: The distributions of $\Delta\phi_{ll}$ and $\Delta\phi_{jj}$, defined as the $\Delta\phi$ of the two leading signal leptons and two leading signal jets respectively, for the different simulated LRSM models.

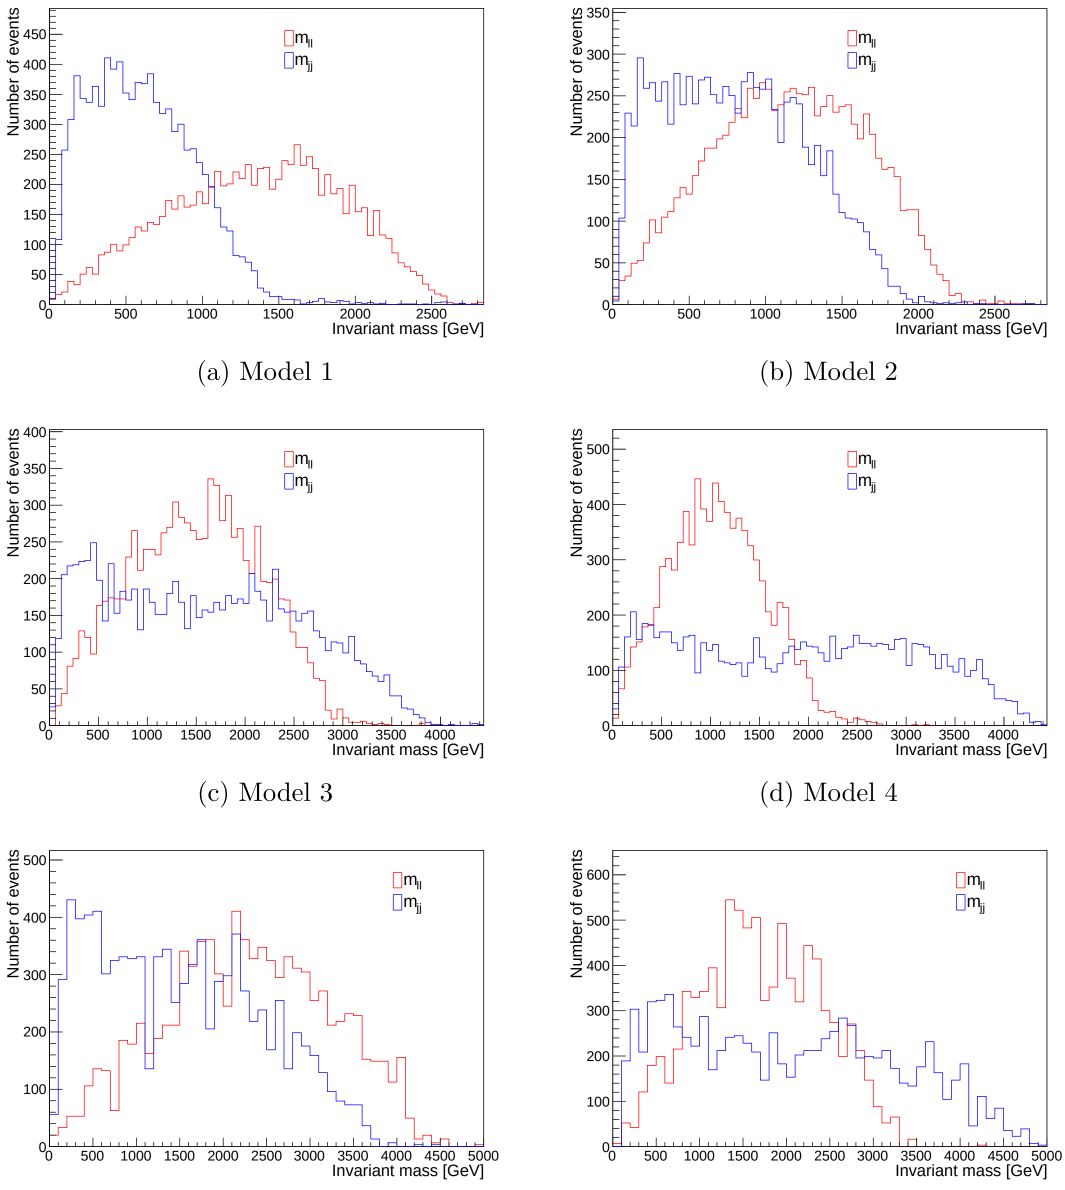

(a) Model 1

(b) Model 2

(c) Model 3

(d) Model 4

(e) Model 5

(f) Model 6

Figure 6.25: The distributions of m_{ll} and m_{jj}, defined as the invariant mass of the two leading signal leptons and two leading signal jets respectively, for the different simulated LRSM models.

also see that for each W_R mass, the model with the larger m_N tends to have smaller m_{ll}. This makes sense, since we know that at least l_1 becomes less energetic when Δm is small.

In figure 6.26, we show the distributions of the invariant mass of the system of the four particles l_1, l_2, j_1 and j_2, which we call m_{lljj}. Since these are all the final state decay products of the initial W_R boson, we know that m_{lljj} is equal to the invariant mass of the W_R, which again is equal to M_{W_R} when it is on-shell. We see that just as expected, all six distributions are sharply peaked around the respective W_R masses in each model. This shows that the W_R is close to on-shell in most events. The distributions do not depend much on the heavy neutrino mass, but there seems to be a tendency for the peaks to be a little bit wider in the plots on the right hand side, with larger m_N. In particular, the tail towards low values of m_{lljj} seems to get larger when m_N is larger. The m_{lljj} distribution from this process can clearly be used to determine the mass of the W_R boson. As we remember from before, the heavy neutrino mass could then be determined from the edges in the lepton transverse momenta distributions.

In figure 6.27, we show the distributions of the missing transverse energy in each event. In all six models, the vast majority of events have very low E_T^{miss}. In models 1 and 2, over half the events have $E_T^{\text{miss}} < 5$ GeV, and in models 5 and 6, most events have $E_T^{\text{miss}} < 40$ GeV. We clearly see that the number of events with large missing transverse energy increases with the mass of W_R (model 6 has around 1 400 events with $E_T^{\text{miss}} > 100$ GeV, while model 1 only has around 800), while the mass of N does not seem to affect the distributions much.

In figure 6.28, we show the distributions of the hadronic activity in each event. The distributions all have relatively similar shapes and the value of m_N does not seem to affect them much. They are however shifted towards larger values when M_{W_R} is increased, and in fact the rightmost edge of the distributions are quite close to M_{W_R} in all models. This makes sense, since M_{W_R} is an approximate measure of the energy in the event, and the total transverse momentum in the event can not be larger than the total energy in the event.

In figure 6.29, we show the distributions of the product of the charges of the leptons l_1 and l_2. Just like in the Majorana models without the W_R boson that we studied in section 6.2, the charge product is expected to be $+1$ and -1 in half the events each, due to the Majorana nature of the heavy neutrino. We can again compute the probability of the observed excess of either OS or SS events in each model, assuming that the number of OS events is binomially distributed. The results are shown in table 6.5. In this case, we see that not all the models have the excess in the same direction (some have more OS events than SS, and some less). Therefore, some of the probabilities in the table are greater than 0.5. In these cases, the interesting probability is actually the complementary probability $P(X \leq x) = 1 - P(X \geq x)$. Replacing each probability that is larger than 0.5 with this complementary probability, we see that the smallest one is 0.05. This is relatively small, but not so small that it is unreasonable, especially since we performed six different simulations. Since there does not seem to be any correlation across the six simulations, we conclude that the simulations work correctly, and that the SS/OS nature of each event is in fact chosen with equal probability.

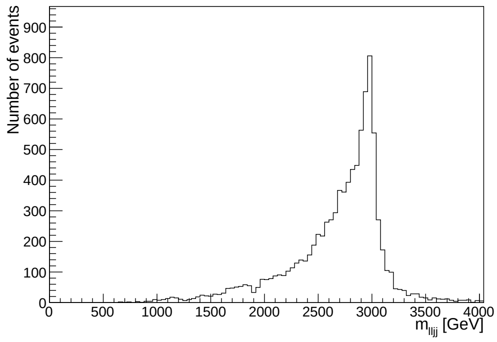

(a) Model 1

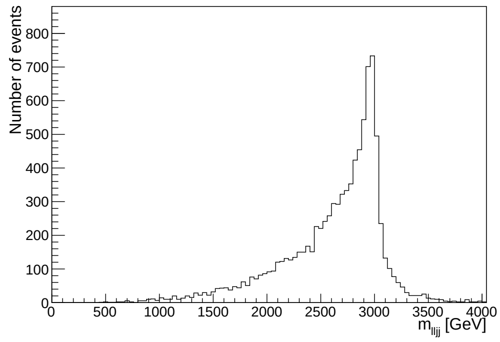

(b) Model 2

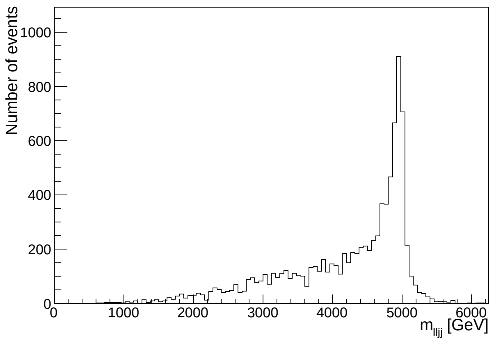

(c) Model 3

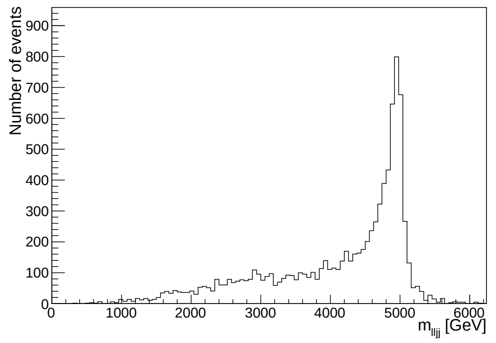

(d) Model 4

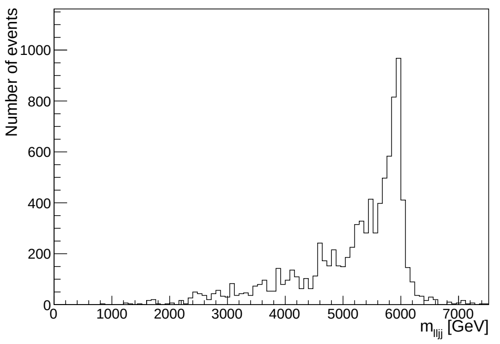

(e) Model 5

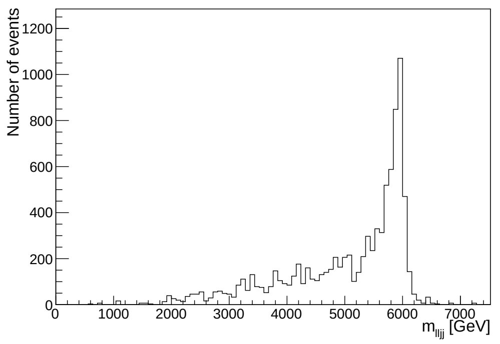

(f) Model 6

Figure 6.26: The distributions of the invariant mass of the two leading signal leptons and two leading signal jets combined in each event, which we call m_{lljj}, for the different simulated LRSM models.

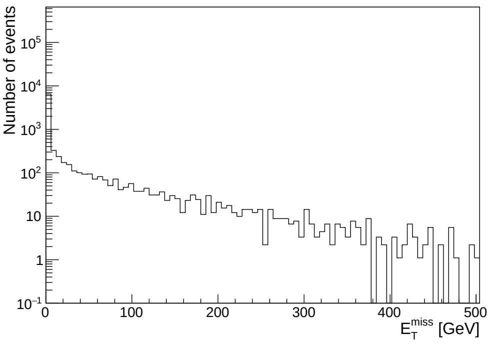

(a) Model 1

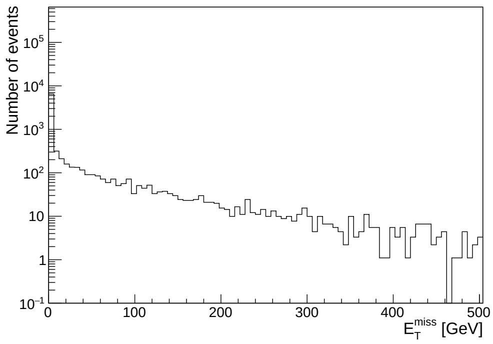

(b) Model 2

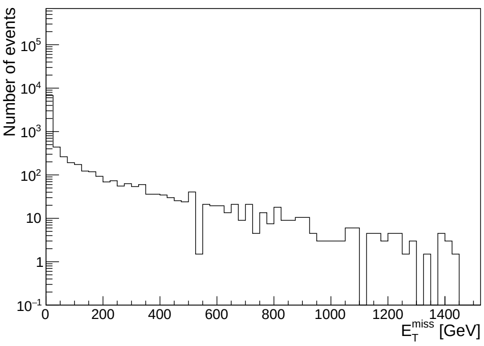

(c) Model 3

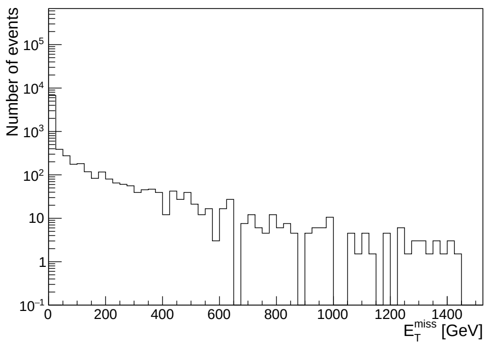

(d) Model 4

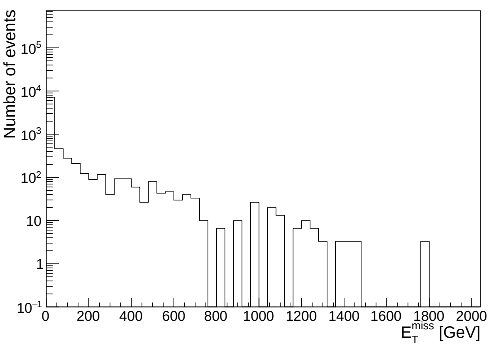

(e) Model 5

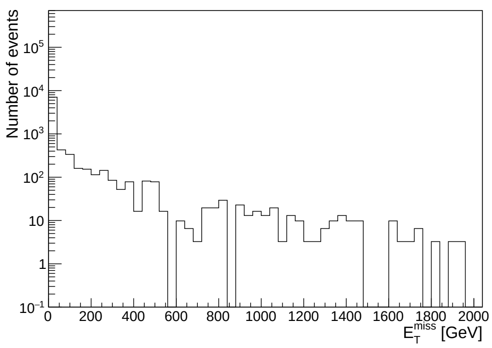

(f) Model 6

Figure 6.27: The distributions of the missing transverse energy E_T^{miss} in each event for the different simulated LRSM models.

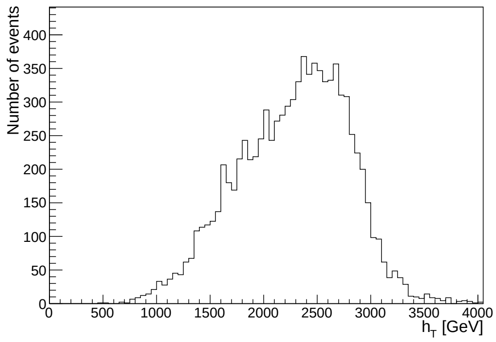

(a) Model 1

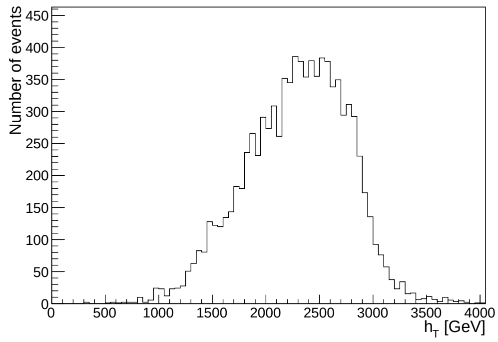

(b) Model 2

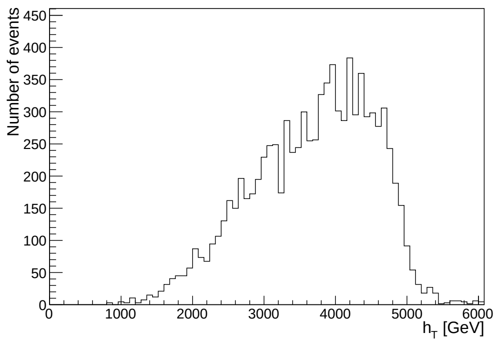

(c) Model 3

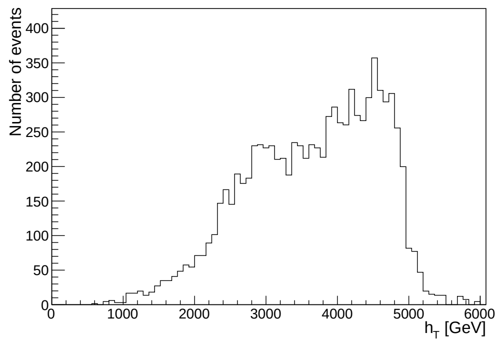

(d) Model 4

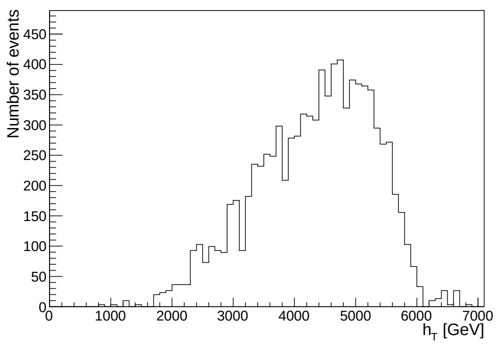

(e) Model 5

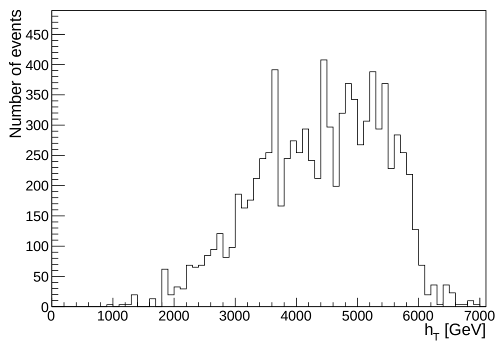

(f) Model 6

Figure 6.28: The distributions of the hadronic activity h_T of all leptons and jets that passed the cuts in each event for the different simulated LRSM models.

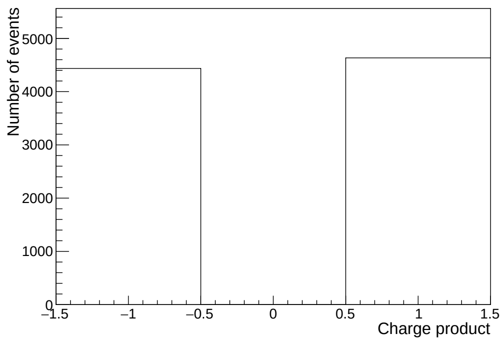

(a) Model 1

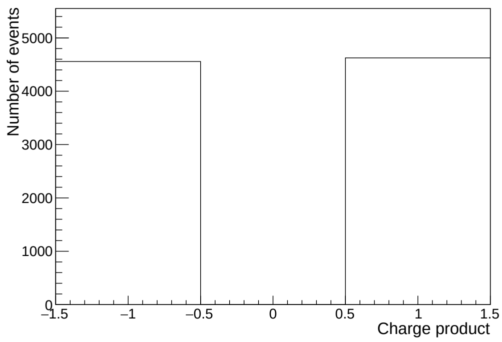

(b) Model 2

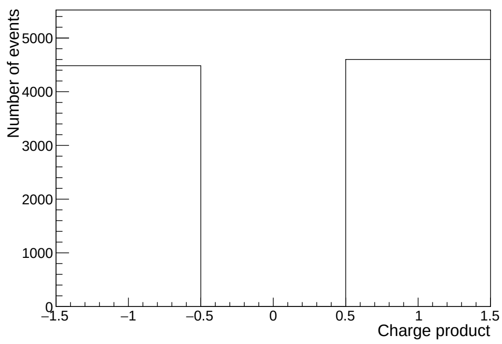

(c) Model 3

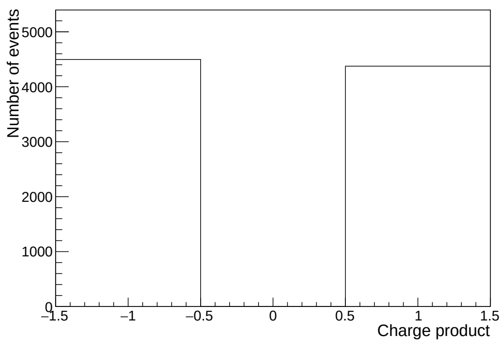

(d) Model 4

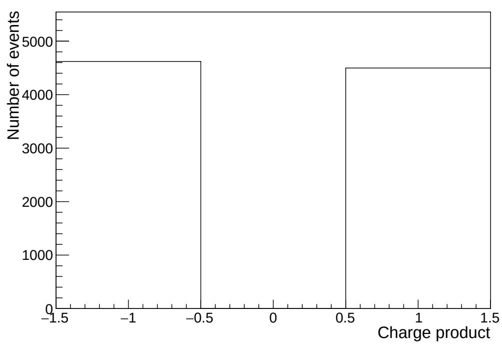

(e) Model 5

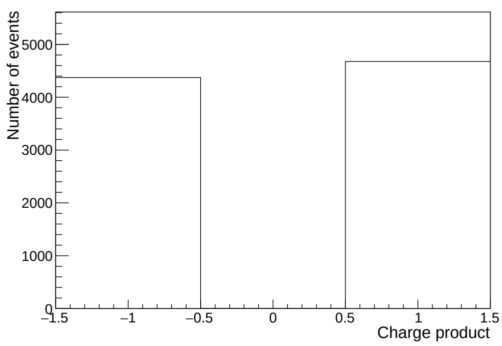

(f) Model 6

Figure 6.29: The distributions of the product of the charges of l_1 and l_2 in each event for the different simulated LRSM models. This shows whether the leptons have the same charge (+1) or the opposite charge (−1).

Table 6.5: The number of events x where the leptons l_1 and l_2 have opposite sign out of the n events that were simulated and passed cuts for each heavy neutrino mass value. The last column gives the probability of seeing this many opposite sign events or more, assuming that x follows the binomial distribution (6.3).

