



UvA-DARE (Digital Academic Repository)

A frequentist analysis of the $n=3$ type-I seesaw model

Krishnamurthy, S.

[Link to publication](#)

Creative Commons License (see <https://creativecommons.org/use-remix/cc-licenses>):
Other

Citation for published version (APA):
Krishnamurthy, S. (2020). A frequentist analysis of the $n=3$ type-I seesaw model.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<http://dare.uva.nl>)

A FREQUENTIST ANALYSIS
OF THE
 $n = 3$ TYPE-I SEESAW MODEL

Suraj Krishnamurthy

A FREQUENTIST ANALYSIS OF THE $n = 3$ TYPE-I SEESAW MODEL

Suraj Krishnamurthy

A Frequentist Analysis of the $n=3$ Type-I Seesaw Model

This work was accomplished at the Gravitation and AstroParticle Physics Amsterdam (GRAPPA) center of excellence of the University of Amsterdam (UvA).

© Suraj Krishnamurthy, 2019

All rights reserved. Without limiting the rights reserved under copyright, no part of this book may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form or by any means (electronic, mechanical, photocopying, recording or otherwise) without the written permission of both the copyright owner and the author of the book.

A Frequentist Analysis of the $n=3$ Type-I Seesaw Model

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K. I. J. Maex
ten overstaan van een door het College voor Promoties
ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op dinsdag 14 januari 2020, te 14 uur

door

SURAJ KRISHNAMURTHY

geboren te Bangalore, Karnataka

PROMOTIECOMMISSIE

PROMOTOR

prof. dr. J. de Boer Universiteit van Amsterdam

COPROMOTOR

dr. C. Weniger Universiteit van Amsterdam

OVERIGE LEDEN

prof. dr. S. Antusch University of Basel

dr. S. Ando Universiteit van Amsterdam

dr. M. Chrzaszcz Institute of Nuclear Physics,
Polish Academy of Sciences

prof. dr. E.L.M.P. Laenen Universiteit van Amsterdam

prof. dr. M.P. Decowski Universiteit van Amsterdam

PUBLICATIONS

This thesis is based on the following preprint:

[1] M. Chruszcz, M. Drewes, T. Gonzalo, J. Harz, S. Krishnamurthy, and C. Weniger. “A frequentist analysis of three right-handed neutrinos with GAMBIT”. in: (2019). arXiv: 1908.02302 [hep-ph]

Author list is in alphabetical order. The project was initiated by CW, MD and SK. SK coded up the constraints related to direct detection, Big Bang nucleosynthesis, and initial versions of all other constraints, along with setting up the necessary connections within the GAMBIT framework. MC contributed to code related to EWPO and CKM constraints, TG contributed to code related to EWPO, LFV and lepton universality and JH contributed code for SM and neutrinoless double-beta decay. CW, TG, MC and SK set up the scans. CW, TG, JH and SK wrote the data analysis and plotting routines used in the study. All authors were involved in discussion of the results. SK drafted the manuscript; all authors had input and were involved in finishing the paper.

Other publications by the author:

[2] R. Bartels, S. Krishnamurthy, and C. Weniger. “Strong support for the millisecond pulsar origin of the Galactic center GeV excess”. In: *Phys. Rev. Lett.* 116 (2016), p. 051102. arXiv: 1506.05104 [astro-ph.HE]

Author list is in alphabetical order. The project was initiated by CW and SK, RB joined later. CW drafted the manuscript and set up the analysis pipeline. RB contributed to the analysis code and performed Monte-Carlo simulations.

Contents

1	Introduction	1
2	The type-I seesaw model	11
2.1	Neutrino oscillation	12
2.1.1	The solar neutrino problem	12
2.1.2	The two-flavour case	14
2.1.3	The three-flavour case	15
2.1.4	Neutrino masses and hierarchies	17
2.2	The seesaw mechanism	22
2.2.1	The type-I seesaw	22
2.2.2	Other seesaw mechanisms	26
2.3	The Casas-Ibarra parameterization	28
2.4	Tuning and symmetries	31
2.4.1	Distinguishing symmetry-protected from tuned parameter choices	33
3	GAMBIT and the statistical analysis framework	37
3.1	Frequentist statistics	37
3.1.1	The profile likelihood ratio	40
3.2	The Bayesian framework	43
3.3	GAMBIT	45
3.3.1	The structure of GAMBIT	46
3.3.2	Models	48

3.3.3	Scanning algorithms	54
4	Observables	59
4.1	Standard Model constraints	60
4.2	Indirect constraints	64
4.2.1	Electroweak precision observables	64
4.2.2	Lepton flavour violation	65
4.2.3	Lepton universality	68
4.2.4	CKM unitarity	71
4.2.5	Neutrinoless double-beta decay	73
4.2.6	Big Bang Nucleosynthesis	75
4.3	Direct searches for right-handed neutrinos	77
4.3.1	PIENU	79
4.3.2	PS-191	80
4.3.3	CHARM	81
4.3.4	DELPHI	83
4.3.5	ATLAS	85
4.3.6	CMS	86
4.3.7	E949	87
4.3.8	NuTeV	88
5	Results	89
5.1	Scanning strategy, parameter ranges and priors	89
5.2	Profile likelihoods	95
5.3	Effect of constraints on the lightest active neutrino mass	102
5.4	Regions of excess likelihood	103
5.5	Flavour mixing pattern and triangle plots	106
5.6	Projected upper limits on parameter space	110

5.7 Comparison with earlier work	111
6 Conclusion and outlook	115
A Decay widths and form factors for LFV observables	119
B Decay widths relevant to Big Bang nucleosynthesis	125
C Additional plots for normal hierarchy	129
D Additional plots for inverted hierarchy	137
E GAMBIT capabilities & module functions	153
E.1 Neutrino models	153
E.2 Right-handed neutrino likelihood functions	154
Bibliography	167
Samenvatting	199
Summary	205
Acknowledgements	211

Chapter 1

Introduction

Nearly everything we know and understand about this planet and this universe at the fundamental level is attributable to the Standard Model (SM) of particle physics and general relativity (GR); one only needs to reflect on the huge leaps forward that have been achieved since the beginning of the twentieth century. A large part of the beauty of the SM+GR framework lies in the fact that it is based on a few extremely powerful underlying principles: quantum mechanics, gauge invariance and general covariance (along with some free parameters). For example, within this framework, all known elementary particles are excitations of a small number of quantum fields and the interactions among them are governed by the gauge symmetries inherent in the Lagrangian [3, 4, 5].

Despite its success in explaining a vast multitude of phenomena, however, it is not a completely satisfactory theory. For one, while matter and all other interactions are described using quantum field theory, gravity is treated classically. At the very least, once we manage to probe energies near the Planck scale, a quantum theory of gravity will be necessary; this is one of the most important gaps in our knowledge of how things work. Within particle physics too, the phenomenon of neutrino oscillations, the origin of baryonic matter and the nature of dark matter cannot be explained within the SM+GR paradigm; the first of these is the only concrete experimental evidence we have for physics beyond the Standard Model (BSM). A common explanation for these three outstanding questions involves the presence of some number of neutral, right-chiral fermions, also dubbed right-handed neutrinos (RHNs). Being neutral with respect to all gauge interactions in the SM (leading to a third alias, “sterile neutrinos”), their interaction is solely via Yukawa couplings [6, 7, 8, 9] (see 2.46 in Chapter 2).

Neutrino oscillations, or flavour violation in the neutrino sector, created a furor in the physics community when they were first observed in the latter half of the 20th century. Since then, there has been a concerted effort to look at it from different perspectives: the disappearance of solar electron neutrinos by the SAGE, GALLEX, GNO and SNO collaborations [10, 11, 12], electron neutrinos in nuclear reactors by KamLAND [13], atmospheric muon neutrinos by Super-Kamiokande [14] and many others, like K2K [15], T2K [16], MINOS [17], OPERA [18], Double Chooz [19], RENO [20] and Daya Bay [21]. It is safe to say that neutrino oscillations are an established fact of reality. The introduction of RHNs allows for the solution of the origin of neutrino masses and hence neutrino oscillations via the type-I seesaw mechanism [8, 22, 23].

The addition of these particles can confer other benefits as well: the Yukawa couplings of the RHNs to the Higgs doublet and the left-handed neutrinos (LHNs) in the SM violates CP symmetry; in the SM, this violation as well as the deviation from equilibrium is too small to account for the baryon asymmetry [24, 25]. Now, the RHNs are gauge singlets, so they can stay out of equilibrium at a temperature where the other particles, experiencing gauge interactions, reach equilibrium. Further, the RHNs' Yukawa couplings contain unconstrained CP-violating phases. Both these facts make baryogenesis via leptogenesis plausible. When one or more right-handed neutrinos are out of equilibrium, lepton asymmetries can arise via various processes like decays, inverse decays, scatterings and flavour oscillations. At temperatures above ~ 130 GeV, sphalerons can convert these asymmetries into a net baryon number [26, 27, 28]. The non-equilibrium conditions can occur during RHN production ("freeze-in" scenario) [29, 30, 31] or once the RHNs "freeze out" and decay [32, 33, 34]. The latter has been shown to require RHNs that are very heavy ($> 10^7$ GeV) [32, 35, 36, 37, 38] (see 2.46 in Chapter 2 for the RHN mass term in the type-I seesaw Lagrangian), while the former can occur with RHNs' in the sub-TeV mass range [29, 39]. This has been studied for the case where two right-handed neutrinos are present (see, for example, [6, 40]), where it was found that such a scenario needs a mass degeneracy between the RHNs of the order of 10^{-3} . This necessity is absent if three RHNs are added to the SM and this has been studied by many groups: for example [41, 42, 43, 44, 45].

Right-handed neutrinos are also dark matter candidates [7, 46, 47]. keV-scale particles are most often considered, for two good reasons. Consider

extremely dense regions, like galactic cores: being fermions, the Pauli exclusion principle prevents RHNs from being packed in an arbitrarily small volume, so they cannot have an arbitrarily small mass [48]. The other reason is due to the fact that the (albeit feeble) mixing between the active and sterile sector (which is a generic case) implies the possibility of a sterile neutrino decaying into an active one and a photon. This decay rate scales as the fifth power of the mass, so its non-observation places an upper limit at the keV scale as well. Recently, a 3.5 keV emission line was observed in the stacked spectrum of galaxy clusters [49, 50], which evoked great excitement because it could be due to the decay of dark matter particles; in fact, 7 keV RHNs could play the part, as argued in that study. However, there are indications that this may not be so; for instance, refer to [51, 52, 53].

The properties of RHNs are hence of interest and the role that RHNs play given existing constraints is strongly dependent on their mass (see, for example, [54]). In any bottom-up approach to the parameter space, there is no limiting factor when it comes to their masses; in principle, even sub-eV masses are possible [55]. Nevertheless, it is possible to exercise some discernment using extant theoretical and experimental knowledge. Constraints from Big Bang nucleosynthesis (BBN), for example, limit their lifetime [56]: the observed abundances of elements, which strongly depend on equilibrium conditions in the primordial environment, require the RHNs to decay sufficiently early so that these limits are satisfied. This indirectly enforces a lower limit on the RHN mass. Along with non-detections from direct searches [57, 58] using decays of mesons and gauge bosons, or the decays of the RHNs themselves, this raises the lower limit on the RHNs' to about 100 MeV. If more than two right-handed neutrinos are present in a model, any value between this and 1-2 orders of magnitude below the Planck mass is experimentally allowed, this upper limit being set by the requirement that the Yukawa interactions be perturbative [59]. So, without further consideration, a mass range 16 orders of magnitude in size is open to study. RHNs with masses near the GUT scale ($\sim 10^{16}$ GeV) can explain neutrino oscillations while simultaneously having sizable Yukawa couplings to the SM neutrinos of $\mathcal{O}(1)$ and leading to standard thermal leptogenesis [59, 60]. The downside to this scenario is that such massive particles will not be observable in any direct search planned for the foreseeable future. A pragmatic route is thus to consider masses up to the TeV scale. Several theoretical arguments have been

made in support of the existence of RHNs in this mass window: for heavier particles, without a mechanism to cancel radiative corrections, the Higgs mass would be destabilized [61], for one. Appeals to Ockham's razor have been put forward in the context of the number of extra particles needed to explain known BSM phenomena [6, 29, 62]. Other arguments for it come from left-right symmetric models [63, 64], the possibility of B-L being a spontaneously broken symmetry which is approximately conserved [65, 66] and the idea that the RHN mass scale and the electroweak scale have a shared origin [67].

Within the 16 orders of magnitude in mass that is open to examination, there are numerous avenues that have been used to look for right-handed neutrinos. Under the umbrella of direct searches, for instance, if it is kinematically allowed, RHNs will take part in any process that involves the SM's "active" neutrinos. To this end, there have been experiments in the past that searched for them in the decay of π and K mesons (where their signature would be missing energy), or even in the RHNs' own subsequent decay (in the form of *invisible* \rightarrow leptons), or even both. Such attempts have been made by PS191 [68], CHARM [69], NuTeV [70], WA66 [71] and NOMAD [72]. Other possible detection methods for RHNs which are 10s of GeV massive include gauge boson decays (this was done at DELPHI [73]), s-channel exchange of W bosons and vector boson fusion [74, 75, 76]. Unfortunately, no experiment has reported a detection of RHNs so far.

In addition, right-handed neutrinos can also be detected via their indirect effects on a range of observables. The RHN Majorana masses violate lepton number, in turn allowing neutrinoless double-beta decay to occur [77]. If they possess masses larger than 100 MeV (which can be argued for, as discussed, based on constraints from BBN and non-detections in experiments), the rate of this process can be large enough to be observable. Further, the Majorana mass also causes flavour violation in the charged lepton sector, allowing processes such as $\mu \rightarrow e\gamma$ [78], lepton universality violation [79] (in leptonic pion/kaon decays, semileptonic decays of B mesons and in tau and W boson decays) and unitarity violation in the PMNS (Pontecorvo-Maki-Nakagawa-Sakata) matrix, i.e. the matrix that parametrizes the mixing among the neutrino flavour and mass eigenstates [80]. Since the addition of RHNs also inevitably modifies the weak currents, multiple electroweak precision observables like the Weinberg angle, mass and decays of the W boson and the

invisible decay width of the Z boson will be affected [81]. Signs of their influence can also be extracted from the considerations of the unitarity of the CKM matrix [80]: the RHNs affect the leptonic decays which have been used to determine the matrix elements. Using the experimental values of these elements and the unitarity of the CKM matrix, the RHNs' mixing can be constrained.

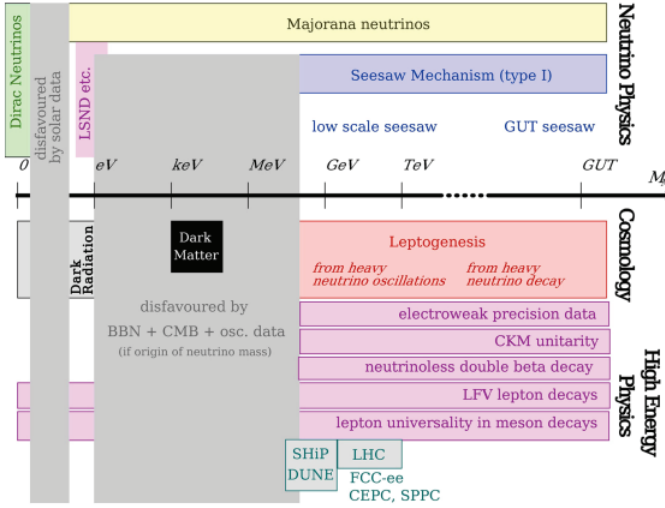


FIGURE 1.1: Taken from [59], this presents a pictorial overview of how various constraints limit the allowed range of RHN masses. Relevant to this study are the constraints coming from the seesaw mechanism (depicted in blue at the top) and all those shown in purple in the “High Energy Physics” section.

Given the complexity of the parameter space and phenomenology, a combined treatment of constraints is necessary to form a realistic picture of the status of RHNs. For instance, the authors of [82] combined constraints from multiple lepton flavour violating processes, while those of [57] used direct detection and BBN. Lepton number violating processes have been used in [83], neutrinoless double-beta decay and direct detection in [84] as well as lepton universality observables in [85]. Further combinations of the observables mentioned above have also been utilized: in [86], for example, lepton universality, the invisible decay width of the Z boson, the W boson mass and three lepton flavour-violating decays were considered. In terms of observables incorporated, some of the most complete studies of the parameter space in the last few years are [80], where multiple electroweak precision observables and flavour-violating decays were included, along with tests of lepton

universality and the unitarity of the CKM matrix; the study conducted in [86] was extended by the inclusion of more electroweak precision observables (the Fermi constant and mass of the Z boson), flavour-violating decays and the unitarity of the CKM matrix in [87]; in [88] as well, lepton universality, flavour-violating decays, the invisible decay width of the Z boson, neutrinoless double-beta decay and direct searches were considered, while a large range of observables was included in [59]: flavour-violating processes, neutrinoless double-beta decay, lepton universality, electroweak precision observables, unitarity of the CKM matrix and the most relevant direct searches at the time.

Many analyses considered the addition of one or two RHNs to the Standard Model. The choice of having $n = 2$ RHNs in a model has been adopted either due to the significant conceptual and computational simplifications that it entails, or because one RHN is taken to be a dark matter candidate, which means it interacts so feebly that its effect on neutrino masses and leptogenesis is negligible (see, for example, [62, 89]). Notwithstanding these reasons, there are a number of motivating factors to extend this to having three RHNs in the model. For one, assuming the so-called “seesaw mechanism” to be the sole generator of neutrino masses, the existence of two non-zero mass differences implies $n \geq 2$, since one RHN is necessary for each observed mass difference. Although it is currently unknown, if the lightest neutrino is discovered to be massive, the preceding logic requires the existence of at least one more RHN ($n \geq 3$). Another reason comes from leptogenesis: if two RHNs are present, their masses have to be quasi-degenerate to generate the BAU [6, 29, 40, 90], whereas the presence of another one allows for relatively more freedom and overcomes both these problems [91, 92].

The overarching goal of this thesis is to conduct a robust analysis of the type-I seesaw model parameter space with $n = 3$ RHNs having masses in the 0.1-500 GeV range. We confirm previous conjectures and extend previous results to provide a solid, statistically sound basis for future work. To do this, we use the “bottom-up” Casas-Ibarra (C-I) parameterization [78]. The parameter space is of high dimensionality, and the various constraints that are included in the analysis have relatively complex forms. This presents numerical problems for “top-down” approaches. As an example, the LHN oscillation parameters have been determined accurately enough that generating parameter points that satisfy these constraints becomes extremely difficult, *when possible*. The C-I scheme thus offers an immense numerical benefit in

comparison. An additional advantage is that it can be cast in to a form where one part encapsulates parameters in the LHN sector, and another contains the remaining parameters of the RHNs. On the other hand, deriving simple relations between the C-I parameters and those that can be motivated by a particular model is non-trivial. To put it another way, if one comes up with a theory-motivated prior on the RHNs' mixings and masses, this would usually acquire a convoluted form in the C-I parameterization. Here, we focus on the frequentist framework to statistically interpret the results; the results obtained are hence prior-independent.

To carry out the scans of the parameter space, the open-source software package GAMBIT (the Global And Modular BSM Inference Tool) is used [93]. It is a flexible global fitting software framework that is very useful for conducting high dimensional analyses. Also, its modular structure allows for the relatively easy addition of observables/constraints, which is very helpful: one “only” needs to code up functions that define the observable (or, more relevantly, likelihoods) to be added. GAMBIT combines these separate likelihoods internally, using the results to drive a scan. It is also “statistically secular”; no particular framework needs to be adhered to. The framework includes an interface to Diver [94], a differential evolution-based scanner that provides improved sampling performance compared to regular methods, especially when it comes to profile likelihood contours (differential evolution has been shown to be efficient at this [95, 96]). The scanner also allows the identification and analysis of small, high-likelihood regions in parameter space. There have been several studies conducted using GAMBIT: the scalar singlet dark matter model [97], Higgs portal dark matter models [98], the MSSM [99], GUT-scale SUSY models [100] and axions [101], for example.

This thesis contains several improvements over earlier analyses:

- In addition to lower masses, the most relevant direct and indirect constraints for RHN masses above that of the W boson have been included in this study. There have been several analyses that have combined subsets of the constraints [86, 87, 88], or have focused on answering a different aspect of the parameter space [59, 83]. Almost all the constraints included here were also in the analysis presented in [59] (for the $n = 3$ case), but only for RHNs below 80 GeV.
- The observed mass differences of the SM neutrinos and the mixing angles contained in the PMNS matrix were fixed at their best-fit values

(such as those presented in [102]). Here, they are encoded as nuisance parameters, using results from the NuFit collaboration [103].

- Electroweak observables such as the Weinberg angle, $\sin \theta_W$, require very precise calculations for their comparison with the equally accurate measurements. To this end, the two-loop order calculations of the SM prediction for $\sin \theta_W^{\text{eff}}$ presented in [104] are utilized.
- Many investigations of lepton flavour violation in neutrino models have focused on the most constraining processes, usually $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$ [59, 80]. This work is inclusive of the most relevant lepton flavour violating processes. In particular, all leptonic τ decays are present, for which the most recent average of experimental results provided by HFLAV [105] is used. Additionally, $\mu - e$ conversion in two nuclei (lead and titanium) is also included as a constraint; this can be a relevant constraint in the near future [106].
- For the constraint exerted by neutrinoless double-beta decay, it is chosen to carry out the analysis conservatively; in addition, the upper limit on the effective Majorana mass (and hence the active-sterile mixing) is encoded in the form of a (one-sided) Gaussian likelihood, not as a strict cut.
- We supplement earlier tests of lepton universality centered on leptonic decays of mesons (K and π), tauons and W bosons with the recently observed semileptonic decays of B mesons [107, 108, 109].
- Triangle plots, showing the distribution of the ratios of flavour-to-total mixing for the RHNs, are presented in the context of the type-I see-saw model with three right-handed neutrinos. These plots have been made for the same model with two RHNs [110]; the plots included here demonstrate how the freedom offered by the extra particle affects the distribution.

In terms of methodology, there are several improvements:

- In [59] and [111], for example, there are regions where full convergence of the scan was not achieved. This is because of the presence of regions that are finely-tuned, making it difficult for scanners to find such regions and explore them. Here, the upper limits on the mixings are fully explored over the entire mass range.

- A differential evolution-based scanning algorithm is used to populate the profile likelihoods constructed here. This class of algorithms has been shown to be much more efficient at such “optimization” problems than the often-used Markov Chain Monte Carlo and nested sampling techniques [94].
- Additionally, existing upper limits were encoded as hard cuts in [59]; they are encoded as likelihoods here. This confers the advantage of being able to consistently study the combined effect of individual experiments, since the limits are often presented at different confidence levels. In the low mass region, in particular, where there are multiple competing constraints, this plays an important role.

The type-I seesaw model which forms the base of the thesis is described in chapter 2, along with a short historical perspective, some important definitions and details about the Casas-Ibarra parameterization. In chapter 3, the various search strategies for right-handed neutrinos that are implemented in GAMBIT have been listed and each described in detail. Chapter 4 provides information on the Bayesian and frequentist frameworks of statistical analysis, outlines some points of comparison and also contains a summary of sorts of GAMBIT itself and its most relevant components. The results are reported and discussed in chapter 5, while a synopsis of the thesis and avenues for future work are presented in chapter 6.

Chapter 2

The type-I seesaw model

Within the Standard Model, neutrinos are weird particles in many aspects. Not only are they several orders of magnitude lighter than all other known fermions (at least 6 orders of magnitude separate neutrino masses from that of the electron, the lightest of the other known fermions), but they also are immune to the electromagnetic and strong forces, having no electric charge or colour; all their interactions are mediated by the vector bosons of the weak sector. To give an idea of how feebly they interact, neutrinos produced in nuclear reactors with energies of about 1 MeV have a cross-section $\sim 10^{-44}$ cm² [112]. This roughly corresponds to a probability of $\sim 10^{-18}$ to interact with a solid detector which is 1 m thick, or of $\sim 10^{-11}$ to interact with the Earth considering a trajectory that passes through the planet's center. Conversely, neutrinos are abundant. The average number density of cosmological neutrinos is $n_\nu \sim 336$ cm⁻³ and approximately 60 billion neutrinos cross each square centimeter of our body every second, having been produced in the core of the Sun by nuclear reactions [113]. A fantastic number ($\sim 10^{58}$) is emitted in the few seconds following the gravitational collapse of a massive star that triggers a supernova.

In the recent past, the Sun and the supernova SN1987A have been the first two astrophysical objects to be observed in neutrinos [114], and the use of neutrinos as messengers in astronomy and cosmology should only increase with time. Their tiny interaction cross section is simultaneously the biggest problem for neutrino astronomy, since massive detectors are required to get reasonable event rates, and also a property that can be exploited, because neutrinos can emerge from deep inside the core of astrophysical objects like stars and blazars [115], giving us information on the processes that occur

there. The IceCube Neutrino Observatory and ANTARES, for instance, currently search for neutrino fluxes of cosmic origin; both are sensitive to neutrinos with energies in the GeV-TeV range [116, 117]. In fact, a few years ago, IceCube presented evidence for the detection of cosmogenic neutrinos [118]. Upcoming neutrino telescopes such as the Cubic Kilometer Neutrino Telescope/KM3NeT [119], SuperNEMO [120] and NESTOR [121] will build on these studies in the future.

2.1 Neutrino oscillation

One of the most revolutionary discoveries in recent memory, the fact that neutrinos can change flavour over distance and time (see [122], for example, for a review), is closely connected to the postulation of the existence of right-chiral neutrinos. A historical exploration of this phenomenon may be beneficial, once again.

2.1.1 The solar neutrino problem

In the 1920s, when solar physics was still in relative infancy, Arthur Eddington suggested that nuclear fusion, proton-proton reactions essentially, powers the Sun. More than a decade later, Hans Bethe pioneered the work that elucidated the details of this process [123], which is now called the proton-proton reaction chain, or pp-chain:

- Two protons form deuteron by one of the two below reactions:

$$p + p \rightarrow d + e^+ + \nu_e \quad \text{or} \quad (2.1)$$

$$p + p + e^- \rightarrow d + \nu_e. \quad (2.2)$$

- The deuteron picks up another proton, forming a helium-3 nucleus, releasing energy in the form of a photon:

$$d + p \rightarrow {}^3\text{He} + \gamma. \quad (2.3)$$

- Helium-3 can join with another proton to make the nucleus of helium-4 (an “alpha particle”), two helium-3s can get together to make an alpha particle and two protons, or the helium-3 can combine with an alpha

particle that was produced in one of the previous reactions to make beryllium-7, with the emission of a photon:

$${}^3\text{He} + p \rightarrow \alpha + e^+ + \nu_e \quad \text{or} \quad (2.4)$$

$${}^3\text{He} + {}^3\text{He} \rightarrow \alpha + p + p \quad \text{or} \quad (2.5)$$

$${}^3\text{He} + \alpha \rightarrow {}^7\text{Be} + \gamma. \quad (2.6)$$

- Finally, beryllium-7 can absorb an electron, making lithium, which then picks up a proton, yielding two alpha particles, or it absorbs a proton, making boron-8, which goes to an excited state of beryllium-8, and from there to two alpha particles:

$${}^7\text{Be} + e^- \rightarrow {}^7\text{Li} + \nu_e, \quad (2.7)$$

$${}^7\text{Li} + p \rightarrow \alpha + \alpha \quad \text{or} \quad (2.8)$$

$${}^7\text{Be} + p \rightarrow {}^8\text{B} + \gamma, \quad (2.9)$$

$${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e, \quad (2.10)$$

$${}^8\text{Be}^* \rightarrow \alpha + \alpha. \quad (2.11)$$

The main point is that there are multiple processes in this chain that produce electron neutrinos. In 1968, Ray Davis et al. [124] reported the results of the first experiments to measure solar neutrinos, to verify the calculations carried out by John Bahcall [125]. He used a massive tank of cleaning fluid (which contains chlorine), because chlorine can absorb a neutrino and convert to argon:

$$\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^-. \quad (2.12)$$

The experiment involved the collection of argon atoms for several months (they were produced at a rate of about one atom *every two days*). The total accumulation, however, was only about a third of what Bahcall predicted, leading to the birth of the infamous “*solar neutrino problem*” [124].

At first, the result was not taken seriously. Bahcall’s calculations were heavily model-dependent and since Davis claimed to have flushed 33 argon atoms out of a tank containing 615 tons of tetrachloroethylene, it could easily have been that he simply missed some. This changed when other detection methods showed the same deficit, and Pontecorvo came up with the simple hypothesis that the neutrinos from the Sun transform, in flight, into a species

which the experiment was not sensitive to: the mechanism now called neutrino oscillation [126].

2.1.2 The two-flavour case

Take two neutrino types, say ν_e and ν_μ . If they can convert between one other, they cannot be eigenfunctions of the Hamiltonian; the true ones (call them ν_1 and ν_2 ; the mass eigenstates) must be some orthogonal, linear combination of them, or, equivalently

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix}, \quad (2.13)$$

θ being the mixing angle. The initial state must be a pure electron neutrino (assume it has a well-defined momentum):

$$|\nu(t=0)\rangle = |\nu_e\rangle = \cos\theta |\nu_1\rangle + \sin\theta |\nu_2\rangle. \quad (2.14)$$

According to the Schrödinger equation, this state evolves in time according to

$$|\nu(t)\rangle = \cos\theta e^{-iE_1 t} |\nu_1\rangle + \sin\theta e^{-iE_2 t} |\nu_2\rangle. \quad (2.15)$$

The energies of the mass eigenstates are¹

$$E_i = \sqrt{p^2 + m_i^2} \approx p + \frac{m_i^2}{2p} \approx E + \frac{m_i^2}{2E}; \quad (2.16)$$

therefore, the probability of finding a muon neutrino at time t comes out to be

$$P(\nu_e \rightarrow \nu_\mu; t) = |\langle \nu_\mu | \nu(t) \rangle|^2 = \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2}{4E} t \right); \Delta m^2 = m_2^2 - m_1^2, \quad (2.17)$$

or, in terms of distance L [127],

$$P(\nu_e \rightarrow \nu_\mu; t) = \sin^2 2\theta \sin^2 \left(1.27 \Delta m^2 \frac{L}{E} \right). \quad (2.18)$$

In the above equation, Δm^2 is in eV^2 , L in km and E in GeV; the factor of 1.27 arises from the combination of the conversion factors (c, \hbar) necessary to

¹Note that here, and in the entirety of this thesis, $c = 1$.

express the equation in natural units.

So, there appear to be 2 key ingredients necessary for oscillation: a non-zero mixing θ and a non-zero mass difference Δm^2 .

Full experimental verification would have to wait until the new millennium, when the results of Super-Kamiokande (Super-K/SK) [128] ($\nu_\mu \rightarrow \nu_\tau$) and the Sudbury Neutrino Observatory (SNO) [129] ($\nu_e \rightarrow \nu_\mu, \nu_\tau$) finally vindicated the hypothesis.