Model number	n	x	$P(X \geq x)$
1	9 070	4 455	0.95
2	9 182	4 569	0.67
3	9 084	4 533	0.57
4	8 871	4 498	0.09
5	9 119	4 597	0.21
6	9 051	4 483	0.81

Chapter 7

Data analysis

In this chapter, we will compare the simulated signals that we studied in the previous chapter to Standard Model backgrounds. For the background we use Monte Carlo generated samples of two lepton events at 13.6 TeV from the official SM background MC production campaign for Run 3 ATLAS analyses. These samples are in the same format as our signal samples, so we can again use `SimpleAnalysis` to perform cuts on signal and background. Now we will also use smearing to mimic the full ATLAS reconstruction, meaning that we need to perform identification of electrons and muons in the MC data, just like we would need in an analysis of real ATLAS data.

7.1 Discovery significance and exclusion

For the different signal models, we want to create a few so-called signal regions. Signal regions are regions in the space of kinematic variables where we expect to have a high sensitivity to a certain signal. This usually means a region where the fraction of signal events to background events is maximised. Once a signal region is chosen, we put all background and signal events through the corresponding cuts on the kinematic variables. In order to say something meaningful about the validity of the signal model, the number of signal and background MC events remaining should be compared to the number of events found in the same signal region in an experiment. However, even in the absence of real experimental data, we can perform a similar analysis by assuming certain outcomes of a hypothetical experiment. In any case, we need to use statistical methods to quantify the agreement or non-agreement between the simulated data and the experimental data. To do this, we call the number of MC simulated background events in the signal region the "expected background" b, and we call the number of simulated signal events in the signal region the "expected signal" s. The number of events in the signal region from the experiment, either real or hypothetical, we call the "number of observed events" N_{obs}. Now, there are two types of statistical analyses one can perform, depending on the (assumed) outcome of the (hypothetical) experiment. If N_{obs} is larger than the expected background b, this indicates a possible discovery of the signal process. Therefore, the analysis one wants to perform in this case is to quantify the significance of the discovery. The other possibility is that N_{obs} is smaller than the expected number of events from signal and background combined, which is $s + b$. This would indicate that the signal does not exist, and in this case statistics is used to quantify the confidence with which the hypothesis of the existence of the signal can be excluded.

7.1.1 Discovery significance

We first look at how the discovery significance can be estimated. The significance of an excess is connected to the probability of seeing that excess or more (this probability is called the p-value p) by random chance under the assumption that the signal does not exist (this assumption is often called the background hypothesis). The less likely it is to see the observed excess of events, the more significant the excess is said to be. Instead of the p-value, significance is commonly measured in standard deviations σ, which refers to the standard deviation of a normal distribution. For a given p-value p, the significance in units of σ will be the number of standard deviations above the mean of a normal distribution that would result in an integral of the normal distribution above this value that is equal to p. If the normal distribution has a mean of $\mu = 0$, the relation between p and a significance of $Z\sigma$ becomes

$$p = \int_{Z\sigma}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} e^{-x^2/(2\sigma^2)}\,\mathrm{d}x = \int_{Z}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2}\,\mathrm{d}y = 1 - \Phi\left(Z\right), \tag{7.1}$$

where we performed a change of variables to remove the standard deviation from the integral, and $\Phi\left(x\right)$ is the cumulative distribution function of the standard normal distribution. Inverting the expression, we get

$$Z = \Phi^{-1}\left(1 - p\right). \tag{7.2}$$

A significance of 5σ is usually required in order to claim a discovery in particle physics, which corresponds to a p-value of $p = 2.87 \times 10^{-7}$. With a significance of 5σ it is therefore highly unlikely that the observed excess would occur under the background hypothesis.

In the absence of an actual experiment, it is common to compute what is called the "expected significance" as the mean or median value of Z when assuming some signal model with an expected number of signal events s. As is shown in the note *Discovery sensitivity for a counting experiment with background uncertainty* by Cowan [66], the expected significance in a particle physics experiment with a large background with negligible systematic uncertainty is approximated by

$$Z = \frac{s}{\sqrt{b}}. \tag{7.3}$$

When the systematic uncertainty in the background, which we call σ_b, is not negligible (which it is usually not in particle physics experiments) it reduces the expected significance of the excess. The systematic uncertainty is added to the statistical uncertainty $\sqrt{b}$ in quadrature, giving

$$Z = \frac{s}{\sqrt{b + \sigma_b^2}}. \tag{7.4}$$

In the note by Cowan, modifications to the expressions in equations (7.3) and (7.4) are also shown, which give more accurate results in the case when s is not much smaller than b. The modified expressions for the expected significance are

$$Z = \sqrt{2\left[\left(s + b\right)\ln\frac{s + b}{b} - s\right]}, \tag{7.5}$$

for the case of negligible systematic uncertainty in the background, and

$$Z = \sqrt{-2\left((s+b)\ln\left[\frac{s+b+\frac{b^2}{\sigma_b^2}}{\left(1+\frac{b}{\sigma_b^2}\right)(s+b)}\right]+\frac{b^2}{\sigma_b^2}\ln\left[\frac{\frac{b}{\sigma_b^2}\left(s+b+\frac{b^2}{\sigma_b^2}\right)}{\left(1+\frac{b}{\sigma_b^2}\right)\frac{b^2}{\sigma_b^2}}\right]\right)} \tag{7.6}$$

for the case with systematic uncertainty σ_b. If an actual experiment is performed, and $N_{\mathrm{obs}} > b$ events are observed, the "observed significance" is given by similar expressions to the expected significance, only with s replaced by $N_{\mathrm{obs}} - b$.

7.1.2 Exclusion

When we want to quantify the certainty with which we can exclude a signal model based on an observation of $N_{\mathrm{obs}} \approx b < s + b$, we could again do this by calculating a p-value. In this case, it would be the probability of seeing N_{obs} events or less in the experiment, given that N_{obs} is drawn from a distribution with an expectation value of $s + b$ (what is called the signal+background hypothesis). This is the so-called "frequentist" approach to exclusion. However, this approach has some drawbacks. The frequentist approach allows us to exclude signals that the experiment is not sensitive to, since any experiment that yields $N_{\mathrm{obs}} = 0$ would give a small p-value for any signal model with a large s. Also, the probability we are most interested in is not the probability of seeing the data given the signal model, which we would write in mathematical notation as $P\,(\mathrm{Data} \mid \mathrm{Model})$. Instead, what we really want to know is the opposite conditional probability, namely the probability of the signal model being correct given the observed data, or $P\,(\mathrm{Model} \mid \mathrm{Data})$. This probability is not defined in the frequentist approach, since the signal model is not the outcome of an experiment that can be repeated many times. We will therefore use so-called "Bayesian" statistics to perform exclusions instead. This is based on Bayes' theorem, which relates the two conditional probabilities in the following way

$$P\,(\mathrm{Model} \mid \mathrm{Data}) = \frac{P\,(\mathrm{Data} \mid \mathrm{Model})\,P\,(\mathrm{Model})}{P\,(\mathrm{Data})}. \tag{7.7}$$

We see that we can use Bayes' theorem to calculate the probability $P\,(\mathrm{Model} \mid \mathrm{Data})$, which in Bayesian statistics is called the "posterior", from the probability $P\,(\mathrm{Data} \mid \mathrm{Model})$, which is called the "likelihood" $\mathcal{L}$. The posterior probability also depends on the probability of the observed data, $P\,(\mathrm{Data})$. Since this is independent of the model, it can be treated as a normalisation constant. Finally, there is the probability $P\,(\mathrm{Model})$, which is called the "prior" probability, as it states what we believe the probability distribution to be before performing the experiment. In the particle physics application, the prior says something about our belief about the signal model before we have seen any data. It is not easy to define such a prior, and even more difficult to justify any particular choice. If we define the signal model in terms of the expected number of signal events s, the most common choice of prior in particle physics is to assume a uniform distribution for $s \geq 0$, and zero probability for $s < 0$. Then the only part of the right hand side of equation (7.7) that is dependent on s is the likelihood. Since a particle physics experiment is essentially a counting experiment, the distribution of n, the random variable from which N_{obs} is sampled, is usually assumed to be the Poisson distribution. In that case, the likelihood of observing N_{obs} events under the signal+background hypothesis becomes

$$\mathcal{L}\,(n = N_{\mathrm{obs}} \mid \lambda = s + b) = \frac{(s+b)^{N_{\mathrm{obs}}}\,e^{-(s+b)}}{N_{\mathrm{obs}}!} \tag{7.8}$$

Then we can use Bayes' theorem to state the probability of the signal model with s expected events given the experiment where only N_{obs} (we think of N_{obs} as being smaller than $s+b$) events were observed as

$$P\left(s \mid N_{\text{obs}}\right) = N\left(s+b\right)^{N_{\text{obs}}} e^{-(s+b)}, \tag{7.9}$$

where N is a normalisation constant that incorporates both the uniform prior probability distribution, the constant in the denominator of Bayes' theorem and the $N_{\text{obs}}!$ constant from the likelihood. It normalises the probability such that

$$\int_0^{\infty} P\left(s \mid N_{\text{obs}}\right) \mathrm{d}s = 1. \tag{7.10}$$

If a signal model gives a very small probability $P\left(s \mid N_{\text{obs}}\right)$, we want to exclude the existence of that signal based on the observation. To define the threshold for when a signal is said to be excluded, we calculate the probability $P\left(s \geq s_{\text{up}} \mid N_{\text{obs}}\right)$ of a signal model giving more expected events than some "exclusion limit" s_{up}. Instead of a lower threshold on the probability $P\left(s \geq s_{\text{up}} \mid N_{\text{obs}}\right)$, the exclusion threshold is usually stated in terms of a Confidence Level (CL), defined as

$$\text{CL} = 1 - P\left(s \geq s_{\text{up}} \mid N_{\text{obs}}\right) = \int_0^{s_{\text{up}}} P\left(s \mid N_{\text{obs}}\right) \mathrm{d}s\,. \tag{7.11}$$

Any signal model that gives $s > s_{\text{up}}$ is then said to be excluded at the given confidence level. We see that a large CL corresponds to a small probability of a signal with as many as s_{up} expected events. A commonly used threshold is CL$= 0.95$, and signals with $s > s_{\text{up}}$ are then said to be excluded at 95% CL.

To be more precise, the exclusion limit s_{up} in equation (7.11) is what we call the "observed exclusion limit", since it is based on a real experiment that observed N_{obs} events. Just like when we calculated significances, we can also define the "expected exclusion limit", which is the exclusion limit we would get from a hypothetical experiment with $N_{\text{obs}} = b$. The expected exclusion limit at 95% CL is therefore defined by

$$\int_0^{s_{\text{up}}} P\left(s \mid b\right) \mathrm{d}s = 0.95. \tag{7.12}$$

In particle physics, it is often more useful to put exclusion limits on cross sections rather than the number of signal events, since s will depend on specifics of the analysis, like the cuts applied to events. In order to translate the exclusion limit to the cross section, we simply rewrite the expected value of the Poisson distributed likelihood of equation (7.8) as

$$\lambda = L\epsilon_s\sigma + b, \tag{7.13}$$

where L is the integrated luminosity of the (hypothetical) experiment, ϵ_s is the signal efficiency defined as the fraction of the total number of signal events that pass all the cuts and end up in the signal region, and σ is the signal cross section. Then the 95% CL expected exclusion limit σ_{up} on the cross section is given by

$$N \int_0^{\sigma_{\text{up}}} \left(L\epsilon_s\sigma + b\right)^b e^{-(L\epsilon_s\sigma+b)} \mathrm{d}\sigma = 0.95. \tag{7.14}$$

Finally, we need to take into account the fact that the values of L, ϵ_s and b have uncertainties. Incorporating these uncertainties is slightly more complicated than in the frequentist approach. Here, we will treat the uncertain variables as additional parameters, often called "nuisance parameters", of the likelihood with their own probability distributions. In order to remove the posterior's dependence on the nuisance parameters, we need to integrate over their probability distributions, giving the so-called marginalised probability

$$P(\sigma \mid b) = N \int \mathcal{L}(\sigma, L, \epsilon_s, b)\, P(L)\, P(\epsilon_s)\, P(b)\, \mathrm{d}L\, \mathrm{d}\epsilon_s\, \mathrm{d}b\,, \tag{7.15}$$

and the 95% CL expected exclusion limit on the cross section σ_{up} is given by

$$N \int_0^{\sigma_{\text{up}}} \left[\int (L\epsilon_s\sigma + b)^b\, e^{-(L\epsilon_s\sigma+b)} P(L)\, P(\epsilon_s)\, P(b)\, \mathrm{d}L\, \mathrm{d}\epsilon_s\, \mathrm{d}b \right] \mathrm{d}\sigma = 0.95. \tag{7.16}$$

The probability distributions $P(L)$, $P(\epsilon_s)$ and $P(b)$ of the nuisance parameters are typically taken to be Gaussian distributions with means equal to the values found in the (MC) experiment, and standard deviations equal to the estimated uncertainties in each parameter. The integral in equation (7.16) is typically computed using Monte Carlo integration. This involves sampling values for each nuisance parameter from their respective probability distributions using a pseudo-random number generator. For each set of nuisance parameter values, the expected exclusion limit can easily be calculated, since the integral reduces to a one-dimensional one. We call each such calculation of an expected limit using a set of nuisance parameter values a "pseudo-experiment". The final value for the expected exclusion limit σ_{up} can then be computed by taking the mean of all the expected limit results from each pseudo-experiment, since taking the mean is equivalent to integrating over the probability distributions. A benefit of computing the limit in this way is that it allows us to quantify the uncertainty in the result from the spread in the results from all pseudo-experiments. Like before, we want to express the uncertainty in terms of the standard deviations of the normal distribution. For the standard normal distribution, the 1σ deviations from the mean are defined by $\Phi(-1) \approx 0.159$ and $\Phi(1) \approx 0.841$. We therefore call the expected limits that are larger than approximately 15.9% and 84.1% of all the limits from the other pseudo-experiments the $\pm 1\sigma$ deviations from the mean expected exclusion limit. Similarly, the expected limits that are larger than 2.28% and 97.7% of all the other limits are called the $\pm 2\sigma$ deviations of the expected exclusion limit.

7.2 Majorana and pseudo-Dirac model

In this section, we compare the MC background to the Majorana and pseudo-Dirac models, using the three-lepton final state that we studied extensively in chapter 6. For our signal models, we choose the same type of models as in section 6.2, meaning models with one Majorana or pseudo-Dirac neutrino with masses in the range 200 GeV to 1200 GeV, and mixing matrix elements

$$|V_{e1}|^2 = 10|V_{\mu 1}|^2 = 100|V_{\tau 1}|^2. \tag{7.17}$$

We now choose $V_{e1} = 0.1$, which is actually slightly optimistic considering the limits given in equation (3.51). Unfortunately, we did not manage to implement smearing on these signal models, which might be due to some incompatibility between the versions of `SimpleAnalysis` that we use and the `Athena` software used to generate truth files.

Table 7.1: The requirements used to define leptons and jets in the analysis of signal and background with the three lepton final state, without smearing.

	Electrons	Muons	Jets
p_T	> 5	> 5	> 20
$\lvert\eta\rvert$	$(0, 1.37]$ or $[1.52, 2.47]$	< 2.5	< 2.5

Table 7.2: The cuts made on kinematic variables to define the signal regions for the Majorana and Dirac models with a three lepton final state. The cuts are given in terms of the requirement that must be fulfilled for an event to pass. See the main text for the definitions of the variables.

Variable	Cut	Variable	Cut
n_{leptons}	$= 3$	$m'_{ll,1}$ [GeV]	> 100
$p'_{T,1}$ [GeV]	> 70	$m'_{ll,2}$ [GeV]	> 95
$p'_{T,2}$ [GeV]	> 50	$m'_{ll,3}$ [GeV]	> 50
$p'_{T,3}$ [GeV]	> 15	$\left\lvert m'_{ll,3} - m_Z \right\rvert$ [GeV]	> 3
$m'_{T,1}$ [GeV]	> 60	$\Delta\phi'_{ij}$	> 0.1
$m'_{T,23}$ [GeV]	> 80	h_T [GeV]	> 250
E_T^{miss} [GeV]	> 20		

The object definitions are therefore very simple, since we effectively work directly on the truth information. Nevertheless, the definitions for this analysis are given in table 7.1. Since tau leptons would require a completely different type of analysis in an experiment, we will not consider models in which the heavy neutrino couples most strongly to the tau. Therefore we will also not count taus as leptons for this analysis, as is common practice in experimental particle physics[1]. Although we are using truth objects for this analysis, we implemented overlap removal using the same parameters that we presented in section 6.1. This makes the analysis somewhat more realistic.

The next, and maybe most important, step is to choose the cuts used to define the signal regions. The ability of our cuts to remove as many background events as possible while keeping as many signal events as possible will decide the significance with which we can discover a signal and the confidence with which we can exclude one. We will use only what are known as rectangular cuts, meaning cuts that can be stated as ”some variable must be smaller/larger than some value”. We make two signal regions for the analysis of the three lepton final state. The only difference between the two regions is that we require the charges of the leptons that are identified as l_1 and l_2 in each event to have either the same or opposite signs. We call the two regions Same Sign (SS) and Opposite Sign (OS) respectively. All the other cuts defining signal regions SS and OS are identical, and these are listed in table 7.2. To choose the cuts in table 7.2, we plotted

[1]There are usually dedicated τ analyses in parallel with electron and muon analyses, but for this case we will leave this for future research.

the kinematic distributions of the model with a 200 GeV Majorana neutrino and the model with a 800 GeV Majorana neutrino along with the SM background before passing any other cuts than requiring at least three leptons (electrons and muons) according to the object definitions in table 7.1. All of these plots are shown in appendix A.1. The variable n_{leptons} in table 7.2 is the number of leptons in each event. The variables $p'_{T,i}$ are as before the transverse momenta of the leptons l'_i, where the prime indicates that i refers to the ordering by p_T. The same applies to the variables $m'_{T,1}$, $m'_{T,23}$ and $m'_{ll,i}$. These all have the same definitions as in section 6.2, but with the ordering by vertex replaced by the ordering by p_T, which is the only ordering we can know with confidence in an experiment. The variables $\Delta\phi'_{ij}$ with $i, j = 1, \ldots, 3$ and $i \neq j$ are the $\Delta\phi$ between the charged leptons number i and j when ordered by p_T.

After applying all the cuts of table 7.2, we create the SS and OS signal regions in order to increase sensitivity to the Majorana and pseudo-Dirac models respectively. For this we need to choose which two leptons in each event to use in the sign comparison. To get the optimal sensitivity to the pseudo-Dirac model and to possibly separate the Majorana and pseudo-Dirac models in data containing a signal of unknown type, we want to choose the leptons l_1 and l_2 from the Feynman diagram in figure 5.4 in each event, since it is this pair of leptons that always have opposite signed charges in the pseudo-Dirac model. To identify this pair in each event we can use the method we developed in section 6.4. We expect this to give the best results for the models we look at, but the method is also highly specialised to these specific models. A more general approach, which we also expect to give reasonable results is to simply use the pair of leptons with the largest p_T in each event. We will therefore perform the same analysis for both methods, and we will call the p_T method the "baseline" method, which defines the baseline OS and SS signal regions, and our identification method using the product $p_T \times \Delta\phi$ we call the "specialised" method defining the specialised OS and SS signal regions. To see the difference between the two methods, we plot the distributions of the product of charges of the pairs of leptons chosen by both the baseline and specialised methods for signal and background processes. The plots are shown in figures 7.1 and 7.2. The signal models shown in these plots are the Majorana and pseudo-Dirac models with 1200 GeV, which we expect to give the largest improvement for the specialised method over the baseline. We see from figure 7.1 that the SS region increases the ratio of the Majorana signal to the background significantly for both methods, but the improvement is slightly larger for the specialised SS region. Another important thing to note is that the specialised SS region contains more signal events than the baseline SS region. In figure 7.2, we see a similar situation for the pseudo-Dirac signal. Both methods give increased s/b ratio in the OS region, but the specialised method gives a slightly higher ratio. Since the number of signal events in the baseline OS region is already quite large, the relative improvement in this region is small. In the SS region, on the other hand, we see that the specialised method reduces the number of pseudo-Dirac signal events by a factor of around 10, while the background remains approximately the same.

To quantify the effectiveness of the cuts we made for the different OS and SS signal regions, we look at the number of background and signal events that passed the cuts for each signal region. We call the fraction of events that passed the cuts (including the requirements of the object definitions) the efficiency ϵ of the cuts. The ideal scenario has a very small efficiency for the background events, which we call ϵ_b, and an efficiency close to 1 for signal events, which we call ϵ_s. Both the number of events remaining in each signal region when scaled to an integrated luminosity of 300 fb^{-1} and the efficiencies of

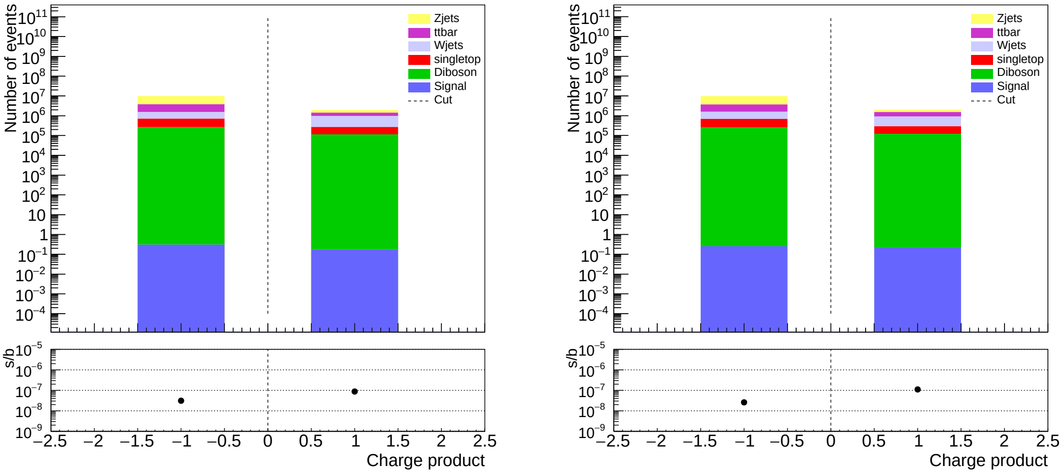

(a) Lepton pair identified by p_T (b) Lepton pair identified by $p_T \times \Delta\phi$

Figure 7.1: The distributions of the charge product of pairs of leptons in each event for the full SM background, and a signal with a 1200 GeV Majorana type heavy neutrino. The difference between the two plots is that in the left plot, the lepton pair was chosen as the two leading leptons in each event, while in the plot on the right, the pair was chosen using the vertex identification method we have developed.

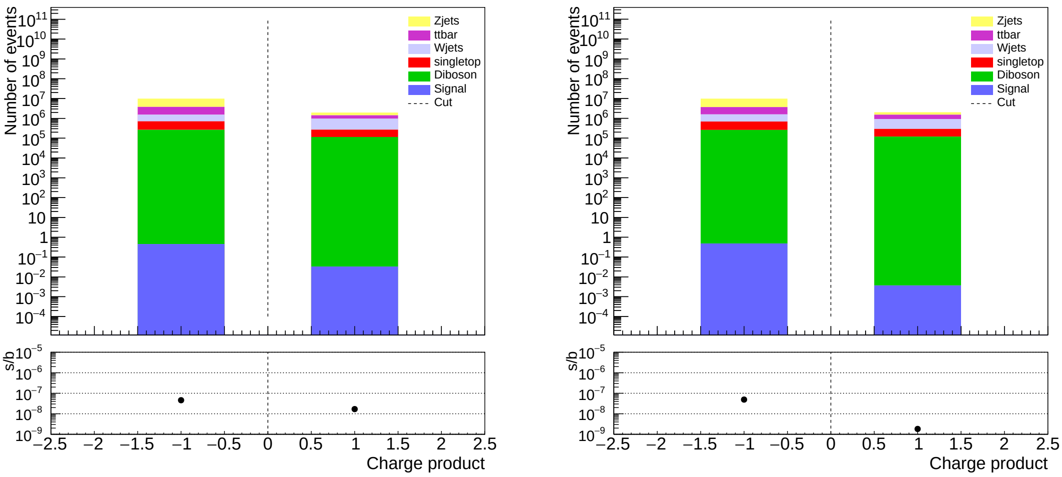

(a) Lepton pair identified by p_T (b) Lepton pair identified by $p_T \times \Delta\phi$

Figure 7.2: The distributions of the charge product of pairs of leptons in each event for the full SM background, and a signal with a 1200 GeV pseudo-Dirac type heavy neutrino. The difference between the two plots is that in the left plot, the lepton pair was chosen as the two leading leptons in each event, while in the plot on the right, the pair was chosen using the vertex identification method we have developed.

Table 7.3: The cut efficiencies on the background and signal events for the signal regions for the three lepton final state. The number of events remaining in each region for an integrated luminosity of 300 fb^{-1} is also shown.

Signal region	Type of events	ϵ	n_{events} (300 fb^{-1})
Baseline SS	Background	1.7×10^{-7}	3 225
	Majorana, $m_N = 200$ GeV	0.057	50
	Majorana, $m_N = 400$ GeV	0.16	12
	Majorana, $m_N = 800$ GeV	0.20	1
	Majorana, $m_N = 1200$ GeV	0.20	≈ 0.1
Baseline OS	Background	7.2×10^{-7}	13 966
	pseudo-Dirac, $m_N = 200$ GeV	0.087	75
	pseudo-Dirac, $m_N = 400$ GeV	0.36	25
	pseudo-Dirac, $m_N = 800$ GeV	0.49	2
	pseudo-Dirac, $m_N = 1200$ GeV	0.53	≈ 0.4
Specialised SS	Background	1.6×10^{-7}	3 172
	Majorana, $m_N = 200$ GeV	0.053	46
	Majorana, $m_N = 400$ GeV	0.18	13
	Majorana, $m_N = 800$ GeV	0.24	1
	Majorana, $m_N = 1200$ GeV	0.26	≈ 0.2
Specialised OS	Background	7.3×10^{-7}	14 019
	pseudo-Dirac, $m_N = 200$ GeV	0.091	79
	pseudo-Dirac, $m_N = 400$ GeV	0.38	27
	pseudo-Dirac, $m_N = 800$ GeV	0.52	2
	pseudo-Dirac, $m_N = 1200$ GeV	0.56	≈ 0.4

the cuts for both signal and background are listed in table 7.3. The integrated luminosity of 300 fb^{-1} corresponds to the expected luminosity of Run 3 of the LHC, and the values of s and b in table 7.3 therefore give the number of signal and background events we would expect to see in the respective signal regions at the end of this run. We see that the efficiencies on the background are between 10^{-6} and 10^{-7}. There are more than four times more background events in the two OS regions than in the SS regions, while the difference between the baseline regions and the specialised regions is very small for the background. For the signals, we see that the efficiency of the cuts is always smallest for the smallest neutrino masses, as expected. There are approximately half as many Majorana signal events in the SS regions as there are pseudo-Dirac signal events in the OS regions, since the cut on the charge product of the leptons removes around half of the Majorana events, while ideally keeping all of the pseudo-Dirac events. The Majorana models still have better signal to background ratios, since the background is so much smaller in the SS regions. Comparing the baseline and specialised regions, we see that there is not much difference for the Majorana models, which is expected since choosing the correct lepton pair is not as important when the pair can have either product of signs. For the $m_N = 200$ GeV Majorana model, there are even three fewer signal events in the specialised region than in the baseline. Also for the pseudo-Dirac models, the improvement in the specialised region is disappointingly small. We saw in section 6.4 that our lepton identification method can have up to 50% higher efficiency than only using p_T, but the cut efficiencies on the pseudo-Dirac models are increased by at most 10%. This must mean that relatively many of the events in which our identification method is better than the p_T method have already been removed by the other cuts. But a 10% increase in signal events is also useful, so we will keep using the specialised OS region for the pseudo-Dirac signals.

Now we will study the distributions of m_{3l} for both the background and signal in each signal region. We did not use this variable to define the signal regions, so that we can instead study its distributions, and perform statistical analyses by counting the number of events either above some threshold value or in some selected bins. The plots of m_{3l} in the specialised signal regions are shown in figures 7.3 and 7.4. In the upper subplots, the signal and background are plotted in separate histograms. In the lower subplots, the ratio s/b for each bin where $s \neq 0$ and $b \neq 0$ is shown. We see that in all the plots, the ratio of signal to background events tends to increase for larger m_{3l}. In the specialised SS region, the largest value of m_{3l} in any background event is around 1600 GeV, while all the signal models have events with larger values than this. In the specialised OS region, there are some background events that have significantly larger m_{3l}, up to around 2 700 GeV. This causes the signal to background ratio to be lower for the pseudo-Dirac models with low m_N, where it is of the order of 10^{-2} in most bins. For the two pseudo-Dirac models with largest m_N, the signal distributions are shifted more towards larger values, and there are a few bins with s/b of the order of 10^{-1}. However, these bins have much less than 1 event each. For the Majorana models in the SS region, on the other hand, we see that the tail of the signal distribution of the $m_N = 200$ GeV model is almost as large as the tail of the background distribution. For the two Majorana models with the largest mass, the signal distributions extend far beyond the background distribution, but the number of events in this tail is very small. The maximum s/b value is also smaller for the Majorana models with large m_N, compared to the one with $m_N = 200$ GeV. With the particular binning that we chose, there are no bins that have $s/b \geq 1$, but in most of the plots there are some bins where the ratio is of the order of 10^{-1}.

Now we want to calculate the expected significance of each signal at Run 3 of the LHC, by counting the events in some bins of the m_{3l} distributions. We have seen that the ratio of signal to background tends to increase towards larger values. We therefore choose to sum the signal and background events from some lower limit for our statistical analysis. This is equivalent to making a new signal region, since the lower limits on m_{3l} are essentially new cuts. The specific lower limits of m_{3l} that give the best expected significance will depend on the neutrino mass and the distribution of the background. But we want to avoid making too many signal regions, as this will make the analysis very model dependent. As a compromise, we only make two different lower limits; one for the Majorana models in the SS regions, and one for the pseudo-Dirac models in the OS regions. Looking at the background distributions with finer binning than in figures 7.3 and 7.4 we see that the background in the SS region becomes very small at around $m_{3l} = 1100$ GeV, while the background in the OS region falls off at around 1300 GeV.

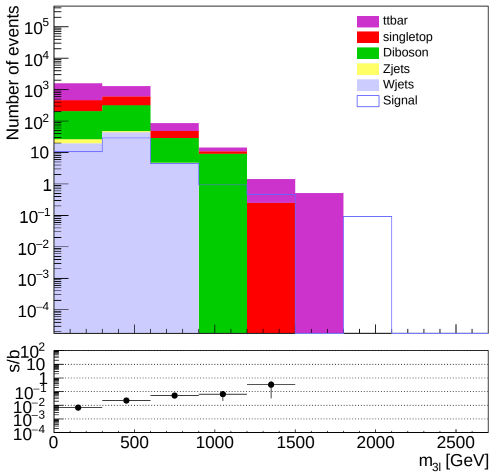

(a) Majorana model, $m_N = 200$ GeV, Region Specialised SS

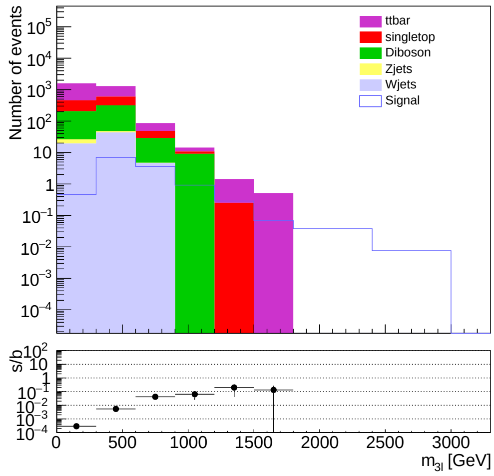

(b) Majorana model, $m_N = 400$ GeV, Region Specialised SS

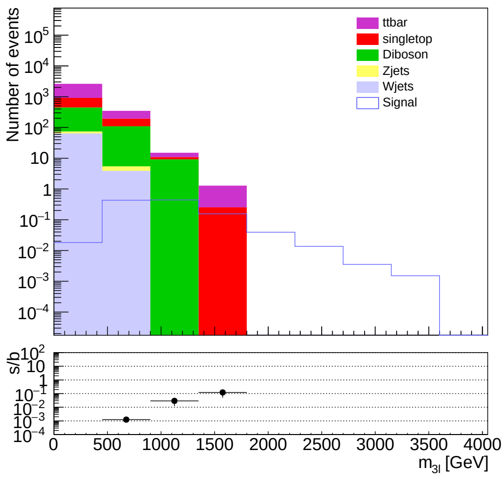

(c) Majorana model, $m_N = 800$ GeV, Region Specialised SS

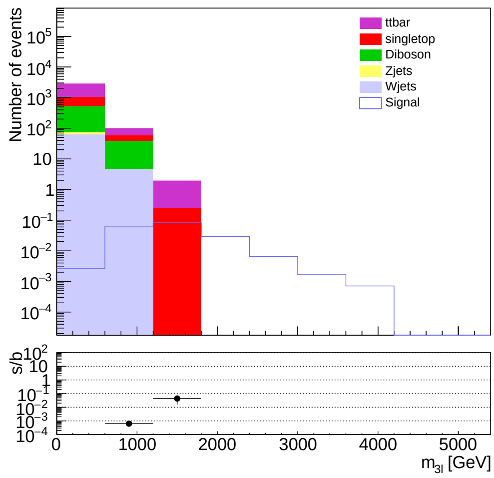

(d) Majorana model, $m_N = 1200$ GeV, Region Specialised SS

Figure 7.3: The distributions of the invariant mass of the three leading leptons in each event for the full SM background and four different Majorana models, after we applied cuts. The top plot in each subfigure shows the distributions for signal and backgrounds, where the backgrounds are stacked on top of each other, while the signal is plotted in a separate histogram. The bottom plot shows the ratio of the number of signal events s to the number of background events b in each bin where it is well-defined.

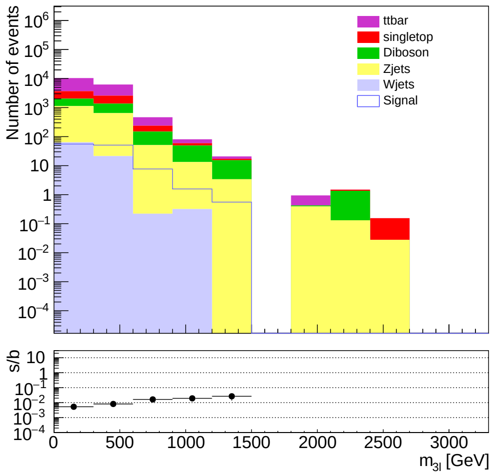

(a) Pseudo-Dirac model, $m_N = 200$ GeV, Region Specialised OS

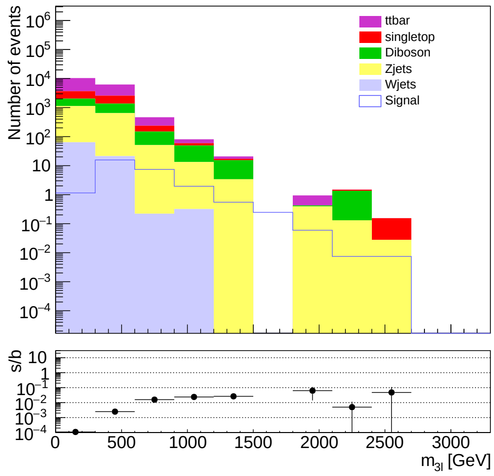

(b) Pseudo-Dirac model, $m_N = 400$ GeV, Region Specialised OS

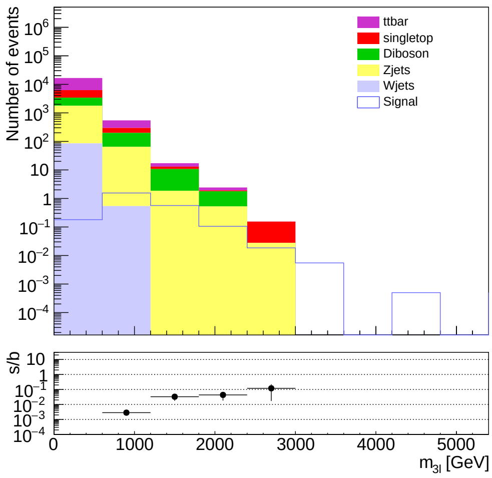

(c) Pseudo-Dirac model, $m_N = 800$ GeV, Region Specialised OS

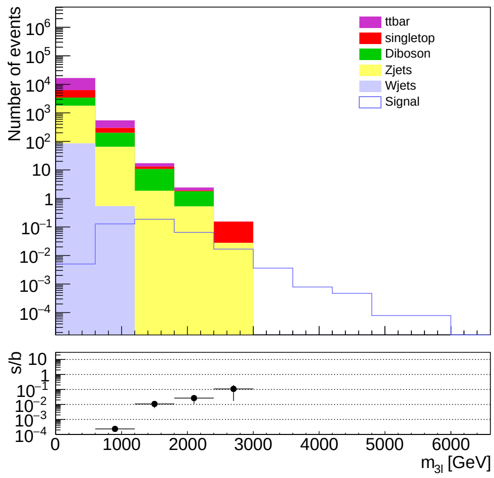

(d) Pseudo-Dirac model, $m_N = 1200$ GeV, Region Specialised OS

Figure 7.4: The distributions of the invariant mass of the three leading leptons in each event for the full SM background and four different pseudo-Dirac models, after we applied cuts. The top plot in each subfigure shows the distributions for signal and backgrounds, where the backgrounds are stacked on top of each other, while the signal is plotted in a separate histogram. The bottom plot shows the ratio of the number of signal events s to the number of background events b in each bin where it is well-defined.

We therefore choose these values for our cuts on m_{3l}, defining new SS and OS regions. In table 7.4, we list the number of background and signal events in the new signal regions with an integrated luminosity of 300 fb^{-1}, and the expected discovery significances. The expected significances were computed using the expressions from equations (7.5) and (7.6), and we compute one value assuming no systematic uncertainty in the background, and one value assuming a 20% systematic uncertainty in the background. Although we do not know the exact systematic uncertainty in the background, the systematic uncertainty in MC generated backgrounds are usually relatively large. The recent ATLAS study on the Keung-Senjanović process performed a thorough estimation of the total uncertainty in the background of the two lepton plus two jet final state, and it is found to be less than 20% [61] (see figure 4(a)). We therefore expect that 20% systematic uncertainty is a conservative estimate also for the three lepton final state. In any case, the significance computed with systematic uncertainty almost certainly gives a more accurate representation of the true expected significance than the value computed without systematic uncertainty.

From the values in table 7.4, we see first of all that none of the expected significances are very large, the largest when background uncertainty is taken into account being 0.61σ for the $m_N = 200$ GeV Majorana model in the baseline SS region. This shows that, at least with this method, these signals can not be discovered during Run 3 of the LHC. The reason for the low significances is the small number of expected signal events, which is caused by both the low cross sections of the signals, and the hard cuts required to remove most of the background. The significances tend to be largest for the models with the smallest neutrino masses, except for the pseudo-Dirac model with $m_N = 200$ GeV, which has relatively small expected significance. Both models with $m_N = 1200$ GeV have significances of 0.1σ or less, so these seem very difficult to discover. We notice that the difference between the significances with and without systematic uncertainty in the background is very small, which means that it is the statistical uncertainty of b that dominates. We also notice that the expected significances of the pseudo-Dirac models are larger in the specialised region than the baseline region, both because the number of signal events is slightly larger and because the background is smaller. For the Majorana models on the other hand, the results are better in the baseline SS region than the specialised one for all masses. From now on we will therefore use only the baseline SS region for the Majorana signals and the specialised OS region for the pseudo-Dirac signals, and call these simply the SS region and OS region. The only potential issue with this choice is that the two signal regions are no longer orthogonal, meaning that the same event can enter both regions. This would make it more challenging to determine the type of an unknown heavy neutrino, but as long as we only care about discovery, regardless of type, the orthogonality of signal regions is not important[2].

Since discovery during Run 3 of the LHC proved difficult, we want to investigate the possibilities at the future runs known as the High Luminosity LHC or HL-LHC, with an expected final integrated luminosity of 3000 fb^{-1} at $\sqrt{s} = 14$ TeV. For this analysis, we need to make a list of assumptions, and it will therefore only be a rough estimate of the possible results at the HL-LHC. For the signal processes, we can easily calculate the cross sections at 14 TeV using `MadGraph`, as we did in section 5.3. However, we do not

[2]The non-orthogonality of the signal regions would be relevant if we wanted to statistically combine results from the two regions, since this would introduce some over-counting. We will however not combine the regions, as this has little benefit for both pseudo-Dirac (most of the signal is in the OS region) and Majorana (much more background in the OS region) signals.

Table 7.4: The number of background and signal events remaining in the different signal regions after performing cuts on m_{3l} for an integrated luminosity of 300 fb^{-1}. The expected discovery significances both when assuming 20% systematic uncertainty in the background and when neglecting the systematic uncertainty are also listed. These correspond to the expected significances of Run 3 of the LHC.

Signal region	Type of events ("Background" or N type, m_N [GeV])	n_{events} (300 fb^{-1})	Z ($\sigma_b = 0.2b$)	Z ($\sigma_b = 0$)
Baseline SS, $m_{3l} > 1100$ GeV	Background	≈ 2	-	-
	Majorana, 200	≈ 0.8	0.61	0.63
	Majorana, 400	≈ 0.6	0.44	0.45
	Majorana, 800	≈ 0.4	0.30	0.31
	Majorana, 1200	≈ 0.1	0.087	0.089
Baseline OS, $m_{3l} > 1300$ GeV	Background	≈ 6	-	-
	Pseudo-Dirac, 200	≈ 0.2	0.067	0.075
	Pseudo-Dirac, 400	≈ 0.5	0.18	0.20
	Pseudo-Dirac, 800	≈ 0.5	0.16	0.18
	Pseudo-Dirac, 1200	≈ 0.2	0.078	0.086
Specialised SS, $m_{3l} > 1100$ GeV	Background	≈ 3	-	-
	Majorana, 200	≈ 0.7	0.37	0.39
	Majorana, 400	≈ 0.7	0.40	0.42
	Majorana, 800	≈ 0.4	0.24	0.25
	Majorana, 1200	≈ 0.1	0.085	0.090
Specialised OS, $m_{3l} > 1300$ GeV	Background	≈ 5	-	-
	Pseudo-Dirac, 200	≈ 0.3	0.12	0.13
	Pseudo-Dirac, 400	≈ 0.6	0.24	0.26
	Pseudo-Dirac, 800	≈ 0.5	0.22	0.24
	Pseudo-Dirac, 1200	≈ 0.2	0.10	0.11

Table 7.5: The expected discovery significances at the HL-LHC ($L = 3000$ fb^{-1}, $\sqrt{s} = 14$ TeV) for the Majorana and pseudo-Dirac signals, both when assuming 20% systematic uncertainty in the background, and when neglecting the systematic uncertainty.

Signal region	Signal (type, m_N [GeV])	Z ($\sigma_b = 0.2b$)	Z ($\sigma_b = 0$)
SS, $m_{3l} > 1100$ GeV	Majorana, 200	1.6	2.0
	Majorana, 400	1.2	1.5
	Majorana, 800	0.85	1.1
	Majorana, 1200	0.23	0.29
OS, $m_{3l} > 1300$ GeV	pseudo-Dirac, 200	0.25	0.42
	pseudo-Dirac, 400	0.53	0.89
	pseudo-Dirac, 800	0.50	0.84
	pseudo-Dirac, 1200	0.21	0.36

have MC samples of the full SM background for 14 TeV, and we will therefore use the 13.6 TeV background without any modifications. This will be a slight underestimation of the background, as some background processes will have larger cross sections at 14 TeV. Additionally, we assume that both the cut efficiencies on the background and on the signal events remain the same as for Run 3. Using these assumptions, and the same cuts of $m_{3l} > 1100$ GeV and $m_{3l} > 1300$ GeV in the SS and OS regions respectively, we get the expected significances for Run 4 of the LHC shown in table 7.5. We see that the expected significances for the HL-LHC are much larger than for Run 3. The main reason for this is the factor 10 increase of the integrated luminosity, which when the statistical uncertainty in the background dominates increases the expected significance with a factor close to $\sqrt{10} \approx 3$. The difference between the significances with and without σ_b is now slightly larger, since the background has ten times more events. The signal with the highest significance is still the Majorana model with $m_N = 200$ GeV, with a significance of 1.6σ with systematic uncertainty in the background. This is still far from large enough to result in a discovery, but it is large enough that it will likely give a noticeable anomaly in the data. The Majorana model with $m_N = 400$ GeV also gets an expected significance of more than 1σ, while all the pseudo-Dirac models still have very small significances.