2.1.3 The three-flavour case

Extending this to 3 generations, the relation between the flavour and mass eigenstates can be written as

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (2.19)$$

labeling the flavour eigenstates with Greek, and the mass eigenstates with Latin indices. U is the PMNS matrix [130]. It is unitary; this follows from the fact that the flavour (and mass) eigenstates are orthogonal to each other. Now, the most general $N \times N$ complex matrix requires $2N^2$ independent real parameters to describe it. With a unitary condition $\sum_i U_{\alpha i} U_{\beta i}^* = \delta_{\alpha\beta}$, N^2 of these can be eliminated. These remaining parameters can be divided into $N(N-1)/2$ angles and $N(N+1)/2$ phases. One can then make any matrix element of U real and positive by redefining a flavour eigenstate; for example, $U_{e1} = |U_{e1}| e^{i\phi_{e1}}$ can be made real by redefining $|\nu_e\rangle \rightarrow |\nu_e\rangle e^{i\phi_{e1}}$. This can be done for a column, eliminating N phases.

If neutrinos are Dirac particles, redefining mass eigenstates is also allowed, and $N-1$ more phases can be eliminated, so $2N-1$ phases can be made unphysical, leaving $(N-1)(N-2)/2$ physically relevant phases.

If neutrinos are Majorana particles, on the other hand, the phase of the mass eigenstates cannot be arbitrarily changed, because of the form of the corresponding mass term. Two “Majorana” phases will also be physical, in this case. These phases, however, have no effect on flavour transitions, since they are common to a column in the mixing matrix.

The neutrino oscillation probability does not, in general, respect CP or T symmetry. Consider the oscillation probability

$$P(\nu_\alpha \rightarrow \nu_\beta) = \left| \sum_j U_{\beta j} U_{\alpha j}^* e^{-im_j^2 \frac{L}{2E_\nu}} \right| \quad (2.20)$$

$$= \sum_{j=1,3} |U_{\beta j}|^2 |U_{\alpha j}|^2 \quad (2.21)$$

$$+ \sum_{j < k} 2\text{Re}|U_{\beta j} U_{\beta k}^* U_{\alpha j}^* U_{\alpha k}| \cos\left(\frac{\Delta m_{jk}^2 L}{2E}\right) \quad (2.22)$$

$$+ \sum_{j < k} 2\text{Im}|U_{\beta j} U_{\beta k}^* U_{\alpha j}^* U_{\alpha k}| \sin\left(\frac{\Delta m_{jk}^2 L}{2E}\right), \quad (2.23)$$

where Δm_{jk} is the difference between the masses of the j th and k th neutrino, i.e. $\Delta m_{jk} = m_j - m_k$.

With three generations of neutrinos, as in the SM, the preceding discussion implies that there will be three mixing angles, call them θ_{12} , θ_{23} and θ_{13} , and, assuming neutrinos to be Dirac particles for now, one physically relevant phase δ . The mixing angles are defined through the elements of the matrix U [131]:

$$\cos^2 \theta_{12} = \frac{|U_{e1}|^2}{1 - |U_{e3}|^2}, \quad \sin^2 \theta_{12} = \frac{|U_{e2}|^2}{1 - |U_{e3}|^2}, \quad (2.24)$$

$$\cos^2 \theta_{23} = \frac{|U_{\tau 3}|^2}{1 - |U_{e3}|^2}, \quad \sin^2 \theta_{23} = \frac{|U_{\mu 3}|^2}{1 - |U_{e3}|^2}, \quad (2.25)$$

$$\sin^2 \theta_{13} = |U_{e3}|^2. \quad (2.26)$$

- A CP transformation is equivalent to replacing neutrinos with anti-neutrinos, meaning we have to make the change $U \rightarrow U^*$. This will cause the “sin” terms to switch sign; the “cos” terms stay the same. CP violation thus requires the mixing matrix to have a non-zero imaginary component, which is possible only if δ is different from 0 or π .
- A time reversal can be achieved by interchanging α and β . Once again, the “sin” terms change sign, while the “cos” terms do not, leading to the same condition on δ as above.
- A CPT transformation ($U \rightarrow U^* + \alpha \leftrightarrow \beta$) leaves the probability invariant. This is expected, since the CPT theorem [132] requires, in vacuum, that $P(\nu_\alpha \rightarrow \nu_\beta) = P(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha)$.

- Apart from the condition on δ , all 3 mixing angles must be non-zero; $\theta_{12}, \theta_{23}, \theta_{13} \neq 0$.

The study of, and search for such violation is a very interesting and ongoing topic of research, since it might play a vital role in explaining the baryon asymmetry of the universe [133].

2.1.4 Neutrino masses and hierarchies

Various experimental techniques and approaches have been utilized to determine the oscillation parameters: solar neutrino experiments like SAGE [10], GALLEX/GNO [11] (using gallium) and the original Homestake chlorine experiment [124], SNO [12] (heavy water detector), KamLAND [13] (liquid scintillators), Kamiokande/Super-Kamiokande [14] (radiochemical/light water Cherenkov detectors) and Borexino [134, 135] have all collected large amounts of data that have led to the determination of the so-called solar parameters, Δm_{sol}^2 and $\sin^2 \theta_{\text{sol}}$ (due to the trappings of convention, the mass eigenstates $|\nu_i\rangle$ are usually labeled such that $|U_{e1}|^2 > |U_{e2}|^2 > |U_{e3}|^2$, implying that the ν_1 component of ν_e is larger than the ν_2 component, which is in turn larger than the ν_3 component. In this light, since ν_1 and ν_2 would be “electron neutrino rich”, the oscillations observed in the solar neutrino problem are governed by $\Delta m_{21}^2 \equiv m_2^2 - m_1^2$; Δm_{sol}^2 and $\sin^2 \theta_{\text{sol}}$ are thus synonymous with Δm_{21}^2 and $\sin^2 \theta_{21}$).

Atmospheric neutrino experiments, on the other hand, such as Super-K [14], MACRO [136] and IceCube [137], along with long baseline experiments like K2K [15] and MINOS [17] have measured the corresponding “atmospheric neutrino parameters”, $|\Delta m_{\text{atm}}^2|$ and $\sin^2 \theta_{\text{atm}}$ (atmospheric neutrino oscillations are, analogous to the solar case, more dependent on Δm_{31}^2 and Δm_{32}^2 , making the “atm” label synonymous with these quantities).

Other experiments like Daya Bay, RENO and Double Chooz have measured the third mixing angle, θ_{13} . Now, SNO determined the ordering of the solar neutrino pair, showing that $m_2^2 > m_1^2$, so $\Delta m_{21}^2 > 0$. Unfortunately, this leaves the position of ν_3 ambiguous. It could be heavier than the other two, or much lighter (hence the modulus sign for this mass splitting). If $m_3^2 > m_2^2$, the structure is referred to as the “normal hierarchy/ordering” (NH/NO), while the other case is called the “inverted hierarchy/ordering” (IH/IO). Which

one occurs in reality is yet to be determined. The rest of the parameters (the CP-violating phases) are also undetermined.

Further, the oscillation results also indicate that one splitting is small, while the other is relatively larger; this seems like a similar structure as that seen in the leptons (e and μ closer together, τ the loner) as well as in the quark sector (u and c relatively close, t higher and d and s close, b higher).

Unfortunately, oscillations are only sensitive to differences in the squares of the neutrino masses, and, of course, one would like to measure the individual neutrino masses directly, which turns out to be very hard. Beta decay is the standard method used to directly determine the mass scale of neutrinos, particularly since it is a model-independent approach [138]. An experiment typically yields the electron energy spectrum; a non-zero neutrino mass would cause slight changes in the slope of this curve and also affect the maximum permitted energy (or endpoint) for the electron and these changes are examined. However, a number of conditions restrict the accuracy of an experiment: the statistical error due to the small number of electrons in the endpoint of the spectrum, the energy resolution of the detector, the nature of the source (it should be thin enough to prevent energy loss of the electron post-decay) and its endpoint energy (the sensitivity improves with smaller values of this quantity) and background atomic and nuclear effects. Measurements of the effective Majorana mass from neutrinoless double-beta decay is another method used, where the amplitude of the peak in the electron energy spectrum gives information on the mass. The effective mass enters the amplitude through nuclear matrix elements, which suffer from quite large uncertainties; this also makes the calculations model-dependent [139]. Due to these significant difficulties, only placing limits has been possible so far, with the strongest ones coming from neutrinoless double-beta decay experiments [140, 141, 142] (the latter method is applicable specifically to Majorana neutrinos, while the former is applicable to both Dirac and Majorana neutrinos).

However, although the mass of the lightest neutrino is unknown, the sum of the neutrino masses has been bounded from above using cosmological data. Massive neutrinos are initially relativistic in the Universe, becoming non-relativistic after a transition when their rest masses begin to dominate. Imprints have been left on the cosmic microwave background (via the Sachs-Wolfe effect [143], which shifts the amplitude and location of the CMB acoustic peaks [144]) and the large scale structure (through suppression of the

clustering of matter because of the large free-streaming velocity of neutrinos [145]). Assuming the Lambda cold dark matter (Λ CDM) model and combining this information with that coming from baryon acoustic oscillations (BAO), the PLANCK collaboration has placed tight constraints on the sum of neutrino masses: $\sum m_\nu < 0.12$ eV [146] (the limit can vary depending on the cosmological datasets used, as is evident in their analysis). The sum is also bounded from below, since the mass differences are known: for NH, $\sum m_\nu > 0.06$ eV; for IH, $\sum m_\nu > 0.1$ eV [146]. The upper limit on the sum of the neutrino masses has thus begun to put pressure on the inverted mass hierarchy.

Additionally, the nature of neutrinos, i.e. whether they are Dirac or Majorana particles, is also an unresolved issue: all mass terms in the SM involve a left- and right-handed field. Only left-handed neutrinos have been observed to exist, so the idea, coming from Ettore Majorana [147], is then to form a right-handed field from a left-handed one and make a mass term this way. Here,

$$n_L^C = C \bar{n}_L^T \quad (2.27)$$

would be a right-handed field; C is the charge conjugation matrix, which can be defined in a representation-independent way [148] through

$$C^{-1} \gamma_\mu C = -\gamma_\mu^T. \quad (2.28)$$

A field defined by $n_L + n_L^C$ would be self-conjugate, i.e. the particle and anti-particle are identical. Since this would violate charge conservation, neutrinos are currently the only possible candidates within the Standard Model for having a Majorana nature. It also turns out that one cannot form a mass term with the fields in the SM: the left-handed neutrinos have a spin of $1/2$ and a weak hypercharge of -1 . A term such as $\bar{n}_L^C n_L$, on the other hand, would have a spin of 1 and a weak hypercharge of -2 , meaning that their coupling requires a Higgs with weak hypercharge $+2$ and spin -1 (aka a Higgs triplet with hypercharge $+2$); such a field is non-existent in the SM and the simplest way to construct such Majorana fields is to postulate the existence of right-handed neutrinos.

Dirac and Majorana mass terms

The SM confers two common properties on all fermions *except* neutrinos: they are Dirac particles and they get their masses through the Higgs mechanism. As discussed above, neutrinos could very well be Majorana particles; it may also be that the origin of their masses does not involve the Higgs mechanism.

Suppose n neutral, right-handed fields are postulated, in juxtaposition to the left-handed neutrinos. Like the other right-handed fields, they are usually taken to be $SU(2)_L$ (the symmetry group associated with weak isospin) singlets; with the definition of electric charge, this implies that the weak hypercharge $Y = 2(Q - T_3) = 0$, Q being the electric charge and T_3 the third component of weak isospin. The right-handed fields are thus singlets of the whole gauge group of the Standard Model.

Nevertheless, they will have non-trivial effects, the first of which is new interactions in the Yukawa sector, like

$$- \sum_{l,l'} F_{ll'} \bar{\ell}_{lL} \tilde{\Phi} n_{l'R} + \text{h.c.}, \quad (2.29)$$

where $F_{ll'}$ are coupling constants, ℓ_L denotes the lepton doublet and $\tilde{\Phi} = \epsilon \Phi^*$, where ϵ is the Levi-Civita tensor and Φ is the SM Higgs doublet. Identifying the newly introduced fields n_R as precisely the right-handed components of the left-handed SM neutrinos², this gives rise to the Dirac mass terms

$$- \mathcal{L}_{\text{mass}} = \sum_{l,l'} F_{ll'} \frac{v}{\sqrt{2}} \bar{n}_{lL} n_{l'R} + \text{h.c.} \quad (2.30)$$

In flavor space, this corresponds to a mass matrix

$$M_{ll'} = \frac{v}{\sqrt{2}} F_{ll'}. \quad (2.31)$$

Generally, this matrix is not diagonal, and the fields are not connected to the physical fermion fields; to get the physical fields, one must find the eigenvectors of the mass matrix, which is usually achieved by diagonalizing it using a (bi-unitary) transformation [149]:

$$U^\dagger M V = M_D. \quad (2.32)$$

²The invariance of equation 2.29 necessitates the assignment of lepton number to the right-handed fields.

Defining new states by

$$n_L \equiv \sum_{\alpha} U_{\alpha} \nu_{\alpha L}, \quad (2.33)$$

$$n_R \equiv \sum_{\alpha} V_{\alpha} \nu_{\alpha R}, \quad (2.34)$$

Now, equation 2.30 can be recast as

$$- \mathcal{L}_{\text{mass}} = \sum_{\alpha} \bar{\nu}_{\alpha L} M_{D\alpha} \nu_{\alpha R} + \text{h.c.} \quad (2.35)$$

Here, $M_{D\alpha}$ is the α th diagonal element of M_D . The ν_{α} thus come out to be fields with definite masses and are physical particles. The mass term in equation 2.35 is a *Dirac mass*, which contains particles of opposite chirality.

With this, neutrinos are treated on the same footing as the other known fermions. Unfortunately, the model leaves something to be desired; it does not provide any information on the eigenvalues of $F_{ll'}$, which determine the neutrino masses. Neither does it say anything about the size of their mixings. Of course, if the coupling constants are small compared to those that influence the lepton and quark masses, then the observed mass differences may come out, but a priori, there is no cause for suspecting that this is so. Further, the model is incomplete; there are other gauge-invariant terms that can still be written.

Case in point, a *Majorana mass* term, which involves fields of the same chirality, can be written:

$$- \frac{1}{2} M_M \bar{\nu}_R^C \nu_R + \text{h.c.} \quad (2.36)$$

Although the introduction of such a term is not forbidden by the Standard Model gauge group, it violates lepton number (the difference between the number of leptons and anti-leptons in a process) by two units. Left-chiral neutrinos are usually not given a Majorana mass term, since such a term would break hypercharge [150]. Further, since right-chiral neutrinos do not participate in weak processes, nothing prevents them from having Majorana mass terms. The introduction of both kinds of mass terms is the most widely used route to the “seesaw mechanism”, which generates neutrino masses and can explain why the left-chiral ones are light.

2.2 The seesaw mechanism

2.2.1 The type-I seesaw

The type-I seesaw mechanism is characterized by the following:

- There is no Majorana mass term for left-chiral neutrinos.
- The Higgs sector of the SM is unchanged.
- $M_M \gg M_D$, i.e. the Majorana masses for the right-chiral neutrinos are much larger than the Dirac masses.

One-flavour case

As an example to demonstrate the seesaw mechanism, consider one generation of neutrinos. The possible mass terms that can be formed, based on the details outlined earlier are

$$-M_D \bar{\nu}_L \nu_R - \frac{1}{2} M_M \bar{\nu}_R^C \nu_R + \text{h.c.} \quad (2.37)$$

For clarification and expression in terms of a mass matrix, introduce

$$\nu = \begin{pmatrix} \nu_L \\ \nu_R^C \end{pmatrix}, \quad (2.38)$$

which reduces the mass terms to

$$-\frac{1}{2} \bar{\nu} \mathcal{M} \nu + \text{h.c.}, \quad (2.39)$$

with the mass matrix

$$\mathcal{M} = \begin{pmatrix} 0 & M_D \\ M_D & M_M \end{pmatrix}. \quad (2.40)$$

Using an orthogonal transformation such as $\mathcal{M} = \mathcal{O} m \mathcal{O}^T$, this can be diagonalized. The diagonal elements turn out to be

$$m_{22(11)} = \frac{1}{2} M_M \pm \sqrt{M_M^2 + 4M_D^2}. \quad (2.41)$$

These can be negative; to ensure that the elements of m are positive, let us introduce $m'_{ij}\eta_{ij} = m_{ij}$, where $\eta_{ij} = \pm 1$. A re-definition of the fields:

$$\nu_M = (\mathcal{O}\sqrt{\eta})^\dagger \nu + [(\mathcal{O}\sqrt{\eta})^\dagger \nu]^C \quad (2.42)$$

allows a mass term similar to equation 2.39:

$$-\frac{1}{2}\bar{\nu}_M m \nu_M + \text{h.c.} \quad (2.43)$$

Notice that $\nu_M^C = \nu_M$. We now have 2 Majorana fields with masses m_{11} and m_{22} .

In the limit of $M_M \gg M_D$, the mass terms for the 2 Majorana fields become

$$m_{11} \cong \frac{M_D^2}{M_M} \quad \text{and} \quad (2.44)$$

$$m_{22} \cong M_M, \quad (2.45)$$

and we end up with a heavy neutrino and a light(er) one. The interplay between M_M , M_D and the resulting masses is the origin of the “seesaw” terminology [8].

The three-flavour case

Adding 3 RHNs to the Standard Model’s particle content introduces 18 new real parameters (3 RHN masses, 3 LHN masses, 3 mixing parameters for the LHNs as well as 3 CP-violating phases and 3 complex or 6 real angles for the mixing among LHNs and RHNs). In order to establish notation, we briefly go over the neutrino mixing matrices and masses, following the conventions in [59].

With the addition of the new particles, the Lagrangian has the following form:

$$\mathcal{L} = \mathcal{L}_{SM} + i\bar{\nu}_R \not{\partial} \nu_R - \bar{\ell}_L F \nu_R \tilde{\Phi} - \tilde{\Phi}^\dagger \bar{\nu}_R F^\dagger \ell_L - \frac{1}{2} \left(\bar{\nu}_R^C M_M \nu_R + \bar{\nu}_R M_M^\dagger \nu_R^C \right). \quad (2.46)$$

Here, $\ell_L = (\nu_L, e_L)^T$ are the left-handed leptons³ in the SM and $\tilde{\Phi} = \epsilon \Phi^*$, where ϵ is the Levi-Civita tensor and Φ is the Higgs doublet. M_M is the

³Four-component spinor notation is used here, though the chiral spinors ν_R and ℓ_L have only two non-zero components ($P_R \nu_R = \nu_R$ and $P_L \ell_L = \ell_L$).

Majorana mass matrix for ν_R and F is the Yukawa coupling matrix. We work in a flavour basis where $M_M = \text{diag}(M_1, M_2, M_3)$.

After electroweak symmetry breaking (EWSB), the complete neutrino mass term reads

$$-\frac{1}{2}(\bar{\nu}_L \bar{\nu}_R^C) \mathcal{M} \begin{pmatrix} \nu_L^C \\ \nu_R \end{pmatrix} + \text{h.c.}, \quad (2.47)$$

with

$$\mathcal{M} = \begin{pmatrix} \delta m_\nu^{1\text{-loop}} & M_D \\ M_D^T & M_M \end{pmatrix}, \quad (2.48)$$

where $M_D = Fv$, v being the Higgs field expectation value ($v = 174$ GeV at $T = 0$). The 1-loop correction $\delta m_\nu^{1\text{-loop}}$ has been included, since it allows the performance of an analysis that is consistent at second order in the Yukawa couplings F . It is given by [151]

$$\left(\delta m_\nu^{1\text{-loop}} \right)_{\alpha\beta} = \sum_I F_{\alpha I} M_I F_{I\beta}^T l(M_I), \quad (2.49)$$

where $l(M_I)$ is a correction factor given by

$$l(M_I) = \frac{1}{(4\pi)^2} \left[\left(\frac{3\ln[(M_I/m_Z)^2]}{(M_I/m_Z)^2 - 1} \right) + \left(\frac{\ln[(M_I/m_H)^2]}{(M_I/m_H)^2 - 1} \right) \right]. \quad (2.50)$$

The mass matrix (equation 2.48) can be diagonalized by a matrix having the form [86]

$$U = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta^\dagger) & \cos(\theta^\dagger) \end{pmatrix} \begin{pmatrix} U_\nu & \mathbf{0} \\ \mathbf{0} & U_N^* \end{pmatrix}, \quad (2.51)$$

with

$$\cos(\theta) = \sum_{n=0}^{\infty} \frac{(-\theta\theta^\dagger)^n}{(2n)!} \quad \text{and} \quad (2.52)$$

$$\sin(\theta) = \sum_{n=0}^{\infty} \frac{(-\theta\theta^\dagger)^n \theta}{(2n+1)!}. \quad (2.53)$$

Here, θ is the matrix that mediates the mixing between the active neutrinos ν_L and the sterile neutrinos ν_R . Therefore, in general,

$$U^\dagger \mathcal{M} U^* = \begin{pmatrix} m_\nu^{\text{diag}} & \mathbf{0} \\ \mathbf{0} & M_N^{\text{diag}} \end{pmatrix}, \quad (2.54)$$

or, in other words,

$$M_N^{\text{diag}} = U_N^T M_N U_N = \text{diag}(M_1, M_2, M_3) \quad \text{and} \quad (2.55)$$

$$m_v^{\text{diag}} = U_v^\dagger m_v U_v^* = \text{diag}(m_1, m_2, m_3), \quad (2.56)$$

with

$$m_v = m_v^{\text{tree}} + \delta m_v^{\text{1-loop}}, \quad (2.57)$$

$$m_v^{\text{tree}} = -M_D M_M^{-1} M_D^T = -\theta M_M \theta^T = -v^2 F M_M^{-1} F^T \quad \text{and} \quad (2.58)$$

$$M_N = M_M + \frac{1}{2} \left(\theta^\dagger \theta M_M + M_M^T \theta^T \theta^* \right). \quad (2.59)$$

Equation 2.58 shows the 3-generation analogue of the seesaw behavior that was seen in equation 2.44. The additional complex conjugation of U_N ensures that the relation among mass and flavour eigenstates will be analogous for LHNs and RHNs within the notation.

In the second part of equation 2.55, the difference between the eigenvalues of M_M and M_N , which is of second order in θ , is neglected. This is justified because of the experimental constraints on the magnitude of the elements $\theta_{\alpha I}$. The limit of small $\theta_{\alpha I}$ is usually referred to as the *seesaw limit* because it corresponds to $M_D \ll M_M$ (in terms of eigenvalues). It allows the approximation

$$\theta = M_D M_M^{-1} \quad (2.60)$$

and

$$U = \left[\begin{pmatrix} \mathbf{1} - \frac{1}{2} \theta \theta^\dagger & \theta \\ -\theta^\dagger & \mathbf{1} - \frac{1}{2} \theta^\dagger \theta \end{pmatrix} + \mathcal{O}(\theta^3) \right] \begin{pmatrix} U_v & \mathbf{0} \\ \mathbf{0} & U_N^* \end{pmatrix}. \quad (2.61)$$

The light and heavy neutrino *mass eigenstates* are given by

$$\nu = V_v^\dagger \nu_L - U_v^\dagger \theta \nu_R^c + V_v^T \nu_L^c - U_v^T \theta^* \nu_R, \quad (2.62)$$

and

$$N = V_N^\dagger \nu_R + \Theta^T \nu_L^c + V_N^T \nu_R^c + \Theta^\dagger \nu_L, \quad (2.63)$$

respectively. The matrices V_ν and V_N (which represent the mixing between mass and interaction eigenstates in the respective sectors) can be defined as

$$V_\nu \equiv \left(\mathbf{1} - \frac{1}{2} \theta \theta^\dagger \right) U_\nu, \quad (2.64)$$

$$V_N^* \equiv \left(\mathbf{1} - \frac{1}{2} \theta^T \theta^* \right) U_N, \quad (2.65)$$

while mixing between the two sectors is encoded in the matrix

$$\Theta = \theta U_N^*. \quad (2.66)$$

This quantity is of primary interest because it controls the interactions of the heavy neutrinos with the physical Higgs field h and the W and Z bosons, as shown by the couplings

$$\begin{aligned} & -\frac{g}{\sqrt{2}} \bar{N}_I \Theta_{I\alpha}^\dagger \gamma^\mu e_{L\alpha} W_\mu^+ - \frac{g}{\sqrt{2}} \bar{e}_{L\alpha} \gamma^\mu \Theta_{\alpha I} N_I W_\mu^- \\ & -\frac{g}{2 \cos \theta_W} \bar{N}_I \Theta_{I\alpha}^\dagger \gamma^\mu \nu_{L\alpha} Z_\mu - \frac{g}{2 \cos \theta_W} \bar{\nu}_{L\alpha} \gamma^\mu \Theta_{\alpha I} N_I Z_\mu \\ & -\frac{g}{\sqrt{2}} \frac{M_I}{m_W} \Theta_{\alpha I} h \bar{\nu}_{L\alpha} N_I - \frac{g}{\sqrt{2}} \frac{M_I}{m_W} \Theta_{I\alpha}^\dagger h \bar{N}_I \nu_{L\alpha}. \end{aligned} \quad (2.67)$$

Here, g is the weak gauge coupling constant and θ_W the Weinberg angle. For convenience, we use introduce the notation

$$U_{\alpha I}^2 \equiv |\Theta_{\alpha I}|^2, \quad (2.68)$$

$$U_I^2 \equiv U_{eI}^2 + U_{\mu I}^2 + U_{\tau I}^2 \quad \text{and} \quad (2.69)$$

$$U_\alpha^2 \equiv \sum_I U_{\alpha I}^2. \quad (2.70)$$

Here, $U_{\alpha I}^2$ represents the coupling of right-handed neutrino n_I to flavour α , U_I^2 the total coupling of RHN n_I to all three flavours $\alpha = e, \mu, \tau$ and U_α^2 the total strength of the coupling of all three RHNs ($I = 1, 2, 3$) to a single flavour α .

2.2.2 Other seesaw mechanisms

Aside from what has been discussed above, other variants of the seesaw mechanism have been worked out. One of them was briefly touched upon in

section 2.1.4, while describing how Majorana mass terms can be constructed. There are two variants of the above “standard” seesaw mechanism.

The type-II seesaw

Here, a Higgs triplet Δ with weak hypercharge $Y = +2$ plays the role of the RHNs [23, 152, 153]. In this case, the active neutrino masses are given by

$$m_\nu = \frac{y_\Delta \mu v^2}{M_\Delta^2}, \quad (2.71)$$

and under the expectation that $\mu \sim M_\delta$, if $M_\Delta \gg v$, the active neutrinos are naturally light. This model has been studied in [154], for example, in connection to the baryon asymmetry of the universe and in [155, 156] to test detectability.

Aside from the caveat that small active neutrino masses require the mass scale of the extra particles to be very high, the re-formulation of the Higgs sector makes this model less minimal than the type-I variant, which we focus on in this study.

The type-III seesaw

In this variant, a *fermion* triplet is introduced into the mix [157]; the triplet has zero hypercharge. This leads to neutrino masses given by

$$m_\nu = -y_T^T \frac{1}{M_T} y_T v^2, \quad (2.72)$$

when $M_T \gg v$ (as is required to generate small active neutrino masses). Mirroring the situation in the type-I scenario, at least two triplets are needed to explain the two observed mass splittings. Unlike the RHN singlets in the type-I variant, however, the triplet(s) are capable of having gauge interactions as well as inducing mixing among the charged leptons, implying that it is possible to produce them in colliders and also observe rare decays [158]. The lightest neutral component of a triplet could also play the role of a dark matter candidate [159], and can also successfully lead to leptogenesis [160]. Nevertheless, as mentioned, we focus on the minimal type-I seesaw model here.

The inverse seesaw

In addition to the variants of the “standard” mechanism, there also exists a class of *inverse seesaw* mechanisms, where both right-handed neutrinos and sterile fermions (the sterile states are left-handed, but singlets of the Standard Model gauge group) are added to the Standard Model [161, 162, 163]. An extra (dimensionful) parameter offers the chance to obtain small active neutrino masses, while still having sizeable couplings among the LHNs and RHNs.

In [164], it was shown that two sub-classes of the inverse seesaw (with minimal additions) can account for oscillation data. The first extends the SM by two RHNs and two sterile states. It leads to compliance with all constraints, but is quite fine tuned and strictly leads to a normal hierarchy for the light neutrinos [164]. The second class also extends the SM by two RHNs, but by three sterile states. This class allows the existence of both normal and inverted mass hierarchies for the light neutrinos (IH only marginally so). The mass of the lightest sterile neutrino can vary over a large interval and, depending on its regime, can explain reactor anomalies [164], or provide a dark matter candidate (for a mass of the lightest sterile state around a keV) [164].

Our choice of using the type-I seesaw model is motivated by multiple reasons: it is the most parsimonious of all the variants with respect to the addition of new fields, and is also the most widely studied. Additionally, it effectively describes the inverse seesaw in the limit of conservation of B-L (the difference between baryon and lepton number).

2.3 The Casas-Ibarra parameterization

The bottom-up Casas-Ibarra (C-I) scheme [78] is used in this study. It offers many advantages for our purpose: since the parameter space is of high dimensionality, any simplification offered would be a boon. This scheme, by construction, satisfies neutrino oscillation data, making it easier to sample the space while still obeying these constraints (something which can be hard to do otherwise, since the oscillation parameters have been determined quite accurately). The numerical benefit this confers was a strong reason for this choice.