We can also plot the expected significance we get from summing all signal and background events in the signal regions as a function of the lower limit on m_{3l}. We again look at the expected significances for the HL-LHC, using the same assumptions as before. The plots of the expected discovery significances as functions of the applied cut on m_{3l} for the Majorana models in the SS signal region and the pseudo-Dirac models in the OS signal regions are shown in figure 7.5. Here we show only the significances with $\sigma_b = 0.2b$. The highest values for the cut on m_{3l} that are used are decided by the amount of signal and background events above that value. In the SS region, the number of background events becomes too small to compute a reasonable significance at around $m_{3l} = 1200$ GeV, while in the OS region there are very few signal and background events above $m_{3l} = 1400$ GeV. We see that the peaks in the significance functions for the Majorana models are close to the cut value of 1100 GeV that we used before, and

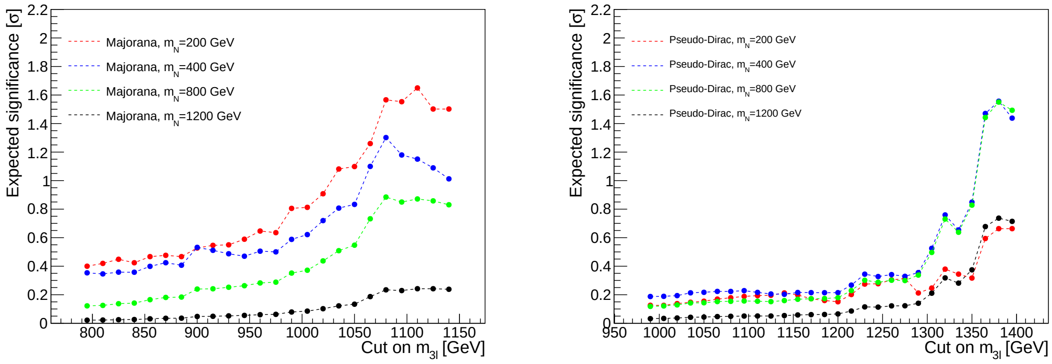

(a) Majorana models

(b) Pseudo-Dirac models

Figure 7.5: The expected significances of the different Majorana and pseudo-Dirac signal models at the HL-LHC as a function of the lower limit cut on m_{3l} performed on signal and background events. For the Majorana models the SS signal region was used, while for the pseudo-Dirac models the OS region was used. The significances were computed using equation (7.6) when assuming a 20% systematic uncertainty in the background.

the maximum values in figure 7.5a are therefore close to the values in table 7.5. For the pseudo-Dirac models, on the other hand, the expected significance increases a lot between the previously used cut value of 1300 GeV and 1400 GeV. The main reason for this is that the number of background events decreases from $b = 45$ for $m_{3l} > 1300$ GeV to $b = 6$ for $m_{3l} > 1400$ GeV. The pseudo-Dirac signals with $m_N = 400$ GeV and 800 GeV, meanwhile, have $s \approx 4$ and $s \approx 5$ respectively for $m_{3l} > 1400$ GeV, which give expected significances of $\approx 1.4\sigma$ and $\approx 1.5\sigma$. Interestingly, among the pseudo-Dirac models, the one with $m_N = 200$ GeV gets the lowest expected significance, while in the Majorana case, $m_N = 200$ GeV has the highest significance. In the pseudo-Dirac case, the reason for the low significance of the $m_N = 200$ GeV model at $m_{3l} > 1400$ GeV is that there are only around 2 signal events in this region. However, as we see in the plot, the significance is not larger for lower cuts on m_{3l} either, due to the larger background.

From figure 7.5, it is clear that both the Majorana model with $m_N = 200$ GeV and the pseudo-Dirac models with $m_N = 400$ and $m_N = 800$ GeV could be observed with significances of around 1.6σ after the entire HL-LHC run, with very specific cuts. The very high sensitivity of the expected significance to the cut value for the pseudo-Dirac models might however indicate that the result for these models is only due to some lucky bins in our analysis, and would disappear if repeated with slightly different signal or background modelling. For the Majorana signals, the significances are slightly more stable, and therefore seem more trustworthy. In this rather optimistic scenario, with mixing matrix elements slightly outside the experimental constraints, it might therefore be possible to detect the signal of a Majorana neutrino with a mass of around 200 GeV using this method with the ATLAS detector during the HL-LHC. But a discovery at 5σ seems definitely out of reach for all the models we looked at with this method.

One goal we had in our analysis was to be able to distinguish between Majorana and pseudo-Dirac neutrinos in data. Unfortunately, this is very difficult when even the detection of the heavy neutrino is almost impossible, especially in the case of a pseudo-Dirac neutrino. It seems like the only hope would be to see a signal in the SS region, in which case it could only be a Majorana neutrino, although as we have seen, this can most likely not be determined with the required confidence.

We now want to calculate the 95% CL expected exclusion limits on the cross sections of the Majorana and pseudo-Dirac signal processes using the method described in section 7.1.2. To calculate the expected limits, we use Gaussian distributions for the nuisance parameters ϵ_s and b. For the integrated luminosity L, we expect the uncertainty to be so small that we can neglect it[3], so we use fixed values of L. For the uncertainty in b, we again assume the conservative 20% systematic uncertainty, which when combined with the statistical uncertainty gives a total uncertainty of

$$\delta_b = \sqrt{b + (0.2b)^2}. \tag{7.18}$$

For the uncertainty in ϵ_s, we need to look at how ϵ_s is calculated. In our analysis, we find ϵ_s from the following expression

$$\epsilon_s = \frac{s}{L\sigma}, \tag{7.19}$$

where σ is the cross section of the signal process. Then we can use error propagation to get the error in ϵ_s (again neglecting the error in L)

$$\delta_{\epsilon_s} = \epsilon_s \sqrt{\left(\frac{\delta_s}{s}\right)^2 + \left(\frac{\delta_\sigma}{\sigma}\right)^2} = \epsilon_s \sqrt{\frac{1}{s} + \left(\frac{\delta_\sigma}{\sigma}\right)^2}, \tag{7.20}$$

where δ_σ is the uncertainty in the signal cross section, and we assumed a statistical error $\delta_s = \sqrt{s}$ (not to be confused with centre of mass energy) in the expected number of signal events. We see that the uncertainty in ϵ_s becomes large when there are few expected signal events. We find that this has a large effect on the resulting expected exclusion limit, and to get reasonable results we therefore have to make slightly different signal regions from the ones we made in the previous section, at least for the Run 3 analysis. We find that for the Majorana models in the baseline SS region, a cut of $m_{3l} > 600$ GeV gives relatively good[4] exclusion limits. For the pseudo-Dirac models in the specialised SS region, we perform a cut of $m_{3l} > 900$ GeV. With a luminosity of $L = 300$ fb^{-1} and centre of mass energy $\sqrt{s} = 13.6$ TeV, corresponding to the entire Run 3 of the LHC, we get the expected exclusion limits shown in figure 7.6. In the figure, the dashed lines show the 95% CL expected exclusion limit on the cross section, and the green and yellow band show the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty regions for the expected limit. This means that any signal model that is above the dashed line in the plots is excluded at 95% CL. We see that in this case, none of the models can be excluded, not even the extremely unrealistic models with $V_{e1} = 1$. For both the Majorana and pseudo-Dirac models, the model with $m_N = 200$ GeV has a relatively large expected limit. This is because these models have by far the smallest ϵ_s, and so a large cross section is required to get a significant number of signal events. The expected limit reaches a minimum for the $m_N = 400$ and $m_N = 800$ GeV models in the Majorana and pseudo-Dirac case respectively, since these have larger

[3]As an example of the order of magnitude of the uncertainty in L, the uncertainty in the integrated luminosity of Run 2 of the LHC was estimated to be 1.7% [67], which is much smaller than the uncertainty in b, and in most cases also in ϵ_s.

[4]Although, as we will see, still not very good.

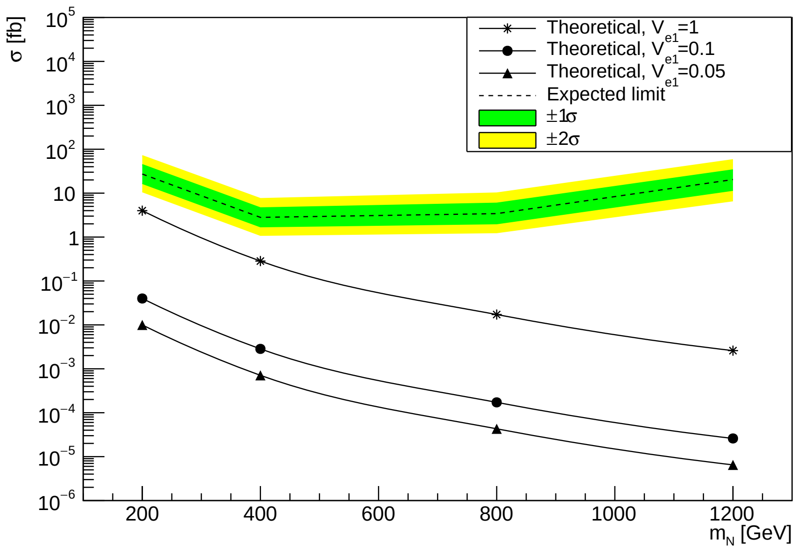

(a) Majorana models

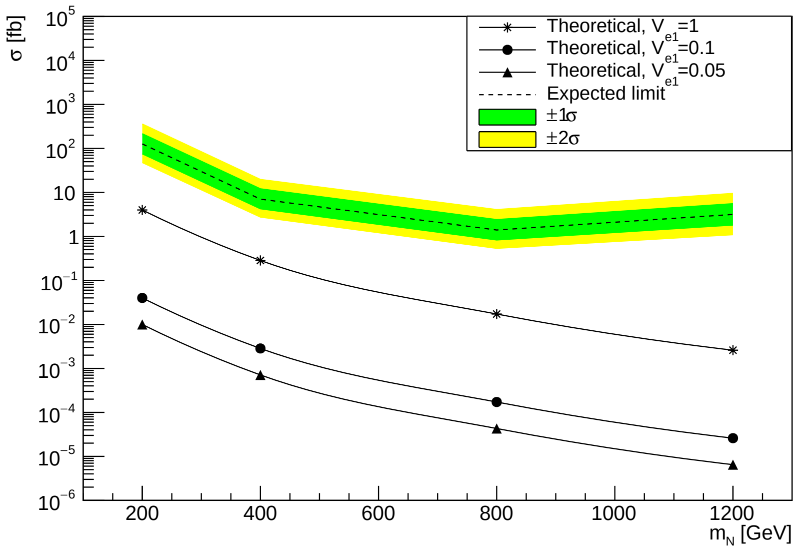

(b) Pseudo-Dirac models

Figure 7.6: The 95% CL expected exclusion limits on the cross sections of the Majorana and pseudo-Dirac signal processes for Run 3 of the LHC. The lines labelled "Theoretical" show the signal cross sections of the processes for m_N = 200, 400, 800 and 1200 GeV computed in `MadGraph` for the case where the mixing parameter V_{e1} is equal to the experimental constraint of 0.05, for the optimistic case $V_{e1} = 0.1$ and the completely unrealistic case of $V_{e1} = 1$.

ϵ_s, and not too small s. Then for $m_N = 1200$ GeV the limits become worse again in both plots, which is mainly because s is less than 1 for these models, so the relative uncertainty in ϵ_s becomes very large. Comparing the Majorana and pseudo-Dirac limits, the Majorana limits are lower for the two lower neutrino masses, while the pseudo-Dirac limits are lower for the two higher masses. The reason for this is that even though the m_{3l} cut is much less strict on the Majorana models, there are nevertheless fewer background events in the SS signal region than in the OS region, and for $m_N = 200$ and $m_N = 400$ GeV, there are more signal events in the SS region than in the OS region. For $m_N = 800$ and $m_N = 1200$ GeV there are more signal events in the OS region (something similar can be seen in table 7.4, although the m_{3l} cuts are slightly different), which gives slightly better exclusion limits for these models. In conclusion, it is safe to say that none of the models we have studied can be excluded after Run 3 of the LHC with this method.

Finally, we compute the 95% CL expected exclusion limits for the HL-LHC. We use the same assumptions as in the previous section to translate our data to the higher luminosity and centre of mass energy. In this case, the number of expected signal events is large enough that we can use the stricter cuts on m_{3l} that gave the best results for significance, namely $m_{3l} > 1100$ GeV for the SS signal region and $m_{3l} > 1400$ GeV for the OS signal region. With these cuts, the expected exclusion limits for the HL-LHC become as shown in figure 7.7. We see that the increase in luminosity and centre of mass energy moved the expected limits significantly lower. However, it is still only the completely unrealistic $V_{e1} = 1$ model with $m_N = 200$ GeV that is on the verge of being excluded at 95% CL, while a few more of the unrealistic models are touching the lower 2σ band. We also notice that the higher neutrino masses get the lowest expected limits in this case, which is due to the stricter cuts favouring the model with higher masses. Unfortunately, we have to conclude that even at the HL-LHC, none of the realistic Majorana or pseudo-Dirac models can be excluded.

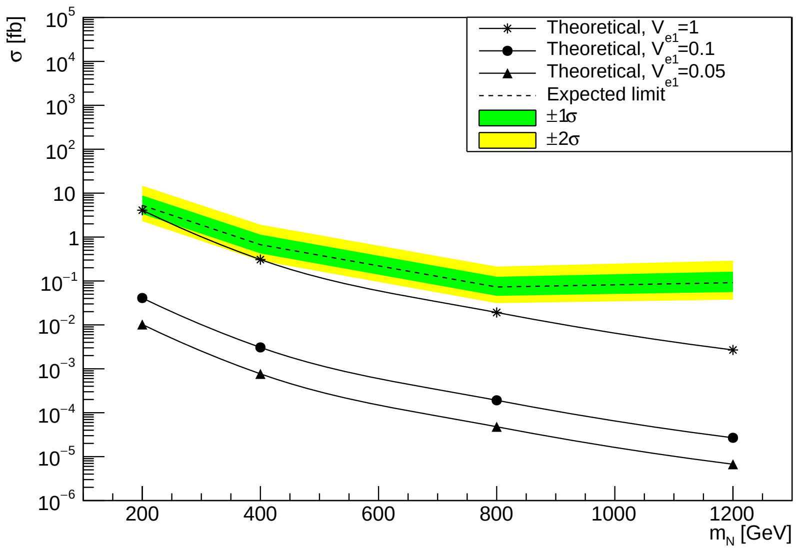

(a) Majorana models

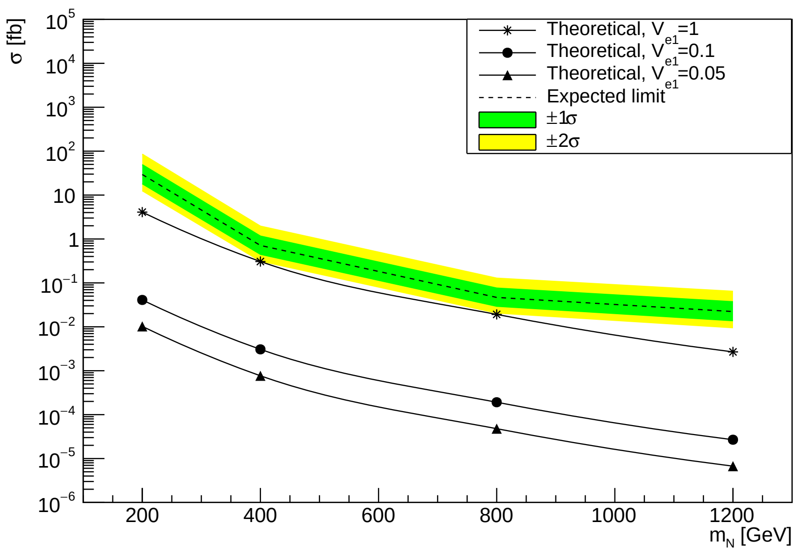

(b) Pseudo-Dirac models

Figure 7.7: The 95% CL expected exclusion limits on the cross sections of the Majorana and pseudo-Dirac signal processes for the HL-LHC. The lines labelled "Theoretical" show the signal cross sections of the processes for $m_N = 200$, 400, 800 and 1200 GeV computed in `MadGraph` for the case where the mixing parameter V_{e1} is equal to the experimental constraint of 0.05, for the optimistic case $V_{e1} = 0.1$ and the completely unrealistic case of $V_{e1} = 1$.

Table 7.6: The requirements that were used to define electrons, muons and jets in the analysis of signal and background with the two leptons and two jets final state. For electrons and muons, the working points are chosen from the ones available in `SimpleAnalysis`, and are based on the ones made in the recent ATLAS paper [61]. The names of the working points are the names they have in `SimpleAnalysis`.

	Electrons	Muons	Jets
Quality	`ETightLH`	`MuMedium` if $p_T < 300$ GeV else `MuHighPt`	-
Isolation	`EIsoLoose`	`MuIsoPflowLoose_VarRad`	-
p_T [GeV]	> 25	> 25	> 20
$\|\eta\|$	$(0, 1.37]$ or $[1.52, 2.47]$	< 2.5	< 2.5

7.3 LRSM model

In this section we perform a similar analysis as in the previous section, but for the LRSM models with the final state of two leptons and two jets (2l+2j). We will use the same masses for the W_R boson and the heavy neutrino N as we did in section 6.5, and we continue to refer to these as models 1–6. Now we will also for each set of masses consider the case where the heavy neutrino is the one that couples to the muon. This gives us a total of 12 signal models, and we refer to them as Model ie and Model $i$mu, $i = 1, \ldots, 6$ when the neutrino couples to the electron and muon respectively. Since we will not study the case where the neutrino couples to the tau, we again only count electrons and muons as leptons in this analysis. For models 1 and 2, we know that they have been excluded under the assumption that the right-handed weak coupling constant is identical to the weak coupling in the SM [61]. For these models our goal will therefore be to first confirm that we can also exclude them with the same assumption, and then to attempt to put an upper limit on the right-handed weak coupling in these models. Models 3–6, on the other hand, are just outside of the mass region that was excluded in the most recent study by the ATLAS collaboration [61], and our goal will therefore be to either exclude these using more specialised signal regions, or to say something about when these models can be expected to be discovered or excluded in the future of the LHC.

For this analysis, we succeeded in implementing smearing on the MC truth electrons and muons. The analysis will therefore better reflect what can be achieved in real experiments. It also means that we need to choose parameters and working points for object identification and reconstruction. For fake rates we use the default values provided in `SimpleAnalysis`. For overlap removal, we use the same parameters that we presented in section 6.1. Finally, we need to choose working points for quality and isolation of reconstructed electrons and muons, as described in section 4.2.1. We choose the working points based on a combination of what is available in `SimpleAnalysis`, the recommended working points for analyses for Run 3 of the LHC and release 22 of `AnalysisBase`, and the working points used in the ATLAS paper [61] (see Table 2 in the paper for their object definitions). In our case, the object definitions are the ones listed in table 7.6.

We again begin by finding the expected significances of our signal models for the entire Run 3 of the LHC, in some signal regions. We will not make separate signal regions

Table 7.7: The cut efficiencies on the background and each of the simulated LRSM models in the signal regions rSRSS2e and rSRSS2mu from the ATLAS paper [61]. The number of events remaining in the respective signal region for an integrated luminosity of 300 fb^{-1} and the expected discovery significances for each signal both when assuming 20% systematic uncertainty in the background and when neglecting the systematic uncertainty are also listed. These correspond to the expected significances of Run 3 of the LHC.

	Type of events	ϵ	n_{events} (300 fb^{-1})	Z ($\sigma_b = 0.2b$)	Z ($\sigma_b = 0$)
Region rSRSS2e	Background	1.4×10^{-9}	27	-	-
	1e	0.16	580	25	51
	2e	0.17	355	19	36
	3e	0.18	3	0.35	0.51
	4e	0.16	1	0.097	0.14
	5e	0.18	≈ 0.5	0.066	0.095
	6e	0.17	≈ 0.1	0.020	0.029
Region rSRSS2mu	Background	1.0×10^{-9}	19	-	-
	1mu	0.19	684	31	61
	2mu	0.19	398	23	42
	3mu	0.20	3	0.51	0.69
	4mu	0.19	1	0.14	0.19
	5mu	0.20	≈ 0.6	0.094	0.13
	6mu	0.20	≈ 0.2	0.029	0.039

for different masses, because the signal models we study are relatively similar, and we want to be as model independent as possible. We also only consider the case where the heavy neutrino is of Majorana type, so there is no need to separate into SS and OS regions. We will however split our signal region in two based on the flavours of the two leading leptons. One region requires the two leptons to be electrons, and we call this the electron channel. The other region requires the two leptons to be muons, and we call this the muon channel. Before we choose the other cuts to define our two signal regions, we will get some baseline results by using the signal regions rSRSS2e and rSRSS2mu defined in the ATLAS paper [61] (see Table 4), which we can later compare to our own electron and muon channels. We apply these cuts on the 13.6 TeV MC background and each of the signal models. The efficiency and the number of remaining events for background and the signal models with a luminosity of 300 fb^{-1} in both signal regions are listed in table 7.7. The expected significances with and without systematic uncertainty in the background are also given for each signal model in the respective signal region. We see that as expected, the signal regions from the ATLAS paper give very large expected significances for the excluded models 1 and 2. For the other models, none have expected significances above 1σ, which means that they can not be discovered at Run 3 of the LHC with these signal regions. We notice that the efficiency of the cuts is almost the

Table 7.8: The cuts made on kinematic variables to define our signal region for the 2l+2j final state, which is subsequently split into two depending on the flavour of the two leading leptons. The cuts are given in terms of the requirement that must be fulfilled for the specific event to pass. See the main text for the definitions of the variables.

Variable	Cut	Variable	Cut
n_{leptons}	< 4	n_{jets}	> 2
m_{jj} [GeV]	> 800	m_{ll} [GeV]	> 400
$p'_{T,l1}$ [GeV]	> 400	$p'_{T,l2}$ [GeV]	> 150
$p'_{T,j1}$ [GeV]	> 500	$p'_{T,j2}$ [GeV]	> 250
E'_{j1} [GeV]	> 800	E'_{j2} [GeV]	> 400
$\text{Sign}(l'_1)\times\text{Sign}(l'_2)$	$= +1$	h_T [GeV]	> 2000

same for all the signals. The efficiency is slightly lower for models 1 and 2 and 4, which is probably due to the cuts on p_T and m_{ll}. We expect that it is possible to make cuts with higher efficiency, especially on models 3–6, since these have very energetic events that are relatively easy to separate from the background.

We will now attempt to make cuts that give better signal to background ratios than the cuts made in the ATLAS paper. For this we again plot the most relevant kinematic distributions for both background and a few signal models to choose our cuts. The plots are shown in appendix A.2, and the cuts we chose based on them are listed in table 7.8. Here n_{leptons} and n_{jets} are the number of leptons (electrons and muons only) and jets in an event that fulfill the objects definitions in table 7.6, and m_{ll} and m_{jj} are the invariant masses of the leading pair of leptons and jets respectively. The primed transverse momenta $p'_{T,li}$ are the transverse momenta of lepton number i when ordered by p_T, and $p'_{T,ji}$ is the same for jets, and E'_{ji} are the energies of jet number i when ordered by p_T. Finally, $\text{Sign}(l'_i)$ refers to the sign of the charge of lepton i when ordered by p_T. In addition to the cuts in table 7.8, we require the flavour of the two leading leptons to be either electrons or muons, defining the electron and muon channel respectively. The distributions of the invariant mass of the two leading leptons and two leading jets, m_{lljj}, for the background and signal models 3 and 4 after applying all these cuts with an integrated luminosity of 300 fb^{-1} are shown in figure 7.8. In the upper subplots we again plot the signal and background in separate histograms. We see that the total number of background events in each signal region is of the order of 1. Also, the peaks of the signal distributions for the two models are at higher m_{lljj} values than any of the background events. In fact, for model 3, the two bins around $m_{lljj} = 5$ TeV contain approximately three signal events and no background events, meaning that if the signal can be discovered, an approximate value for M_{W_R} could also be deduced from these distributions. Since we use a bin width of 1 TeV in these plots (the large bin width is required in order to have any significant number of signal events in each bin), we expect the precision of such a measurement to be of the order of 0.5 TeV. For model 4, there is only around 1 signal event in the two bins, which makes the discovery slightly more difficult.