Let us start simple and begin the derivation at tree level. Also, let's work in a basis where the charged lepton mass matrix is diagonal. With the type-I seesaw model (containing 3 RHNs) in mind, the full neutrino mass matrix after EWSB is

$$\mathcal{M} = \begin{pmatrix} \mathbf{0} & M_D \\ M_D^T & M_M \end{pmatrix}, \quad (2.73)$$

which, as in equation 2.54, can be diagonalized by a unitary matrix U via a rotation between neutrino flavour and mass eigenstates. As before, since the "active" block of U is unitary to a good degree of approximation, this mixing matrix can be expanded:

$$U = \begin{pmatrix} \mathbf{1} - \frac{\theta\theta^\dagger}{2} & \theta \\ -\theta^\dagger & \mathbf{1} - \frac{\theta\theta^\dagger}{2} \end{pmatrix} \begin{pmatrix} U_\nu & \mathbf{0} \\ \mathbf{0} & U_N^* \end{pmatrix} + \mathcal{O}(\theta^3) = \begin{pmatrix} U_\nu & \theta U_N^* \\ -\theta^\dagger U_\nu & U_N^* \end{pmatrix} + \mathcal{O}(\theta^2). \quad (2.74)$$

Also, from the diagonalization of \mathcal{M} , the following relations can be obtained at leading order in θ :

$$\theta^* M_M \theta^\dagger \approx -U_\nu^* m_\nu^{\text{diag}} U_\nu^\dagger \quad \text{and} \quad (2.75)$$

$$M_M \approx U_N M_N^{\text{diag}} U_N^T. \quad (2.76)$$

These two relations imply that

$$(\theta U_N^*)^* M_N^{\text{diag}} (\theta U_N^*)^\dagger \approx -U_\nu^* m_\nu^{\text{diag}} U_\nu^\dagger. \quad (2.77)$$

The Casas-Ibarra parameterization [78] takes the LHN masses and angles and the PMNS matrix phases as input, ensuring that the constraints in this sector are always satisfied. The first step is to manipulate equation 2.75 and re-write it in the form

$$\left(i \frac{1}{\sqrt{m_\nu^{\text{diag}}}} U_\nu^\dagger \theta U_N^* \sqrt{M_N^{\text{diag}}} \right) \left(i \frac{1}{\sqrt{m_\nu^{\text{diag}}}} U_\nu^\dagger \theta U_N^* \sqrt{M_N^{\text{diag}}} \right)^T = 1. \quad (2.78)$$

The general solution to this is $\mathcal{R}\mathcal{R}^T = 1$, where \mathcal{R} is constrained to be a 3×3 complex, orthogonal matrix which encodes the degrees of freedom associated with the RHN sector. With this, $\Theta \equiv \theta U_N^*$, which represents the mixing among the LHNs and RHNs, can be written in terms of \mathcal{R} as:

$$\Theta = i U_\nu \sqrt{m_\nu^{\text{diag}}} \mathcal{R} \sqrt{M_N^{\text{diag}}}. \quad (2.79)$$

Following [165] and inserting $\delta m_\nu^{1\text{-loop}}$ from equation 2.49 into equation 2.57 allows us to re-write the latter as

$$m_\nu = -\theta \tilde{M} \theta^T, \quad (2.80)$$

analogous to equation 2.58, where

$$\tilde{M} = U_N \tilde{M}^{\text{diag}} U_N^T; \quad (2.81)$$

$$\tilde{M}^{\text{diag}} = M_I \delta_{IJ} \left[\mathbf{1} - \frac{M_I^2}{v^2} I(M_I) \right]. \quad (2.82)$$

Under these conditions, the matrix U which diagonalized the full neutrino mass matrix is not modified by radiative corrections, apart from the change in m_ν . Without such corrections, \tilde{M} would be identical to M_N at second order in θ . The one-loop corrected Casas-Ibarra scheme is now slightly modified, becoming

$$\Theta = i U_\nu \sqrt{m_\nu^{\text{diag}}} \mathcal{R} \sqrt{\tilde{M}^{\text{diag}}}^{-1}. \quad (2.83)$$

The parameterization of \mathcal{R}

The matrix \mathcal{R} is, in turn, parameterized by 3 complex angles ω_{ij} in this way:

$$\mathcal{R} = \mathcal{R}^{23} \mathcal{R}^{13} \mathcal{R}^{12}, \quad (2.84)$$

with \mathcal{R}^{ij} having the non-zero elements

$$\mathcal{R}_{ii}^{ij} = \mathcal{R}_{jj}^{ij} = \cos \omega_{ij}, \quad (2.85)$$

$$\mathcal{R}_{ij}^{ij} = -\mathcal{R}_{ji}^{ij} = \sin \omega_{ij}, \quad (2.86)$$

$$\mathcal{R}_{kk}^{ij} = 1; k \neq i, j. \quad (2.87)$$

Since we work in the flavor basis where the Yukawa couplings of the charged leptons are diagonal, U_ν can be parameterized as

$$U_\nu = V^{23} U_\delta V^{13} U_{-\delta} V^{12} \text{diag}(e^{i\alpha_1/2}, e^{i\alpha_2/2}, 1), \quad (2.88)$$

where $U_{\pm\delta} = \text{diag}(e^{\mp i\delta/2}, 1, e^{\pm i\delta/2})$. The matrix V^{ij} has the non-zero entries

$$V_{ii}^{ij} = V_{jj}^{ij} = \cos \theta_{ij}, \quad (2.89)$$

$$V_{ij}^{ij} = -V_{ji}^{ij} = \sin \theta_{ij}, \quad (2.90)$$

$$V_{kk}^{ij} = 1; k \neq i, j, \quad (2.91)$$

where θ_{ij} are the LHN mixing angles. The parameters α_1, α_2 and δ are CP-violating phases.

Note that Θ has an exponential dependence on the imaginary parts of ω_{ij} : $\Theta^2 \sim \frac{\exp(2\text{Im}(\omega))}{M}$. Large values of $\text{Im}(\omega)$ produce mixings that are too large to pass any constraints, so we pre-emptively disallow choices that lead to $|\Theta|_{ij}^2 > 1$.

2.4 Tuning and symmetries

An easy way to generate small neutrino masses m_i is the limit $M_I \gg v$, as motivated by Grand Unified Theories ($M_I \sim 10^{16}$ GeV) [166]. However, such heavy RHNs would be hard to detect experimentally.

The case can be made for the naturalness of lower values of M_I ($M_I \sim \text{eV}$), since lepton number conservation (broken by the Majorana masses M_I) is restored in the limit of all $M_I \rightarrow 0$. However, then, the above suppression by v/M_I that explains the smallness of the m_i in the conventional seesaw scenario becomes insufficient to generate active neutrino masses that agree with experimental values [54]. By bringing the RHN masses to the scale of experimentally accessible values and ignoring the mixing between left-handed neutrinos (which contributes $\mathcal{O}(1)$ corrections), the largest coupling and mixing that a right-handed neutrino can have turns out to be [59]

$$F_{\alpha I} \sim f_I^2 \equiv \frac{M_I}{v^2} \sqrt{\Delta m_{\text{atm}}^2 + m_{\nu_0}^2}; \quad (2.92)$$

$$U_{\alpha I} \sim u_I^2 \equiv \frac{1}{M_I} \sqrt{\Delta m_{\text{atm}}^2 + m_{\nu_0}^2}, \quad (2.93)$$

with m_{ν_0} being the mass of the lightest left-handed neutrino. The u_I^2 defined above can be very small; for an RHN of mass 1 GeV, assuming a lightest neutrino mass of about 0.2 eV (the approximate limit KATRIN is expected to place in the near future [167]), $f_I \approx 9 \times 10^{-8}$ and $u_I^2 \approx 2.3 \times 10^{-10}$.

This is smaller than the expected sensitivity of upcoming experiments like NA62 [110], SHiP [168] and T2K [169].

There is a way out, though, and it exploits the fact that the estimate in equation 2.92 assumes that there are no cancellations in the seesaw relation. Such cancellations would allow for larger individual $U_{\alpha I}^2$ while keeping the eigenvalues m_i^2 of $m_\nu m_\nu^\dagger$ small. One way for this to happen is for the Lagrangian (equation 2.46) to (approximately) respect the B-L symmetry of the SM [65, 66, 170]. Following [111], this can be realized if the Yukawa coupling and mass matrix can be brought into the form

$$M_M = \overline{M} \begin{pmatrix} 1 - \mu & 0 & 0 \\ 0 & 1 + \mu & 0 \\ 0 & 0 & \mu' \end{pmatrix} \quad \text{and} \quad (2.94)$$

$$F = \begin{pmatrix} F_e(1 + \epsilon_e) & iF_e(1 - \epsilon_e) & F_e\epsilon'_e \\ F_\mu(1 + \epsilon_\mu) & iF_\mu(1 - \epsilon_\mu) & F_\mu\epsilon'_\mu \\ F_\tau(1 + \epsilon_\tau) & iF_\tau(1 - \epsilon_\tau) & F_\tau\epsilon'_\tau \end{pmatrix} \quad (2.95)$$

by a rotation in flavour space. Here $\epsilon'_\alpha, \epsilon_\alpha, \mu, \mu' \ll 1$ are small parameters that violate lepton number. In the RHN mass basis, this implies that

$$F_{\alpha 1} \simeq \frac{1}{\sqrt{2}} (F_\alpha + \epsilon_\alpha) \quad , \quad M_1 \simeq \overline{M} - \mathcal{O}[\mu_i] \quad (2.96)$$

$$F_{\alpha 2} \simeq \frac{i}{\sqrt{2}} (F_\alpha - \epsilon_\alpha) \quad , \quad M_2 \simeq \overline{M} + \mathcal{O}[\mu_i] \quad (2.97)$$

$$F_{\alpha 3} \simeq \epsilon'_\alpha \quad , \quad M_3 \simeq M' \quad (2.98)$$

$$(2.99)$$

i.e. two form a Dirac spinor with mass \overline{M} , while the other is left lighter, with a mass M' , and has a feeble coupling strength (in fact, this is precisely the pattern in the ν MSM; see [66])⁴.

This scenario predicts that the two RHNs that are mass degenerate will have approximately equal mixings in all three flavours: $U_{\alpha I}^2 \cong U_{\alpha J}^2$. Further,

⁴An exception is when $\mu \ll 1$ and $M' \simeq \overline{M}$, in which case even small off-diagonal elements μ_{ij} can lead to a large misalignment between the basis in which F has the form in 2.95 and the RHN mass basis. In this scenario, all RHNs will have unsuppressed Yukawa couplings $\sim F_\alpha$, despite $\epsilon'_\alpha \ll 1$, as discussed in [171]. Further, all three mass eigenstate N_I will have approximately the same mass \overline{M} , precluding a kinematic distinction. Then, the experimentally relevant mixing is $U_{\alpha I}^2$, whose magnitude is controlled by the (large) entries F_α . Heavy neutrino oscillations in the detector [172, 173, 174, 175, 176, 177, 178, 179, 180, 181] could facilitate an indirect way to access the small mass splitting and phenomenologically study this case.

there is no upper limit on $U_{\alpha I}^2$ from neutrino oscillation data or neutrinoless double-beta decay. Although such lepton number violating observables exert strong constraints on $U_{\alpha I}^2$, the tiny parameters $\epsilon'_\alpha, \epsilon_\alpha$ and μ/\bar{M} in equations 2.94 and 2.95 suppress these signals. In the mass range considered here, the actual upper limit then comes from other experimental constraints, like direct searches. For larger masses there is a theoretical bound $U_{\alpha I}^2 < 4\pi(n-1)(v/\bar{M})^2$ from the requirement that the Yukawa couplings remain perturbative [182], n being the number of right-handed neutrinos. Hence, there is theoretical motivation for a low-scale seesaw with experimentally accessible $U_{\alpha I}^2 \gg u_1^2$. Specific examples that motivate this limit include “inverse seesaw”-type scenarios [64, 161, 183, 184], a “linear seesaw” [185, 186], scale-invariant models [42], some technicolour-type models [187, 188] or the ν MSM [66].

Within the C-I parameterization, this approximate B-L symmetry protected scenario can be realized by making two RHNs are degenerate in mass, say $M_1 = M_2$, and the corresponding complex angle in \mathcal{R} is large, while the other two are small or zero, i.e. $\omega_{12} \rightarrow \pm i\infty, \omega_{13}, \omega_{23} = 0$. These settings allow one to obtain the highest mixing values that are allowed by constraints, while still bypassing the strong limits on U_{eI}^2 coming from neutrinoless double-beta decay; see section 4.2.5 in chapter 4.

On the other hand, the Casas-Ibarra parameterization, by construction, respects the constraints from light neutrino oscillation data, and it is easy to obtain very large values of $U_{\alpha I}^2$ by simply choosing $|\text{Im}(\omega_{ij})| > 1$. However, most parameter points obtained this way will not respect (approximate) B-L symmetry, but are instead due to an unmotivated “fine-tuning”; the smallness of the m_i is the result of accidental cancellations in m_ν .

2.4.1 Distinguishing symmetry-protected from tuned parameter choices

The Casas-Ibarra-parameterization (2.83) is, as previously mentioned, a “bottom up” scheme, and it is not easy to see, directly from the values of its fundamental parameters, whether these parameters are symmetry-protected. An analytic exploration of all possible solutions and their classification between symmetric and fine-tuned would be useful, but this is outside the scope of this work. In our scan, we take a pragmatic approach: first, a large amount

of parameter choices is generated by randomizing the parameter values and the order of the matrices \mathcal{R}^{ij} in equation 2.84 to ensure maximal coverage of the parameter space. Then, the following cut is applied to distinguish the symmetry protected points:

$$\begin{aligned} \frac{|M_2 - M_1|}{M_2 + M_1} &< \epsilon, \quad \frac{m_{\nu_0}}{\mu\text{eV}} < 1, \\ |F_{\alpha 3}| &< \epsilon, \quad \frac{|F_{\alpha 1} + iF_{\alpha 2}|}{|F_{\alpha 1}| + |F_{\alpha 2}|} < \epsilon. \end{aligned} \quad (2.100)$$

This cut practically enforces the structure in equations 2.94 and 2.95 on the masses and couplings.

Using the above cut (2.100) requires some care for two reasons. First, equations 2.94 and 2.95 do not encapsulate all symmetry-protected points; see footnote 4. Some symmetry-protected points may inadvertently be misidentified as tuned. Nevertheless, we find that the number of such points is small. Second, when using the Casas-Ibarra parameterization, it is possible to generate points that mimic the forms of equations 2.94 and 2.95 (and hence pass the cut in equation 2.100), but in reality exhibit a significant amount of tuning.

As a demonstration of the second point, let us work at tree level and approximate $\tilde{M}^{\text{diag}} \simeq M_M$, yielding

$$F \approx iU_\nu \sqrt{m_\nu^{\text{diag}}} \mathcal{R} \sqrt{M_M} / v. \quad (2.101)$$

For the inverted hierarchy, one gets a pseudo-Dirac pair of heavy neutrinos with the choices

$$M_1 = M_2 = \overline{M}, \quad (\omega_{12}, \omega_{13}, \omega_{23}) = (\omega, 0, 0), \quad (2.102)$$

with $|\text{Im}\omega| \gg 1$. Additionally, one has to set $m_{\nu_0} = 0$ to find the symmetry-protected region, since a non-zero lightest neutrino mass is not consistent with $\epsilon'_\alpha \rightarrow 0$. With these choices, the upper left block of the matrix $\mathcal{R}\sqrt{M_M}$ in equation 2.101 is large, while the third row, which multiplies $m_3 = m_{\nu_0}$, is small, reproducing the structure in equations 2.94 and 2.95. Thus, the decoupling $F_{\alpha 3} = 0$ and the vanishing mass of the lightest neutrino can be interpreted as results of the symmetry.

Now, for normal ordering, making the choices in 2.102 will also yield a structure that passes the cut in 2.100, but this is, in fact, a tuned solution mimicking the structure. In this case, $m_{\nu_0} = m_1$ multiplies the large components of $\mathcal{R}\sqrt{M_M}^{-1}$ in equation 2.101 and hence, the approximate symmetry makes the wrong left-handed neutrino mass small (m_3 instead of $m_{\nu_0} = m_1$). Although $m_1 = 0$ can be enforced by hand in 2.101, this choice cannot be justified by the symmetry. Therefore, despite the choices in 2.102 leading to a pseudo-Dirac structure amongst the ν_R (as predicted by the symmetry), the vanishing mass of the lightest neutrino is not a result of that symmetry.

The problem is that the Casas-Ibarra parameterization allows the enforcement of $m_{\nu_0} = 0$ by hand, but gives no indication if this leads to accidental cancellations. In order for the symmetry to be realized, the eigenvalues of M_M and the non-zero ω_{ij} must be chosen so that the large block in the matrix $\mathcal{R}\sqrt{M_M}^{-1}$ in equation 2.101 multiplies the two non-zero light neutrino masses. For normal ordering, this can be done with the choices

$$M_2 = M_3 = \bar{M}, (\omega_{12}, \omega_{13}, \omega_{23}) = (0, 0, \omega), \quad (2.103)$$

again with $m_{\nu_0} = 0$. Note that one cannot choose ν_{R3} to be the particle that decouples. This is, however, not a fundamental problem, because the labels of the ν_{RI} have no physical meaning, but it does mean that the labeling and order of the matrices \mathcal{R}^{ij} have to be carefully considered when applying a cut to identify symmetry-protected points in the data.

As mentioned earlier, we randomize the order of the three matrices \mathcal{R}^{ij} in equation 2.84 to generate more points. Making the choices in 2.103 (2.102) for normal (inverted) neutrino mass ordering, reproduces the approximate $B - L$ conserving limit, regardless of the ordering of the \mathcal{R}^{ij} . However, in light of the cuts we perform (2.100), the effect of small perturbations around these limits must be considered, since they have a strong dependence on this ordering. The effect is smallest if the matrices \mathcal{R}^{ij} are ordered in a way that the one with large entries (controlled by ω) directly multiplies $\sqrt{M_M}^{-1}$ in equation 2.101. For normal ordering, this is the case with $\mathcal{R} = \mathcal{R}^{23}\mathcal{R}^{13}\mathcal{R}^{12}$; for inverted ordering, this is independent of permutations, since two of the ω_{ij} are zero.

One last point which is worth noting is that the choices in 2.102 and 2.103 are not the only ones that yield the symmetry-protected scenario, but just the simplest. The non-trivial structure of the complex rotation matrix \mathcal{R} allows

for many solutions to the required layout described above. Case in point, we take advantage of this fact by taking, for normal ordering,

$$M_1 = M_2 = \overline{M}, (\omega_{12}, \omega_{13}, \omega_{23}) = (0, \pi/2, \omega), \quad (2.104)$$

since, in this work, we focus on the case where M_1 and M_2 are almost degenerate.

Chapter 3

GAMBIT and the statistical analysis framework

As has been mentioned before, the issue of prior independence as well as the difficulty in setting up reasonable priors in the Casas-Ibarra parameterization scheme made us adopt a frequentist approach to analyze the parameter space. Before going into a description of GAMBIT (the Global And Modular BSM Inference Tool), its structure and inner machinery, some material regarding the two prevalent statistical approaches, frequentist and Bayesian, will be presented. The debate over which statistical framework is more logically and philosophically consistent is an old and contentious one and one that is still ongoing [189]. The Bayesian approach dates back about two and a half centuries ago, to Thomas Bayes, whose paper was posthumously published in 1763 [190]. It was further developed by Pierre-Simon Laplace [191] and axiomatized by Harold Jeffreys [192]. The Polish statistician Jerzy Neyman played a pivotal role in the development of frequentist statistics [193].

The contentious nature of the debate is not without good reason: if the data is not very constraining, the interpretation of experimental results can vary appreciably depending on which methodology is employed. In these cases, a frequentist approach arguably provides a more objective and inclusive perspective on the allowed parameter space.

3.1 Frequentist statistics

The differences between Bayesians and frequentists has its root in how the concept of probability is interpreted everywhere by each. For frequentists,

the probability p of something is defined in terms of a large number N of essentially identical, independent trials: if this “something” happens in s trials, p is defined as the limit of the ratio s/N , as $N \rightarrow \infty$. Since a repeated series of trials is required, assigning probabilities to single events is a meaningless endeavour in frequentist statistics. For example, if the question of whether it snowed in Oslo yesterday is posed, the frequentist position would be that even though this may not be known, in actual fact it either was snowing or not, so this is not a matter for assigning a probability to. There is also the, let’s say, technical problem of being unable to set up a large number of yesterdays in order to define a sensible probability in this framework. Another relevant example is that of physical constants. In the frequentist approach, one cannot assign probabilities to statements involving the numerical values of physical parameters, since this (again) is a situation which cannot be checked by replication. Hypotheses are either true or false, and not directly suitable for frequentist probabilities. A similar argument applies to statements about theories: the framework will not allow probability assignments as to whether, for example, the Higgs boson exists.

Consider a simple problem: an experiment attempts to measure the mass m of a particle and gets the result m_0 ; the experimentalists understand the apparatus well and know that it yields values in a Gaussian distribution around the true mass m_t with variance σ_m^2 . Assuming the measured value m_0 is much above 0, the *probability density function* (pdf) $P(m|m_t)$ for getting out a mass m from the result, given the true mass m_t is

$$P(m|m_t) = N(m_t, \sigma_m) = \frac{1}{\sqrt{2\pi\sigma_m^2}} e^{-\frac{(m-m_t)^2}{2\sigma_m^2}}. \quad (3.1)$$

Gaussian/normal distributions like the one depicted in equation 3.1 are ubiquitous in statistical analyses. Part of their ubiquity lies in the central limit theorem (CLT), which states that if one has n independent, identically distributed random variables X_i ($i : 1, \dots, n$), with means μ_i and variances σ_i^2 , the sample average $\bar{X} = \sum_i X_i/n$ will have a Gaussian distribution characterized by mean $\mu = \sum_i \mu_i/n$ and variance $\sigma^2 = \sum_i \sigma_i^2/n$, regardless of the underlying distribution (in fact, this is true of any distribution with a finite variance [194]).

The *likelihood* $\mathcal{L}(m_t|m_0)$ has the same form as equation 3.1, with an important change of perspective: it encodes the behavior of how the function would

change for different values of m_t , given the observed value m_0 .

$$\mathcal{L}(m_t|m_0) = \frac{1}{\sqrt{2\pi\sigma_m^2}} e^{-\frac{(m_0-m_t)^2}{2\sigma_m^2}}. \quad (3.2)$$

It should not be confused with equation 3.1, not only for this reason, but also because likelihoods must not be interpreted as probability densities of model parameters. They do not behave as pdfs do under a transformation of parameters (a probability density function of a random variable is a function whose integral across an interval gives the probability that the value of the variable lies within the interval).

Consider equation 3.1. Neyman first developed the technique [193] to construct what are now called “classical” confidence intervals. To construct a central 68% CL confidence interval for m_t , one finds the value of $m_l < m_0$, such that $(100\%-68\%)/2 = 16\%$ of the area under $N(m_l, \sigma_m)$ is covered for values of $m > m_0 > m_l$ and $m_u > m_0$ such that 16% of the area under $N(m_u, \sigma_m)$ is covered for values of $m_u < m < m_0$. This leads to $m_l = m_0 - \sigma_m$ and $m_u = m_0 + \sigma_m$, forming the conventional statement that the measured mass is $m_0 \pm \sigma_m$.

The canonical method to construct, say 1σ CL confidence intervals from \mathcal{L} for arbitrary distributions is to find m_l and m_u such that at these values, $\ln \mathcal{L}$ is less than its maximum value by 0.5, or equivalently, those points at which $-2 \ln \mathcal{L}$ increases by 1. This is related to Wilks’ theorem [195], which is discussed below.

Central to the construction of confidence intervals is the concept of coverage. In terms of similar experiments aiming to measure some quantity m and compute confidence intervals for m_t at say, 68% CL, the classical construction provides the guarantee that if the ensemble of experiments is large enough, 68% of these intervals will contain the (unknown) true value m_t , or “cover” m_t . This definition of coverage is indicative of the frequency with which the statement “ m_t is within (m_l, m_u) ” is true. In such a statement, the variables are m_l and m_u ; m_t is fixed, but unknown. It is *not* a statement that is meant to tell us how strongly we should believe that m_t lies within the interval set by any particular experiment. The concept of degree of belief is absent in frequentist confidence intervals, which are constructed to be independent of any probabilistic statements about m_t (a restatement of the earlier fact that one cannot treat parameters probabilistically in the frequentist framework).

However, as Neyman himself said [193], coverage, as a concept, is not restricted in applicability to large ensembles of near-identical experiments. Frequentist confidence intervals have a very powerful property: suppose that in an ensemble of different experiments, each measures different observables, constructing 68% CL intervals for them. In the long run, 68% of the intervals will cover the true value of their respective observables.

Techniques which lead to a coverage that is equal to the nominal value (e.g. 68%) for all values of the parameter are said to have *exact coverage*. If the coverage drops below the nominal value, the method is said to *under-cover*. This is regarded as a bad quality: if the actual coverage for determining the parameter is, say, 25% and not the nominal 68%, quoting the range for the parameter as determined by that method can confuse a reader into believing that the result is more reliable than it is. Over-coverage will not suffer from this problem, of course, but such confidence intervals are more *conservative* (wider) than they need be.

3.1.1 The profile likelihood ratio

As an example, consider the Reines and Cowan experiment [196], where a detector sensitive to neutrinos interacting in it was built close to a nuclear reactor. There were background processes which mimic the interactions of the reactor neutrinos. The observed number of counts n would have been Poisson-distributed with mean $b + \mu$:

$$P(n) = \frac{e^{-(b+\mu)}(b+\mu)^n}{n!}, \quad (3.3)$$

b being the expected background, and μ the signal rate. If, however, b is not known or well-understood, this will lead to a systematic uncertainty; b is then referred to as a “nuisance parameter”.

Likelihood functions can often be dependent on quantities/parameters that are not the focus of the analysis. In more complicated cases, the likelihood function can be dependent on many parameters, while we might be interested in only a subset of them; for example, in this study, we are primarily interested in seeing the behavior of the likelihood function as the right-handed neutrino mass and coupling to left-handed neutrinos vary. Other quantities

like those contained in the PMNS matrix are not relevant in this context. Nuisance parameters are commonly handled using maximum/profile likelihood ratios, which we discuss next.

In practice, to construct the profile likelihood contours, we make use of Wilks' theorem [195]. The theorem provides an asymptotic distribution that can be used for the log-likelihood statistic, similar to how the asymptotic limit prescribed by the central limit theorem is the normal distribution. Consider a set of parameters \vec{x} that is modeled by the distribution of the likelihood function for the true values of the model parameters $\theta_1, \dots, \theta_n$. A statistic called the *maximum likelihood ratio* can be constructed:

$$\Lambda(\theta_1, \dots, \theta_k | \vec{x}) = -2 \ln \frac{\mathcal{L}(\theta_1, \dots, \theta_k, \hat{\theta}_{k+1}, \dots, \hat{\theta}_n | \vec{x})}{\mathcal{L}(\hat{\theta}_1, \dots, \hat{\theta}_n | \vec{x})}, \quad (3.4)$$

where the $\hat{\theta}_i$ are so-called maximum likelihood estimators (MLEs) for the function \mathcal{L} ; they are values of the parameters θ_i which maximize \mathcal{L} . The quantities $\hat{\theta}_i$ denote the profiled values of the parameters, i.e. those values of each parameter θ_i that maximize \mathcal{L} , keeping all other parameters $\theta_{j \neq i}$ fixed. Wilks' theorem states that in the asymptotic limit, this statistic will follow a χ^2 distribution with k degrees of freedom, i.e. in the limit of a large amount of data,

$$\Lambda(\theta_1, \dots, \theta_k | \vec{x}) \sim \chi_k^2. \quad (3.5)$$

For example, take the case where the likelihood function is a normal distribution with one degree of freedom, the mean. Then,

$$\Lambda = -2 \ln \frac{e^{-(\vec{x} - \vec{\mu})^2 / 2\sigma^2}}{e^{-(\vec{x} - \vec{\mu})^2 / 2\sigma^2}}. \quad (3.6)$$

Setting the derivatives of the log-likelihood with respect to $\vec{\mu}$ to zero, one can obtain the MLEs, which turn out to be $\vec{\mu} = \vec{x}$, leaving

$$\Lambda = -2 \ln e^{-(\vec{x} - \vec{\mu})^2 / 2\sigma^2} = \frac{(\vec{x} - \vec{\mu})^2}{\sigma^2}, \quad (3.7)$$

which is the χ^2 distribution with one degree of freedom.

With these ingredients, the procedure for constructing confidence intervals is as follows: as an example, let one model parameter θ_1 hold our interest and let the others be nuisance parameters. The (one-dimensional) maximum

log-likelihood ratio is

$$\Lambda(\theta_1|\vec{x}) = -2 \ln \frac{\mathcal{L}(\theta_1, \hat{\theta}_2, \dots, \hat{\theta}_n|x)}{\mathcal{L}(\hat{\theta}_1, \dots, \hat{\theta}_n|x)}. \quad (3.8)$$

First, $-2 \ln \mathcal{L}(\theta_1, \dots, \theta_n|x)$ is *minimized* with respect to all model parameters. Then, the parameter of interest (θ_1) is varied until $-2 \ln \mathcal{L}$ increases by the desired confidence level c , while simultaneously fitting the nuisance parameters. The values of θ_1 obtained then define the required confidence interval:

$$I(\vec{x}_{\text{obs}}) = \{\theta_1 | \Lambda(\theta_1|\vec{x}_{\text{obs}}) < c\}. \quad (3.9)$$

This can be generalized to higher dimensional confidence intervals as

$$I(\vec{x}_{\text{obs}}) = \{\vec{\theta} \in \mathbb{R}^k | \Lambda(\theta_1, \theta_2, \dots | \vec{x}_{\text{obs}}) < c\}. \quad (3.10)$$

Note that the value of c that defines the threshold(s) of the confidence interval depends on the interval's dimensionality: the cumulative distribution function (the cdf of a random variable, evaluated at a particular value c , gives the probability that the variable will take a value less than or equal to c) of a χ^2 distribution with k degrees of freedom is

$$F(c; k) = \frac{\gamma\left(\frac{k}{2}, \frac{c}{2}\right)}{\Gamma\left(\frac{k}{2}\right)}, \quad (3.11)$$

where Γ is the function defined by

$$\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx, \quad \Re(z) > 0, \quad (3.12)$$

and the function γ is the lower incomplete gamma function, defined by

$$\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt, \quad s \in \mathbb{C}, \Re(s) > 0. \quad (3.13)$$

For a χ^2 distribution with one degree of freedom ($k = 1$), $F(c) = \text{erf}(\sqrt{c/2})$. A 68.3% (1σ) confidence interval means one must solve $F(c) = \alpha (= 0.683)$. Therefore, in this case,

$$\text{erf}(\sqrt{c/2}) = \alpha \implies c = 2(\text{erf}^{-1}(\alpha))^2 = 2(\text{erf}^{-1}(0.683))^2 = 1.003. \quad (3.14)$$

On the other hand, if the distribution has two degrees of freedom ($k = 2$), $F(c) = 1 - e^{-c/2}$ for $c > 0$, so in this instance,

$$1 - e^{-c/2} = \alpha \implies c = -2\ln(1 - \alpha), \quad (3.15)$$

and for $\alpha = 0.683$, $c = -2\ln(0.317) = 2.298$.

A couple of caveats should be mentioned here regarding our usage of the asymptotic distributions prescribed by Wilks' theorem and the central limit theorem. We construct profile likelihood contours in the mixing-mass plane, and hence use a χ^2 distribution with two degrees of freedom while extracting the intervals. However, the limits from, for instance, direct detection experiments are in the form of confidence intervals for the mixing alone (using a single degree of freedom). Since a relatively larger region of the parameter space will be covered by a χ^2 distribution with two degrees of freedom, it must be kept in mind that the contours we obtain can be somewhat weaker than what the limits from the experiments would suggest. Additionally, the large sample limit that is implicitly assumed in the CLT is not necessarily satisfied by the individual experiments considered here, since they often report non-detection, or 1-2 candidate events at most. Nevertheless, keeping these general caveats in mind, the frequentist approach provides a largely objective view on the validity of different parts of the RHN parameter space.

3.2 The Bayesian framework

For Bayesians, probability is a personal assessment of how likely it is that something is true [197]. It depends on personal biases and/or previous knowledge about the situation, and can differ from person to person. For example, suppose I toss a coin, and ask you what the probability is of the result being heads, a reasonable guess would be 50%. But maybe I am dishonest, looked at the coin and saw that it was tails. For me, the probability of heads is definitely 0%. Or perhaps I just managed a quick glance, and think (but am not certain) that it was tails, so I assign a probability of 25% to heads. This personal view of probability means that it is possible to give numerical estimates for one-off situations (who wins this year's election), for parameter values (eg. fraction of dark matter), or regarding theories (existence of Higgs boson). Again, the numerical assessments could vary from person to person. In fact, it sounds as if this is not conducive to numerical estimates. But the

Bayesian assessment of probability *must* be consistent with the “fair bet” concept: if a Bayesian believes that a certain statement has a 10% probability of being true, they should be prepared to offer 9 to 1 odds (or 1 to 9) to someone who is prepared to bet against them (or with them, respectively) [198, 199].

The probability of two independent events A and B both happening is given by well-known result [190]

$$P(A \text{ and } B) = P(A)P(B). \quad (3.16)$$

However, if they are not independent, this changes to

$$P(A \text{ and } B) = P(A|B)P(B), \quad (3.17)$$

where $P(A|B)$ is the *conditional probability* of A occurring, taking into account the fact that B has occurred. For example, let A be the occurrence of snowfall in Oslo and B be the occurrence of a day in December. Picking a random day of the year, the probability of it being a snowy December day in Oslo, $P(A \text{ and } B)$, is the probability of snowfall in Oslo in December $P(A|B)$ multiplied by the probability of the day being in December $P(B) = 31/365$.

Further, $P(A \text{ and } B)$ is symmetric in A and B , so

$$P(A \text{ and } B) = P(A|B)P(B) = P(B|A)P(A). \quad (3.18)$$

From the second equality, one gets Bayes’ theorem:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \quad (3.19)$$

which relates $P(A|B)$ and $P(B|A)$.

Once A is taken to be a theoretical parameter and B to be the experimental data; the theorem then states:

$$P(\text{param}|\text{data}) \propto P(\text{data}|\text{param})P(\text{param}). \quad (3.20)$$

$P(\text{data}|\text{param})$ is the likelihood function from earlier; $P(\text{param})$ is the *prior density*. The prior expresses what was known about the parameter before the measurement and encodes the biases or “personal” aspect of Bayesian analysis. $P(\text{param}|\text{data})$ is the *posterior probability density* for the parameter, and contains the information about the parameter obtained by combining

prior information with that from measurement. However, due to the complications that can arise when choosing an appropriate probability density for the prior; conclusions drawn from the data can be heavily dependent on these choices. This has been seen, for instance, in other scans done using GAMBIT: while performing a global fit of the MSSM [99] and while studying axions [101].¹ As I have mentioned before, these disadvantages informed our decision to perform a frequentist analysis of the parameter space.

3.3 GAMBIT

It is very often the case, especially in the case of BSM (Beyond the Standard Model) theories, that multiple experiments, using different approaches, provide data that is pertinent to a theory. Global fits need to correctly weigh the sum total of all the data and make statistical statements about the theory. This has been attempted in several cases; for example, in the context of electroweak physics [200, 201, 202], to CKM fits [203], neutrinos [102, 204, 205, 206, 207], as well as supersymmetry (SUSY) [99, 100, 208, 209, 210, 211, 212, 213, 214, 215, 216, 217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 230, 231, 232]. In fact, BSM global fits have been largely focused on SUSY, particularly on lower-dimensional subspaces of the minimally supersymmetric standard model (MSSM) [233], or the next-to-minimal variant (NMSSM) [234]. The software frameworks that formed the base of these analyses [200, 202, 220, 222, 235] have long development cycles, even more so when new physics needs to be added to them.

In addition, carrying out such analyses in many-dimensional parameter spaces is technically challenging, requiring, among other things, a comprehensive understanding of many disparate theory calculations and experiments, coding experience, large amounts of computing time, and careful attention to statistical and numerical methods [236].

¹Nevertheless, in some cases, the “data overwhelms the prior”, and the result becomes insensitive to the choice of prior. For example, the mass of the intermediate vector boson, the Z^0 , was measured at LEP at CERN. Having found that the likelihood function was essentially a Gaussian at $91.188 \text{ MeV}/c^2$, with a width of $2 \text{ MeV}/c^2$. Any reasonable choice of prior for the Z^0 mass will vary negligibly over the range (a few parts in 10^5), and so in this case the posterior is essentially independent of the prior.

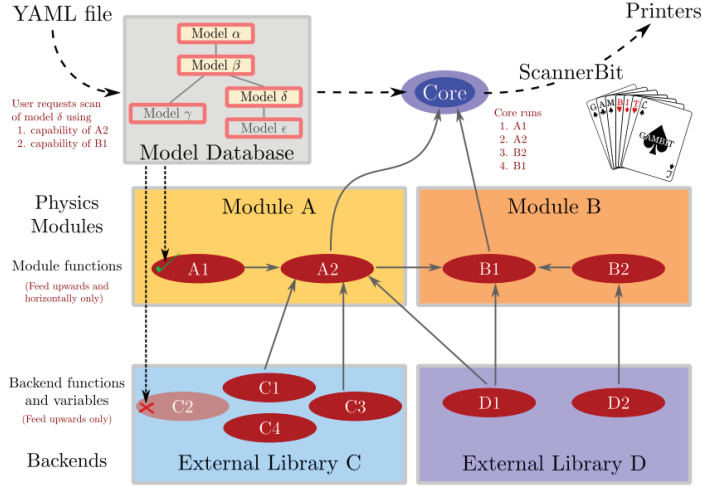


FIGURE 3.1: Taken from [93], the figure shows a schematic representation of the structure of a GAMBIT scan. Once provided with a YAML input file (see www.yaml.org), which specifies a model to scan and observables or likelihoods to calculate, the requested model δ and its ancestor models β and α are activated by GAMBIT. All model-dependent module and backend functions/variables are checked for compatibility; anything that is not is disabled (C2 above). Module functions that can provide the requested quantities A2 and B1 are found, and other module functions are identified to fulfill their dependencies. Once all dependencies are resolved, the Core determines the correct function evaluation order, which is passed to ScannerBit. This module chooses parameter combinations to sample, running the module functions in the determined order for each parameter combination. The requested quantities are output by the printer system for each tested combination.

GAMBIT (the Global And Modular BSM Inference Tool) [93] is a global fitting software framework designed to address the needs listed above: flexibility, relatively easy extension to new observables and interfaces, public code availability, statistical secularism and computational speed.

3.3.1 The structure of GAMBIT

GAMBIT consists of a number of modules, referred to as “Bits”, in addition to a number of core components. Aside from one, the scanning module, all are physics-related.

ColliderBit calculates particle collider observables and likelihoods, including implementations of LEP, ATLAS and CMS searches for particle production, and measurements of the Higgs boson [237].

FlavBit calculates observables and likelihoods from flavour physics, in particular B, D and K meson decays as observed by LHCb, including angular observables and correlations [238].

DarkBit calculates observables and likelihoods related to dark matter, like the relic abundance, direct and indirect searches [239].

SpecBit provides interfaces external spectrum generators to calculate pole masses and running parameters. In addition, it can provide these to the rest of GAMBIT in a standardized spectrum container format [240].

DecayBit calculates decay rates of relevant particles in a BSM theory and also contains decay data for all SM particles [240].

PrecisionBit calculates model-dependent precision corrections to masses, couplings and other observables, as well as precision nuisance parameter likelihoods [240].

ScannerBit is the module that handles statistics and sampling. It provides the user with an interface to different sampling algorithms; one can then compute profile likelihoods, Bayesian posteriors, and other statistical quantities [94].

In this study, most observables and likelihoods do not fall into any of these existing modules, and a new physics module has been created: **NeutrinoBit**, which contains all calculations relevant to both left- and right-handed neutrino physics. All of the module functions used in this study are in this Bit, unless otherwise stated.

Physics modules are collections of module functions, each of which can compute a physical or mathematical quantity (an observable, likelihood or any intermediate quantity). Every function is tagged with a **capability**. Along with the specified datatype, this fully describes the quantity it calculates. Module functions are the basic building blocks of GAMBIT. The intermediate quantities, observables and likelihoods may be **dependencies**, quantities which are required by a function for its calculations. All calculations are either done within GAMBIT, or by external tools interfaced with GAMBIT, which are called backends.

Once a scan is set to run, the dependency resolver in GAMBIT searches for module functions to fulfill each dependency by matching the declared capabilities of module functions with the declared dependencies of other module functions. It also confirms that the datatypes do not have any conflict. An example is the likelihood that is calculated for one of the LHN mass splittings ($\Delta m_{3\ell}^2$), identified by the capability `md3l_1nL` and an identically named module function. To do this, it requires the calculated splitting and ordering:

```
#define CAPABILITY md3l_1nL
START_CAPABILITY
    #define FUNCTION md3l_1nL
    START_FUNCTION(double)
    DEPENDENCY(ordering, bool)
    DEPENDENCY(md31, double)
    DEPENDENCY(md32, double)
    #undef FUNCTION
#undef CAPABILITY
```

The dependencies are explicitly declared; the splitting is to be provided as a double, which the function accesses via pointers (`ordering` and `md31` or `md32`) that are in a namespace reserved for dependencies (so-called “Pipes”), for example `Pipes::md3l_1nL::Dep`.

Given the large amounts of computation time inherent in almost any scan, GAMBIT is constructed to be parallelized. This is done at two levels: at the scanner level via MPI [241] and at the level of module functions with OpenMP [242], facilitating the scaling to many thousands of cores, as most major external sampling packages use MPI.

3.3.2 Models

GAMBIT achieves the activation of appropriate module functions by utilizing a hierarchical model database. A model is defined as a set of free parameters and a series of relations to other models. It is thus possible to have a parent-child relation between models; the child model’s parameter space can be easily mapped to a subspace of that of the parent model with a translation function (such a relation exists for our RHN model, which will be described shortly). Additionally, models can also have translations defined to so-called “friend models” outside their family tree. Multiple models can be scanned over simultaneously, and their parameter values given to a module function

that needs them. In this way, SM parameters can be treated like nuisance parameters in a BSM scan. For a simple example of GAMBIT's activation process, see figure 3.1.

Anyway, once the requirements of all module functions are found to have been met, a directed graph of dependencies (with no internal closed loops) is formed – a directed acyclic graph, which has an implicit topology. GAMBIT uses this ordering to ensure that module functions which satisfy dependencies are computed prior to the functions that depend on them. It also optimizes the order beyond this by considering how long a function takes to be evaluated and how important the particular function is in excluding previous parameter choices. ScannerBit is then called to carry out the actual sampling; it uses whichever scanning algorithm is specified in the input file, and at the end of it all, passes the output to the printer system to be made available in the user-specified format.

GAMBIT ships with pre-defined models; new ones can also be defined easily. All the SM and active neutrino parameters are defined in a model called **StandardModel_mNudiff**, a daughter model of **StandardModel_SLHA2**. The latter includes SM parameters written in the SLHA2 convention [243]. The former contains the parameters `mNu_light` (m_{ν_0}), `dmNu21` (Δm_{21}^2) and `dmNu31` ($\Delta m_{3\ell}^2$), which give the lightest left-handed neutrino mass and mass splittings respectively. Other parameters in this model that are relevant for this study and are scanned over are `alpha1` (α_1), `alpha2` (α_2), `delta13` (δ), `theta12` (θ_{12}), `theta23` (θ_{23}) and `theta13` (θ_{13}). A translation function connects the parent and daughter models.

The daughter model **StandardModel_mNudiff** is defined as shown below.

```
#define MODEL StandardModel_mNudiff
#define PARENT StandardModel_SLHA2

START_MODEL
DEFINEPARS(alphaInv, GF, alphaS)
DEFINEPARS(mZ)
DEFINEPARS(mE, mMu, mTau)
DEFINEPARS(mNu_light, dmNu21, dmNu31)
DEFINEPARS(mD, mU)
DEFINEPARS(mS, mCmC)
DEFINEPARS(mBmB, mT)

// CKM parameters
DEFINEPARS(CKM_lambda, CKM_A, CKM_rhoBar, CKM_etaBar)
```

```

// PMNS parameters
DEFINEPARS(theta12, theta23, theta13)
DEFINEPARS(delta13, alpha1, alpha2)

INTERPRET_AS_PARENT_FUNCTION(StandardModel_mNudiff_to_StandardModel
_SLHA2)
#undef PARENT
#undef MODEL

```

The definition of the daughter model differs only in description of the LHN masses, and a line to specify the translation function, which contains

```

#define MODEL StandardModel_mNudiff
void MODEL_NAMESPACE::StandardModel_mNudiff_to_StandardModel_SLHA2
(const ModelParameters &myP, ModelParameters &targetP)
{
    logger()<<"Running interpret_as_parent calculations for
StandardModel_mNudiff --> StandardModel_SLHA2."<<LogTags::info<<EOM;
    if (myP["dmNu3l"] >0.) // normal hierarchy, l = 2
    {
        targetP.setValue("mNu1", myP["mNu_light"]*1e-9);
        targetP.setValue("mNu2",
            pow(myP["mNu_light"]*myP["mNu_light"]+myP["dmNu21"], 0.5)*1e-9);
        targetP.setValue("mNu3",
            pow(myP["mNu_light"]*myP["mNu_light"]+myP["dmNu3l"]+myP["dmNu21"],
                0.5)*1e-9);
    }
    else // inverted hierarchy, l = 1
    {
        targetP.setValue("mNu3", myP["mNu_light"]*1e-9);
        targetP.setValue("mNu1",
            pow(myP["mNu_light"]*myP["mNu_light"]-myP["dmNu3l"], 0.5)*1e-9);
        targetP.setValue("mNu2",
            pow(myP["mNu_light"]*myP["mNu_light"]-myP["dmNu3l"]+myP["dmNu21"],
                0.5)*1e-9);
    }
}
#undef MODEL

```

The differential model thus stores the left-handed neutrino masses in eV.

The right-handed neutrino sector is defined in a separate model, **RightHandedNeutrinos**, containing the RHN masses M_I and the real and imaginary parts of the ω_{ij} parameters in the C-I parameterization (section 2.3). To aid in

the exploration of the symmetry-protected region (section 2.4), a differential model inherited from this is also defined: **RightHandedNeutrinos_diff**.

```
#define MODEL RightHandedNeutrinos
START_MODEL
DEFINEPARS(M_1, M_2, M_3, ReOm23, ImOm23, ReOm13, ImOm13, ReOm12,
ImOm12, Rorder)
#undef MODEL
```

```
#define MODEL RightHandedNeutrinos_diff
#define PARENT RightHandedNeutrinos
START_MODEL
DEFINEPARS(M_1, delta_M21, M_3, ReOm23, ImOm23, ReOm13, ImOm13,
ReOm12, ImOm12, Rorder)
INTERPRET_AS_PARENT_FUNCTION(RightHandedNeutrinos_diff
_to_RightHandedNeutrinos)
#undef PARENT
#undef MODEL
```

```
#define MODEL RightHandedNeutrinos_diff
void MODEL_NAMESPACE::RightHandedNeutrinos_diff_to_RightHandedNeutrinos
(const ModelParameters &myP, ModelParameters &targetP)
{
    logger()<<"Running interpret_as_parent calculations for
    RightHandedNeutrinos_diff -->
    RightHandedNeutrinos."<<LogTags::info<<EOM;
    targetP.setValues(myP,false);
    targetP.setValue("M_2", myP["M_1"]+myP["delta_M21"]);
}
#undef MODEL
```

In the differential model, the parameter M_2 is replaced by δM_{21} , and a translation function takes the parameters of the daughter model to the parent model, analogous to what is done for the LHNs, i.e.

$$M_2 = M_1 + \delta M_{21}. \quad (3.21)$$

`Rorder` encodes the ordering of the matrices \mathcal{R}^{ij} in equation 2.84, which allows us to fully cover all the parameter space with the C-I parameterization.

The YAML initialization file

The master initialization file is written in the YAML format. YAML is a ‘human-friendly, cross-language, Unicode-based data serialization language’.

Its format is quite similar to that of Python. The top node of the master initialization file is a dictionary that contains eight entries.

`Parameters` contains the scan parameters for different models.

`Priors` describes the priors to be placed on the scan parameters.

`ObsLikes` describes observables and likelihoods that the user would like to be calculated in a scan.

`Rules` specifies additional rules to guide the resolution of dependencies and backend requirements.

`Printer` provides details about how and where to store the results of the scan.

`Scanner` provides information about the scanning algorithm to be utilized in a scan.

`Logger` presents options for logging messages during the scan.

`KeyValues` is an additional global option section.

Multiple models can be scanned simultaneously. In this context of sampling parameter values, a prior specifies how a parameter should be sampled before any data is brought to bear. A simple option would be to assign independent, flat distributions for each parameter: “flat priors”. When paired with a naive random scanner, this would lead to uniform sampling of the parameter values.

In it’s most basic form, the fast prior section can just specify a value to fix a parameter to during a scan (“delta-function prior”):

```
model_1:
  parameter_1: 5.0
  parameter_2: 25.0
```

or, equivalently

```
model_1:
  parameter_1:
    fixed_value: 5.0
  parameter_2:
    fixed_value: 25.0
```

This can be trivially extended to a set of points to be picked in a particular order.

```
model_1:
  parameter_1: [5.0, 25.0, 100.0]
  parameter_2:
    fixed_value: [75.0, 205.0, 350.0]
```

In case parameters are very simply related, whether they are from the same model or not, they can be combined using the `same_as` keyword.

```
model_1:
  parameter_1:
    same_as: model_2::parameter_2
    scale: scale
    shift: shift
```

`model_1::parameter_1` will then be set from the value assigned to `model_2::parameter_2` at each point in the scan. The keywords `scale` and `shift` can also be optionally specified; thus, in the above example, `model_1 ::parameter_1 = shift + (scale * model_2::parameter_2`.

Other fast priors can be chosen via the `prior_type` keyword, which can be set to flat, log (uniform in the log of the parameter value), or various trigonometric functions (cos, sin, tan or cot):

```
model_1:
  parameter_1:
    prior_type: chosen_prior
    range: [low, high]
  parameter_2:
    prior_type: log
    range: [5, 75]
```

The `scale` and `shift` parameters work with `prior_type`, in the same way as presented for `same_as`.

If no fixed value is specified for a parameter, and both `prior_type` and `same_as` are absent, but `range` is given, a flat prior is assumed. Additional custom priors can be written as plugins for ScannerBit.

The `ObsLikes` section determines what is calculated during a scan. Each entry typically corresponds to a likelihood contribution or an observable. Likelihood functions and observables are largely the same within GAMBIT; likelihoods are used to inform and drive a scan, observable are just recorded and may be used in the analysis of data. The minimal allowed entry has the form

```
ObsLikes:
- capability: example_capability
  purpose: example_purpose
- ...
```

In GAMBIT, by convention, all likelihoods are given in terms of the log-likelihood.

Supposing different module functions provide a single capability requested in the `ObsLikes` section, or several module functions provide the necessary capability-type pair required by a function through its dependencies, specifications in the `Rules` section required to fully define the scan and avoid any conflict.

3.3.3 Scanning algorithms

Many of the techniques used as algorithms for scanning the parameter space were, towards the beginning, Markov Chain Monte Carlos (MCMCs). In cosmology, for example, they have been frequently used because of their theoretical near-linear scalability with parameter dimensionality [244, 245, 246] and also proved popular in particle physics for the exploration of moderately complex supersymmetric model parameter spaces [209, 212, 220]. In a few years, nested sampling [247, 248] also came into use; the method is good at mapping posterior distributions and calculating the Bayesian evidence.

It was mentioned briefly in section 3.1.1 that approximate methods like the profile likelihood are often used instead of fully frequentist Neyman constructions. A pragmatic reason for this is that calculation of full likelihood functions that are usually involved are computationally expensive. Still, the previously mentioned MCMC and nested sampling methods may not be optimal [95] for profile likelihood calculations. Estimating the Bayesian posterior requires integrating the likelihood in directions of the parameter space (so, to some degree, it is an integration problem from the perspective of numerical analysis), whereas the profile likelihood needs the maximization of the likelihood in these directions (making it an optimization problem). Recently, genetic algorithms and differential evolution have been shown to be more efficient than these methods [95, 96] and a variant of the latter is used for this analysis; it is described a little later in section 3.3.3.

In GAMBIT, ScannerBit splits the sampling problem into three steps:

- Choosing n values in the interval $[0, 1]$. These values are a “point” in the n -dimensional “unit hypercube”.
- Transformation of the point in the unit hypercube into a point in the physical n -dimensional parameter space.
- Passing the values of the physical parameters to a user-specified function, which may calculate any quantity.

The steps repeat until some convergence criterion is satisfied, with the results of third step used to choose the next point in the unit hypercube in the first step.

Differential evolution

In this technique, successive generations of points in the parameter space evolve through a form of vector addition between members of the current population, like a random walk with a step size provided by the population. In its simplest form, DE has three controlling parameters.² These 3 parameters are connected to the three main steps that any DE variant goes through: mutation, crossover, and selection, controlled by the population size `NP`, the mutation scale factor `F`, and the crossover rate `Cr`. The simplest form of DE, “rand/1/bin” [249], is widely used. The first two parts of the name refer to the strategy for mutation, the third to crossover. The zeroth step is to pick a population of points (“target vectors”) $\{\mathbf{X}_i^g\}$ randomly from the allowed parameter space; i indexes members of the population, g the generation. Each subsequent generation of the population is chosen by performing mutation, crossover and selection on the previous generation. It has a computational advantage in being easy to parallelize, since each member of the population can be simultaneously and independently evaluated against a replacement candidate. The population-based mutation also leads to contour matching, which is especially useful to map likelihood contours. This also means that members of a population have a good chance of moving among local minima, which hastens convergence towards the global minimum.

Mutation will produce donor vectors \mathbf{D}_i from the current population of target vectors \mathbf{X}_i^0 . In the rand/1 mutation scheme, a random vector (hence the

²In variants that allow self-adaptation of parameters, this can be less.

“rand”) is combined with one difference vector scaled by the mutation scale factor F . Consider three random vectors \mathbf{X}_{r1} , \mathbf{X}_{r2} and \mathbf{X}_{r3} from the current population (none are the same, and none match the current target vector \mathbf{X}_i). The donor vector is then

$$\mathbf{V}_i = \mathbf{X}_{r1} + F(\mathbf{X}_{r2} - \mathbf{X}_{r3}). \quad (3.22)$$

Therefore, the size and shape of the population (which reflect the same properties of the likelihood contours) influence the scan. F controls how broad the search is; for convergence to be possible, it must be less than 1, but not too low, since then convergence can happen too quickly without sufficient exploration of the parameter space.

In the next step, *crossover* (or *recombination*), the donor vectors from the mutation step are combined with the original population of target vectors to produce trial vectors \mathbf{U}_i . The trial vectors are potential candidates for the next generation. The level of mixing between the donor and target vectors in the trial is controlled by the Cr , which takes a value between 0 and 1. In binomial crossover (the “bin” of rand/1/bin), the trial vector is chosen according to the following procedure:

- For the k th component of the trial vector \mathbf{U}_i , $\mathbf{U}_{i,k}$, a random number r_k is chosen between 0 and 1.
- If $r_k \leq Cr$, the component is taken from the donor vector: $\mathbf{U}_{i,k} = \mathbf{V}_{i,k}$.
- If, instead, $r_k \geq Cr$, it is taken from the target vector: $\mathbf{U}_{i,k} = \mathbf{X}_{i,k}$.
- Once all components of \mathbf{U}_i are chosen, one component is reassigned ensuring that trial vectors are always different from their parent target vectors. A dimension l is chosen at random for each member of the population. The corresponding component of the donor vector is assigned to the target vector: $\mathbf{U}_{i,l} = \mathbf{V}_{i,l}$.

High values of Cr therefore lead to increased exploration, as the trial vectors will differ from the target vectors along many dimensions. Low values of Cr are effective when the likelihood is a separable function of the parameters; if it is non-separable, Cr should be kept high to allow better exploration.

The final step, *selection*, gives the next generation of vectors. Here, the value of the objective function (typically the likelihood) for each target vector \mathbf{X}_i^g (the previous population) is compared with the trial vector \mathbf{U}_i that has been

constructed. The one with a higher likelihood moves on to the next generation and becomes \mathbf{X}_i^{g+1} . If both have the same likelihood, the trial vector \mathbf{U}_i is taken so that the population does not get stuck on flat surfaces.

Many variations of rand/1/bin exist, involving a different choice of base vector, or a different method to calculate the difference vector. One strategy that encompasses several possible options is [250]

$$\mathbf{V}_i = \lambda \mathbf{X}_{\text{best}} + (1 - \lambda) \mathbf{X}_1 + \sum_{q=0}^Q \mathbf{F}_q (\mathbf{X}_{2q} - \mathbf{X}_{3q}), \quad (3.23)$$

where \mathbf{X}_1 is the current vector and $\mathbf{X}_{2,3}$ are chosen randomly from the population. No vectors may be used twice. This form allows rand base vectors ($\mathbf{X}_1 = \mathbf{X}_{\text{rand}}$ and $\lambda = 0$), current base vectors ($\mathbf{X}_1 = \mathbf{X}_i$ and $\lambda = 0$), best base vectors ($\lambda = 1$), rand-to-best base vectors ($\mathbf{X}_1 = \mathbf{X}_{\text{rand}}$ and $0 < \lambda < 1$), and current-to-best base vectors ($\mathbf{X}_1 = \mathbf{X}_i$ and $0 < \lambda < 1$).

What is used here is based on self-adaptive differential evolution, and is called jDE [251]; it is based on rand/1/bin, but adapts \mathbf{F} and \mathbf{Cr} throughout a run.

Diver

Diver can perform λ jDE optimization (dynamically varies λ in equation 3.23), regular jDE, rand/1/bin and all strategies in between available. In either of the jDE option, the value of \mathbf{F} in each generation is controlled by a value τ_1 . So that evolution can take place, but also ensure convergence, \mathbf{F} can be between $\mathbf{F}_l = 0.1$ to $\mathbf{F}_u = 0.9$. The initial value of for each vector is drawn randomly from a uniform distribution in $[\mathbf{F}_l, \mathbf{F}_u]$. Prior to mutation, a random number is compared to τ_1 ; if it smaller, \mathbf{F} takes a random value in the allowed range, which is used for mutation. During selection, if the trial vector is accepted, the new value for is retained; if it is rejected, the previous value is kept. Similarly, \mathbf{Cr} is controlled by τ_2 . \mathbf{Cr} has no limitation, since no pathological behaviour occurs in either of the limits. So, in this case, for each member of the population, \mathbf{Cr} is set to a random value between 0 and 1. For each generation, before crossover, a trial value for is chosen, a uniform random number is compared to τ_2 and so on.

In the case of λ jDE, λ is also varied in this way. However, after some tests, we concluded that this does not provide an improvement on the sampling of

the parameter space (because the best-fit region is away from the boundaries, as we will see later).

Duplicate vectors can be troublesome for differential evolution and they can naturally arise in rand/ or best/ mutations. Premature convergence can result from this, but there is a worse issue: if a pair of duplicates is chosen to create the difference vector (in mutation), the donor vector will match the third vector chosen, possibly creating another duplicate. In best/ mutation, this can lead to an infection of the population by the best vector. Fortunately, it is unlikely to happen in current/ mutation (`current = true`), jDE (`jDE = true`), and λ jDE (`lambdajDE = true`).

Chapter 4

Observables

In this chapter, the observables that constrain the parameter space of right-handed neutrinos most strongly are listed and discussed.

Firstly, *neutrino oscillation data* (section 4.1) has determined the mass splittings and mixing angles to a good degree of precision and any points that are generated by the scan must satisfy them. Instead of fixing the left-handed neutrino parameters at their experimentally-determined best-fit values, we use the χ^2 distributions that the NuFit collaboration obtained for these quantities [102].

In the context of *direct detection* (section 4.3), kinematically accessible right-handed neutrinos can, of course, be produced in a well-designed experiment, or their decay products looked for. For example, PIENU [252] has looked at the process

$$\pi^+ \rightarrow e^+ \nu \tag{4.1}$$

and searched for missing momentum that would be carried by the neutrino.

On the other hand, RHNs can leave indirect signatures, leading to deviations from the Standard Model prediction of quantities. Since the presence of RHNs in the model modifies the weak currents, *electroweak precision observables* (section 4.2.1) such as the Weinberg angle and the Fermi constant are relevant; contributions from RHNs cannot lead to results counter to the precision measurements of these quantities. This has been observed in [80] and [253], for example. [80] and [253] also observed the relevance of *lepton universality* (section 4.2.3) constraints coming from both, fully leptonic and semileptonic decays, which we include.

The existence of *lepton flavour violation* (section 4.2.2) among neutrinos carries over to the charged leptons as well. Hence, flavour-violating decays like

$\mu \rightarrow e\gamma$ provide strong bounds that can restrict the parameter space. We include all the relevant processes, in addition to flavour conversion in nuclei.

The *unitarity of the CKM matrix* (section 4.2.4) and the precision with which the elements have been experimentally determined offers another avenue through which the RHN parameter space can be studied. The matrix elements have been determined from processes (kaon and tau decays) within the framework of the Standard Model; the modification to these processes coming from the RHNs is used to constrain LHN-RHN couplings.

Another well-studied process that we include is *neutrinoless double-beta decay* (section 4.2.5). The Majorana mass possessed by the neutrinos violates lepton number conservation, hence mediating this decay, which is forbidden in the SM.

Lastly, *Big Bang nucleosynthesis* (section 4.2.6) also forces the RHNs' properties to conform to certain limits; the restriction comes from the fact that they have to couple strongly enough to decay within a timeframe that leaves the (experimentally well-observed) elemental abundances of deuterium, helium-3, helium-4 and lithium-7 within observational limits.

To summarize, included are those direct detection experiments which currently exert the strongest constraints over the considered mass range; these are beam dump and peak search experiments which searched for RHNs in meson, tau and gauge boson decays. On the indirect side, the constraints implemented are those coming from the LHN mass differences and mixings, electroweak precision observables (EWPO), lepton flavor violating (LFV) processes, lepton universality, CKM unitarity, neutrinoless double-beta decay ($0\nu\beta\beta$) and Big Bang nucleosynthesis (BBN). As a central point of reference, a list of the constraints that are implemented in the current analysis and the type of likelihood used to encode each is presented in table 4.1.¹

4.1 Standard Model constraints

In the last two decades, solar, atmospheric, reactor and accelerator neutrino experiments have conclusively shown the existence of neutrino oscillations,

¹Aside from these constraints, two other “theoretical” limits are enforced: firstly, the unitarity of U in equation 2.61 implies the relation $(\delta m_\nu^{1\text{-loop}})_{\alpha\alpha} = \sum_i m_i (V_\nu)_{\alpha i}^2 + \sum_l M_l \Theta_{\alpha l}^2$. All points in the scan must satisfy this condition. The second condition enforced on all points is that the Yukawa couplings must be perturbative, i.e. $F^2 < 4\pi$ [182].