Since it does not really make sense to calculate the expected significance in a signal

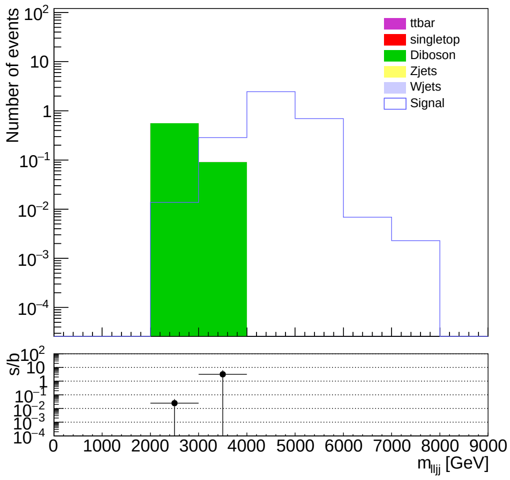

(a) Model 3e, Electron channel

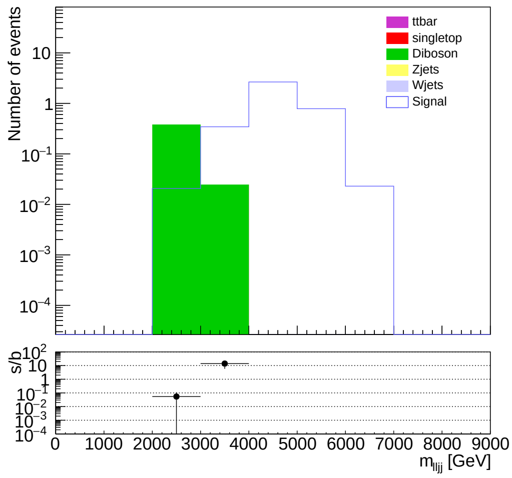

(b) Model 3mu, Muon channel

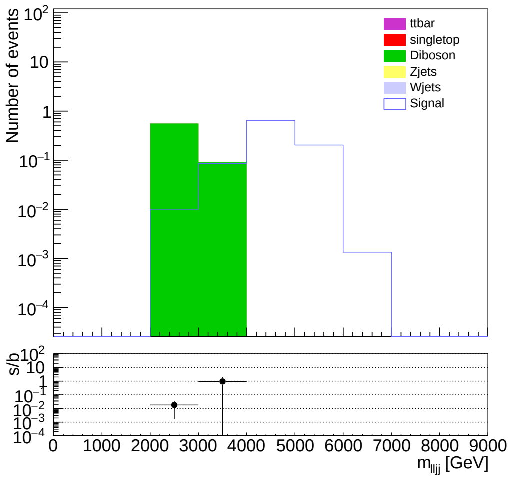

(c) Model 4e, Electron channel

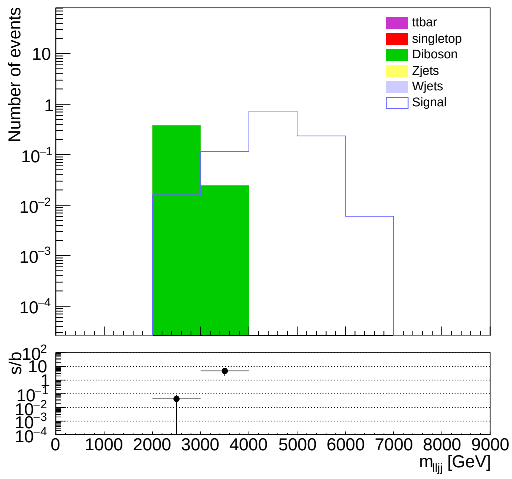

(d) Model 4mu, Muon channel

Figure 7.8: The distributions of m_{lljj} in each event for the full SM background and four different LRSM models, after we applied the cuts in table 7.8 and separated into electron and muon channels. The top plot in each subfigure shows the distributions for signal and backgrounds, where the backgrounds are stacked on top of each other, while the signal is plotted in a separate histogram. The bottom plot shows the ratio of the number of signal events s to the number of background events b in each bin where it is well-defined.

Table 7.9: The cut efficiencies on the background and signal events for each of the simulated LRSM models in the electron and muon channels defined by our own cuts. The number of background and signal events remaining in the respective signal region for an integrated luminosity of 300 fb^{-1} and the expected discovery significances both when assuming 20% systematic uncertainty in the background, and when neglecting the systematic uncertainty are also listed. These correspond to the expected significances of Run 3 of the LHC.

	Type of event	ϵ	n_{events} (300 fb^{-1})	Z ($\sigma_b = 0.2b$)	Z ($\sigma_b = 0$)
	Background	3.4×10^{-11}	≈ 0.7	-	-
	1e	0.026	93	24	27
Electron	2e	0.062	129	29	33
Channel	3e	0.23	3	2.8	2.9
	4e	0.21	≈ 1	0.97	0.99
	5e	0.19	≈ 0.5	0.56	0.57
	6e	0.26	≈ 0.2	0.26	0.26
	Background	2.1×10^{-11}	≈ 0.4	-	-
	1mu	0.027	97	26	29
Muon	2mu	0.067	140	32	37
Channel	3mu	0.25	4	3.4	3.5
	4mu	0.25	≈ 1	1.3	1.3
	5mu	0.25	≈ 0.7	0.89	0.90
	6mu	0.29	≈ 0.2	0.35	0.36

region with zero background events, we will not perform an additional cut on m_{lljj} like we did on m_{3l} in the three lepton final state. The final electron and muon channel signal regions are therefore the ones defined by the cuts in table 7.8 and the requirement on the flavour of the lepton pair. The efficiencies on the background and signal models for these two signal regions along with the expected significances for each of the signal models are given in table 7.9. We see that our cuts have almost 50 times smaller efficiency on the background than the cuts from the ATLAS paper, and that there is only of the order of one background event left in our signal regions at 300 fb^{-1}. For the signal models 1 and 2, our cuts have much smaller efficiency than the cuts from the ATLAS paper. This is expected, since our cuts are stricter, and we focused on optimising the cuts for models 3–6. Nevertheless, we still get extremely large expected significances for models 1 and 2, both in the electron and muon channels. We see that we succeeded in adapting the cuts to models 3–6, as we have approximately the same efficiencies or larger than the ATLAS paper for these signals. Since we have cut much more of the background, this leads to much larger expected significances. We get a particularly large significance for model 3, with almost 3σ in the electron channel, and more than 3σ in the muon channel. This means that this model, with $M_{W_R} = 5$ TeV and $m_N = 4$ TeV, could almost be discovered

Table 7.10: The number of background and LRSM signal events remaining in the two signal regions defined by our own cuts for $\sqrt{s} = 14$ TeV and an integrated luminosity of 3000 fb^{-1}. The expected discovery significances both when assuming 20% systematic uncertainty in the background, and when neglecting the systematic uncertainty are also listed. These correspond to the expected significances of the HL-LHC.

	Type of event	n_{events} (3000 fb^{-1})	Z ($\sigma_b = 0.2b$)	Z ($\sigma_b = 0$)
	Background	7	-	-
	1e	1066	56	94
Electron	2e	1436	66	113
Channel	3e	47	9.0	11
	4e	13	3.4	4.1
	5e	8	2.3	2.7
	6e	3	1.1	1.3
	Background	4	-	-
	1mu	1102	64	101
Muon	2mu	1564	77	125
Channel	3mu	52	11	14
	4mu	15	4.7	5.4
	5mu	11	3.7	4.1
	6mu	4	1.5	1.7

during Run 3 of the LHC. Also for model 4 we get significances of approximately 1σ in the electron channel and 1.3σ in the muon channel, which is far from a discovery, but might be enough to be noticeable in experimental data. Models 5 and 6 have significances smaller than 1σ in both channels, and would therefore be difficult to discover. For all models we see that the muon channel has larger ϵ_s than the electron channel, and since ϵ_b is smaller in the muon channel, the muon channel gives larger expected significances. We also notice that since the number of background events is so small, the statistical uncertainty dominates over the systematic one, such that the expected uncertainty with $\sigma_b = 0.2$ is almost identical to the one with no systematic uncertainty.

Next, we will compute the expected discovery significances of the twelve models for the complete HL-LHC. We use the same assumptions as in section 7.2; that the cross section of all backgrounds and the cut efficiencies on signals and background remain the same. For the signal models, we use the $\sqrt{s} = 14$ TeV cross sections given in table 5.2[5]. We use the same electron and muon channel signal regions, and do not cut on m_{lljj}. The expected significances become as shown in table 7.10. Again, models 1 and 2 would clearly be very easy to discover at the HL-LHC (which is expected, since they are

[5]The case when the heavy neutrino couples to the muon gives essentially identical cross sections to the case when it couples to the electron, since the energy of the process is so large that the mass difference does not matter.

already excluded). But with the increase in luminosity and signal cross section, model 3 also gets expected significances above 5σ, meaning that this signal can be discovered during the HL-LHC run, both in the case of the neutrino coupling to the electron and the case when it couples to the muon. The expected significances in the muon channel are generally higher, with model 4 also getting very close to 5σ significance, and model 5 getting 3.7σ, which is definitely enough to be noticeable. In the electron channel, on the other hand, models 4 and 5 only get 3.4σ and 2.3σ respectively. Model 6 still gets relatively low significance, on the order of 1σ, in both channels.

Now we want to find the expected exclusion limits for our six models after Run 3 and the HL-LHC respectively. In addition to upper limits on the cross sections, we also want to find upper limits on the right-handed weak coupling, parametrised by the relative strength κ_R as follows

$$g_R = g\kappa_R, \tag{7.21}$$

where g is the SM weak coupling constant. Translating the upper limit σ_{up} on the cross section of the process to an upper limit κ_R^{up} on the relative right-handed coupling is straight forward. As shown in section 5.3.2, the NWA, which is an even better approximation when $\kappa_R < 1$, gives

$$\sigma = \sigma' \kappa_R^2, \tag{7.22}$$

which immediately gives the following relation between σ_{up} and κ_R^{up}

$$\begin{aligned} \sigma_{\text{up}} &= \sigma' \left(\kappa_R^{\text{up}}\right)^2 \\ \kappa_R^{\text{up}} &= \sqrt{\frac{\sigma_{\text{up}}}{\sigma'}}. \end{aligned} \tag{7.23}$$

For the LRSM models, we have two unknown mass parameters, and since we have only six points in the (M_{W_R}, m_N) plane, we will list the expected exclusion limits instead of plotting them. To compute the exclusion limits, we use the same cuts on signal and background for the electron and muon channels as before. We again assume that the uncertainty in the integrated luminosity is negligible, and use the expressions from equations (7.18) and (7.20) for the uncertainties in b and ϵ_s. The 95% CL expected exclusion limits on the cross sections and relative right-handed coupling with $\pm 1\sigma$ deviations for Run 3 and the HL-LHC are given in table 7.11. In general, it must be noted that when κ_R^{up} is larger than 1 the limit should not be trusted, since the NWA breaks down for any κ_R significantly larger than 1, and so the assumption that $\sigma \propto \kappa_R^2$ no longer holds. We show these values in the table nevertheless, mostly to indicate the relative difference between σ_{up} and the signal cross section σ' for the most natural case of $\kappa_R = 1$. Any model that gives $\kappa_R^{\text{up}} < 1$ is obviously excluded at 95% CL if we demand $\kappa_R = 1$.

Looking first at the expected limits after Run 3, we see that only models 1 and 2 are expected to be excluded at 95% CL if no signal is seen. The expected 95% CL upper limit that can be put on κ_R for these two models is around 0.2 if the heavy neutrino couples to the electron, and a slightly lower if it couples to the muon. For model 3, the upper limits of $\kappa_R^{\text{up}} = 1.5$ and $\kappa_R^{\text{up}} = 1.2$ in the electron and muon channels respectively are probably small enough to be valid under the NWA assumption. For models 4, 5 and 6, there is no hope of excluding them after Run 3 of the LHC, as the upper limits on the cross sections are at least ten times larger than σ'.

Table 7.11: The 95% CL expected upper limits on the cross sections and the relative right-handed coupling strength for each LRSM signal model for both the complete Run 3 of the LHC and the complete HL-LHC run. The columns marked σ' give the cross sections of each signal under the assumption $\kappa_R = 1$, calculated using `MadGraph`.

	Run 3			HL-LHC		
Signal model	σ' [fb]	$(\sigma_{\rm up})^{+1\sigma}_{-1\sigma}$ [fb]	$\left(\kappa_R^{\rm up}\right)^{+1\sigma}_{-1\sigma}$	σ' [fb]	$(\sigma_{\rm up})^{+1\sigma}_{-1\sigma}$ [fb]	$\left(\kappa_R^{\rm up}\right)^{+1\sigma}_{-1\sigma}$
1e	12.2	$0.58^{0.76}_{0.40}$	$0.22^{0.25}_{0.18}$	13.9	$0.098^{0.17}_{0.064}$	$0.084^{0.11}_{0.068}$
2e	7.1	$0.24^{0.32}_{0.17}$	$0.19^{0.21}_{0.15}$	7.9	$0.041^{0.094}_{0.027}$	$0.072^{0.094}_{0.059}$
3e	0.051	$0.12^{0.17}_{0.080}$	$1.5^{1.8}_{1.3}$	0.069	$0.014^{0.024}_{0.0088}$	$0.44^{0.59}_{0.36}$
4e	0.015	$0.28^{0.40}_{0.18}$	$4.4^{5.2}_{3.5}$	0.020	$0.016^{0.029}_{0.011}$	$0.90^{1.2}_{0.72}$
5e	0.009	$0.61^{0.87}_{0.39}$	$8.1^{9.6}_{6.4}$	0.015	$0.026^{0.047}_{0.016}$	$1.3^{1.8}_{1.1}$
6e	0.0029	$1.1^{1.6}_{0.68}$	19^{23}_{15}	0.0046	$0.022^{0.040}_{0.014}$	$2.2^{3.0}_{1.7}$
1mu	12.2	$0.39^{0.76}_{0.39}$	$0.18^{0.25}_{0.18}$	13.9	$0.082^{0.14}_{0.058}$	$0.077^{0.10}_{0.064}$
2mu	7.1	$0.15^{0.30}_{0.15}$	$0.15^{0.21}_{0.15}$	7.9	$0.033^{0.055}_{0.023}$	$0.064^{0.084}_{0.054}$
3mu	0.051	$0.070^{0.15}_{0.070}$	$1.2^{1.7}_{1.2}$	0.069	$0.011^{0.019}_{0.0074}$	$0.39^{0.52}_{0.33}$
4mu	0.015	$0.14^{0.32}_{0.14}$	$3.1^{4.6}_{3.1}$	0.020	$0.012^{0.019}_{0.0083}$	$0.77^{0.96}_{0.64}$
5mu	0.009	$0.22^{0.50}_{0.22}$	$4.8^{7.3}_{4.8}$	0.015	$0.016^{0.028}_{0.011}$	$1.0^{1.4}_{0.85}$
6mu	0.0029	$0.54^{1.3}_{0.54}$	14^{21}_{14}	0.0046	$0.016^{0.029}_{0.011}$	$1.9^{2.5}_{1.5}$

From the expected limits after the HL-LHC, we see that models 1, 2, 3 and 4 are all expected to be excluded at 95% CL if no signal is seen. The exclusion limits again tend to be slightly stricter in the muon channel, which is due to the smaller background and higher number of signal events, as seen in table 7.10. For models 1 and 2, the 95% CL upper limits on the relative coupling are smaller than 0.1, and as small as 0.064 for model 2mu. For model 3, κ_R^{up} is around 0.4, while for model 4 it is 0.90 for the neutrino coupling to the electron and 0.77 for the neutrino coupling to the muon. Since the current limits exclude most models with $M_{W_R} = 5$ TeV and $m_N < 4$ TeV, we can conclude that any model with a W_R of 5 TeV and a Majorana neutrino with $m_N \leq 4.5$ TeV can be excluded after the HL-LHC if no evidence is seen for this process[6]. Model 5 is very close to being possible to exclude at 95% CL. In the electron channel, the upper limit on the relative coupling for this model is $\kappa_R^{\mathrm{up}} = 1.3$, while in the muon channel it is $\kappa_R^{\mathrm{up}} = 1.03$. If we relax the confidence requirement, we find that model 5mu is expected to be excluded at 90% CL with an upper limit of $\kappa_R^{\mathrm{up}} = 0.89$, while model 5e is excluded at 80% CL with $\kappa_R^{\mathrm{up}} = 0.94$. Finally, model 6 is not expected to be excluded. We can therefore expect that after the HL-LHC, the 95% CL excluded models with $M_{W_R} = 6$ TeV cover all heavy Majorana neutrino masses up to around 4 TeV, while m_N slightly above 4 TeV can be excluded at 90% CL, under the assumption $\kappa_R = 1$.

[6]The case where m_N and M_{W_R} are nearly degenerate might require special treatment, so we can not necessarily exclude all models with $M_{W_R} = 5$ TeV and $m_N < M_{W_R}$.

Chapter 8

Conclusion and outlook

The seesaw mechanism provides a natural explanation for the small masses of the Standard Model (SM) neutrinos. We have covered the theory behind three versions of this mechanism that include new heavy neutrinos that could interact with the SM particles. In two of these versions, the heavy neutrinos couple to the SM W, Z and Higgs bosons, and the heavy neutrinos are either Majorana or pseudo-Dirac fermions. We call these the Majorana and pseudo-Dirac model respectively. The third version of the seesaw mechanism we studied is contained in the framework of the Left-Right Symmetric Model (LRSM), and includes heavy Majorana neutrinos that couple only to a new right handed gauge boson W_R and SM charged leptons of a specific flavour, while the W_R also couples to SM quarks. For each of the three models we chose one particular interaction of the heavy neutrinos to study extensively using Monte Carlo (MC) simulations at next-to leading order in QCD. These interactions were chosen because they are well suited for detection in a proton-proton collider experiment like the LHC, and the goal of studying the interactions was to discover characteristics of the interactions that could be used to aid in their discovery or give information about the underlying model if they were discovered. For each model we also chose masses for the heavy neutrinos based on the requirements for the seesaw mechanism, previous exclusion limits and accessibility at the LHC. For the Majorana and pseudo-Dirac models we studied masses in the range $200 - 1200$ GeV, while for the LRSM we studied W_R masses from 3 TeV to 6 TeV and neutrino masses from 1.5 TeV to 5 TeV.

For the Majorana and pseudo-Dirac models, we chose to study the production process of a heavy neutrino from an initial state of two quarks through a W boson, and with decay into a final state with three charged leptons and a SM neutrino. When studying the kinematics of this process, we discovered that the transverse momenta p_T and the transverse masses m_T of the charged leptons, the azimuthal angle difference $\Delta\phi$ between the charged leptons and the SM neutrino, and the invariant mass of the three charged leptons m_{3l} are all sensitive to the mass of the heavy neutrino. These could therefore potentially be used to estimate the mass of the heavy neutrino if this process is detected. We also found that, as expected, the two leptons originating from the production and decay vertex of the heavy neutrino will always have charges with opposite signs in the pseudo-Dirac model, while in the Majorana model the signs are opposite or the same with equal probability. This means that the signs of the charges of these leptons can be used to distinguish between the two models in an experiment, if the correct pair of leptons can be identified. We therefore used our study of the process kinematics to develop a method for identifying this pair of leptons. We found that the best method is to choose the pair among the three leptons with the largest p_T that have the largest value of the product $p_T \times \Delta\phi$. This method has a success rate of around 1/3 for the lowest neutrino mass we considered, of 200 GeV, and above 90% for neutrino masses above 800 GeV.

For the LRSM, we studied the famous Keung-Senjanović process for the case of $m_N < M_{W_R}$, where m_N and M_{W_R} are the masses of the heavy neutrino and W_R boson respectively. This process has a final state of two charged leptons and two hadronic jets. We found that this allows relatively precise determination of M_{W_R} in experiments from the invariant mass of all four final state objects, m_{lljj}. We also discovered that the transverse momenta of the two charged leptons can be used to determine the mass difference between the W_R boson and the heavy Majorana neutrino. This is especially the case when this mass difference is small, in which case we showed that one of the charged leptons has an upper limit on its transverse momentum approximately equal to the mass difference.

In the final part of our analysis we used the signal processes we simulated to study the sensitivity of the ATLAS experiment at the LHC to these processes for the current Run 3 and the future High Luminosity LHC (HL-LHC). For this study, we used exclusively MC simulated background that was generated for Run 3 ATLAS analyses. For the analysis of the LRSM process we implemented smearing on the signal and background MC events, which mimics the full ATLAS reconstruction algorithms. We created signal regions for the Majorana and pseudo-Dirac signals by performing cuts on the variables that are most efficient in separating the signal from background, and separated them into same sign (SS) and opposite sign (OS) regions using the lepton identification method we developed. Even with these highly specialised signal regions, we find that the strict limits on the mixing matrix elements coupling the heavy neutrinos to the SM W boson makes it impossible to discover these signals during Run 3. After the entire HL-LHC run with an integrated luminosity of 3000 fb^{-1}, we find that even for an optimistic choice of mixing matrix elements, the signal with the highest expected significance is the Majorana model with $m_N = 200$ GeV, with 1.6σ. With reasonable cuts, none of the pseudo-Dirac signals have expected significances above 1σ, which is mainly due to the larger background in the OS region. We also computed the expected exclusion limits on the cross sections of the Majorana and pseudo-Dirac models using Bayesian statistics. We find that none of the models we studied can be excluded at 95% CL at the HL-LHC.

For the LRSM signals we created a signal region that performs better than the one used in the most recent study on the Keung-Senjanović process from the ATLAS collaboration [61], in the sense that it gives higher expected significances for Run 3. For two models with $M_{W_R} = 5$ TeV and $m_N = 4$ TeV or $m_N = 4.5$ TeV, we increase the expected significance from their 0.35σ (0.51σ) and 0.097σ (0.14σ) to 2.8σ (3.4σ) and 0.97σ (1.3σ) for the case when the heavy neutrino couples to the electron (muon). This means that using our method, we could expect to see evidence of the model with $m_N = 4$ TeV, if it exists. In that case we also show that the distribution of m_{lljj} could be used to estimate M_{W_R}, at least with a precision of around ± 0.5 TeV. After the entire HL-LHC run, we find that our method is expected to give a significance above the discovery threshold of 5σ for the $M_{W_R} = 5$ TeV, $m_N = 4$ TeV model if it exists and under the assumption that the coupling constant g_R of W_R is equal to the coupling constant g of the SM W boson. Models with $M_{W_R} = 5$ TeV, $m_N = 4.5$ TeV and $M_{W_R} = 6$ TeV, $m_N = 4$ TeV have expected significances of 3.4σ (4.7σ) and 2.3σ (3.7σ) for coupling to electrons (muons), which means they will likely be noticeable, although discovery will be challenging after the HL-LHC. Finally, we determine the expected 95% CL upper limits on the ratio $\kappa_R = g_R/g$ for the six different models after both Run 3 and HL-LHC. We find that after Run 3, models with $M_{W_R} = 3$ TeV, and m_N between 1.5 TeV and 2 TeV, which have already been excluded under the natural assumption $\kappa_R = 1$, will get upper limits on κ_R of around 0.2 if no signal is seen. None of the signals we studied with $M_{W_R} = 5$ TeV and $M_{W_R} = 6$ TeV and $m_N > 4$ TeV can be excluded after Run 3. After the HL-LHC meanwhile, the expected limits on κ_R for the models with $M_{W_R} = 3$ TeV drop below 0.1, and models with $M_{W_R} = 5$ TeV and m_N between 4 and 4.5 TeV can be excluded at 95% CL for neutrinos coupling to either electrons or muons. This would expand the current excluded region from the ATLAS collaboration [61]. Finally, the case with $M_{W_R} = 6$ TeV, $m_N = 4$ TeV is expected to be excluded at 90% (80%) CL if the neutrino couples to electrons (muons).

Although we have found that the discovery or exclusion of Majorana and pseudo-Dirac neutrinos that couple to the W boson with masses on the order of 100 GeV to 1 TeV is very difficult, our analysis is far from exhaustive. We focused on one specific signal process, and there are several other options that could give better results, in particular if one were to combine searches in multiple channels. Among the possible processes to study there is the process analogous to the Keung-Senjanović process, with the W_R replaced with a SM W, giving a final state of two leptons and two jets. There are also the processes where the heavy neutrino decays to either a Z boson or a Higgs, giving even more possible final states with new kinematics. In addition to these, there are other next-to leading order production processes that can be important at the LHC, such as W-photon and W-Z fusion [59]. For the LRSM case, we have found promising signs that current and future runs of the LHC will provide sensitivity to new regions of the mass parameter space. However, when extrapolating our analysis to the HL-LHC we made some simplifying assumptions that most likely made our results slightly better than if we performed a proper analysis with background simulated for the HL-LHC. In addition to this, when simulating the production processes for large M_{W_R} we found that we got large uncertainties in the cross section and a large fraction of negatively weighted events, which may have made our results less reliable. In future studies, it may be desirable to develop new techniques to better model the production of W_R bosons with large masses.

In addition to the seesaw types we have looked at, there are also the so-called type 2 [68] and type 3 [69] seesaw, with different particle contents and therefore completely different phenomenology. There has also recently been done studies on the possibilities of generating neutrino masses at 1-loop order [70], which gives an essentially endless number of possible models each with their own new particles.

All in all, it is an exciting time for neutrino physics, as both theoretical and experimental research is closing in on the mysterious properties of the neutrinos, including the origin of their masses and their Dirac or Majorana nature. Hopefully, it will not be long before these mysteries are uncovered, and with it new doors to physics beyond the Standard Model could be opened.

Appendices

Appendix A

Plots used for making cuts

In this appendix we show the plots used to choose the cuts that define the signal regions for both the three lepton final state and the two leptons and two jets final state. Both the signals and background are Monte Carlo generated at $\sqrt{s} = 13.6$ TeV, and the number of events is scaled to an integrated luminosity of 300 fb^{-1}. Each plot shows the distributions of a kinematic variable for both the SM background and a specific signal. The events have not passed any cuts, apart from the object definitions in tables 7.1 and 7.6 for the three lepton and $2l + 2j$ final states respectively. Then for the three lepton final state, events are required to have at least three leptons in each events, and for the $2l + 2j$ final state, events are required to have at least two leptons and at least two jets following the object definitions.

Each plot contains two subplots. The top subplot shows the distributions for both signal and background, which are stacked with the background on top of the signal. The bottom subplot shows the ratio between the number of signal events s and the number of background events b in each bin of the histograms above. The cuts made based on the distributions are also marked with dashed lines in the plots. In some cases, the cut keeps only events with higher values of the kinematic variable, while in other cases only events with lower values are kept. The precise statements of the cuts are also given in tables 7.2 and 7.8. The cuts are chosen in an attempt to simultaneously maximize the amount of signal event that pass the cuts, while also maximizing the ratio s/b in the signal region.

A.1 Three lepton final state

In this section, we look at the signal processes for three lepton final state, meaning the simple Majorana or pseudo-Dirac heavy neutrino extensions of the Standard Model, with no new gauge bosons. First we find the cuts that are shared between the signal regions SS and OS. For this we do not care about the Majorana/pseudo-Dirac nature of the heavy neutrino, since we expect their kinematics to be similar. We therefore choose two versions of the extension of the SM by a single heavy Majorana neutrino N_1, one model with $m_N = 200$ GeV, and one with $m_N = 800$ GeV. In both cases the heavy neutrino is assumed to couple most strongly to the electron, with $V_{e1} = 0.1$, while $V_{\mu 1} = 0.1/\sqrt{10}$ and $V_{\tau 1} = 0.01$. The strictness of the cuts are mainly limited by the model with $m_N = 200$ GeV, since this is the lowest value of m_N we will study. We see that almost all the cuts remove a larger fraction of signal events from the 200 GeV model than the 800 GeV model. The reason for this is that the signals with lower m_N have

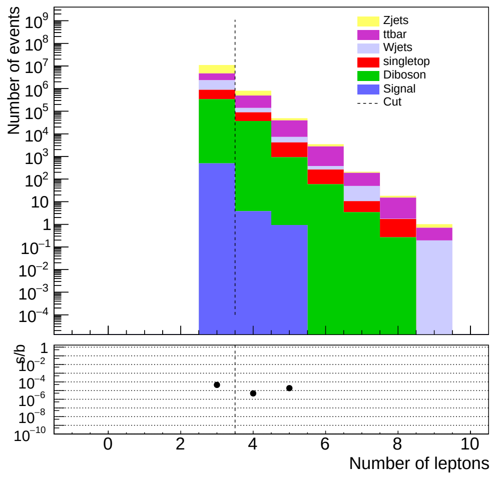

(a) $m_N = 200$ GeV

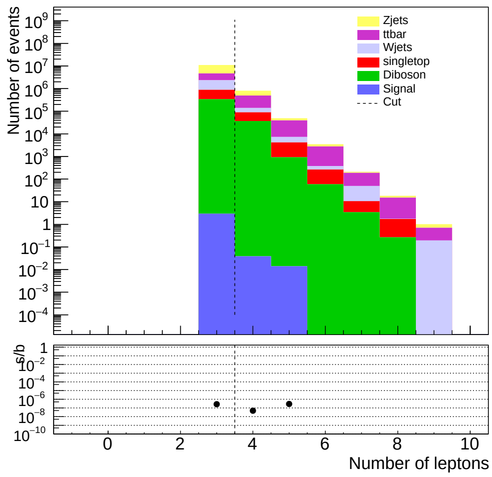

(b) $m_N = 800$ GeV

Figure A.1: The distributions of the number of leptons in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

kinematic distributions that are more similar to the background than the signals with larger m_N. But since models with lower m_N also give larger cross sections, we can afford to remove a larger fraction of signal events and still have enough events remaining. The plots of the two signals along with the SM background and the cuts made are shown in figures A.1–A.14.

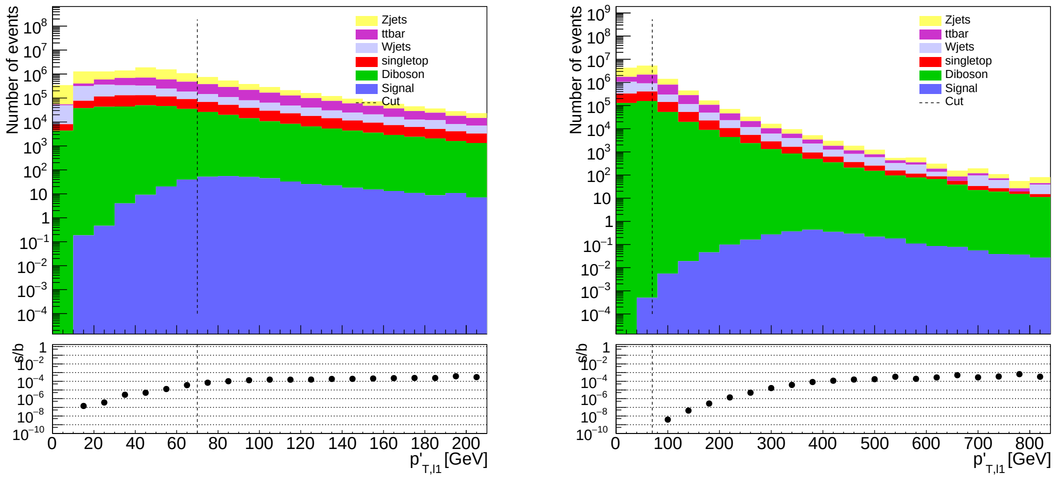

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.2: The distributions of the p_T of the leading lepton in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

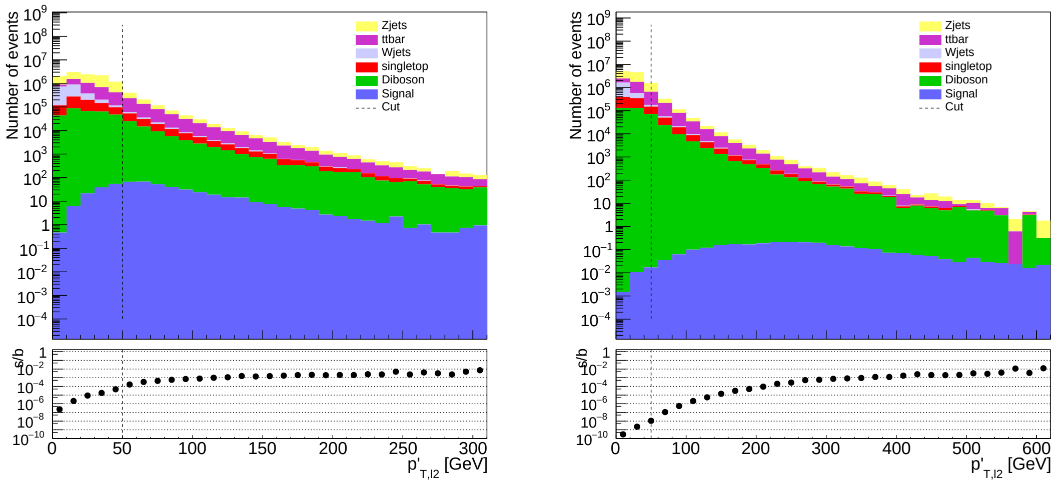

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.3: The distributions of the p_T of the second leading lepton in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

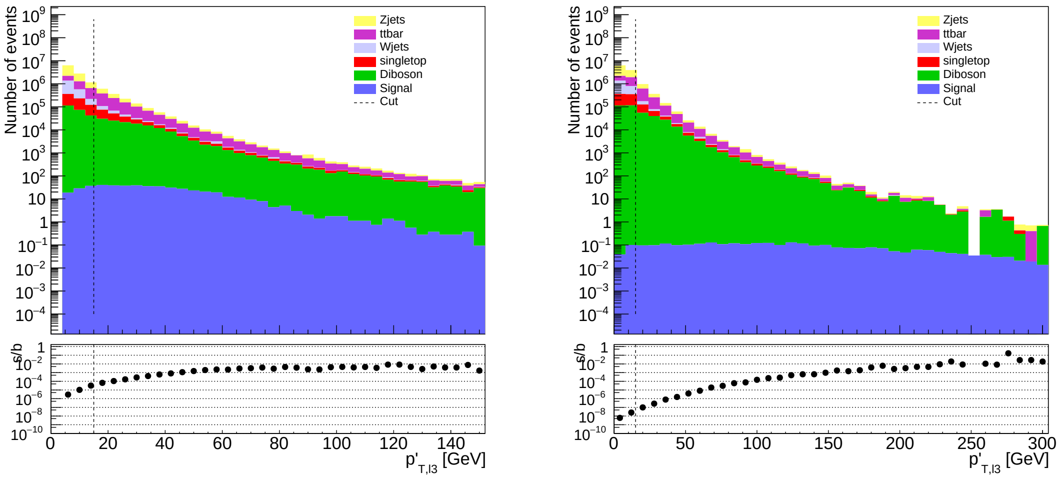

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.4: The distributions of the p_T of the third leading lepton in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

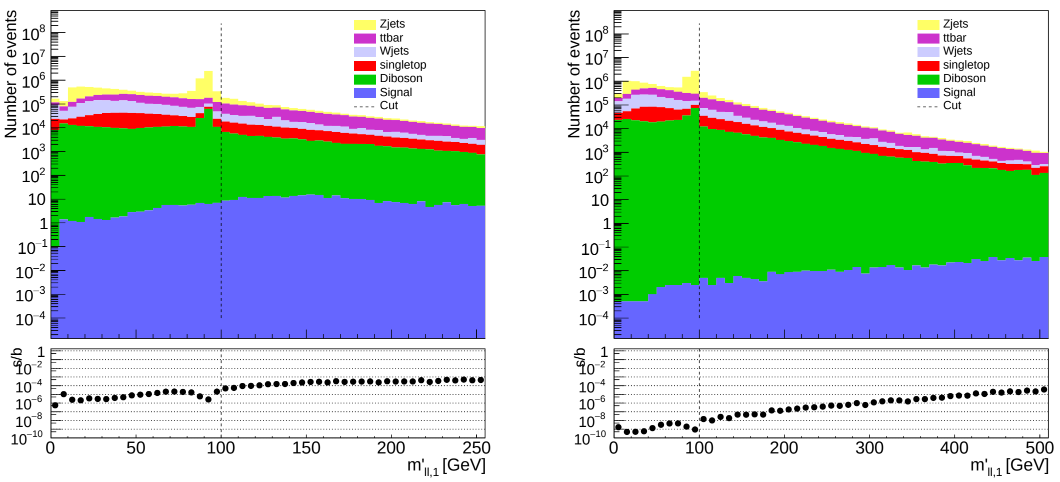

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.5: The distributions of the invariant mass of lepton 1 and 2 when ordered by p_T in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

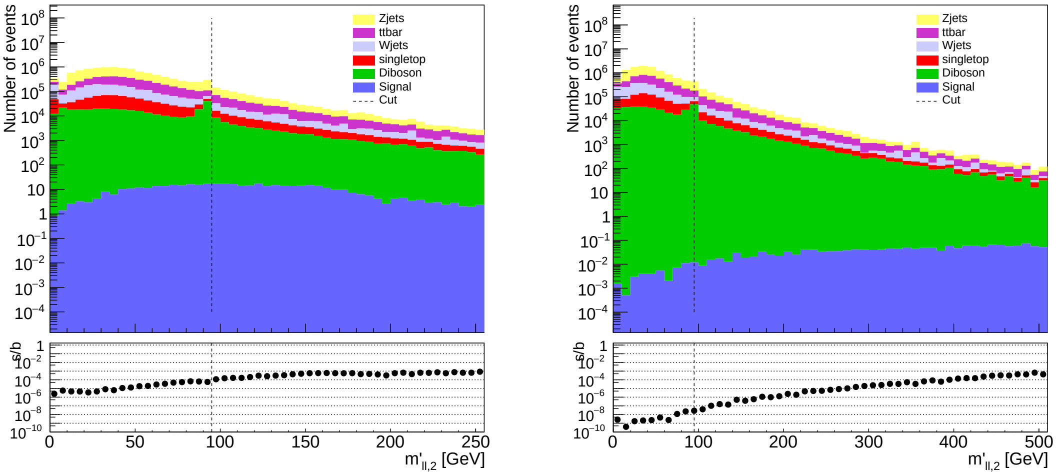

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.6: The distributions of the invariant mass of lepton 1 and 3 when ordered by p_T in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

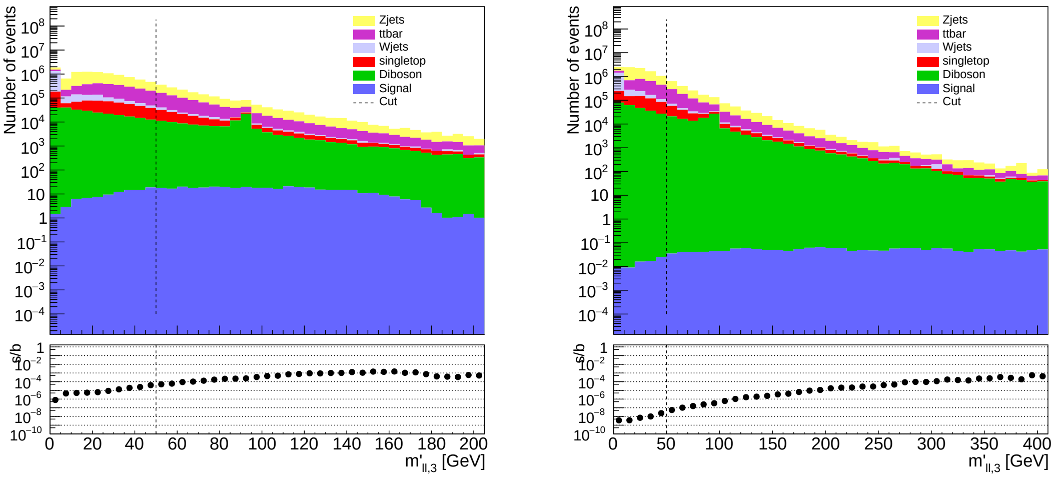

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.7: The distributions of the invariant mass of lepton 2 and 3 when ordered by p_T in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

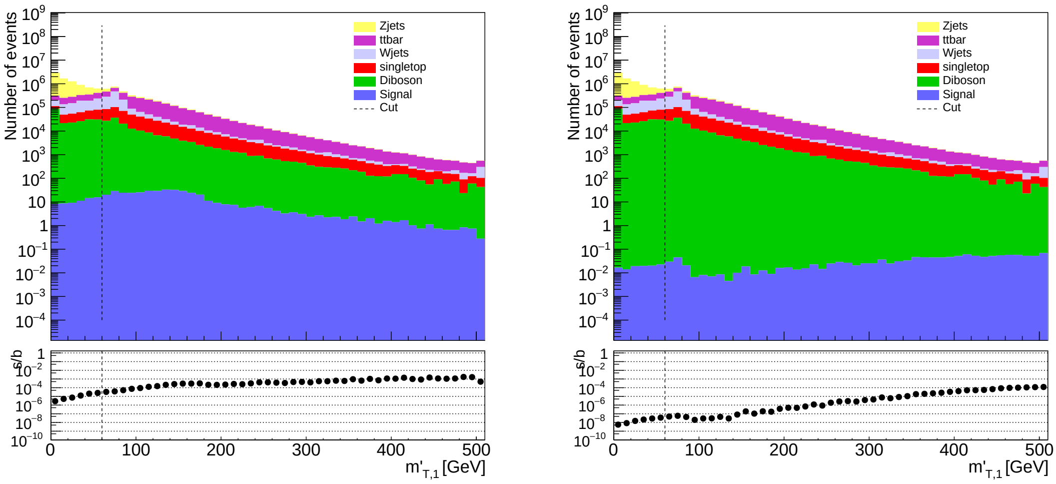

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.8: The distributions of the transverse mass of the leading charged lepton and the final state neutrino in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

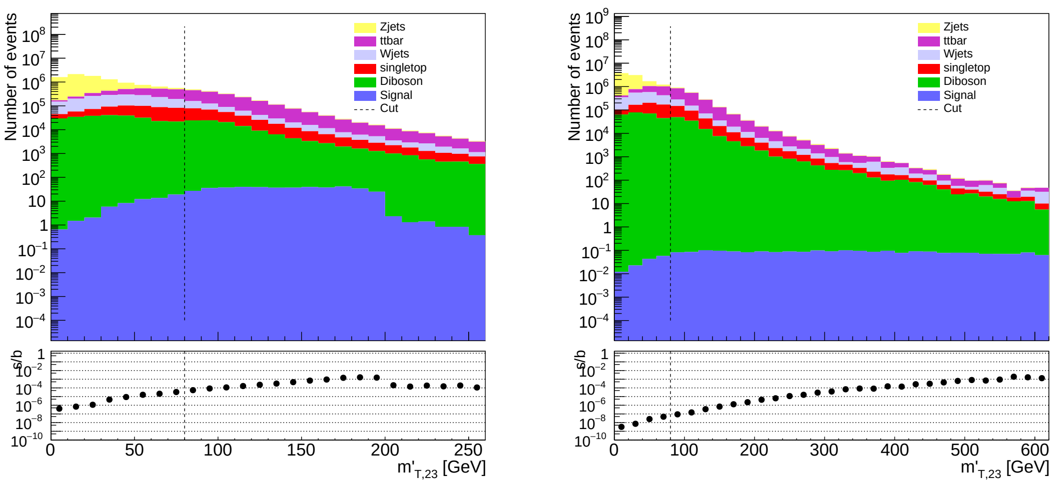

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.9: The distributions of the transverse mass of the second leading charged lepton, third leading charged lepton and the neutrino in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

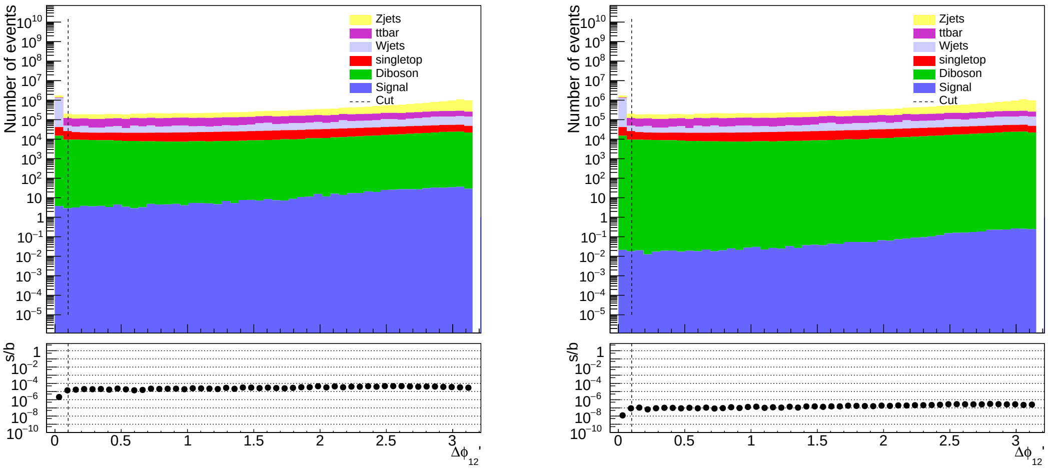

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.10: The distributions of the $\Delta\phi$ of the leading and second leading charged leptons in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