Observables	Likelihood	References
<i>Standard model constraints</i>		
Active neutrino likelihoods	Gaussian	[102, 103]
<i>Indirect constraints</i>		
EW precision observables	Gaussian	[254]
Lepton flavour violation	Half-Gaussian	[105, 255, 256, 257, 258, 259] [260, 261, 262, 263, 264]
Lepton universality	Gaussian	[105, 107, 109, 265, 266, 267, 268]
CKM unitarity	Gaussian	[269, 270, 271, 272, 273]
Neutrinoless double-beta decay	Half-Gaussian	[274, 275]
BBN	Step function	[56]
<i>Direct searches</i>		
PIENU	Half-Gaussian	[252]
PS-191 (e and μ)	Poissonian	[57, 68]
CHARM (e and μ)	Poissonian	[69]
DELPHI	Poissonian	[73]
ATLAS (e and μ)	Half-Gaussian	[276]
CMS (e and μ)	Half-Gaussian	[277]
E949	Half-Gaussian	[278]
NuTeV	Poissonian	[70]
CHARM (τ)	Half-Gaussian	[279]

TABLE 4.1: Likelihood and observable overview. The active neutrino observables are discussed in section 4.1, EWPO in section 4.2.1, LFV in section 4.2.2, lepton universality in section 4.2.3, CKM unitarity in section 4.2.4, neutrinoless double-beta decay in section 4.2.5 and BBN in section 4.2.6. The direct searches are outlined in section 4.3. The entries in the “Likelihood” column specify the type of function used to encode each observable in this study.

driven by non-zero neutrino masses and mixings. The “disappearance” of solar ν_e , reactor $\bar{\nu}_\mu$ and atmospheric ν_μ and $\bar{\nu}_\mu$ has been observed in multiple solar neutrino experiments [10, 11, 12], KamLAND [13] and Super-Kamiokande [14] respectively. The accelerator experiments K2K [280], T2K [281] and MINOS [282] have also observed ν_μ disappearance, while OPERA [18] and Super-K [283] have reported on $\nu_\mu \rightarrow \nu_\tau$ oscillations. All this work culminated in the 2015 Nobel Prize being awarded to Takaaki Kajita of the Super-Kamiokande collaboration and Arthur McDonald of the SNO collaboration for their contribution.

These efforts have led to the determination of the mass differences between the two pairs of neutrinos, Δm_{21}^2 and $\Delta m_{3\ell}^2$ ($\ell = 1$ for NH, $\ell = 2$ for IH), as

well as of the three mixing angles θ_{ij} and the CP-violating phase δ .

The simplest constraint that is enforced on the masses of the left-handed neutrinos in this study comes from PLANCK [284], in the form of their sum. As described in chapter 2, the transition of massive neutrinos from being initially relativistic in the Universe to becoming non-relativistic leaves detectable signatures on the cosmic microwave background [143] and the large scale structure [145]). Assuming the Lambda cold dark matter (Λ CDM) model and considering the power spectrum of the CMB, PLANCK has constrained the sum of the left-handed neutrino masses to be $\sum m_\nu \leq 0.23$ eV [284]. The limit has been updated with time and also changes depending on which cosmological datasets are included; for instance, by combining the information from CMB with that from baryon acoustic oscillations (BAO), the constraint becomes tighter: $\sum m_\nu < 0.12$ eV [146]. We use the former limit of $\sum m_\nu \leq 0.23$ in our analysis, enforced it as a hard cut: the parameter point is flagged as being invalid if $\sum m_\nu > 0.23$ eV. This slightly older value was included in the analysis before the updated limits were published; in the course of the work, we found that the constraint on the sum of active neutrino masses has a negligible effect on the RHN parameter space, and so chose to keep this value for the limit. The constraint's insignificant effect on the parameter space is also why it is modeled as a hard cut off and not a more sophisticated likelihood function.

In addition to the sum being bound from above, it is also bounded from below, since the mass differences are known: for NH, $\sum m_\nu > 0.06$ eV; for IH, $\sum m_\nu > 0.1$ eV [146]. In the future, if it is found that $\sum m_\nu \lesssim 0.1$ eV, the inverted hierarchy would be in trouble.

Additionally, the individual mass splittings Δm_{21}^2 and $\Delta m_{3\ell}^2$, the three mixing angles θ_{12} , θ_{13} , θ_{23} and the CP-violating phase δ_{CP} are encoded in likelihood form using the results from the NuFIT collaboration [102, 103], who performed a global analysis of solar, atmospheric, reactor and accelerator neutrino data in the framework of three neutrino oscillations using

- SAGE [10], the Homestake chlorine experiment [285], Gallex/GNO [286], SNO [287], the four phases of Super-Kamiokande [288, 289, 290] and two phases of Borexino [291, 292, 293]
- the atmospheric experiments IceCube/DeepCore [294]

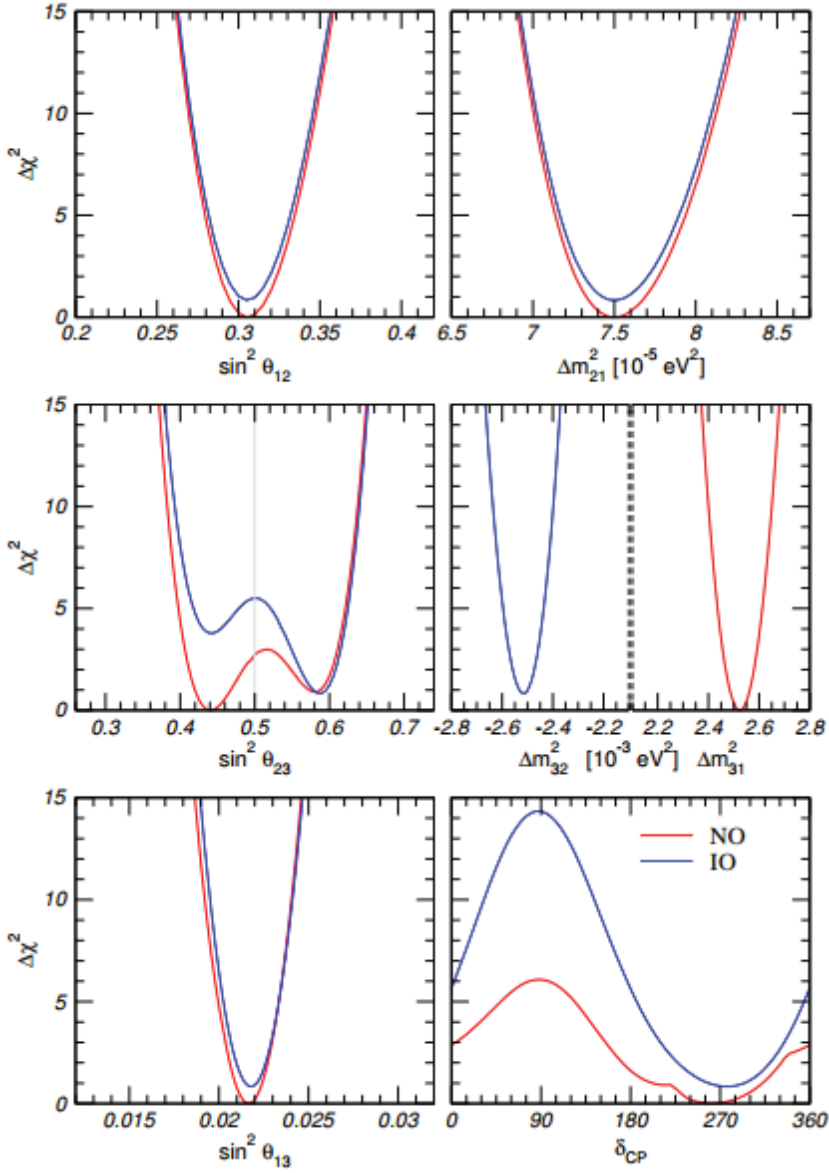


FIGURE 4.1: 1-dimensional χ^2 distributions determined by the NuFit collaboration, which are used in this study. Red curves are for NH, blue for IH; plots taken from [102].

- the reactor experiments KamLAND [295], Double-Chooz [296], Daya-Bay [297] and Reno [298]
- the accelerator experiments MINOS [299, 300], T2K [301] and NO ν A [302].

The collaboration provides one- and two-dimensional $\Delta\chi^2$ tables for these quantities for both orderings. Since we scan over the LHN parameters individually in our analysis, neglecting any correlations between them, we use the one-dimensional distributions. This should have a negligible effect on the results.

4.2 Indirect constraints

4.2.1 Electroweak precision observables

From the relevant terms in equation 2.67, one can see that the presence of RHNs modifies the weak current and the interaction strength of the LHNs. Now, the Fermi constant is usually determined using muon decay, specifically the process $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$, through the equation [303]:

$$G_F = \sqrt{\frac{192\pi^3}{\tau_\mu m_\mu^5}}; \quad (4.2)$$

radiative corrections not being included; here, τ_μ is the lifetime of the muon and m_μ is its mass. Due to the muon being lighter than the RHNs considered here, the latter can be integrated out in the analysis, leaving only the states defined in equation 2.62 in the end. The weak interaction strength of ν_L is thus suppressed, in turn suppressing the decay rate of the muon, leading to a Fermi constant that will be different from that determined *including* BSM physics. The relation between the constant determined from muon decay “ G_μ ” and the true G_F can be written as [59]

$$G_\mu^2 = G_F^2 [1 - (\theta\theta^\dagger)_{\mu\mu} - (\theta\theta^\dagger)_{ee}]. \quad (4.3)$$

It is caused by the non-unitarity of the flavour mixing matrix V_ν (see equation 2.65) which leads to the slight suppression of the muon decay.

Since the weak mixing angle θ_W and the mass of the W boson m_W depend on the Fermi constant at one loop, they get a correction given by [80]

$$\sin^2 2\theta_W = [\sin^2 2\theta_W]_{SM} \sqrt{1 - (\theta\theta^\dagger)_{\mu\mu} - (\theta\theta^\dagger)_{ee}}, \quad (4.4)$$

$$\frac{m_W^2}{[m_W^2]_{SM}} = \frac{[\sin^2 \theta_W]_{SM}}{\sin^2 \theta_W} \sqrt{1 - (\theta\theta^\dagger)_{\mu\mu} - (\theta\theta^\dagger)_{ee}}. \quad (4.5)$$

Experiments typically measure the effective Weinberg angle $\sin^2 \theta_{W,\text{eff}}$, so, for the SM prediction in the above equation, the (highly accurate) calculation from [104] which includes corrections up to two loops is used.

$$[\sin^2 \theta_{W,\text{eff}}]_{\text{SM}} = 0.23152 \pm 0.00010, \quad (4.6)$$

$$[m_W]_{\text{SM}} = 80.361 \pm 0.010 \text{ GeV}. \quad (4.7)$$

Other quantities are also affected by the presence of the heavy neutrinos: the invisible decay width of the Z boson, Γ_{inv} , and leptonic decays of the W. Under the assumption that radiative corrections factorize from the heavy neutrino contribution, at least up to order θ^2 [86, 304], one can write the invisible decay width of the Z as [253]

$$\Gamma_{\text{inv}} = \sum_{i,j} \Gamma_{Z \rightarrow \nu_i \nu_j} |_{\text{SM}} \left[|V_\nu^\dagger V_\nu|_{ij}^2 + |V_\nu^\dagger \Theta|_{ij}^2 \left(1 - \frac{m_{N_j}^2}{m_Z^2} \right)^2 \left(1 + \frac{1}{2} \frac{m_{N_j}^2}{m_Z^2} \right) \right]. \quad (4.8)$$

The contribution from $Z \rightarrow N_i N_j$ has been ignored due to it being $\mathcal{O}(\theta^4)$. For the process $Z \rightarrow \nu_i \nu_j$, the 2-loop calculation from [305] is used.

The RHN contribution to the decay widths of the W boson to leptons can be written in the form [80]

$$\Gamma_{W \rightarrow l_\alpha \bar{\nu}} = \frac{G_\mu m_W^3}{6\sqrt{2}\pi} \frac{((1 - \frac{1}{2}\theta\theta^\dagger)_{\alpha\alpha})(1 - x_\alpha)^2(1 + x_\alpha)}{\sqrt{1 - (\theta\theta^\dagger)_{\mu\mu} - (\theta\theta^\dagger)_{ee}}}, \quad (4.9)$$

where $x_\alpha \equiv m_{l_\alpha}^2 / m_W^2$.

Gaussian likelihoods are constructed for these observables using the experimental measurements x_0 and uncertainties σ displayed in table 4.2, in the form:

$$\ln \mathcal{L} = -\frac{1}{2} \log(2\pi\sigma^2) - \frac{1}{2} \frac{(x - x_0)^2}{\sigma^2}, \quad (4.10)$$

where x is the value of the relevant observable that comes out of our calculation.

4.2.2 Lepton flavour violation

A well-established fact about flavour-changing neutral processes, such as lepton flavour violation (LFV) is that they are suppressed in the Standard Model due to the Glashow-Iliopoulos-Maiani (GIM) mechanism [306] (in

Observable	Value
<i>Input parameters</i>	
G_μ [GeV ⁻²]	$1.1663787(6) \times 10^{-5}$
m_Z [GeV]	91.1875(21)
<i>Constraints</i>	
m_W [GeV]	80.361(10)
s_{eff}^2	0.23152 ± 0.00010
Γ_{inv} [MeV]	499.0 ± 1.6
$\Gamma_{W \rightarrow e \bar{\nu}_e}$ [MeV]	223 ± 6
$\Gamma_{W \rightarrow \mu \bar{\nu}_\mu}$ [MeV]	222 ± 5
$\Gamma_{W \rightarrow \tau \bar{\nu}_\tau}$ [MeV]	237 ± 6

TABLE 4.2: Electroweak precision observable measurements and uncertainties [254].

fact, lepton flavour is exactly conserved in the SM). Because of this, if a BSM theory involves contributions to such processes, they would dominate any contribution coming from the SM. Experimentally, the potential observation of such behaviour is firmly considered to be an indication of new physics, and experiments have attempted to measure LFV processes. This has been done with great precision [255, 258, 261, 262], unfortunately with no positive measurements; this does enable the enforcement of upper limits on the branching ratios involved.

Table 4.3 contains a list of the most significant of these observables, along with the experimental upper bound on their branching ratios and the experiment(s) that provided it.

Note that, in the table, the upper bounds for μ and τ decays are given as branching fractions with respect to the total decay width of the lepton involved. The total decay widths used are [254, 307]

$$\Gamma_\mu = (2.995984 \pm 0.000003) \times 10^{-19} \quad \text{and} \quad (4.11)$$

$$\Gamma_\tau = (2.2670 \pm 0.0039) \times 10^{-12}. \quad (4.12)$$

In the model considered here, with three RHNs, the leading contributions to these observables are from dipole and box diagrams. The expressions used for the calculation of the decay widths is in appendix A; for example, in the

Process	BR	Reference
<i>LFV decay</i>		
$\mu^- \rightarrow e^- \gamma$	4.2×10^{-13}	MEG [255]
$\tau^- \rightarrow e^- \gamma$	5.4×10^{-8}	BaBar [256], Belle [257]
$\tau^- \rightarrow \mu^- \gamma$	5.0×10^{-8}	BaBar [256], Belle [257]
$\mu^- \rightarrow e^- e^- e^+$	1.0×10^{-12}	SINDRUM [260]
$\tau^- \rightarrow e^- e^- e^+$	1.4×10^{-8}	Belle [261], BaBar [262]
$\tau^- \rightarrow \mu^- \mu^- \mu^+$	1.2×10^{-8}	ATLAS [258], LHCb [259], Belle [261], BaBar [262]
$\tau^- \rightarrow \mu^- e^- e^+$	1.1×10^{-8}	Belle [261], BaBar [262]
$\tau^- \rightarrow e^- e^- \mu^+$	0.84×10^{-8}	Belle [261], BaBar [262]
$\tau^- \rightarrow e^- \mu^- \mu^+$	1.6×10^{-8}	Belle [261], BaBar [262]
$\tau^- \rightarrow \mu^- \mu^- e^+$	0.98×10^{-8}	Belle [261], BaBar [262]
<i>LFV conversion</i>		
$\mu - e$ (Ti)	1.7×10^{-12}	SINDRUM II [263]
$\mu - e$ (Pb)	4.6×10^{-11}	SINDRUM II [264]

TABLE 4.3: Experimental upper bounds on LFV processes, along with the experiments that provided the bound. When more than one experiment is cited, the HFLAV average is used [105]. All upper bounds are given at 90% C.L.

case of $l \rightarrow l \gamma$ processes, the decay width is given by

$$\Gamma_{l_\alpha^- \rightarrow l_\beta^- \gamma} = \frac{\alpha_{\text{em}} m_{l_\alpha}^5}{4} \left(|K_2^L|^2 + |K_2^R|^2 \right), \quad (4.13)$$

where

$$K_2^L = \frac{G_F}{4\sqrt{2}\pi^2} \frac{m_{l_\beta}}{m_{l_\alpha}} \sum_{a=1}^6 \Theta_{\alpha a} \Theta_{\beta a}^* G \left(\frac{m_{\nu_a}^2}{m_W^2} \right) \quad \text{and} \quad (4.14)$$

$$K_2^R = \frac{G_F}{4\sqrt{2}\pi^2} \sum_{a=1}^6 \Theta_{\alpha a} \Theta_{\beta a}^* G \left(\frac{m_{\nu_a}^2}{m_W^2} \right). \quad (4.15)$$

In the above equations, α_{em} is the fine structure constant, $m_{l_{\alpha/\beta}}$ is the mass of the relevant lepton, G_F is the Fermi constant and the loop function $G(x)$ is defined by [308]:

$$G(x) = \frac{-7 + 33x - 57x^2 + 31x^3 + 6x^2(1 - 3x) \log(x)}{12(-1 + x)^4}. \quad (4.16)$$

Aside from decays, LFV processes can also give rise to a neutrinoless $\mu - e$ conversion inside a nucleus. Now, a muon cannot “convert” into an electron

in vacuum, simply by considering energy and momentum conservation (in the reference frame where the muon is at rest, the electron also has to be at rest from momentum conservation. Since the rest energy is simply the mass, energy will not be conserved.). In the presence of a nucleus that can account for this extra energy, though, this process can happen. Muons captured by a nucleus typically decay in orbit, resulting in a continuous spectrum of energy for the electron in the final state. However, in coherent flavour violating conversion, by preparing “muonic” atoms in a particular state, the final state electrons will have a discrete energy spectrum. Experiments then attempt to measure the rate at which this conversion happens with respect to the rate at which the nucleus captures electrons:

$$R_{\mu-e} = \frac{\Gamma_{\text{conv}}}{\Gamma_{\text{capt}}}. \quad (4.17)$$

The expressions for the conversion ratio and the nuclear parameters for the two nuclei studied (^{48}Ti and ^{208}Pb) are also in appendix A.

The likelihoods for all the observables mentioned are modeled as Gaussian upper limits. They are computed as

$$\ln \mathcal{L} = \begin{cases} -\frac{1}{2} \log(2\pi\sigma^2), & x > x_0 \\ -\frac{1}{2} \log(2\pi\sigma^2) - \frac{1}{2} \frac{(x-x_0)^2}{\sigma^2}, & x < x_0, \end{cases} \quad (4.18)$$

using the experimental data from table 4.3. The measured value x_0 is assumed to be zero in case an experiment does not provide a value. The standard deviation is set as $\sigma = v/1.28$ (a 90% CL upper limit corresponds to exclusion beyond 1.28σ from the mean [309]), v being the quoted upper limit.

4.2.3 Lepton universality

Lepton universality, or the paradigm in which SM interactions treat all three charged leptons identically, has been quite extensively tested. Experiments performed in the past at LEP [310] and SLC [311], for example, used lifetime measurements of the tauon and muon, as well as the partial decay widths of the Z boson and confirmed this, but more recent measurements from meson decays [107, 108, 109] have brought this into question, leading to the formulation of many BSM theories attempting to explain this deviation [253, 312].

The fact that RHNs modify the leptonic currents means they will contribute to these tests, most importantly, in the case of leptonic decays of charged mesons, $X^+ \rightarrow l^+ \nu$, and semileptonic decays of B mesons $B^{0/\pm} \rightarrow X^{0/\pm} l^+ l^-$. When calculating the decay widths of pseudoscalar mesons, the trouble from hadronic uncertainties must be dealt with. To do this, tests of universality are most often formulated as ratios between different species, and the uncertainties are then assumed to cancel out. For the cases of leptonic and semileptonic decays of mesons, these ratios are expressed as

$$R_{\alpha\beta}^X = \frac{\Gamma(X^+ \rightarrow l_\alpha^+ \nu_\alpha)}{\Gamma(X^+ \rightarrow l_\beta^+ \nu_\beta)} \quad \text{and} \quad (4.19)$$

$$R_X = \frac{\Gamma(B^{0/\pm} \rightarrow X^{0/\pm} l_\alpha^+ l_\alpha^-)}{\Gamma(B^{0/\pm} \rightarrow X^{0/\pm} l_\beta^+ l_\beta^-)} \quad (4.20)$$

respectively.

Fully leptonic decays

In this case, the test of lepton universality can be neatly expressed in terms of deviations from the SM prediction:

$$R_{\alpha\beta}^X = R_{\alpha\beta,SM}^X (1 + \Delta r_{\alpha\beta}^X), \quad (4.21)$$

where the sterile neutrino contribution coming from $\Delta r_{\alpha\beta}^X$ can be calculated using the active-sterile mixing matrix Θ as [59, 313]

$$\Delta r_{\alpha\beta}^X = \frac{1 + \sum_I |\Theta_{\alpha I}|^2 [G_{\alpha I} - 1]}{1 + \sum_I |\Theta_{\beta I}|^2 [G_{\beta I} - 1]} - 1, \quad (4.22)$$

where

$$G_{\alpha I} = \vartheta(m_X - m_{l_\alpha} - M_I) \frac{r_\alpha + r_I + (r_\alpha - r_I)^2}{r_\alpha (1 - r_\alpha)^2} \sqrt{1 - 2(r_\alpha + r_I) + (r_\alpha - r_I)^2}, \quad (4.23)$$

with ϑ being the Heaviside step function, $r_\alpha \equiv m_{l_\alpha}^2 / m_X^2$ and $r_I \equiv M_I^2 / m_X^2$. The SM predictions used in equation 4.21 for the tests of lepton universality for pions and kaons are $R_{e\mu,SM}^\pi = 1.235 \times 10^{-4}$ and $R_{e\mu,SM}^K = 2.488 \times 10^{-5}$ respectively [314].

Semileptonic decays

The contribution from right-handed neutrinos to semileptonic decays of B mesons is less significant than to leptonic decays. In fact, the effect on the decay of B to charmed mesons, $B^\pm \rightarrow D l \nu$, is negligible, as argued in [253]. Semileptonic decays to K mesons show a larger effect, particularly $B^+ \rightarrow K^+ l^+ l^-$ and $B^0 \rightarrow K^{*0} l^+ l^-$. Assuming that $m_l \ll m_{K^{(*)}}$ and that the Wilson coefficient $C_7 \ll C_9, C_{10}$, the ratios R_K and R_{K^*} can be approximated by [315]

$$\begin{aligned} R_{K^{(*)}} &= \frac{\Gamma(B^{\pm/0} \rightarrow K^{\pm/0} \mu^+ \mu^-)}{\Gamma(B^{\pm/0} \rightarrow K^{\pm/0} e^+ e^-)} \\ &\approx \frac{|C_{10}^{SM} + \Delta C_{10}^\mu|^2 + |C_9^{SM} + \Delta C_9^\mu|^2}{|C_{10}^{SM} + \Delta C_{10}^e|^2 + |C_9^{SM} + \Delta C_9^e|^2}. \end{aligned} \quad (4.24)$$

The BSM contribution to the Wilson coefficients ΔC_9^α & ΔC_{10}^α can be expressed as [316]

$$\Delta C_9^\alpha = -\Delta C_{10}^\alpha = -\frac{1}{4 \sin^2 \theta_W^2} \sum_I |\Theta_{\alpha I}|^2 E(x_t, x_I), \quad (4.25)$$

with $x_t = m_t^2/m_W^2$, $x_I = M_I^2/m_W^2$ and the loop function

$$\begin{aligned} E(x, y) = xy &\left\{ -\frac{3}{4} \frac{1}{(1-x)(1-y)} \right. \\ &+ \left[\frac{1}{4} - \frac{3}{2(x-1)} - \frac{3}{4(x-1)^2} \right] \frac{\log x}{x-y} \\ &\left. + \left[\frac{1}{4} - \frac{3}{2(y-1)} - \frac{3}{4(y-1)^2} \right] \frac{\log y}{y-x} \right\}. \end{aligned} \quad (4.26)$$

Next-to-next-to-leading order (NNLO) calculations for the SM contribution to the Wilson coefficients C_9 and C_{10} from [317, 318] are used in equation 4.24. The values are $C_9^{SM} = 4.211$ and $C_{10}^{SM} = -4.103$.

Apart from meson decays, other common tests of lepton universality are the decays of the W boson to leptons and τ decays.

- The ratio of the decay widths of the W to charged leptons l_α and l_β can be written as [80]

$$R_{\alpha\beta}^W = \frac{\Gamma(W^+ \rightarrow l_\alpha^+ \nu_\alpha)}{\Gamma(W^+ \rightarrow l_\beta^+ \nu_\beta)} = \sqrt{\frac{1 - (\theta\theta^+)_{\alpha\alpha}}{1 - (\theta\theta^+)_{\beta\beta}}}. \quad (4.27)$$

- The test for τ decays follows the same form as in equation 4.22; the SM prediction is $R_{\mu e, SM}^\tau = 0.973$ [319]. In this particular case, we only calculate the correction if $M_I > m_\tau$; if not, no contribution from BSM physics is assumed. This is a pragmatic decision; it is possible to calculate the contribution when $M_I < m_\tau$ too [253], but as is demonstrated in the referred paper, the resulting deviation from the SM is very small and will have a negligible influence on any results. Further, this calculation involves numerical integration, which would also increase computation time.

All the tests mentioned are implemented as Gaussian likelihoods centered on the experimentally measured value. The experimental measurements of these observables as well as the corresponding uncertainties² are shown in table 4.4. There, two measurements are shown for $R_{K^*}^B$ corresponding to two regions of the dilepton invariant mass: $0.045 < q^2 < 1.1 \text{ GeV}^2/c^4$ for (1) and $1.1 < q^2 < 6.0 \text{ GeV}^2/c^4$ for (2). The form of the likelihood is identical to equation 4.10.

Observable	Measured value
$R_{e\mu}^\pi$	$(1.235 \pm 0.005) \times 10^{-4}$ [265]
$R_{e\mu}^K$	$(2.488 \pm 0.010) \times 10^{-5}$ [266]
$R_{\mu e}^\tau$	0.9762 ± 0.0028 [105]
$R_{\mu e}^W$	0.980 ± 0.018 [267]
$R_{\tau e}^W$	1.063 ± 0.027 [268]
$R_{\tau\mu}^W$	1.070 ± 0.026 [268]
R_K^B	0.745 ± 0.089 [109]
$R_{K^*}^B$ (1)	0.66 ± 0.09 [107]
$R_{K^*}^B$ (2)	0.69 ± 0.10 [107]

TABLE 4.4: Experimental measurements for all tests of lepton universality.

4.2.4 CKM unitarity

The determination of CKM matrix elements $(V_{CKM}^{\text{exp}})_{ab}^i$ is usually done under the implicit assumption of $\Theta = 0$. Within the type-I seesaw model, therefore, these measurements have to be adjusted to take into account the effect of the RHNs. In this analysis, the smallest element of the CKM matrix, $(V_{CKM})_{ub}$,

²The experimental uncertainties for $R_{K^{(*)}}^B$ are obtained as the sum in quadrature of the statistical and systematic uncertainties provided by [107] and [109].

is neglected, since its absolute value $|(V_{\text{CKM}})_{ub}|^2 \sim 10^{-5}$ is much smaller than our sensitivity to Θ . Under the assumption of the unitarity of the CKM matrix, one then has the relation

$$|(V_{\text{CKM}})_{ud}|^2 + |(V_{\text{CKM}})_{us}|^2 = 1. \quad (4.28)$$

Various experimental measurements of these two matrix elements [269, 270, 271, 272] can thus be utilized to simultaneously constrain the true value of $|(V_{\text{CKM}})_{us}|$ and the active-sterile mixing matrix Θ . Note that because of equation 4.28, only one of the two matrix elements is independent.

Unitarity constraints are most relevant for $q\bar{q} \rightarrow W^* \rightarrow lN$; leptonic electroweak decays are the most sensitive, given that if one considered $q\bar{q} \rightarrow W^* \rightarrow q\bar{q}$, terms of the form $|V_{qq}| * |V_{qq}|$ would be involved. Instead of focusing on constraints coming from V_{us} (essentially kaon decays), pion decays, where one has the $u\bar{d}$ vertex are utilized in this study. They are more relevant since V_{ud} is larger and better measured, i.e. with a smaller error, than V_{us} .

Following [59, 80], the experimental measurements and true values of the CKM matrix elements can be related by

$$|(V_{\text{CKM}}^{\text{exp}})_{us,ud}^i|^2 = |(V_{\text{CKM}})_{us,ud}|^2 [1 + f^i(\Theta)], \quad (4.29)$$

where the functions f^i encode the contribution from RHNs to the process considered in each experiment. $|(V_{\text{CKM}}^{\text{exp}})_{us}|$ has been extracted using various decays, and the functions are given by [80]

$$K_L \rightarrow \pi^+ e^- \bar{\nu}_e : 1 + f^1(\Theta) = \frac{G_F^2}{G_\mu^2} [1 - (\theta\theta^\dagger)_{ee}], \quad (4.30)$$

$$K_S \rightarrow \pi^+ e^- \bar{\nu}_e : f^2(\Theta) = f^1(\Theta), \quad (4.31)$$

$$K^- \rightarrow \pi^0 e^- \bar{\nu}_e : f^3(\Theta) = f^1(\Theta), \quad (4.32)$$

$$K_L \rightarrow \pi^+ \mu^- \bar{\nu}_\mu : 1 + f^4(\Theta) = \frac{G_F^2}{G_\mu^2} [1 - (\theta\theta^\dagger)_{\mu\mu}], \quad (4.33)$$

$$K^- \rightarrow \pi^0 \mu^- \bar{\nu}_\mu : f^5(\Theta) = f^4(\Theta), \quad (4.34)$$

$$\frac{\tau^- \rightarrow K^- \nu_\tau}{\tau^- \rightarrow \pi^- \nu_\tau} : 1 + f^6(\Theta) = 1 + (\theta\theta^\dagger)_{\mu\mu} \quad \text{and} \quad (4.35)$$

$$\tau^- \rightarrow \pi^- \bar{\nu}_\tau : 1 + f^7(\Theta) = 1 + (\theta\theta^\dagger)_{ee} + (\theta\theta^\dagger)_{\mu\mu} - (\theta\theta^\dagger)_{\tau\tau}. \quad (4.36)$$

For $|(V_{\text{CKM}}^{\text{exp}})_{ud}|$, the uncertainty is dominated by measurements of superallowed $0^+ \rightarrow 0^+$ nuclear beta transitions, and it is only this which needs to

be modified; the modification is identical to that for the $K \rightarrow \pi e \nu$ processes above [80]:

$$0^+ \rightarrow 0^+ : 1 + f^1(\Theta) = \frac{G_F^2}{G_\mu^2} [1 - (\theta\theta^\dagger)_{ee}]. \quad (4.37)$$

We then treat V_{us} as a nuisance parameter and optimize its value to remove its Θ -dependence. The value $(V_{\text{CKM}})_{ud}$ is calculated with equation 4.28 for a given value of $(V_{\text{CKM}})_{us}$. The values of the matrix elements that come from these various processes and that of $|(V_{\text{CKM}}^{\text{exp}})_{ud}|$ are listed in table 4.5. The constraint is modeled as a χ^2 distribution in the current analysis:

$$\chi^2 = \sum_{i=1}^7 \frac{[(V_{\text{CKM}}^{\text{exp}})_i - (V_{\text{CKM}})_{us} \cdot (1 + f^i(\Theta))]^2}{\sigma_i^2} + \frac{[(V_{\text{CKM}}^{\text{exp}})_{ud} - (V_{\text{CKM}})_{ud} \cdot (1 + f^1(\Theta))]^2}{\sigma^2}. \quad (4.38)$$

Parameter	Process	Value	Reference
$(V_{\text{CKM}}^{\text{exp}})_{us} f_+(0)$	$K_L \rightarrow \pi e \nu$	0.2163(6)	[269, 273]
	$K_L \rightarrow \pi \mu \nu$	0.2166(6)	
	$K_S \rightarrow \pi e \nu$	0.2155(13)	
	$K^\pm \rightarrow \pi^0 e \nu$	0.2160(11)	
	$K^\pm \rightarrow \pi^0 \mu \nu$	0.2158(14)	
$(V_{\text{CKM}}^{\text{exp}})_{us}$	$\frac{\text{BR}(\tau \rightarrow K \nu)}{\text{BR}(\tau \rightarrow \pi \nu)}$	0.2262(13)	[270, 271]
	$\tau \rightarrow K \nu$	0.2214(22)	
	$\tau \rightarrow l, \tau \rightarrow s$	0.2173(22)	
$(V_{\text{CKM}}^{\text{exp}})_{ud}$	Average	0.97417(21)	[272]

TABLE 4.5: Experimental values of $(V_{\text{CKM}})_{us}$ and the average value of $(V_{\text{CKM}})_{ud}$ used in the calculation of the CKM likelihood. The value of the factor $f_+(0)$ is taken from [273].