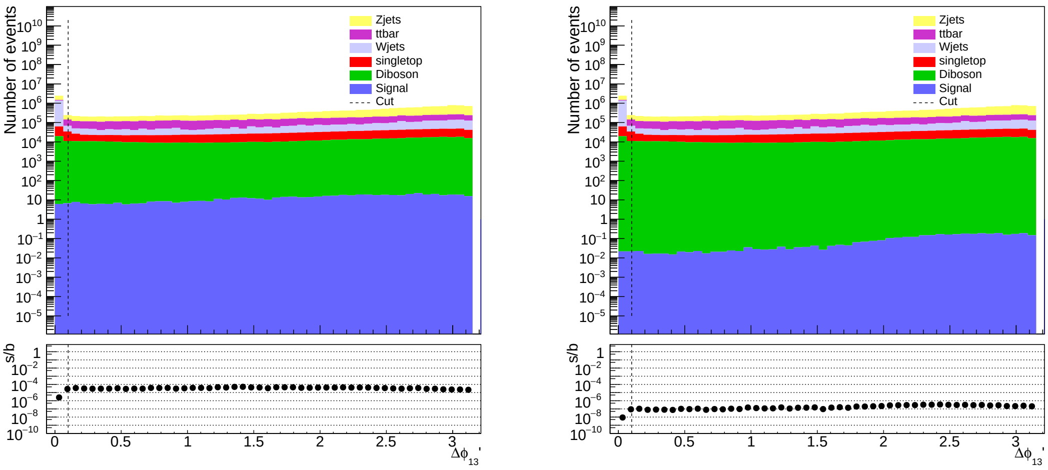

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.11: The distributions of the $\Delta\phi$ of the leading and third leading charged leptons in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

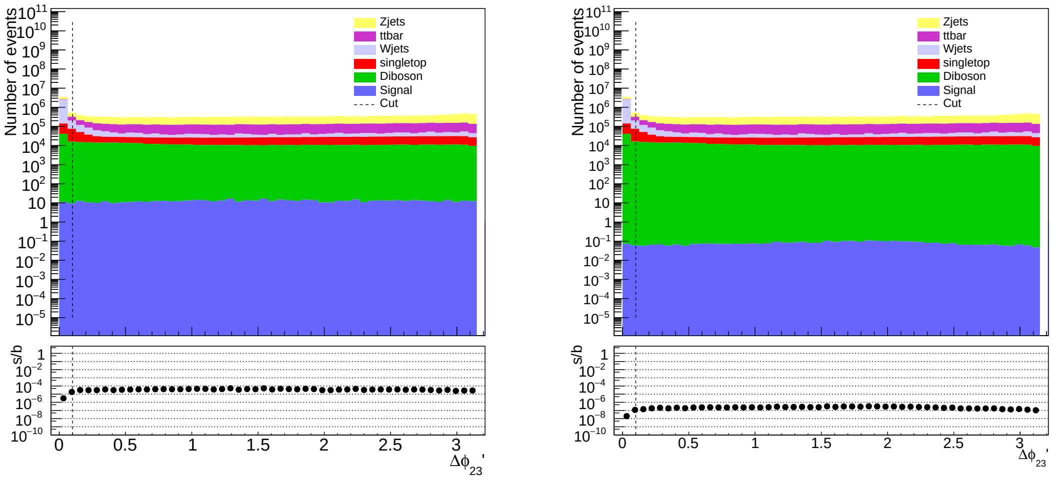

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.12: The distributions of the $\Delta\phi$ of the second and third leading charged leptons in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

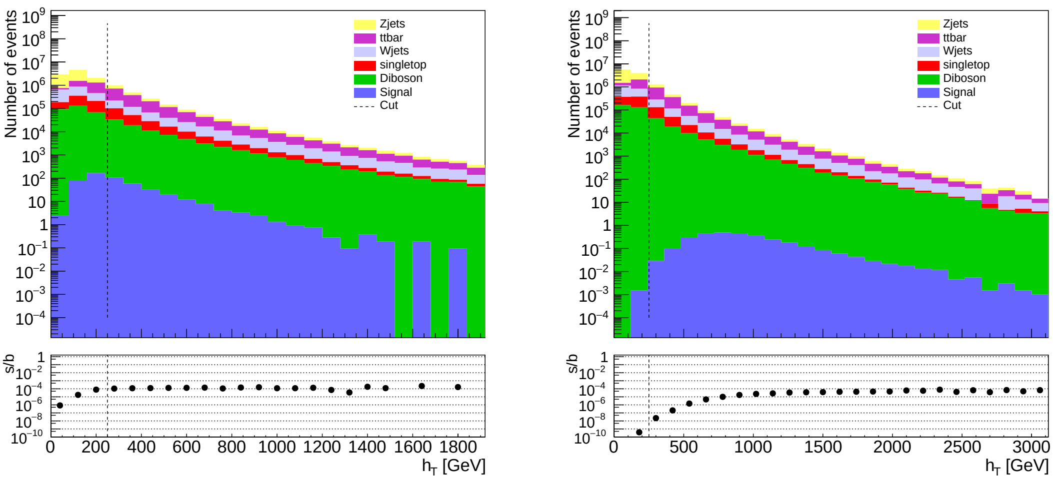

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.13: The distributions of the hadronic activity in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

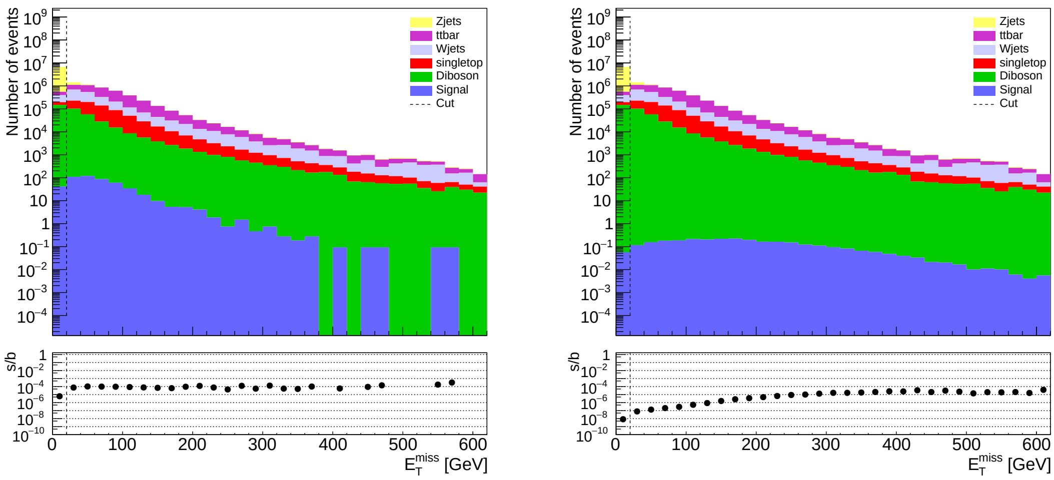

(a) $m_N = 200$ GeV (b) $m_N = 800$ GeV

Figure A.14: The distributions of the missing transverse energy E_T^{miss} in each event for the full SM background and the two different Majorana signal models, and the cut that was made based on these.

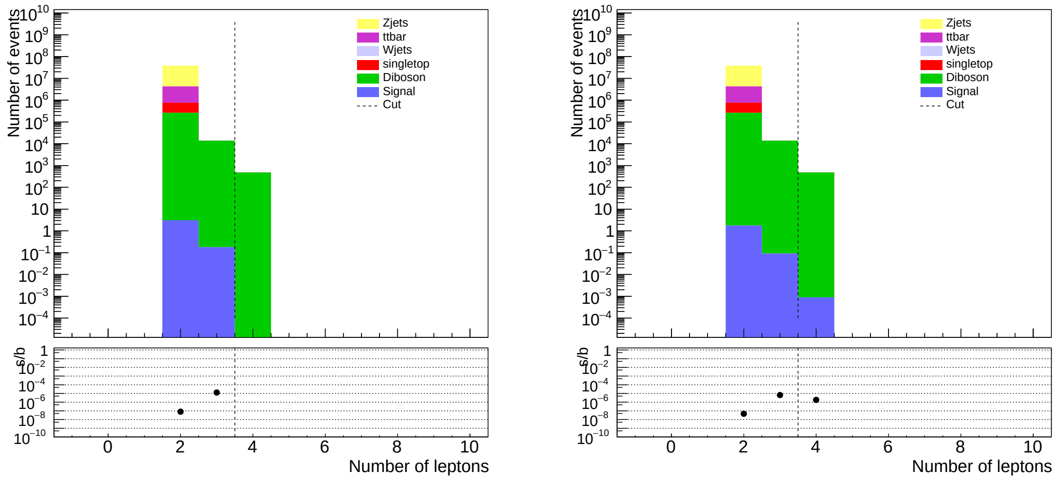

(a) Model 4e

(b) Model 5e

Figure A.15: The distributions of the number of leptons in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

A.2 2l+2j final state

In this section, we look at the signal processes for the LRSM model with the two lepton and two jet final state. We want to choose the cuts that define our single signal region, which is subsequently split into a electron channel and a muon channel. As we saw in section 6.5, the relative mass difference is one of the most important factors for the kinematics of the signal. We therefore use models 4 and 5 for the kinematical cuts, since model 4 has the smallest relative difference between m_N and M_{WR}, while model 5 has the largest relative difference of the models 3–6, which are the main focus of our analysis. The plots of these two signals along with the SM background and the cuts made are shown in figures A.15–A.26.

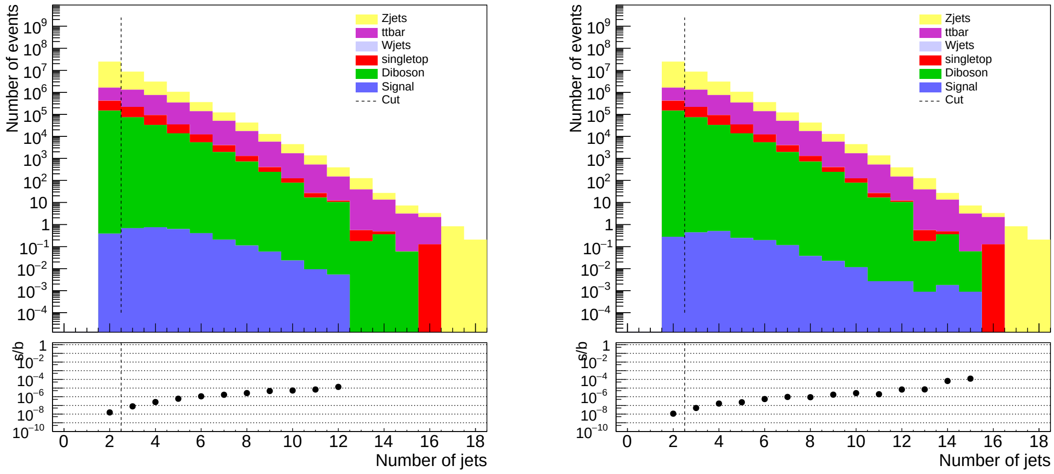

(a) Model 4e (b) Model 5e

Figure A.16: The distributions of the number of jets in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

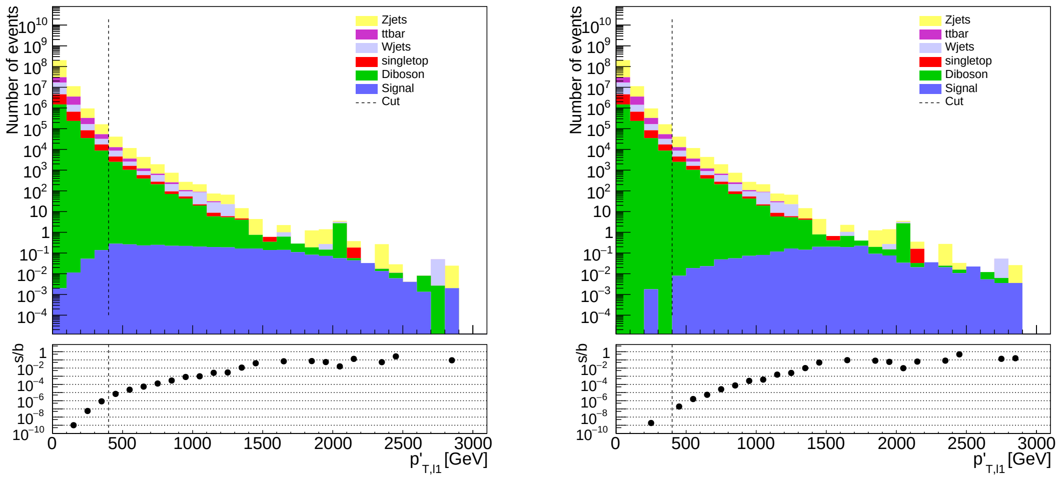

(a) Model 4e (b) Model 5e

Figure A.17: The distributions of the transverse momentum of the leading lepton in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

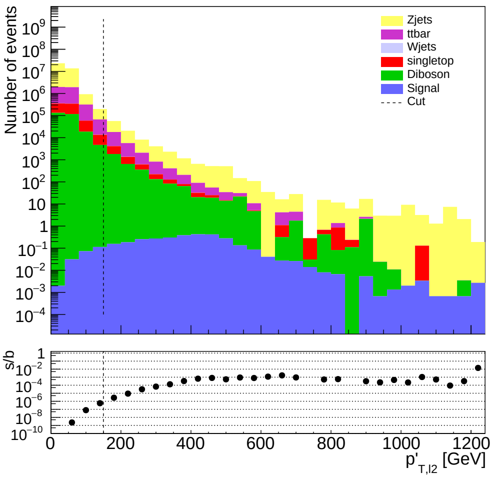

(a) Model 4e

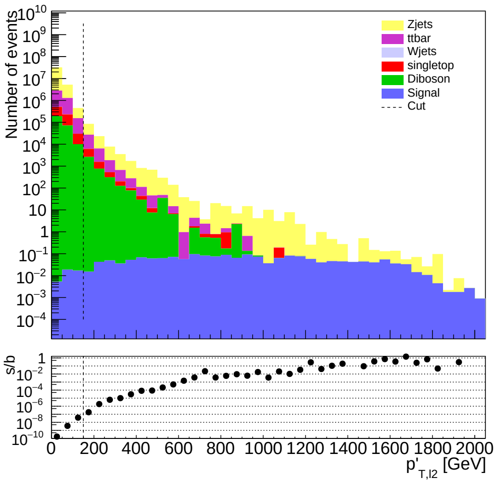

(b) Model 5e

Figure A.18: The distributions of the transverse momentum of the second leading lepton in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

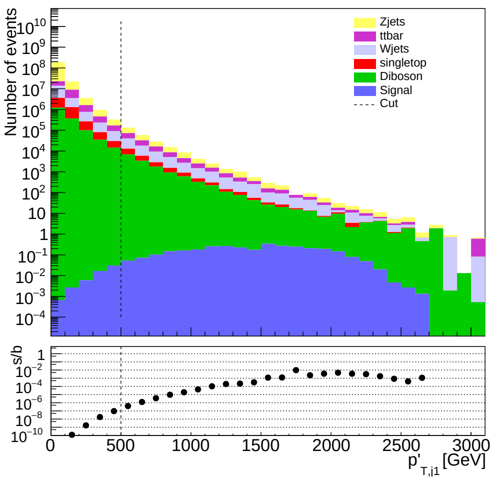

(a) Model 4e

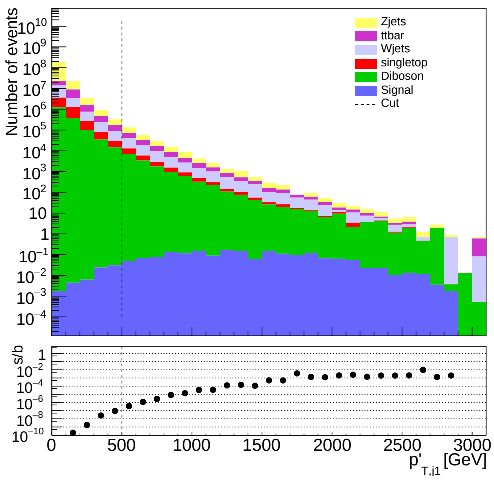

(b) Model 5e

Figure A.19: The distributions of the transverse momentum of the leading jet in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

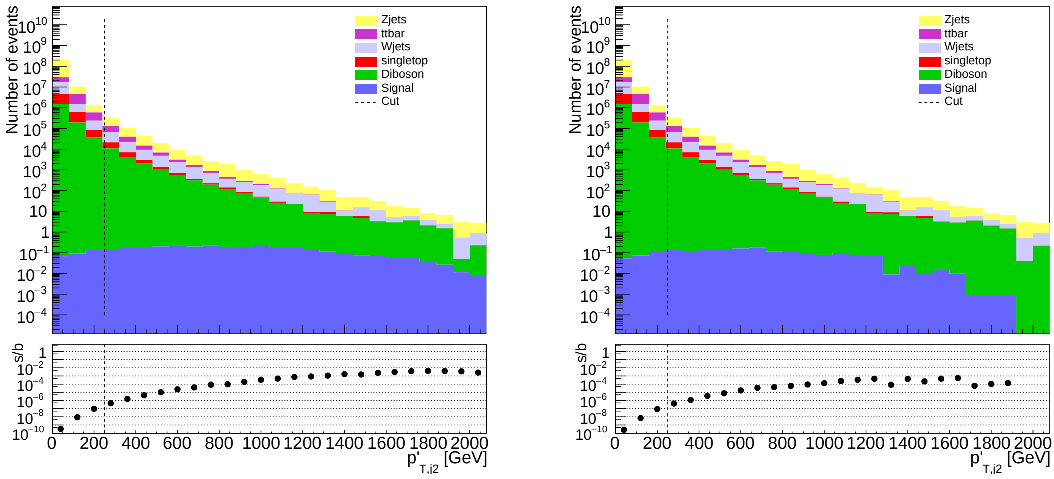

(a) Model 4e (b) Model 5e

Figure A.20: The distributions of the transverse momentum of the second leading jet in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

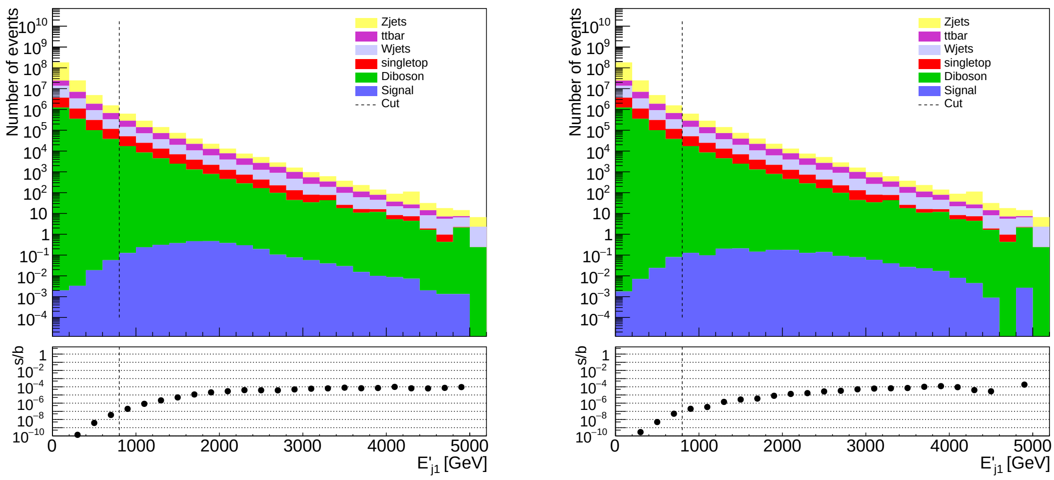

(a) Model 4e (b) Model 5e

Figure A.21: The distributions of the energy of the leading jet in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

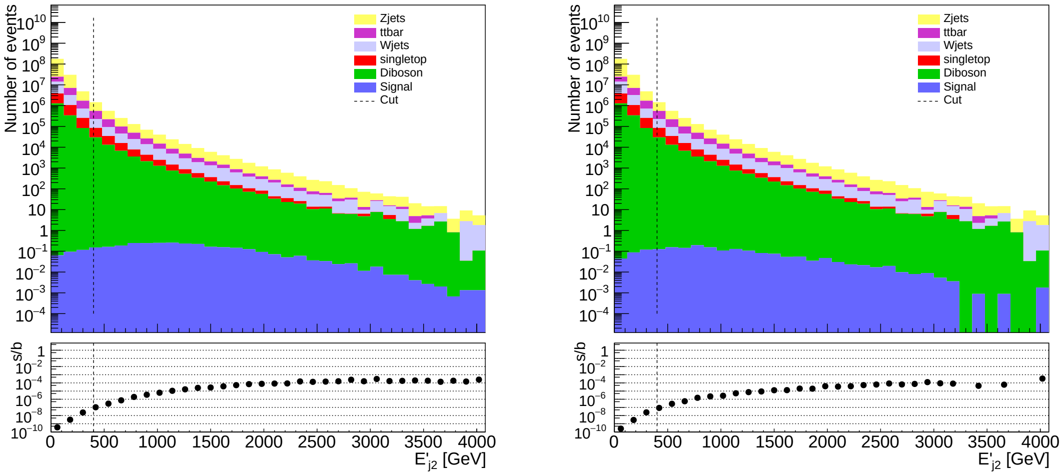

(a) Model 4e

(b) Model 5e

Figure A.22: The distributions of the energy of the second leading jet in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

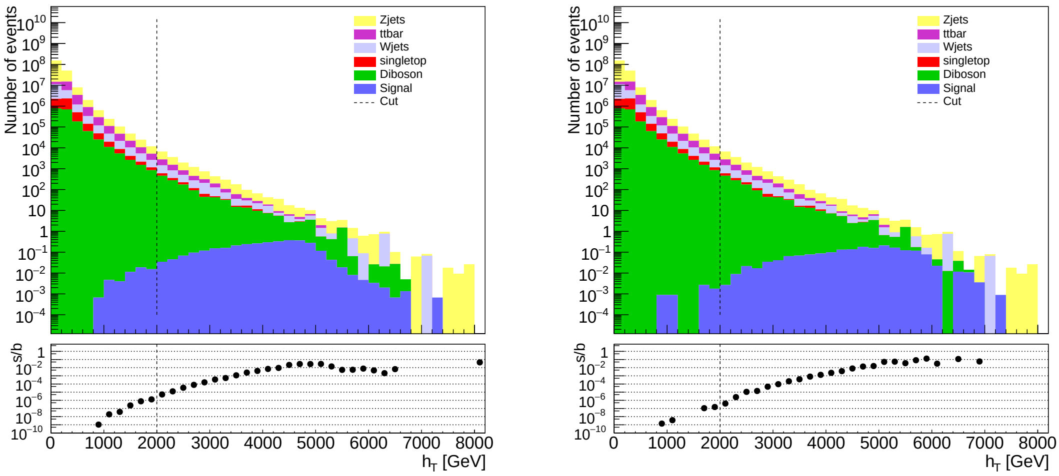

(a) Model 4e

(b) Model 5e

Figure A.23: The distributions of the hadronic activity in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