4.2.5 Neutrinoless double-beta decay

During double-beta decay, two neutrons are transformed into two protons; two electrons and two anti-neutrinos are released due to lepton number conservation. Majorana neutrino masses violate lepton number and can hence induce neutrinoless double-beta decay. Here, the experimental constraints

are enforced in terms of $m_{\beta\beta}$, specifically in the form [320]

$$m_{\beta\beta} = \left| \sum_i (U_{\nu})_{ei}^2 m_i + \sum_I \Theta_{eI}^2 M_I f_A(M_I) \right|. \quad (4.39)$$

The first term gives the contribution from LHNs, the second, from RHNs; the term $f_A(M_I)$ parameterizes the suppression arising out of the fact that the RHNs are virtual if M_I is larger than the typical momentum exchange in $0\nu\beta\beta$ (about 100 MeV), as is the case here. If contributions come solely from the light neutrinos, one cannot constrain any neutrino properties using neutrinoless double-beta decay, since the limit on $\sum m_\nu$ precludes high values of $m_{\beta\beta}$, i.e. values of $m_{\beta\beta}$ that are consistent with the cosmological limit already lie below experimental bounds on the former. Additionally, the suppression of the RHN contribution due to f_A has, in the past, been thought to render RHN influence negligible. However, recent studies [59, 165, 320, 321] have shown that this suppression is not enough at higher masses (\sim GeV) and, in fact, right-handed neutrinos can dominantly contribute to the decay.

Specifically, accidental cancellations in the active neutrino mass matrix m_ν can lead to small LHN masses while still allowing large mixings. If these cancellations are truly by accident, or if they protected by a symmetry like approximate lepton number (B-L) conservation with the B-L breaking parameters being relatively large³, then $m_{\beta\beta}$ can exceed the standard prediction. As observed in [320] this can occur for mass splittings $\Delta M/M$ of $\mathcal{O}(10^{-3})$ and $|\text{Im}(\omega)| > 2$. These values are covered in the ranges of our scan, which leads us to expect that the limits on U_{eI}^2 in [139], for example, will be stronger than what we will obtain.

A central issue with respect to this observable has to do with the nuclear matrix elements (NMEs) (which are implicitly part of f_A); the amplitude for neutrinoless double-beta decay is proportional to the NMEs and these quantities suffer from fairly large uncertainties. This is partly due to the fact that they can only be inferred indirectly. This uncertainty in the elements that determine the exchanged momentum naturally leads to uncertainty in the function f_A , and we use the approximation given in [320], i.e.

$$f_A(M) \approx \frac{p^2}{p^2 + M^2}. \quad (4.40)$$

³Note that in the symmetric limit, the rate for neutrinoless double-beta decay is always small.

The experimental constraints from GERDA [274] and KamLAND-Zen [275] are implemented in this analysis. In the above equation, for p^2 , we adopt the values from [139], which are listed there for multiple elements. To be specific, the “Argonne” model is used along with the lower of the two values cited for p^2 : $\sqrt{\langle p^2 \rangle} = 178$ MeV for xenon (used in KamLAND-Zen), and $\sqrt{\langle p^2 \rangle} = 159$ MeV for germanium (used in GERDA); this yields the most conservative constraints. Since a frequentist interpretation of the results is focused on here, this approach is equivalent to profiling over systematic uncertainties, assuming a flat prior that spans the entire range of $\langle p^2 \rangle$ values in [139]. The constraints on $m_{\beta\beta}$ from GERDA - $m_{\beta\beta} < 0.15 - 0.33$ eV (90% CL) - and KamLAND-Zen - $m_{\beta\beta} < 0.061 - 0.165$ eV (also 90% CL) - are encoded as one-sided Gaussian likelihoods, choosing the higher of the two values to stay conservative. The form of the log-likelihood is the same as in equation 4.18.

4.2.6 Big Bang Nucleosynthesis

BBN is a comprehensive phenomenon, in the sense that three of the four known fundamental interactions play different, but important roles in it. First, consider the electromagnetic force, which all charged particles and photons are subject to. For processes involving these particles and electromagnetic interactions, the timescale for reaching equilibrium is much less compared to the expansion time. These particles thus reach thermal equilibrium quickly and remain equilibrated at some (common) temperature T .

The weak interactions lead to a sharp change in the equilibrium behaviour: at $T \gtrsim \text{MeV}$, the neutral particles it affects (neutrons, neutrinos) stay in equilibrium, but as the expansion of spacetime reduces the temperature, this balance is lost, and the particles are said to “freeze out”.

At high temperatures, processes like $n + \nu_e \rightarrow p + e^-$ are in chemical equilibrium, i.e. the forward and reverse processes occur at the same rate. Chemical and thermal equilibria are interconnected, however, so they are lost simultaneously, when the neutron-to-proton ratio freezes out at lower temperatures. The strong interactions are also responsible for the production of nuclei consisting of more than one nucleon.

With the addition of right-handed neutrinos to the SM’s particle content, the predictions of BBN also change, making it a valuable cosmological testing

ground for BSM models. The RHNs can decay, but can have relatively very long lifetimes (on the order of a second, even) because their coupling strength is so weak. Their decay products are then injected into the primordial environment, leading to an increase in temperature and a shift in the chemical equilibrium.

If RHNs decay shortly before or during BBN, the typical energy of the decay products ($\sim M_I \geq 50$ MeV) will be significantly higher than the plasma temperature at that time (~ 100 keV). Therefore, either by dissociating formed nuclei or by causing deviations from thermal equilibrium, they will affect the abundances of primordial elements, which are observationally well constrained. The requirement that RHN decay happens sufficiently early enough implies an upper limit on the lifetime τ_I of RHNs, or equivalently, a *lower* bound on the mixing U_I^2 [322] (in order to decay within some time limit, the RHNs must have a corresponding minimum mixing strength). However, in the presence of multiple RHN species, BBN cannot, in general, constrain individual mixing angles $U_{\alpha I}^2$, only the total mixing U_I^2 .

Here, decay channels of RHNs up until the mass range of B mesons is considered; beyond this mass, the constraint of BBN is considerably weaker. The decays are taken from [6] and [323], as well as the values of the necessary decay constants. The decay width for each topology is listed in appendix B. Once the widths are calculated, a rather simple constraint on the lifetime of each RHN is enforced: that it must be less than 0.1 s [56]. This is implemented as a hard cutoff in the parameter space, i.e. if the lifetime is larger than 0.1 s, the parameter point is given a large negative log-likelihood ($\ln \mathcal{L} = -100$), effectively rendering it invalid. We choose to do this and not flag such points to be invalid so that the effect of BBN could be examined during analysis. In principle, this limit can be weakened if the lightest active neutrino has a mass $< \mathcal{O}(10^{-3})$ eV, since the RHNs do not necessarily thermalize in this case [58], but since BBN is already a weak constraint, this is not incorporated here.⁴

⁴Previous analyses such as [6] did not take into account the color factor of 3 that should enter the decay width of RHNs to jets; [6, 323] also use charged current co-efficients for the decay widths, whereas the relevant neutral current factor would have to be included for higher masses. These additions would result in improvements; however, they are not incorporated here as well since their effect would be observable only at masses greater than about 5 GeV, where the BBN bound becomes irrelevant in our study.

4.3 Direct searches for right-handed neutrinos

Beam dump and peak searches are two of the most frequently used methods that have been employed to detect RHNs. In beam dump experiments, the large background signal that is usually present near the target hinders the detection of charged particles that are produced along with the RHNs (and which are usually used for identifying positive signals). On the other hand, RHNs with mass below the D meson scale can be long-lived enough to travel macroscopic distances. Looking for their charged decay products some distance away from the target leads to (almost) background-free experimental situations.

Furthermore, RHN production can generate a peak in the lepton energy spectrum. If, for example, a meson of mass m_X decays into an RHN of mass M_I and an electron/muon with mass m_{l_α} , this peak will be approximately at

$$E_{\text{peak}} \simeq \frac{m_X^2 + m_{l_\alpha}^2 - M_I^2}{2m_X}. \quad (4.41)$$

Even in situations where backgrounds are sizeable, a peak search can hence be used to impose constraints on the mixing.

To implement the direct detection constraints as likelihoods, we follow two different approaches depending on the information that is provided in each study. In one case, the experiments found no signal events and had no background counts after cuts (DELPHI, CHARM, PS191 and NuTeV). Therefore, since the processes in the experiments are essentially Poissonian, we construct the likelihood (to observe n events) as a Poisson distribution. The number of expected counts, μ , is a function of the RHN masses and mixings, i.e. $\mu = \mu(M_I, U_{\alpha I}^4)$ (assuming the experiment does so as well, the fourth power takes both production and decay of RHNs into account). For an μ expected and b background events, the likelihood is

$$\mathcal{L}(n|\mu) = (\mu + b)^n \frac{e^{-(\mu+b)}}{n!}. \quad (4.42)$$

With no reported detection ($n = 0$) and background cuts reducing b to approximately zero,

$$\ln \mathcal{L}(n = 0|\mu) = -\mu. \quad (4.43)$$

To connect μ with our model parameters, we use the fact that the expected

signal counts are proportional to the LHN-RHN mixing, $\mu \propto U_{\alpha I}^4$. The factor of proportionality is set to reproduce the results from the experimental papers; assuming that these limits are based on the Feldman-Cousins procedure [324], a 95% CL upper limit, for example, would correspond to an expected number of signal counts of $\mu = 3.09$.

On the other hand, for the experiments which either quote non-zero signal events and/or backgrounds, or if this information is ambiguous (CHARM ν_τ re-interpretation, PIENU, ATLAS and E949), we model the constraint likelihood as Gaussian upper limits, i.e. we model them as half-Gaussians with zero mean and error set according to the confidence level at which the results are presented. For example, in the case of an experiment that presents limits at 90% CL, for a half Gaussian, this lies within 1.28σ of the mean.

Lastly, a correction (division of the experimental limits by $\sqrt{2}$) is sometimes necessary if the neutrinos were assumed to be Dirac particles in the analysis of an experiment.

Note that simplifications in the model are often made in collider experiments in order to compute confidence intervals. Given that we utilize these results to construct the likelihoods in this study, any simplifications made are also incorporated here. In practice, this means that the actual bounds will likely be slightly weaker than those presented here or in those papers. There is one way to avoid this: perform a full collider simulation of the signal events for each process and construct likelihoods based on the number of observed events, rather than on the experiments' interpretation of the results based on some model (e.g. the ν MSM in the CMS paper). However, this is costly to do since, firstly, the machinery for such searches is not yet available in GAMBIT, and secondly, this means that, for each parameter point, a full simulation with large ($>100,000$) number of events would have to be performed, which takes a lot of CPU time and would make our scans too time-consuming. We use the benefit conferred by the fact that the simplified model approach is a good approximation in most cases; we, as earlier studies did, take the simplified model limits as the true limits. The difference in our results due to this would be insignificant.

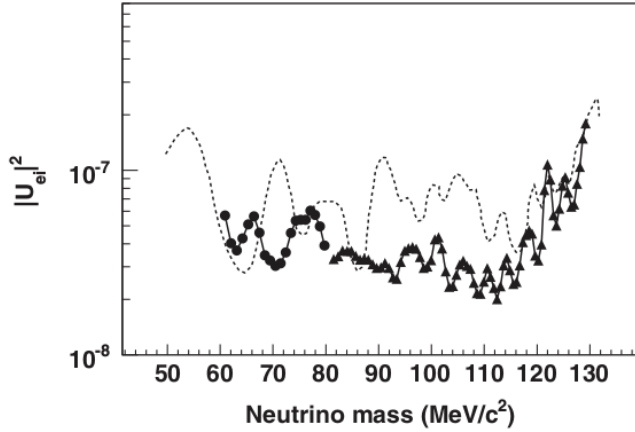


FIGURE 4.2: 90% CL upper limits (solid black line) on U_{eI}^2 placed by PIENU [252]. The dashed line represents older limits which this experiment subsumed.

4.3.1 PIENU

The PIENU experiment [252] sought to detect RHNs in the mass range of 68 – 129 MeV by searching for peaks in the energy spectrum of the decay process

$$\pi^+ \rightarrow e^+ \nu \quad (4.44)$$

(hence the name). This process is helicity-suppressed in the Standard Model, but this suppression is relaxed in the presence of RHNs. There were other attempts at the search using the same process [325], but the background process $\mu^+ \rightarrow e^+ \nu \bar{\nu}$ (the decay of the muon which can be produced in $\pi^+ \rightarrow \mu^+ \nu$) was not fully corrected for. PIENU attempted to improve these results. The difference in the lifetimes of pions and muons ($\tau_\pi = 26$ ns, $\tau_\mu = 2197$ ns) was used for this purpose, by setting up timing cuts based this difference. The spectrum was also cut at an angle to the beam axis to reduce shower leakage and improve the signal-to-noise ratio.

The experiment was sensitive to the mixing $|\Theta_{eI}|^2 \equiv U_{eI}^2$ and μ in equation 4.43 is also taken to scale as U_{eI}^2 in the analysis. Although no peaks were found, exact information on the number of background events is unavailable. Further, production processes in peak searches are, in general, unaffected by the Majorana/Dirac nature of the RHNs; hence, no correction is necessary here.

The constraints on U_{eI}^2 are at 90% CL, so it is implemented as a half-Gaussian with zero mean and error set at 1.28σ . The results of the experiment are shown in figure 4.2.

4.3.2 PS-191

This experiment [68] was designed for the purpose of detecting neutrino decays. RHNs would be produced via either of the following mechanisms:

$$\pi^+ / K^+ \rightarrow e^+ \nu_e \quad \text{or} \quad (4.45)$$

$$\pi^+ / K^+ \rightarrow \mu^+ \nu_\mu. \quad (4.46)$$

The RHN that may be produced via the mixing with the LHNs could then decay via one of the topologies below.

$$N \rightarrow \mu^- e^+ \nu, \quad (4.47)$$

$$N \rightarrow e^- \mu^+ \nu, \quad (4.48)$$

$$N \rightarrow e^- \pi^+, \quad (4.49)$$

$$N \rightarrow \mu^- \mu^+ \nu, \quad (4.50)$$

$$N \rightarrow \mu^- \pi^+ \quad \text{or} \quad (4.51)$$

$$N \rightarrow e^- \pi^+ \pi^0. \quad (4.52)$$

Thus, PS-191 could constrain the quantities U_{eI}^4 and $U_{\mu I}^4$ for RHNs with a mass between 20 – 450 MeV.

Having found no signal or background events, it placed constraints on these parameters at 90% CL, as shown in figure 4.3. There are a couple of deviations from the original analysis. The first is necessitated by the fact that in the original analysis, the constraints were derived under the assumption that the RHNs interact only through the charged current. In the seesaw model, they can interact via the neutral current also, so any of the LHN flavours can appear in the decays. This was recognized a few years ago and the PS191 limits were re-interpreted with the inclusion of neutral current interactions in [57]. The effect of this is to make the signal count proportional to $U_{eI}^2 \times \sum_\alpha c_\alpha U_{\alpha I}^2$. The co-efficients are given by

$$c_e = \frac{1 + 4 \sin^2 \theta_W + 8 \sin^4 \theta_W}{4}, \quad c_\mu, c_\tau = \frac{1 - 4 \sin^2 \theta_W + 8 \sin^4 \theta_W}{4}. \quad (4.53)$$

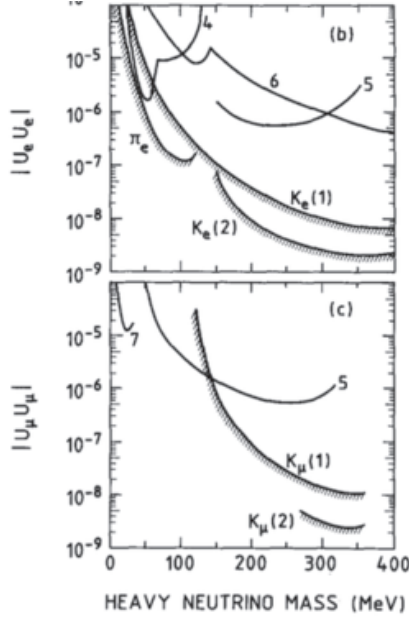


FIGURE 4.3: 90% CL upper limits (shaded lines) on U_{eI}^2 and $U_{\mu I}^2$ from the PS-191 experiment [68]. These are re-interpreted, as mentioned in the text. Other lines represent results from other experiments.

The revised bounds are used here. The limits are encoded in likelihood form as in equation 4.43, with the proportionality factor being 2.44.

4.3.3 CHARM

CHARM [69] sought to detect RHNs using two strategies. One used a neutrino beam from dumping protons on copper (BD), the other utilized a wide-band neutrino beam (WBB) from primary protons.

In the beam dump, the production of RHNs was assumed to occur through the decay of D mesons. They would then decay via

$$N \rightarrow e^+ e^- \nu_e, \quad (4.54)$$

$$N \rightarrow \mu^+ \mu^- \nu_\mu, \quad (4.55)$$

$$N \rightarrow e^+ \mu^- \nu_e \quad \text{or} \quad (4.56)$$

$$N \rightarrow \mu^+ e^- \nu_\mu \quad (4.57)$$

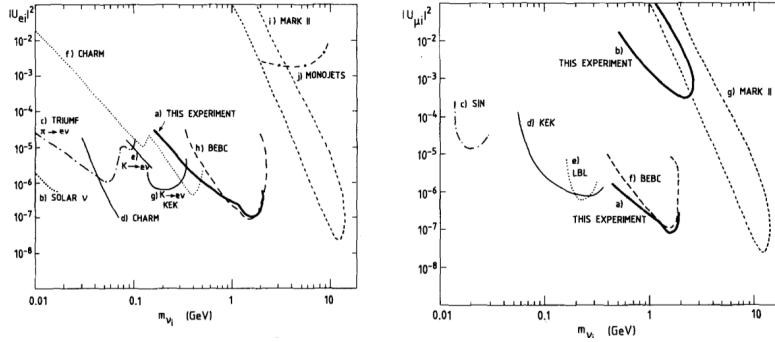


FIGURE 4.4: 90% CL upper limits on U_{eI}^2 and $U_{μI}^2$ from CHARM [69] (“THIS EXPERIMENT” in both plots). These are re-interpreted, as with PS-191.

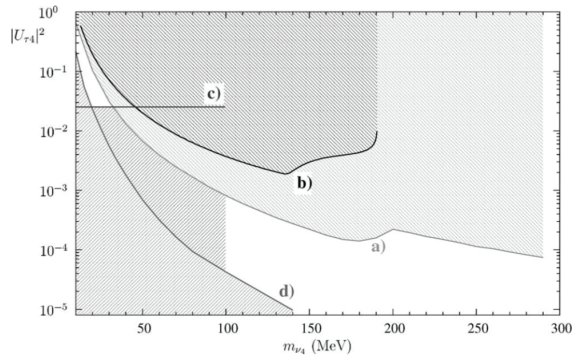


FIGURE 4.5: 90% CL upper limits on $U_{τI}^2$ from CHARM data, assuming RHNs couple only to tau flavour [279] (line a in the plot).

(and the anti-particle counterparts) and the decay products were looked for. The main background here came from hadron showers resulting from the inelastic scattering of electron and muon neutrinos and anti-neutrinos. These events have a wider angular distribution than the decays listed above, which was used to separate signal from background. In addition, the electromagnetic processes occur with more regularity. This was also factored in to suppress the background.

In WBB, RHN production was assumed to occur via neutrino-nucleus neutral current scattering and their subsequent decay into hadrons was searched for:

$$\nu_\mu A \rightarrow XN; \quad N \rightarrow \mu H, \quad (4.58)$$

where A and X are nuclei and H is the hadron from the decay of the RHN. The limits from the WBB analysis are, however, weaker than those exerted by other experiments in the same mass range, and are not considered here.

The BD analysis yielded no candidate events or background and hence placed limits on U_{eI} and $U_{\mu I}$ at 90% CL. The original analysis assumed the possibility of RHNs interacting only via the charged current; the results are re-interpreted accounting for the inclusion of neutral current interactions [57] as discussed in section 4.3.2 and equation 4.43 is once again used to represent the likelihood, with the proportionality factor being 2.44.

The data from CHARM has also been used to constrain the mixing with tau flavour. Using the beam dump information and with the assumption that RHNs mix solely with tau-flavoured leptons, [279] placed limits at 90% CL on $U_{\tau I}$, which we implement as a half-Gaussian with zero mean and error set at 1.28σ .

Dirac RHNs were assumed in both the original and tau-specific analyses, so the limits presented are also re-scaled by dividing them by $\sqrt{2}$.

The experimental limits are displayed in figures 4.4 and 4.5.

4.3.4 DELPHI

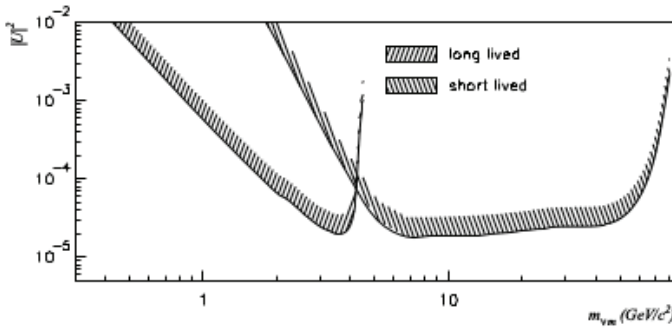


FIGURE 4.6: Upper limits at 95% CL, equally applicable to all 3 mixings, from the DELPHI analysis [73].

At DELPHI [73], the dominant RHN production mechanism was

$$e^+e^- \rightarrow Z^0 \rightarrow N\bar{\nu} \quad (4.59)$$

(the process $Z^0 \rightarrow N\bar{N}$ would be suppressed due to the additional U^2 factor). The products of the RHN decaying via the weak and neutral current were then searched for, according to

$$N \rightarrow \nu Z^*; \quad Z^* \rightarrow \nu\bar{\nu}, l\bar{l}, q\bar{q}, \quad (4.60)$$

or

$$N \rightarrow l'W^*; \quad W^* \rightarrow \nu\bar{l}, q\bar{q}'. \quad (4.61)$$

Since the RHNs could have existed long enough to travel macroscopic distances of upto 100 cm, different signatures had to be considered and the analysis was split to tackle the short- and long-lived cases separately. For short-lived RHNs with mass less than about 30 GeV/c², because of the large boost received by N , the signature would be a monojet. Background came from leptonic Z boson decays or $\gamma\gamma$ processes; angular cuts were made to remove these. Higher masses open the decay channel into $q\bar{q}$ (and a lepton, depending on the particular channel), so the signature would be two acollinear jets which are also acoplanar with respect to the beam axis. Hadronic Z decays with missing energy make up most of the background in this case, and a neural network was used to remove all of them from the final data.

Longer-lived RHNs were looked for using displaced vertices and calorimeter clusters. The former was useful in tracking RHNs with an intermediate lifetime; however, a cluster finding algorithm along with vertex reconstruction did not find any signals in this regime. The latter was used to detect the longest-lived RHNs, whose decay products would interact with the outermost layers/components of DELPHI. The experimental signature would be a cluster of hits in a small angular region, coincident with the beam collision and traceable back to the initial interaction point.

DELPHI could constrain U_{eI} , $U_{\mu I}$ and $U_{\tau I}$ for RHNs having a mass between 0.5 – 80 GeV. Majorana neutrinos were assumed and the analysis yielded one candidate event (from the monojet component of the short-lived RHN analysis) and no background events; the result is shown in figure 4.6. In our analysis, this means the proportionality factor is 3.09 and μ scales as $U_{\alpha I}^4$.

A caveat must be mentioned: DELPHI presented bounds on the mixing in a flavor-independent manner: the limit on U^2 , as presented in the paper, applies equally to U_e^2 , U_μ^2 and U_τ^2 , as they mention. In the mass range under

consideration, the mass of the tauon will, of course, influence the strength of the limit and, as they quote, the presented bounds become weaker for masses below ~ 4 GeV. However, the extent of the kinematic suppression due to the tauon mass is not quantitatively discussed; the limits are used as is, noting that it is highly likely that NA62 will subsume these bounds in the near future [326].

4.3.5 ATLAS

The process relevant for RHN production in ATLAS [276] is

$$pp \rightarrow (W^\pm)^* \rightarrow l^\pm N. \quad (4.62)$$

The RHNs were taken to be heavier than the W boson, allowing it to decay into a lepton and a W boson:

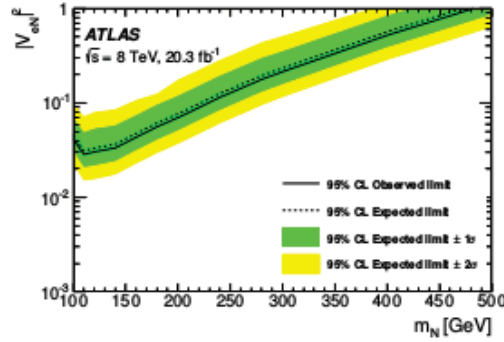
$$N \rightarrow l^\pm W^\mp. \quad (4.63)$$

The W boson would then decay predominantly into a quark-antiquark pair, and the signature of this decay chain was searched for, with either two electrons or muons in the final state.⁵

Within the framework of the SM, one can get the same outgoing particles from Drell-Yan and diboson processes, as well as from decays of top quarks. After modelling all relevant processes through simulation, differentiating between the SM and BSM processes and applying cuts to suppress any additional background (such as a beam-related component, calorimeter noise and cosmic ray muons), the analysis placed 95% CL limits on the two mixing angles U_{eI} and $U_{\mu I}$ in the mass range of 100 – 500 GeV. The experimental limits are displayed in figure 4.7.

Details on the number of observed/expected events and background are available and could be cast into a likelihood function combining Poissonian and Gaussian errors; however, implementing the limits as a half-Gaussian

⁵There is ongoing dispute in the literature concerning the rate of LNV processes at collider experiments, namely whether they are always suppressed and therefore unobservably small [65, 170] or whether coherent flavour oscillations can lead to LNV signatures in spite of the smallness of these parameters [172, 173, 174]. In the range of M_I under consideration here, the strongest direct search constraints do not come from experimental signatures that rely on LNV, and the results are therefore only mildly affected by the outcome of this discussion.



(b)

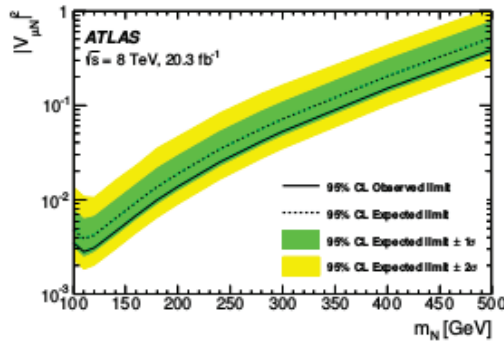


FIGURE 4.7: 95% CL upper limits on electron and muon mixing from the ATLAS analysis [276].

with zero mean and error set at 1.64σ reproduces the experimental limits well enough for the current purpose.

Hence, in our analysis, $\mu \propto U_{\alpha I}^4$, $\alpha = e, \mu$. The original analysis was carried out under the assumption of Majorana RHNs, so no correction is necessary.

4.3.6 CMS

With the LHC having run with a center-of-mass energy of 13 TeV, the CMS detector searched for different event signatures of the same process as ATLAS studied: instead of the W boson decaying into quarks, it decays into a lepton-(anti-)neutrino pair, leaving a final signature of three charged leptons (these cannot have the same charge, but can be any combination of electrons and muons). 95% CL limits were calculated for U_{eI} and $U_{\mu I}$ for RHNs with mass between 1 GeV and 1.2 TeV [277] (shown in figure 4.8).

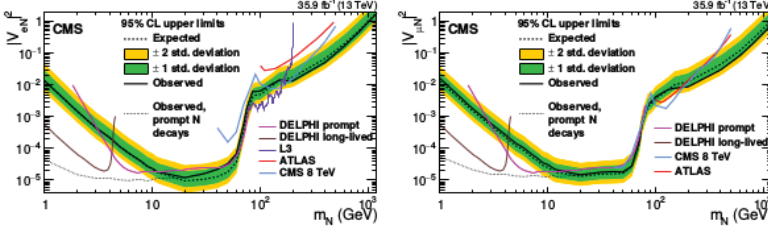


FIGURE 4.8: 95% CL upper limits on electron and muon mixing from the CMS analysis [277].

As before, Majorana RHNs were assumed in the analysis, and the implementation of the limits mirrors that of ATLAS.

Note: Updated bounds from ATLAS [327] and CMS [328] have been released, but are not included, since these papers came out after our scans were completed. Additionally, the new bounds from ATLAS in [327] are comparable to those from DELPHI, while the updated CMS bounds in [328], based on a dilepton search, produces stronger bounds for RHN masses above ~ 500 GeV, beyond our range of study.

4.3.7 E949

In this experiment [278], RHNs were searched for in the decay of kaons produced in a beam dump:

$$K^+ \rightarrow \mu^+ N. \quad (4.64)$$

This is similar to PIENU's and PS191's search strategy, essentially constraining the mixing $U_{\alpha I}^2$ using

$$\Gamma(M^+ \rightarrow l^+ N) = \rho \Gamma(M^+ \rightarrow l^+ \nu_l) U_{\alpha I}^2, \quad (4.65)$$

ρ being a kinematic factor. The main interfering background processes were well-understood and suppressed; these are $K^+ \rightarrow \mu^+ \nu_\mu \gamma$, $K^+ \rightarrow \pi^0 \mu^+ \nu_\mu$ and $K^+ \rightarrow \pi^0 \pi^+ \gamma$.

Constraints on $U_{\mu I}$ were placed at 90% CL in the mass range 175 – 300 MeV; we also divide the limits by a factor of $\sqrt{2}$ to account for the Majorana nature of RHNs in our model. The limits are reproduced in figure 4.9.

The likelihood is modeled as a half-Gaussian with zero mean, error set at 1.28σ and $\mu \propto U_{\mu I}^2$.

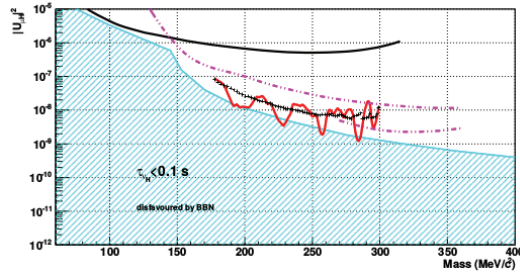


FIGURE 4.9: Upper limits on $U_{\mu I}^2$ at 90% CL from E949 [278].

4.3.8 NuTeV

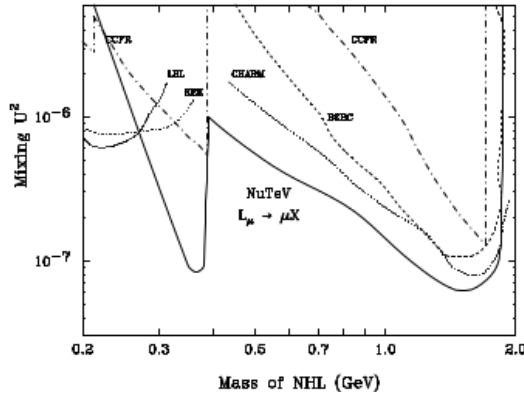


FIGURE 4.10: 90% CL upper limits on $U_{\mu I}^2$ from NuTeV [70].

The NuTeV experiment [70] searched for RHNs through their decay into the following final states: $\mu\nu$, $\mu\mu\nu$, $\mu\pi$ and $\mu\rho$. They were assumed to be produced in the decay of mesons. Thus, the experimental signature was a neutral particle decaying into two charged ones within the detector.

After suppressing major sources of background, which were neutrino interactions in the helium used in the drift chamber (the helium's purpose was to reduce neutrino interactions, but this was achieved imperfectly), interactions with the material beyond this chamber or around it, no signal or background events were detected and 90% CL limits on $U_{\mu I}$ were placed for Majorana RHNs with a mass between 0.25 – 2 GeV, as shown in figure 4.10.

The likelihood is modeled as in equation 4.43, with a proportionality factor of 2.44 and μ scaling as $U_{\alpha I}^4$.

Chapter 5

Results

5.1 Scanning strategy, parameter ranges and priors

One primary goal of this work was to fully map the profile likelihood in the mixing-mass, or $M_I - U_{\alpha I}^2$ plane for RHN masses between 100 MeV and 500 GeV. The lower limit was chosen due to considerations outlined in chapter 1, reiterated here: although there are no limitations on RHN mass a priori, particularly in bottom-up parameterizations, like the Casas-Ibarra scheme, constraints from BBN, for instance, restrict their lifetime [56] (indirectly limiting allowed masses). Coupled with the results of previous direct searches [57, 58] which have not reported any detections, the lower limit on the RHNs' masses can be set to about 100 MeV.

After initial scans, the differential evolution scanner Diver was chosen, since it demonstrably explored the parameter space better. Even so, the parameter space is complex and there are regions where a tuning of the parameters, such as accidental cancellations among terms in the active neutrino mass matrix, is needed to simultaneously satisfy existing constraints and reach the largest allowed mixings. This necessitated the usage of a few strategies in order to better sample the space.

We perform scans using a “differential” model, where the RHN masses are parameterized using the quantities M_1 , $\Delta M_{21} = M_2 - M_1$ and M_3 , as well as using the “full” model. It is known that (an approximate) mass degeneracy between two RHNs can lead to approximate B-L symmetry conservation, in turn leading to larger observable mixings (see, for example, [165]). Considering lepton number-violating processes like neutrinoless double-beta decay, which places one of the tightest constraints on U_{eI}^2 especially, a protecting symmetry suppresses these signals [59], rendering U_{eI}^2 unconstrained by

experimental upper limits on $m_{\beta\beta}$. Hence, a logarithmic prior is placed on the RHN masses M_1 and M_3 , as well as on the mass splitting, ΔM_{21} . This helps with a better sampling of the limits of the most constraining experiments/observables. Further, the Casas-Ibarra parameterization makes it possible to get large values of $U_{\alpha I}^2$ by simply choosing $|\text{Im}(\omega_{ij})| > 1$, since $\left(\Theta^2 \sim \frac{\exp(2\text{Im}(\omega))}{M}\right)$; this is also done by setting appropriate ranges in the configuration file.

Out of the active neutrino parameters, only α_1 and α_2 are unconstrained by oscillation data, hence they are allowed to vary freely from 0 to 2π with flat priors. The ranges for the other neutrino phases and angles are taken from the 3σ ranges from the NuFit collaboration [103], also with flat priors. The mass of the lightest active neutrino has a definite impact on the lower bound of U_I^2 (equation 2.69) [59], so a logarithmic prior is chosen (see table 5.3), which enables us to examine this impact in better detail than a flat prior would allow and keeps the BBN limits relevant [58]. The upper limit on m_{ν_0} is chosen as the broad cosmological bound given by Planck [284], $\sum m_\nu < 0.23$ eV. In order to better fit the active neutrino data, the mass splittings Δm_{21}^2 and Δm_{3l}^2 are chosen as scan parameters, where $l = 1$ and $\Delta m_{3l}^2 > 0$ for normal hierarchy and $l = 2$ and $\Delta m_{3l}^2 < 0$ for inverted hierarchy.

Further, because of the choice of ordering in the definition of R in equation 2.84, the complete possible range of couplings is not covered. Thus, we introduce another parameter R_{order} , which takes discrete values [1, 6] corresponding to each of the possible permutations of the definition of R in terms of R^{ij} . This allows full coverage of the mixing space and, since the likelihood is theoretically independent of this ordering (and confirmed by the data), a uniform distribution of values taken by R_{order} is achieved.

Since the mixing matrix in the C-I parameterization depends on m_H (due to the 1-loop correction), as seen in section 2.3, m_H is taken to be a nuisance parameter with a Gaussian distribution around its averaged measured value [254] and a flat prior. Other SM parameters are fixed to their PDG values [254].

However, we found that for low and high RHN masses ($\lesssim 1$ GeV and $\gtrsim 100$ GeV), there are regions of excess likelihood. This leads to the scanner preferring points in these regions, leaving the rest of the parameter space undersampled, even in other parts of the parameter space with few constraints,

making for less than satisfactory profile likelihood contours. Thus, we remove the likelihood contribution from $R_{e\mu}^K$ from the total likelihood (which drives the scan), and add it back later during postprocessing. Other likelihoods with positive contributions are Γ_{inv} (the invisible decay width of the Z boson), unitarity of the CKM matrix and $R_{e\mu}^\tau$; these push the scan towards large $U_{\tau I}^2$ couplings. This helps in saturating the upper limits of $U_{\tau I}^2$, but leaves regions with low τ coupling undersampled. These features will be discussed in more detail in section 5.4.

In order to isolate these preferred regions and sample regions outside them well, the following strategies were implemented.

- An extra auxiliary likelihood called the *coupling slide* was introduced, which can be tuned to prefer specifiable RHN mass ranges and large mixings. By specifying options in the configuration file: `slope` $\equiv s_C$, `mslope` $\equiv s_M$, `I` (RHN index) and `i` (flavour index), a contribution to the log-likelihood given by

$$\ln \mathcal{L} += (s_C \times \log[\min(U_{iI}, 1)]) + (s_M \times \log M_I) \quad (5.1)$$

can be added. This contribution was then subtracted in the postprocessing stage.

- The full scan was split into three sets, one for each flavour mixing (targeted in *coupling slide*). Each set is in turn comprised of multiple parts, each part working on a different RHN mass range. This ensures that each flavour coupling saturates the relevant experimental upper bound in these ranges. Further, scans were set up to use different values of ΔM_{21} and/or m_{ν_0} to sample fine-tuned regions of the parameter space.
- Since the preferred regions were observed to exist for $U_{\tau I}^2 > 10^{-4}$, the set of scans for this flavour mixing were further divided into a subset where $U_{\tau I}^2$ is restricted to be less than 10^{-4} (so that the scanner can explore the rest of the parameter space at lower coupling values more efficiently) and another where the mixing is unrestricted.

These steps lead to another subtle issue to be considered. Now, the ordering of the RHNs is arbitrary in our scheme; given sufficient computing time, the data (i.e. the samples obtained from the scans) should be agnostic to this. In a profile likelihood plot, therefore, one should not need to specify which RHN's parameter space is being shown; the plot should hold for any

Parameter	Value/Range	Prior
<i>Free parameters</i>		
α_1, α_2 [rad]	$[0, 2\pi]$	flat
Re ω_{ij} [rad]	$[0, 2\pi]$	flat
Im ω_{ij}	$[-15, 15]$	flat
<i>Nuisance parameters</i>		
δ [rad]	$[0, 2\pi]$	flat
θ_{12} [rad]	$[0.547684, 0.628144]$	flat
θ_{23} [rad]	$[0.670206, 0.925025]$	flat
θ_{13} [rad]	$[0.139452, 0.155509]$	flat
Δm_{21}^2 [10^{-5} eV ²]	$[6, 9]$	flat
Δm_{3l}^2 [10^{-3} eV ²]	$[\pm 2, \pm 3]$	flat
m_H [GeV]	$[124.1, 127.3]$	flat
R_{order}	$[1, 6]$	flat

TABLE 5.1: Parameters ranges and priors adopted for the scan.

of them. However, our settings revolve around N_1 : as mentioned, this is to ensure the scan sufficiently explores the relevant upper limits. The scanner thus prefers points where the flavour couplings to this RHN are high, leading to an underexploration of the other, symmetrically allowed possibilities, i.e. N_2 or N_3 playing the role of the RHN with high couplings. When plotting the profile likelihood, which is asking, “Is there an N_I with mass M_I and coupling $U_{\alpha I}^2$ that is within n -sigma (of the best-fit point)?”, there are not enough points for N_2 and N_3 to completely define the likelihood contours. To further circumvent this problem of preferred regions and isolate them, we cap the maximum likelihood to the highest value in the “bulk” of the parameter space, i.e. between masses of about 10-70 GeV.

Table 5.3 shows the list of scans that are common to all three flavour mixings (so, along with the restricted $U_{\tau I}^2$ scans, 4x the number of scans in this table were done). Extra runs for low RHN mass ranges for tau flavour as well as a few dedicated runs for electron and muon flavour couplings also had to be set up to fully explore the profile likelihood; these are listed in tables 5.4 and 5.2. The ranges and priors used for all other parameters are given in table 5.1. The two hierarchies have also been scanned independently; all scans were performed using Diver ver. 1.0.4, with the settings `NP` = 19200, `conv_threshold` = 10^{-10} , and the `ljDE` option in Diver switched on. High flavour couplings of N_1 are favoured through the new likelihood. These settings, including the very low convergence threshold, together with the strategy described above, ensure a thorough exploration of the parameter space,

M_1 [GeV]	ΔM_{21} [GeV]	M_3 [GeV]	m_{ν_0} [eV]	$[s_C, s_M]$
$U_{e\mu I}^2$				
[0.06, 0.14] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]
U_{eI}^2				
[0.14, 0.2] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]
[0.2, 0.4217] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>		$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]
$U_{\mu I}^2$				
[0.14, 0.3162] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]

TABLE 5.2: A list of scans specific to certain flavour mixings, making up part of the full dataset(s), along with parameter ranges and priors. These scans were performed for both mass hierarchies.

M_1 [GeV]	ΔM_{21} [GeV]	M_3 [GeV]	m_{ν_0} [eV]	$[s_C, s_M]$
[0.1, 0.3162] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>	<i>prior: log</i> [0.06, 500]	$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]
[0.1, 0.4217] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>		$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, -0.5]
[0.3162, 2] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>		$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[2, 60] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[10^{-6}, 0.23]$ <i>prior: flat</i>	[0, 5, 0]
[2, 60] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$10^{-4}, 10^{-5}, 10^{-6}$ <i>fixed</i>	[0, 5, 0]
[60, 500] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[10^{-6}, 0.23]$ <i>prior: flat</i>	[0.7, 0.25]
[60, 500] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$10^{-4}, 10^{-5}, 10^{-6}$ <i>fixed</i>	[0.7, 0.25]

TABLE 5.3: A list of scans making up part of the full dataset(s), along with parameter ranges and priors. The settings are common to the three sets for the individual flavour mixings (refer to text for details) as well as to both mass hierarchies. These scans were performed with $U_{\tau I}^2 < 10^{-4}$ and with no restriction.

albeit at the price of computation time. Although none of the observables required heavy computation or simulation, most scans took between 2 and 10 hours of running time on a large number of supercomputer cores varying in number between 250 and 780. All tests and scans were performed across

M_1 [GeV]	ΔM_{21} [GeV]	M_3 [GeV]	m_{ν_0} [eV]	$[s_C, s_M]$
NH+IH				
[0.175, 0.3611] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[10^{-2}, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[0.25, 0.3611] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[10^{-2}, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[0.25, 0.4] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[10^{-2}, 0.23]$ <i>prior: log</i>	[1, 0]
[0.3611, 0.4492] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>		$[10^{-2}, 0.23]$ <i>prior: log</i>	[0.5, −0.5]
[0.3611, 0.4492] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[10^{-2}, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[0.4, 0.5] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>		$[10^{-7}, 0.23]$ <i>prior: log</i>	[0.5, −0.5]
IH				
[1.25, 1.45] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[0.01, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[1.65, 1.85] <i>prior: log</i>	$[10^{-20}, 10^{-10}]$ <i>prior: log</i>		$[0.01, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
IH; $U_{\tau I}^2 < 10^{-3}$				
[0.3162, 1.4] <i>prior: log</i>	$[10^{-10}, 0.1]$ <i>prior: log</i>	[0.06, 500] <i>prior: log</i>	$[0.03, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[1.0, 1.5] <i>prior: log</i>	$[10^{-7}, 0.01]$ <i>prior: log</i>		$[0.03, 0.23]$ <i>prior: log</i>	[0.5, 0.5]
[1.4, 1.78] <i>prior: log</i>	$[10^{-7}, 0.01]$ <i>prior: log</i>		$[0.03, 0.23]$ <i>prior: log</i>	[0.5, 0.5]

TABLE 5.4: A list of scans specific to tau flavour mixing, making up part of the full dataset(s), along with parameter ranges and priors (refer to text for details).

several supercomputer facilities, including the Abel cluster in Oslo, MareNostrum in Barcelona, Marconi in Bologna, Lisa in Amsterdam and Prometheus in Krakow.

Due to the large number of scans and the low convergence threshold used, the size of the full dataset surpassed 1TB for each hierarchy. This made it extremely difficult to use most plotting routines on them. With the objective of making profile likelihood plots in the M_I vs. $U_{\alpha I}^2$ planes, a subset of the data points in each mass-coupling plane was extracted; the combined set was thus optimized for all couplings. A similar optimisation was performed in the m_{ν_0} - $U_{\alpha I}^2$ planes, since we intended to study the effect of m_{ν_0} cuts on the parameter space.

The C-I parametrization is also inherently biased towards the different right-handed neutrinos, in particular through their masses M_I . To remove this bias, after combining the reduced datasets for all the scans, a symmetrisation procedure was conducted on the combined datasets over M_I and $U_{\alpha I}^2$, increasing the size of the datasets sixfold.

Out of these massive datasets, a total of 37.7 million valid data samples were using for plotting, of which 9.4 million correspond to normal hierarchy and 8.7M for inverted hierarchy optimised in the M_I - $U_{\alpha I}^2$ planes, and 9.9M for normal and 9.7M for inverted hierarchy optimised in the m_{ν_0} - $U_{\alpha I}^2$ planes.

In the following sections, separate figures for the inverted hierarchy are presented only if they differ significantly from the corresponding normal hierarchy plots. Plots for the IH that are essentially identical to their NH analogues shown in this chapter are in appendix D.

5.2 Profile likelihoods

The purpose of making profile likelihood plots was to correctly visualize the allowed parameter space of the type-I seesaw model with three RHNs, saturating the upper limits in desired projections of the parameter space (the $U_{\alpha I}^2 - M_I$ plane for most plots, the $U_{\alpha I}^2 - U_{\beta I}^2$ plane for the triangle plots in section 5.5). Phenomenologically, these are the most interesting projections to study. Now, the upper limit in a region may come from an overlap of multiple constraints (placed at different confidence levels) and this approach allows a careful treatment in such instances. This parameter space has been explored before while under the constraint of multiple overlapping observables, but profile likelihood contours in the present mass range with all (or most of) the observables in this study have not been constructed.

Figures 5.1 and 5.2 show 1σ and 2σ (corresponding to 68.3% and 95.45%) likelihood contours in the three $M_I - U_{\alpha I}^2$ planes, as well in the $M_I - U_I^2$ plane. In particular, high mixings for large masses are allowed, which could be interesting for future searches. The drop-off at lowest masses in the electron (figure 5.1a) and muon flavour (figure 5.1b) mixings and the forbidden region in the lower left for tau flavour (figure 5.1c) mixing is due to Big Bang nucleosynthesis constraints: the limit on an RHN's lifetime ($\tau_I < 0.1$ s; see section 4.2.6) implies that an RHN should have sufficiently large (total) mixing U_I^2 for the necessary decay to happen. There are few constraints on $U_{\tau I}^2$

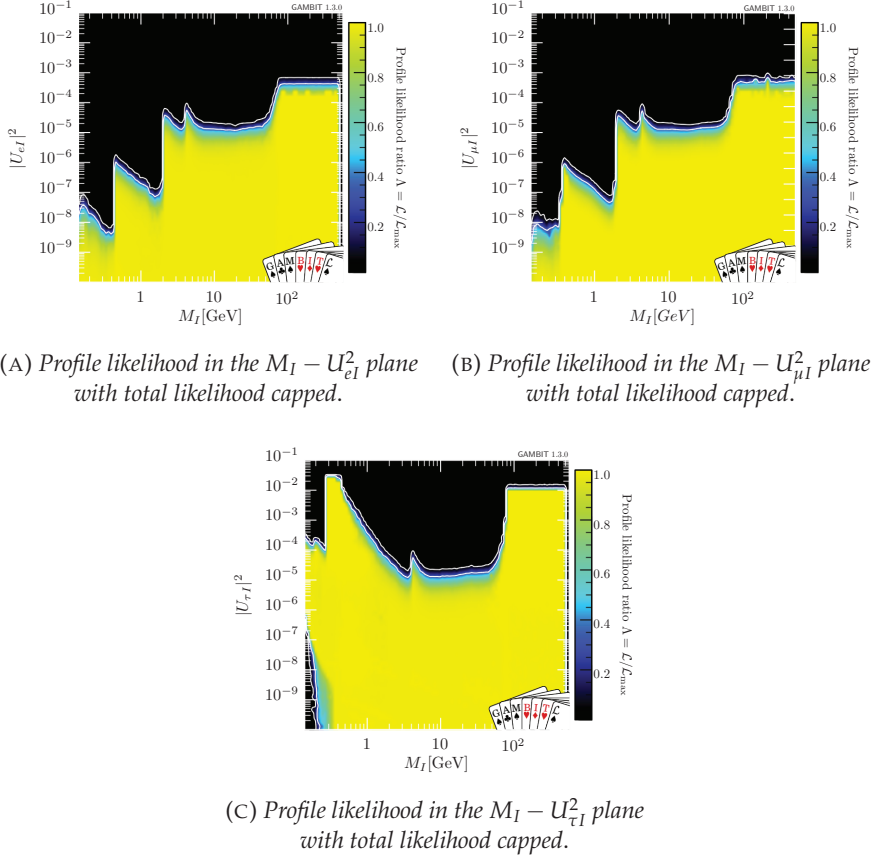


FIGURE 5.1: 1σ and 2σ profile likelihood contours in the electron, muon and tau flavour mixing-mass planes.

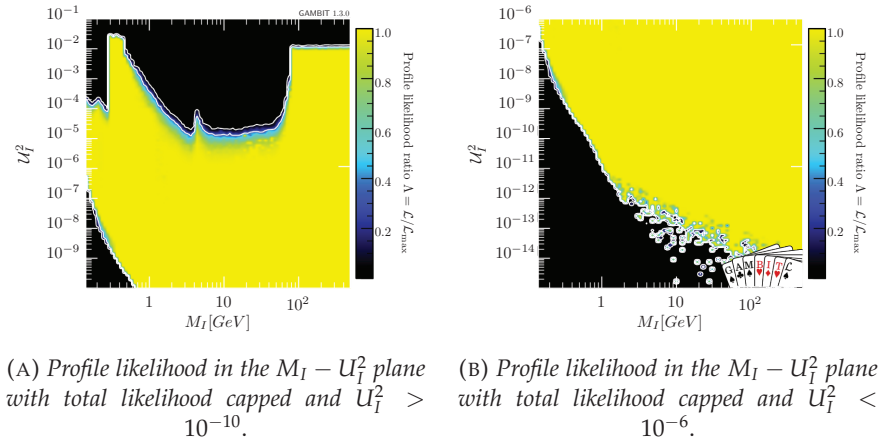


FIGURE 5.2: 1σ and 2σ profile likelihood contours in the total flavour mixing-mass plane.

for these masses, allowing it to go up to $\sim 10^{-4}$, limiting the mixings to the other two flavours. The BBN constraint also means that since the electron and muon flavour mixings are tightly constrained, mixing to tau flavour should have some minimum strength, seen to be close to 10^{-7} at lowest masses.

The lower limits on an RHN's combined mixing to all flavours, U_I^2 , coming from oscillation data and BBN, are more clearly visible in figure 5.2b. Since the lifetime falls off very rapidly with mass ($\tau_I \propto M_I^{-5}$, see appendix B), this limit is most relevant at lower masses and is irrelevant for $M_I > m_K$ [59]. The limit from oscillation data is what provides the lower bound at higher masses (given that it is unlikely that the scanner would explore the low coupling region, due to the reasons detailed in section 5.1, this bound is not well-defined). Note that, for $n = 3$, there is no lower bound on the individual flavour mixings $U_{\alpha I}^2$. This is due to the fact that the mixing $U_{\alpha I}^2$ to any one flavour α can be arbitrarily small, so long as the $U_{\beta \neq \alpha I}^2$ are large enough to make the RHNs decay in time to satisfy the BBN limit, as well as to satisfy the seesaw relation (generating the observed mass splittings among the active neutrinos).

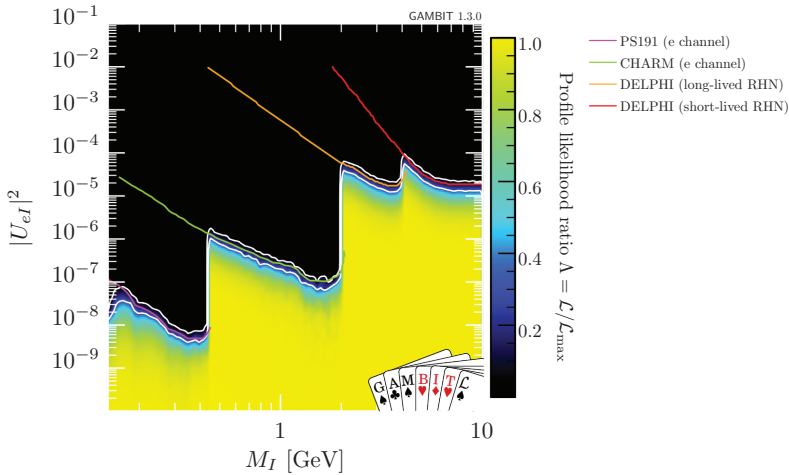


FIGURE 5.3: Profile likelihood in $M_I - U_{eI}^2$ plane, for $M_I < 10$ GeV, with the total likelihood capped and experimental limits overlaid.

For the electron and muon flavour mixings at RHN masses below that of the W boson, $M_I < 80$ GeV, the upper limits come almost entirely from direct detection experiments, as seen in figures 5.3 and 5.4. A few points need to be kept in mind while interpreting these plots. The upper limits from different

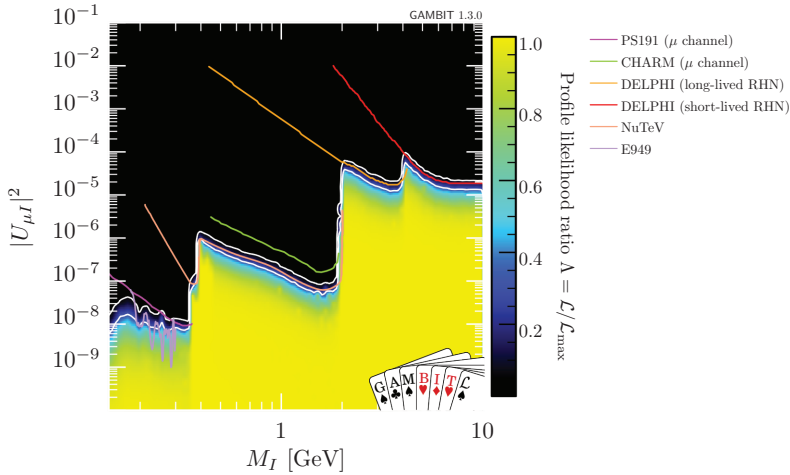


FIGURE 5.4: Profile likelihood in $M_I - U_{\mu I}^2$ plane, for $M_I < 10$ GeV, with the total likelihood capped and experimental limits overlaid.

experiments are at different confidence levels (either 90% or 95% CL); since 1σ ($\equiv 68.3\%$ CL) and 2σ ($\equiv 95.45\%$ CL) likelihood contours are made using the data, there may be a slight mismatch between the contours and the overlaid experimental limits: for example, NuTeV's limits on $U_{\mu I}^2$ were placed at 90% CL, hence the overlaid experimental data is between the 2σ (outer) and 1σ (inner) contours in figure 5.4. Additionally, even limits quoted at 95% CL will be slightly within the 2σ contour since experiments place limits on one parameter (the mixing) as a function of another (the RHN mass), whereas we construct two-dimensional profile likelihoods in the mixing-mass plane.

Another point to be mindful of has to do with the PS-191 and CHARM experiments. As mentioned in sections 4.3.2 and 4.3.3, the limits in the original papers must be re-interpreted [57] to include neutral current interactions of RHNs, which the original papers omitted; the experimental limits are then actually on a linear combination of the flavour mixings, specifically $U_{eI/\mu I}^2 \times \sum_{\alpha} c_{\alpha} U_{\alpha I}^2$ (see section 4.3.3 in chapter 4), not just the individual mixings U_{eI}^2 or $U_{\mu I}^2$. Since the coefficients c_{α} in the combination are all less than one, this means that the limit on any individual flavour is weaker than the limit on the combination, which is reflected in the contours. Lastly, there is the possibility of a mismatch between contours and experimental limits because of binning and interpolation effects while making the profile likelihood contours. This is seen in the plot for the muon flavour mixing (figure 5.4),

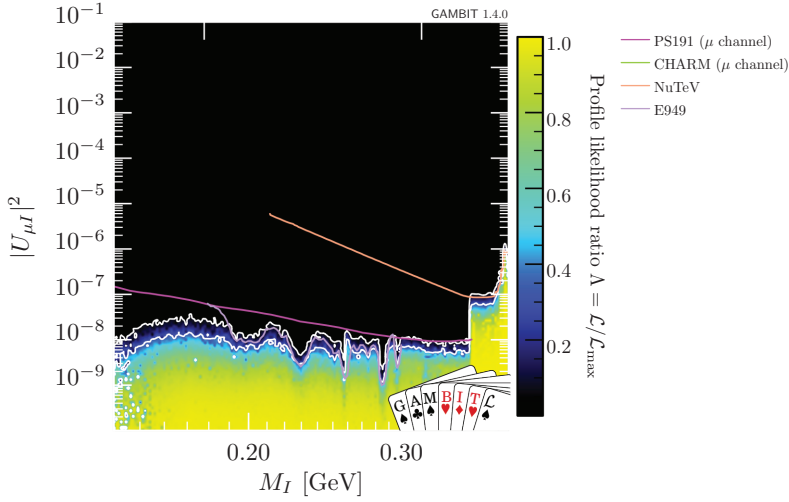


FIGURE 5.5: Profile likelihood in $M_I - U_{\mu I}^2$ plane, for $M_I < 0.4$ GeV, with the total likelihood capped and experimental limits overlayed.

where the E949 upper limits appear to be an order of magnitude stronger than the contours suggest; this is because the experimental limits change sharply in such a small mass region and is not a physical effect. By zooming in on the relevant mass range, as in figure 5.5, the interpolation becomes smoother and one can see that the contours indeed match the experimental limits.

Additionally, at lowest masses, the PS-191/E949 limits are not the most limiting constraint for electron and muon flavour mixings; BBN plays this role, as explained earlier.

The direct searches and/or indirect constraints that form the upper limits on each flavour mixing are shown in appendix C and are listed here for reference (in order of relevance in increasing RHN mass).

- Electron channel: BBN, PS-191, CHARM, DELPHI, CMS, (sine of) the Weinberg angle.
- Muon channel: E949 (stronger than PS-191), NuTeV (stronger than CHARM), DELPHI, CMS, (sine of) the Weinberg angle.
- Tau channel: CHARM (τ re-interpretation bounds), DELPHI, the invisible Z^0 decay width.

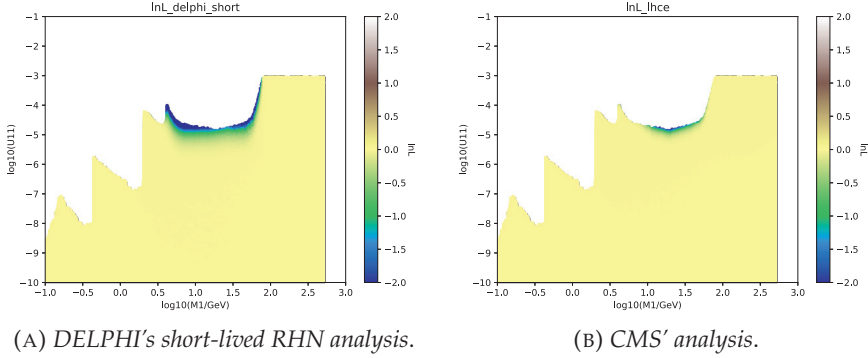


FIGURE 5.6: The effect of the upper limits from DELPHI's short-lived RHN analysis (section 4.3.4) and CMS' analysis (section 4.3.6) on the parameter space.