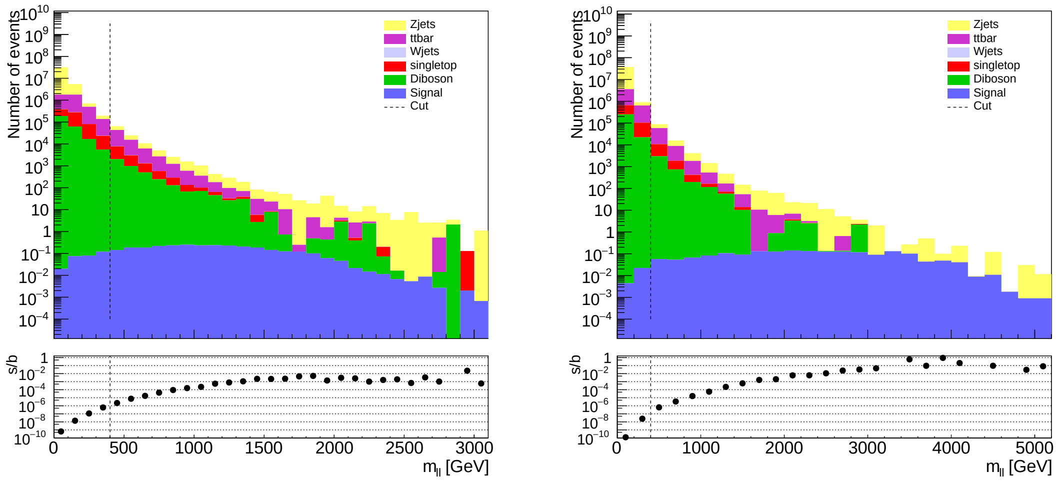

(a) Model 4e

(b) Model 5e

Figure A.24: The distributions of the invariant mass of the two leading leptons in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

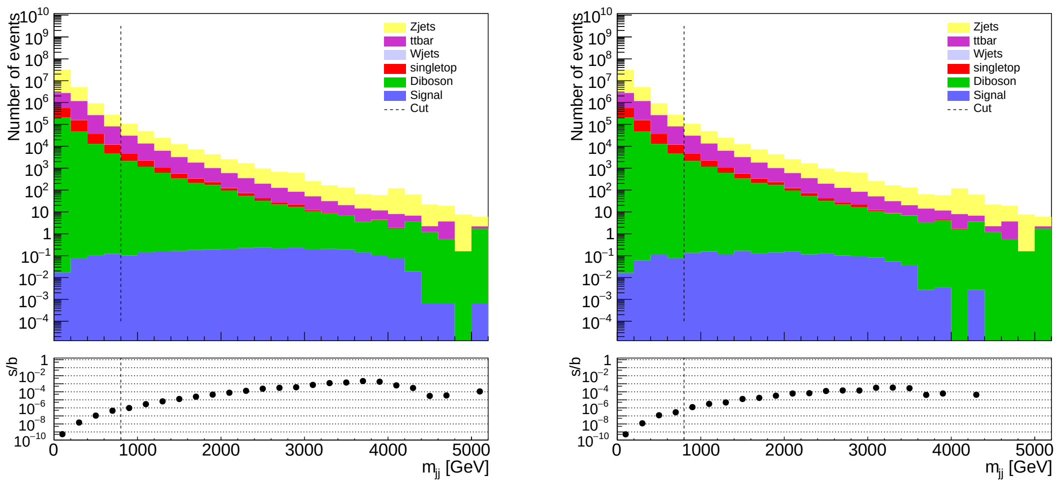

(a) Model 4e

(b) Model 5e

Figure A.25: The distributions of the invariant mass of the two leading jets in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

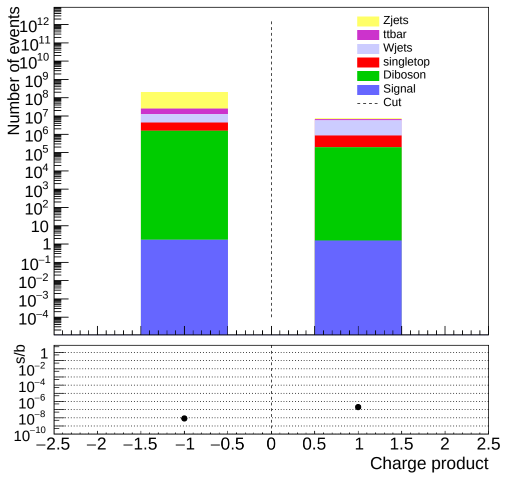

(a) Model 4e

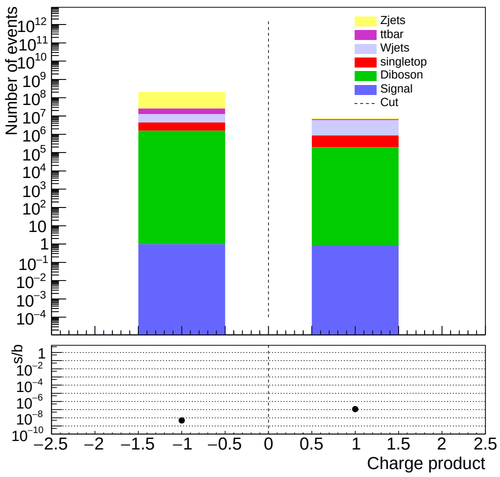

(b) Model 5e

Figure A.26: The distributions of the product of the charges of the two leading leptons in each event for the full SM background and the two different LRSM signal models, and the cut that was made based on these.

Bibliography

[1] Frederick Reines and Clyde L. Cowan. 'The neutrino'. In: *Nature* 178 (1956), pp. 446–449. DOI: 10.1038/178446a0.

[2] Y. Fukuda et al. 'Evidence for Oscillation of Atmospheric Neutrinos'. In: *Physical Review Letters* 81.8 (Aug. 1998), pp. 1562–1567. DOI: 10.1103/physrevlett.81.1562. URL: https://doi.org/10.1103%2Fphysrevlett.81.1562.

[3] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Reading, USA: Addison-Wesley, 1995. ISBN: 978-0-201-50397-5.

[4] Mark Thomson. *Modern particle physics*. New York: Cambridge University Press, 2013. ISBN: 978-1-107-03426-6. DOI: 10.1017/CBO9781139525367.

[5] F. Englert and R. Brout. 'Broken Symmetry and the Mass of Gauge Vector Mesons'. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.321.

[6] Peter W. Higgs. 'Broken Symmetries and the Masses of Gauge Bosons'. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.508.

[7] S. L. Glashow. 'Partial Symmetries of Weak Interactions'. In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.

[8] Marco Drewes. 'The Phenomenology of Right Handed Neutrinos'. In: *Int. J. Mod. Phys. E* 22 (2013), p. 1330019. DOI: 10.1142/S0218301313300191. arXiv: 1303.6912 [hep-ph].

[9] Ta-Pei [0000-0002-1137-0969] Cheng and Ling-Fong [0000-0002-8035-3329] Li. *Gauge Theory of Elementary Particle Physics*. Oxford, UK: Oxford University Press, 1984. ISBN: 978-0-19-851961-4.

[10] J. R. Batley et al. 'A Precision measurement of direct CP violation in the decay of neutral kaons into two pions'. In: *Phys. Lett. B* 544 (2002), pp. 97–112. DOI: 10.1016/S0370-2693(02)02476-0. arXiv: hep-ex/0208009.

[11] B. Pontecorvo. 'Neutrino Experiments and the Problem of Conservation of Leptonic Charge'. In: *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

[12] Evgeny Kh. Akhmedov. 'Do charged leptons oscillate?' In: *JHEP* 09 (2007), p. 116. DOI: 10.1088/1126-6708/2007/09/116. arXiv: 0706.1216 [hep-ph].

[13] R. L. Workman et al. 'Review of Particle Physics'. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

[14] Ubaldo Dore, Pier Loverre and Lucio Ludovici. 'History of accelerator neutrino beams'. In: *The European Physical Journal H* 44.4 (Nov. 2019), pp. 271–305. ISSN: 2102-6467. DOI: 10.1140/epjh/e2019-90032-x. URL: https://doi.org/10.1140/epjh/e2019-90032-x.

[15] L.J. Wen, J. Cao and Y.F. Wang. 'Reactor Neutrino Experiments: Present and Future'. In: *Annual Review of Nuclear and Particle Science* 67.1 (Oct. 2017), pp. 183–211. DOI: 10.1146/annurev-nucl-101916-123318. URL: https://doi.org/10.1146%2Fannurev-nucl-101916-123318.

[16] Yoichiro Suzuki. 'Solar neutrinos'. In: *Int. J. Mod. Phys. A* 15S1 (2000). Ed. by J. Jaros and Michael E. Peskin, pp. 201–228. DOI: 10.1142/S0217751X00005164.

[17] Raymond Davis. 'A review of the homestake solar neutrino experiment'. In: *Progress in Particle and Nuclear Physics* 32 (1994), pp. 13–32. ISSN: 0146-6410. DOI: https://doi.org/10.1016/0146-6410(94)90004-3. URL: https://www.sciencedirect.com/science/article/pii/0146641094900043.

[18] Q. R. Ahmad et al. 'Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory'. In: *Phys. Rev. Lett.* 89 (1 June 2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301. URL: https://link.aps.org/doi/10.1103/PhysRevLett.89.011301.

[19] K. Eguchi et al. 'First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance'. In: *Phys. Rev. Lett.* 90 (2 Jan. 2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.90.021802.

[20] C. Mariani and for K2K Collaboration. *Recent Results from K2K Experiment.* 2005. arXiv: hep-ex/0505019 [hep-ex].

[21] K. Abe et al. 'Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target'. In: *Phys. Rev. D* 91.7 (2015), p. 072010. DOI: 10.1103/PhysRevD.91.072010. arXiv: 1502.01550 [hep-ex].

[22] F. P. An et al. 'New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay'. In: *Physical Review Letters* 115.11 (Sept. 2015). DOI: 10.1103/physrevlett.115.111802. URL: https://doi.org/10.1103%2Fphysrevlett.115.111802.

[23] NuFIT collaboration. *NuFIT 5.2.* Accessed on 2 June 2023. 2022. URL: http://www.nu-fit.org/.

[24] Ivan Esteban et al. 'The fate of hints: updated global analysis of three-flavor neutrino oscillations'. In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). DOI: 10.1007/jhep09(2020)178. URL: https://doi.org/10.1007%2Fjhep09%282020%29178.

[25] B. Abi et al. 'Volume I. Introduction to DUNE'. In: *Journal of Instrumentation* 15.08 (Aug. 2020), T08008. DOI: 10.1088/1748-0221/15/08/T08008. URL: https://dx.doi.org/10.1088/1748-0221/15/08/T08008.

[26] K. Abe et al. 'Hyper-Kamiokande Design Report'. In: (May 2018). arXiv: 1805.04163 [physics.ins-det].

[27] A. Alekou et al. 'The European Spallation Source neutrino super-beam conceptual design report'. In: *Eur. Phys. J. ST* 231.21 (2022), pp. 3779–3955. DOI: 10.1140/epjs/s11734-022-00664-w. arXiv: 2206.01208 [hep-ex].

[28] Joseph A. Formaggio, André Luiz C. de Gouvêa and R.G. Hamish Robertson. 'Direct measurements of neutrino mass'. In: *Physics Reports* 914 (June 2021), pp. 1–54. DOI: 10.1016/j.physrep.2021.02.002. URL: https://doi.org/10.1016%2Fj.physrep.2021.02.002.

[29] M. Aker et al. 'Direct neutrino-mass measurement with sub-electronvolt sensitivity'. In: *Nature Phys.* 18.2 (2022), pp. 160–166. DOI: 10.1038/s41567-021-01463-1. arXiv: 2105.08533 [hep-ex].

[30] Shadab Alam et al. 'Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory'. In: *Physical Review D* 103.8 (Apr. 2021). DOI: 10.1103/physrevd.103.083533. URL: https://doi.org/10.1103%2Fphysrevd.103.083533.

[31] Eleonora Di Valentino, Alessandro Melchiorri and Joseph Silk. 'Cosmological constraints in extended parameter space from the Planck 2018 Legacy release'. In: *Journal of Cosmology and Astroparticle Physics* 2020.01 (Jan. 2020), pp. 013–013. DOI: 10.1088/1475-7516/2020/01/013. URL: https://doi.org/10.1088%2F1475-7516%2F2020%2F01%2F013.

[32] Valentina Biancacci. 'The GERDA Experiment in the Search for Neutrinoless Double-Beta Decay'. In: *Moscow Univ. Phys. Bull.* 77.2 (2022), pp. 359–362. DOI: 10.3103/S0027134922020151.

[33] S. T. Petcov. 'The Nature of Massive Neutrinos'. In: *Adv. High Energy Phys.* 2013 (2013), p. 852987. DOI: 10.1155/2013/852987. arXiv: 1303.5819 [hep-ph].

[34] Evgeny Akhmedov. 'Majorana neutrinos and other Majorana particles:Theory and experiment'. In: Dec. 2014. arXiv: 1412.3320 [hep-ph].

[35] Anupama Atre et al. 'The Search for Heavy Majorana Neutrinos'. In: *JHEP* 05 (2009), p. 030. DOI: 10.1088/1126-6708/2009/05/030. arXiv: 0901.3589 [hep-ph].

[36] S. M. Bilenky and S. T. Petcov. 'Massive neutrinos and neutrino oscillations'. In: *Rev. Mod. Phys.* 59 (3 July 1987), pp. 671–754. DOI: 10.1103/RevModPhys.59.671. URL: https://link.aps.org/doi/10.1103/RevModPhys.59.671.

[37] Peter Minkowski. '$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?' In: *Phys. Lett. B* 67 (1977), pp. 421–428. DOI: 10.1016/0370-2693(77)90435-X.

[38] J. W. F. Valle. 'Gauge theories and the physics of neutrino mass'. In: *Prog. Part. Nucl. Phys.* 26 (1991), pp. 91–171. DOI: 10.1016/0146-6410(91)90010-L.

[39] R. N. Mohapatra and J. W. F. Valle. 'Neutrino Mass and Baryon Number Nonconservation in Superstring Models'. In: *Phys. Rev. D* 34 (1986), p. 1642. DOI: 10.1103/PhysRevD.34.1642.

[40] Ernest Ma and Rahul Srivastava. 'Dirac or inverse seesaw neutrino masses from gauged B–L symmetry'. In: *Mod. Phys. Lett. A* 30.26 (2015), p. 1530020. DOI: 10.1142/S0217732315300207. arXiv: 1504.00111 [hep-ph].

[41] Silvia Pascoli, Richard Ruiz and Cedric Weiland. 'Heavy neutrinos with dynamic jet vetoes: multilepton searches at $\sqrt{s} = 14$, 27, and 100 TeV'. In: *JHEP* 06 (2019), p. 049. DOI: 10.1007/JHEP06(2019)049. arXiv: 1812.08750 [hep-ph].

[42] Gerard 't Hooft. 'Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking'. In: *NATO Sci. Ser. B* 59 (1980). Ed. by Gerard 't Hooft et al., pp. 135–157. DOI: 10.1007/978-1-4684-7571-5_9.

[43] Michael Dine. 'Naturalness Under Stress'. In: *Ann. Rev. Nucl. Part. Sci.* 65 (2015), pp. 43–62. DOI: 10.1146/annurev-nucl-102014-022053. arXiv: 1501.01035 [hep-ph].

[44] Jogesh C. Pati and Abdus Salam. 'Lepton Number as the Fourth Color'. In: *Phys. Rev. D* 10 (1974). [Erratum: Phys.Rev.D 11, 703–703 (1975)], pp. 275–289. DOI: 10.1103/PhysRevD.10.275.

[45] Rabindra N. Mohapatra and Jogesh C. Pati. 'Left-Right Gauge Symmetry and an Isoconjugate Model of CP Violation'. In: *Phys. Rev. D* 11 (1975), pp. 566–571. DOI: 10.1103/PhysRevD.11.566.

[46] G. Senjanovic and Rabindra N. Mohapatra. 'Exact Left-Right Symmetry and Spontaneous Violation of Parity'. In: *Phys. Rev. D* 12 (1975), p. 1502. DOI: 10.1103/PhysRevD.12.1502.

[47] Manimala Mitra et al. 'Neutrino Jets from High-Mass W_R Gauge Bosons in TeV-Scale Left-Right Symmetric Models'. In: *Phys. Rev. D* 94.9 (2016), p. 095016. DOI: 10.1103/PhysRevD.94.095016. arXiv: 1607.03504 [hep-ph].

[48] Mu-Chun Chen, Sally Dawson and C. B. Jackson. 'Higgs Triplets, Decoupling, and Precision Measurements'. In: *Phys. Rev. D* 78 (2008), p. 093001. DOI: 10.1103/PhysRevD.78.093001. arXiv: 0809.4185 [hep-ph].

[49] Olivier Mattelaer, Manimala Mitra and Richard Ruiz. 'Automated Neutrino Jet and Top Jet Predictions at Next-to-Leading-Order with Parton Shower Matching in Effective Left-Right Symmetric Models'. In: (Oct. 2016). arXiv: 1610.08985 [hep-ph].

[50] Rabindra N. Mohapatra and Goran Senjanovic. 'Neutrino Mass and Spontaneous Parity Nonconservation'. In: *Phys. Rev. Lett.* 44 (1980), p. 912. DOI: 10.1103/PhysRevLett.44.912.

[51] 'Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using *pp* collisions at $\sqrt{s} = 13\ TeV$'. In: (May 2017).

[52] Morad Aaboud et al. 'Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV'. In: *Eur. Phys. J. C* 79.8 (2019), p. 639. DOI: 10.1140/epjc/s10052-019-7140-6. arXiv: 1902.04655 [physics.ins-det].

[53] Georges Aad et al. 'Muon reconstruction and identification efficiency in ATLAS using the full Run 2 *pp* collision data set at $\sqrt{s} = 13$ TeV'. In: *Eur. Phys. J. C* 81.7 (2021), p. 578. DOI: 10.1140/epjc/s10052-021-09233-2. arXiv: 2012.00578 [hep-ex].

[54] J. Alwall et al. 'The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations'. In: *Journal of High Energy Physics* 2014.7 (July 2014). DOI: 10.1007/jhep07(2014)079. URL: https://doi.org/10.1007%2Fjhep07%282014%29079.

[55] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. arXiv: 2203.11601 [hep-ph].

[56] Stefano Frixione and Bryan R. Webber. 'Matching NLO QCD computations and parton shower simulations'. In: *JHEP* 06 (2002), p. 029. DOI: 10.1088/1126-6708/2002/06/029. arXiv: hep-ph/0204244.

[57] Katharina Danziger, Stefan Höche and Frank Siegert. 'Reducing negative weights in Monte Carlo event generation with Sherpa'. In: (Oct. 2021). arXiv: 2110.15211 [hep-ph].

[58] Daniel Alva, Tao Han and Richard Ruiz. 'Heavy Majorana neutrinos from $W\gamma$ fusion at hadron colliders'. In: *JHEP* 02 (2015), p. 072. DOI: 10.1007/JHEP02(2015)072. arXiv: 1411.7305 [hep-ph].

[59] Celine Degrande et al. 'Fully-Automated Precision Predictions for Heavy Neutrino Production Mechanisms at Hadron Colliders'. In: *Phys. Rev. D* 94.5 (2016), p. 053002. DOI: 10.1103/PhysRevD.94.053002. arXiv: 1602.06957 [hep-ph].

[60] Wai-Yee Keung and Goran Senjanovic. 'Majorana Neutrinos and the Production of the Right-handed Charged Gauge Boson'. In: *Phys. Rev. Lett.* 50 (1983), p. 1427. DOI: 10.1103/PhysRevLett.50.1427.

[61] Georges Aad et al. 'Search for heavy Majorana or Dirac neutrinos and right-handed W gauge bosons in final states with charged leptons and jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: (Apr. 2023). arXiv: 2304.09553 [hep-ex].

[62] Alessio Maiezza et al. 'Left-Right Symmetry at LHC'. In: *Phys. Rev. D* 82 (2010), p. 055022. DOI: 10.1103/PhysRevD.82.055022. arXiv: 1005.5160 [hep-ph].

[63] D. Berdine, N. Kauer and D. Rainwater. 'Breakdown of the Narrow Width Approximation for New Physics'. In: *Phys. Rev. Lett.* 99 (2007), p. 111601. DOI: 10.1103/PhysRevLett.99.111601. arXiv: hep-ph/0703058.

[64] 'SimpleAnalysis: Truth-level Analysis Framework'. In: (2022). DOI: 10.17181/CERN.R6S3.0QKV.

[65] Kazuya Mochizuki. 'b-jet identification algorithms and performance in the ATLAS experiment'. In: *Nuclear and Particle Physics Proceedings* 273-275 (2016). 37th International Conference on High Energy Physics (ICHEP), pp. 2536–2538. ISSN: 2405-6014. DOI: https://doi.org/10.1016/j.nuclphysbps.2015.09.450. URL: https://www.sciencedirect.com/science/article/pii/S2405601415009396.

[66] Glen Cowan. *Discovery sensitivity for a counting experiment with background uncertainty*. Tech. rep. http://www.pp.rhul.ac.uk/~cowan/stat/medsig/medsigNote.pdf. London: Royal Holloway, 2012. URL: http://www.pp.rhul.ac.uk/~cowan/stat/medsig/medsigNote.pdf.

[67] 'Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC'. In: (June 2019).

[68] T. P. Cheng and Ling-Fong Li. 'Neutrino Masses, Mixings and Oscillations in SU(2) x U(1) Models of Electroweak Interactions'. In: *Phys. Rev. D* 22 (1980), p. 2860. DOI: `10.1103/PhysRevD.22.2860`.

[69] Robert Foot et al. 'Seesaw Neutrino Masses Induced by a Triplet of Leptons'. In: *Z. Phys. C* 44 (1989), p. 441. DOI: `10.1007/BF01415558`.

[70] Carolina Arbeláez et al. 'How many 1-loop neutrino mass models are there?' In: *JHEP* 08 (2022), p. 023. DOI: `10.1007/JHEP08(2022)023`. arXiv: `2205.13063 [hep-ph]`.