As a demonstration of the benefit of performing a likelihood analysis, consider figure 5.6. Both subplots effectively show the influence of the respective partial likelihood on the parameter space: the difference between the value of the partial likelihood and the value of the same partial likelihood at the best-fit point are calculated; the plot is colour-coded according to this difference, with the parameter points having the largest difference plotted last. So, a value of zero (or the dull yellow colour seen in most of the plot) means that the partial likelihood value at that point is the same as the partial likelihood value at the best-fit point. If the colour shifts to green/blue, the partial likelihood values are less than the value at the best-fit point, implying that this likelihood begins to constrain the parameter space in these regions. The shape of the profile likelihood contours in the $U_{eI}^2 - M_I$ plane is thus set by both experiments, but is stronger than if they would simply be overlayed one on top of the other.

Similar plots for the most relevant partial likelihoods at high masses for the three flavour mixings are shown in figure 5.7.

The upper limit coming from EWPO observables (sine of the Weinberg angle, the invisible Z^0 decay width) is in line with the findings of [329]. Further, the small mass splittings between RHNs, favoured due to the logarithmic prior on ΔM_{21} (see tables 5.3 and 5.4), leads to the neutrinoless double-beta decay limit from GERDA to be bypassed, i.e. we get mixings larger than what is suggested in [329].

Refs. [59] and [111] include many of the observables present in this study.

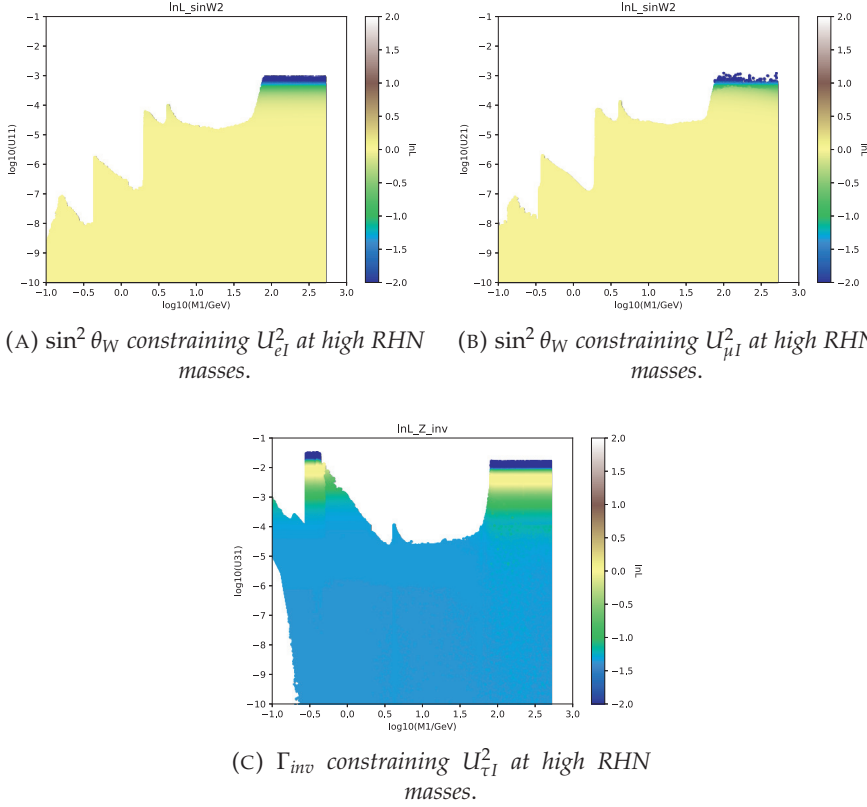


FIGURE 5.7: The tightest constraints exerted on the three flavour mixings at high RHN masses come from the sine of the Weinberg angle (for e and μ) and the invisible decay width of the Z boson (for τ).

However, in both studies, points are absolutely rejected if they fail a constraint, regardless of the level of exclusion with which the constraint was originally stated. As they acknowledge, sampling the entire allowed parameter space was practically impossible, and saturation of the upper limits throughout the entire mass range was thus also not achieved. These are issues that were focused on in this study and, as seen, successfully dealt with.

Since the high mass region ($M_I > 100$ GeV) is essentially flat for all three flavour mixings, we extract 2σ upper limits from the dataset(s). This is done by first trimming the dataset so that all points are within 2σ . Then, the highest value of the respective flavour mixing for $M_I > 100$ GeV are extracted. The upper limits on the three flavour mixings are shown in table 5.5.

Mixing	Upper limit (2σ)
<i>Normal hierarchy</i>	
U_{eI}^2	3.3114×10^{-4}
$U_{\mu I}^2$	5.9554×10^{-4}
$U_{\tau I}^2$	1.0132×10^{-2}
<i>Inverted hierarchy</i>	
U_{eI}^2	5.0551×10^{-4}
$U_{\mu I}^2$	6.8811×10^{-4}
$U_{\tau I}^2$	1.0205×10^{-2}

TABLE 5.5: 2σ (95.45% CL) upper limits on the individual flavour mixings for both mass hierarchies.

5.3 Effect of constraints on the lightest active neutrino mass

The effect of the mass of the lightest active neutrino on the lower limit of the total mixing of an RHN to all flavours (U_I^2) was briefly touched upon earlier. Oscillation data currently gives us information on two mass splittings, but not on the scale of the masses, that is the lightest active neutrino mass. If the lightest active neutrino were to be massless, there is no strict lower bound on the U_I^2 of N_I , since the non-zero mixings of the other two RHNs can explain the observed mass splittings. As the lightest neutrino mass increases, then, this lower bound becomes stronger, since larger and larger minimum mixings are required to generate the masses of the active neutrinos. This can be visualized using the naive estimate [59]

$$U_I^2 \sim \frac{1}{M_I} \sqrt{\Delta m_{\text{atm}}^2 + m_{\nu_0}^2}, \quad (5.2)$$

leading to the overlayed lines in figure 5.8. In the figure, the straight lines visually depict the naive estimate from equation 5.2, while the identically coloured jagged lines are the lowest U_I^2 for the same limiting values of m_{ν_0} from our data. Although undersampling of the low mixing and low m_{ν_0} region in the data yield lower limits on U_I^2 that are not smooth for $m_{\nu_0} < 0.01$ eV, there is no discrepancy with the naive estimate; the higher the mass of the lightest active neutrino, the larger the lower limit on U_I^2 .

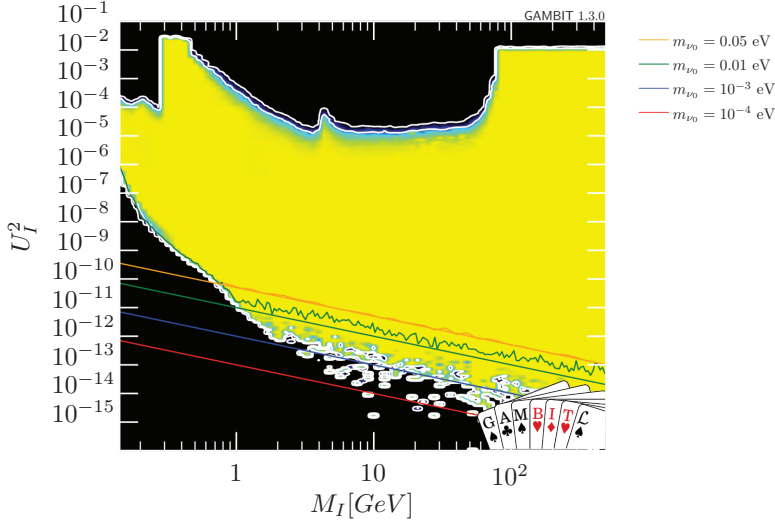


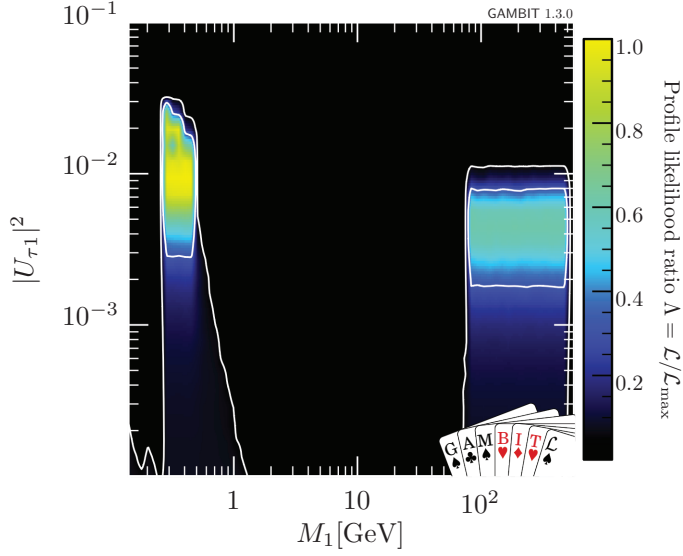
FIGURE 5.8: Profile likelihood in $M_I - U_I^2$ plane, with the total likelihood capped, along with the naive seesaw estimates for a few m_{ν_0} values. The color scheme is as in figure 5.1.

5.4 Regions of excess likelihood

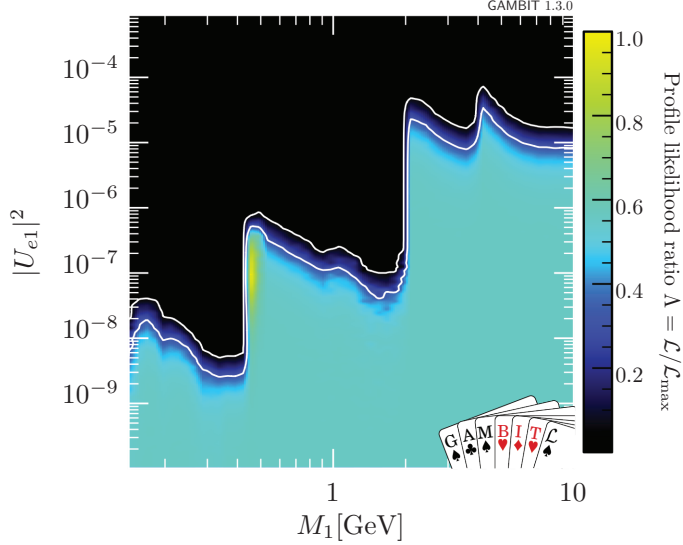
We now come to the regions of excess likelihood. Let us start with the regions in the tau flavour mixing space. In this case, three observables are responsible for this effect: the invisible decay width of the Z boson (section 4.2.1), the lepton universality quantity R_τ (section 4.2.3) and the fit to the entries of the CKM matrix (section 4.2.4). All three quantities show a preference for non-zero (and large) tau flavour mixing to an RHN; this has been observed before for Γ_{inv} and R_τ in [80, 253].

In the plots displaying these preferred regions (figures 5.9a, 5.10a and 5.10b) the labelling has been changed from the earlier $M_I - U_{\tau I}^2$ notation to $M_1 - U_{\tau 1}^2$. Additionally, the cap on the total likelihood is removed, and the other two RHNs are required to be away from these regions: $1 \text{ GeV} < M_{2/3} < 80 \text{ GeV}$. This is done to isolate the regions and highlight the fact that if $M_I - U_{\tau I}^2$ were to be plotted using scans that truly explored the parameter space completely, there would be some $N_{I \neq 1}$ with mixings that would cover the empty regions.

Looking at figure 5.10a, at low masses ($M_1 < 1 \text{ GeV}$), the CKM constraint and that from Γ_{inv} show a likelihood excess more than 1σ away from \mathcal{L}_0 (the



(A) Profile likelihood in $M_1 - U_{\tau 1}^2$ plane, with the total likelihood uncapped.



(B) Profile likelihood in $M_1 - U_{e 1}^2$ plane, with the total likelihood uncapped.

FIGURE 5.9: The regions of excess likelihood in the electron and tau flavour mixing spaces.

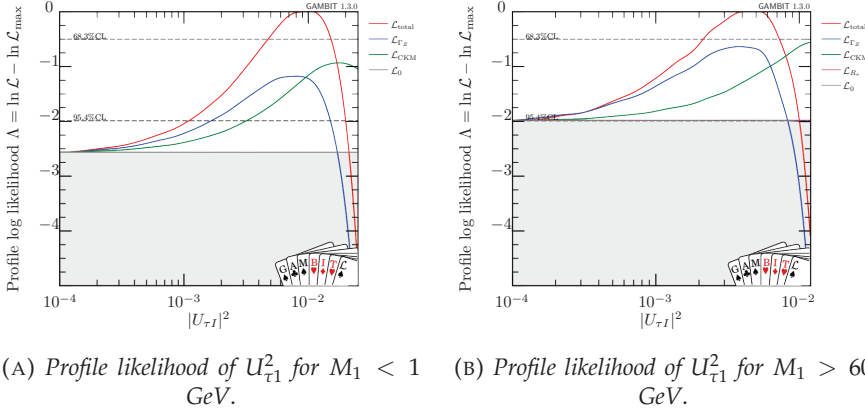


FIGURE 5.10: 1-dimensional profile likelihood(s) of $U_{\tau 1}^2$, along with the partial likelihoods responsible for the excesses at low and high RHN masses. The excesses in the individual likelihoods for Γ_{inv} (blue), CKM unitarity (green) and R_τ (pink in the figure on the right) are shown along with the total likelihood (red). All the likelihoods are normalized so that they show up as excesses over the combination of all other likelihoods (\mathcal{L}_0 in grey).

sum of all other likelihoods); the likelihood peaks at $U_{\tau 1}^2 \sim 10^{-2}$ from the Γ_{inv} observable and at a slightly larger value from the CKM unitarity constraint. Since the preference coming from the invisible decay width of the Z boson drops quickly past $U_{\tau 1}^2 \sim 10^{-2}$, giving *lower* likelihood values than \mathcal{L}_0 for tau flavour mixings above $\sim 2 \times 10^{-2}$, this acts counter to the preference coming from CKM unitarity, eventually leading to an upper limit.

From figure 5.10b, at high masses ($M_1 > 60$ GeV), the story is similar, apart from a preference for a smaller tau flavour mixing ($\sim 3 \times 10^{-3}$ due to Γ_{inv}). The effect of R_τ , although unignorable, is also very small compared to the effects of the other two constraints. This is corroborated by the results in [253], which showed the possibility of a very slight easing of tension between measured and (SM-)predicted values from non-zero mixings of $\sim 10^{-2}$.

In the electron flavour mixing space, a similar occurrence is seen at low masses (see figures 5.9b and 5.11). Note that the labelling change holds here too, for the same reason outlined before. The lepton universality observable R_K prefers non-zero mixings of about 10^{-7} at the 1σ level. Although the excess due to R_K carries over to larger couplings, as seen in figure 5.11), the total likelihood drops because of the CHARM bound.

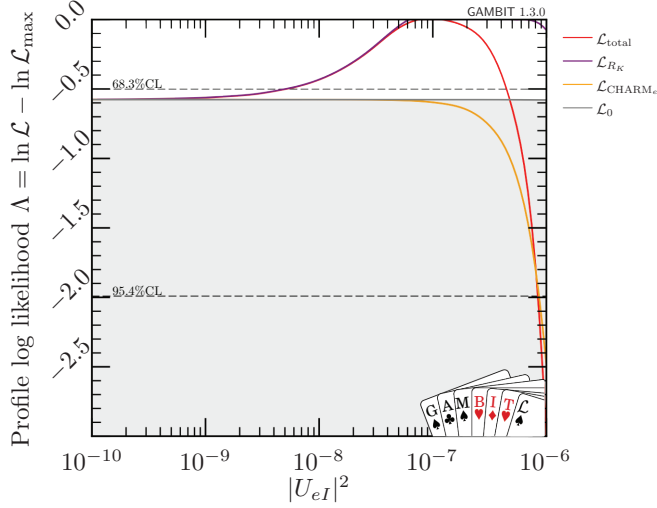


FIGURE 5.11: Profile likelihood of U_{e1}^2 for $M_1 < 10$ GeV. Analogously to figures 5.10a and 5.10b, the one-dimensional likelihoods for R_K (purple) and the reinterpreted CHARM bounds for tau mixing (orange) are shown with respect to the total likelihood (red) over the background of the combination of all other likelihoods (grey).

These excesses are interesting features of the RHN model. However, they are not indicative of a discovery of these particles, since the significance of the deviations is small and probably due to statistical fluctuations. These excesses were also observed in previous studies, coming from EWPO and CKM constraints [80]. Experiments like NA62 [110], SHiP [168], ILC [174] and CEPC [330] are expected to test these excesses in the future, as discussed later, in section 5.6.

5.5 Flavour mixing pattern and triangle plots

The allowed ratios which the three flavour mixings can take can be neatly visualized using triangle plots (for example, see [331]). Consider the limiting case of the lightest left-handed neutrino being massless, $m_{\nu_0} = 0$. This implies that the Yukawa coupling matrix F and therefore $\Theta^\dagger \Theta$ will have vanishing eigenvalues. So, in some direction of the RHN flavour space (which will, in general, be a linear combination of the N_I), there is no coupling to the left-handed neutrinos; the observed mass differences must be explained

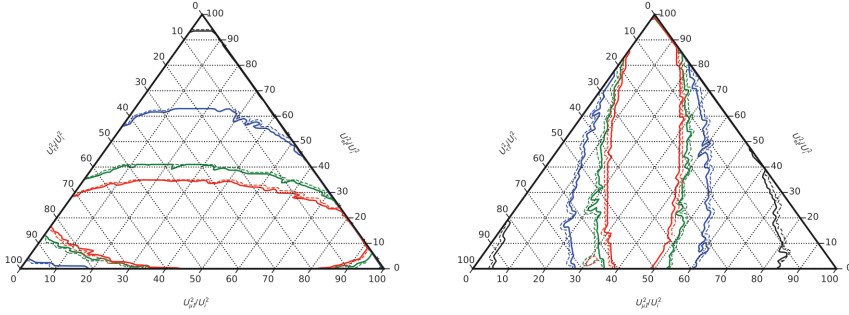


FIGURE 5.12: $U_{\alpha I}^2/U_I^2$ (in percent) for $m_{\nu_0} < 0.01$ meV (red), $m_{\nu_0} < 0.1$ meV (green), $m_{\nu_0} < 1$ meV (blue) and $m_{\nu_0} < 10$ meV (black). Solid (dashed) lines delineate the 1σ (2σ) contours, for normal (left) and inverted hierarchy (right).

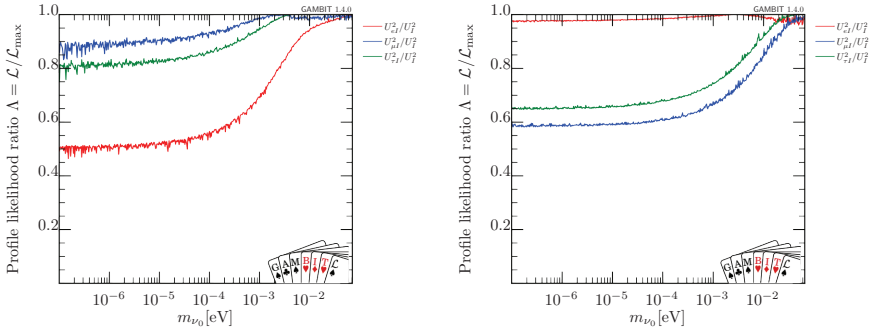


FIGURE 5.13: Upper limits on the coupling ratios $U_{\alpha I}^2/U_I^2$ within 2σ as a function of the lightest left-handed neutrino mass m_{ν_0} , for normal (left) and inverted hierarchy (right).

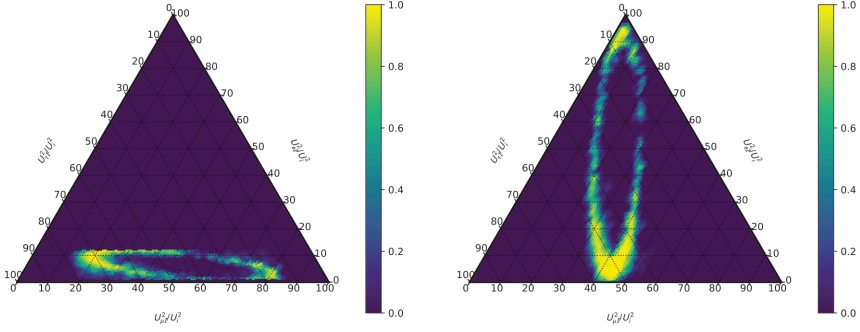


FIGURE 5.14: Profile likelihood for $U_{\alpha I}^2 / U_I^2$ (in percent) in the limit of $n = 2$ in the symmetry protected region for normal (left) and inverted (right) hierarchy. The color scale, or values in the color bar, denote $\mathcal{L} / \mathcal{L}_{\max}$, i.e. the ratio of the likelihood at a point in the space with respect to the maximum likelihood. For details of the cuts performed, we refer the reader to the text.

by two linearly independent RHN flavour eigenstates, leading to constraining relations among the flavour mixings U_{eI}^2 , $U_{\mu I}^2$ and $U_{\tau I}^2$; since not all ratios will be allowed, this should be reflected in a triangle plot as a well-defined region. As m_{ν_0} is allowed to take larger values, the increased freedom should lead to this region becoming bigger, until some limiting value, beyond which the entire triangle is filled (implying very weak relations among the three).

Here, for the first time, triangle plots are constructed for the type-I seesaw model with three right-handed neutrinos. The plots in figure 5.12 are constructed in the following way: for each limit on the mass of the lightest neutrino, those values of the flavour mixings that are within 1σ and 2σ of the global best-fit point are tabulated and their ratios with respect to the total mixing are calculated. These ratios are then plotted in a ternary map. From these plots, we indeed notice the expected behaviour outlined in the previous paragraph: as m_{ν_0} is allowed to take larger values, larger portions of the parameter space is filled. This can be also seen in figure 5.13, where the largest coupling ratio is plotted for each flavour as a function of the lightest LHN mass. In particular, from figure 5.12, there is no visible upper limit on $U_{\mu I}^2 / U_I^2$ or $U_{\tau I}^2 / U_I^2$ for normal hierarchy, while U_{eI}^2 / U_I^2 is constrained to $\lesssim 0.95$. Conversely, for inverted hierarchy there is an upper limit for the μ and τ flavours, but none for the e flavour.

A subtle detail deserves to be mentioned: when constraining m_{ν_0} to very small values, one right-handed neutrino almost decouples. The contribution

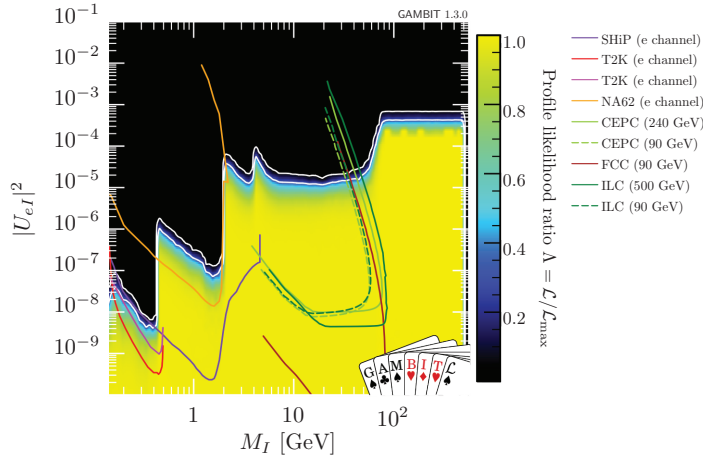
of this (weakly coupled) state to the generation of light neutrino masses is negligible, implying that its properties, such as its flavour mixing pattern, are almost unconstrained by neutrino oscillation data. Due to this, large ratios $U_{\alpha I}^2/U_I^2$ can occur for this particular RHN, although the absolute values of U_I^2 remain negligible, and it has no effect on any near-future experiment. Given that our focus is primarily on right-handed neutrinos that contribute measurably to the generation of left-handed neutrino masses and/or may be discovered in experiments, we applied a cut on $M_I U_I^2 > 10^{-10}$ GeV in figures 5.12 and 5.13 to remove artefacts arising from states that are practically decoupled. The value chosen for the cut is motivated by experimental sensitivities. The NA62 experiment, for example, will only be able to probe up to sensitivities of $M_I U_I^2 \approx \mathcal{O}(10^{-8})$ [110]; under optimistic assumptions, the FCC may test $M_I U_I^2 \approx \mathcal{O}(10^{-11})$ [174] and the LHC, $M_I U_I^2 \approx \mathcal{O}(10^{-8})$ [332].

Triangle plots for the type-I seesaw model with two RHNs have been made before, see [110], for example and it is useful to compare our results with them. For this purpose, however, it is not sufficient to simply insert very small values for m_{ν_0} in the parameterization (2.83), because such values can also be achieved due to accidental cancellations in the left-handed neutrino mass matrix, without the decoupling of any of the right-handed neutrinos; refer to subsection 2.4.1 for a detailed discussion. As mentioned there, in order to remove such fine-tuned points, we impose the cut given in equation 2.100, which is reproduced here:

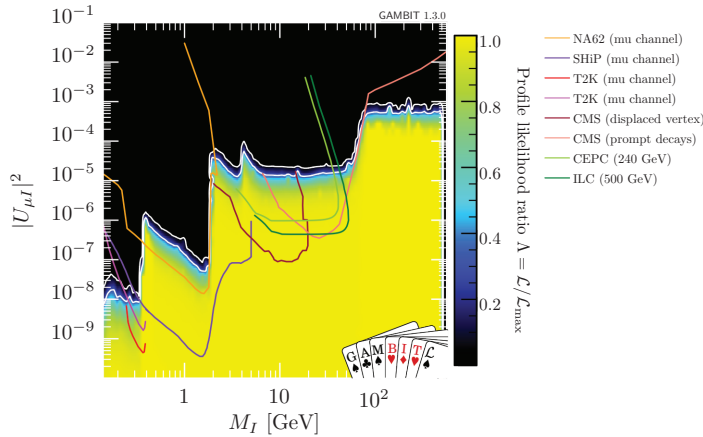
$$\begin{aligned} \frac{|M_2 - M_1|}{M_2 + M_1} &< \epsilon, \quad \frac{m_{\nu_0}}{\mu\text{eV}} < 1, \\ |F_{\alpha 3}| &< \epsilon, \quad \frac{|F_{\alpha 1} + iF_{\alpha 2}|}{|F_{\alpha 1}| + |F_{\alpha 2}|} < \epsilon. \end{aligned} \quad (5.3)$$

Here, ϵ is an arbitrarily small number, which we choose to be $\epsilon = 0.01$ for convenience. In addition, we work in the limits $|\text{Im}\omega_{23}| \gg 1$ and $\text{Re}\omega_{13} \sim \pi/2$ for normal hierarchy, and $|\text{Im}\omega_{12}| \gg 1$ for inverted hierarchy, as detailed in subsection 2.4.1. Since we randomized the order of the matrices \mathcal{R}^{ij} while performing the scan, we can only reproduce the true symmetry-protected regime for the permutation $\mathcal{R} = \mathcal{R}^{23}\mathcal{R}^{13}\mathcal{R}^{12}$. The inverted hierarchy limit is independent of permutations as two of the ω_{ij} are zero. In figure 5.14, we present triangle plots for NH and IH in the symmetry-protected region after applying the aforementioned cuts to remove fine-tuned points. The results are consistent with what was found in [110] for $n = 2$ RHNs.

5.6 Projected upper limits on parameter space



(A) Profile likelihood in $M_I - U_{eI}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlayed.



(B) Profile likelihood in $M_I - U_{\mu I}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlayed.

FIGURE 5.15: Future experimental sensitivities in the electron and muon flavour mixing planes.

There are a number of experiments planned in the next few years that will search for right-handed neutrinos. This may through studying meson decays, as SHiP [168] and T2K [169] will pursue, beam dump experiments like NA62 [110] or colliders using signatures such as displaced vertices, like ILC [174], CMS [174], CEPC [330] and FCC [333].

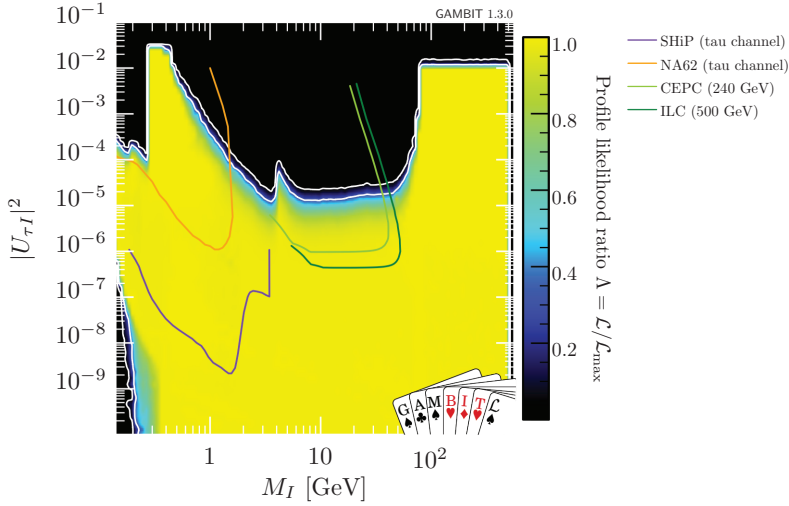


FIGURE 5.16: Profile likelihood in $M_I - U_{\tau I}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlaid.

Figures 5.15a, 5.15b and 5.16 show how the existing allowed parameter region that comes out our scan will be affected/tested by these future experimental efforts.

These experiments will probe the currently allowed parameter space below ~ 100 GeV and are expected to place stronger limits. In particular, between 300 MeV and 5 GeV, SHiP is projected to improve limits by three orders of magnitude for the electron and muon flavour couplings and even more for the tau flavour coupling. In view of the regions with excess likelihood, SHiP and NA62 can explore the one due to R_K , which showed up in figure 5.9b of the $M_1 - U_{e1}^2$ plane, while NA62, CEPC and ILC will test the excesses at low RHN masses in the $M_1 - U_{\tau 1}^2$ plane.

5.7 Comparison with earlier work

One of the highlights of this study is the construction of fully explored profile likelihoods in the RHN mixing-mass plane(s). Many groups and individuals have examined the parameter space of the type-I seesaw model under the constraint of multiple observables; some have focused on a single extra sterile neutrino (eg. [88]), or two (eg. [6, 57, 84]); studies with three RHNs have

included only some of the constraints (eg. [80, 88, 253]) that are in this thesis and/or employed hard cuts to decide on allowed points (eg. [59]). Some have also encountered difficulty in exploring the $n = 3$ parameter space because of its complexity and computation restrictions [59, 111]. The ($n = 3$) profile likelihoods presented in section 5.2, in contrast, saturate the upper limits in the entire mass range of $0.1 < M_I < 500$ GeV.

Aside from the direct detection experiments, which almost completely comprise the upper limits on the parameter space below 10 GeV (see the overlay plots in section 5.2), the most relevant indirect observables that act to restrict the parameter space are the fits to CKM matrix elements, considering CKM unitarity, the invisible decay width of the Z boson and the (sine of the) Weinberg angle. Comparing our results with those presented in [80] corroborates what is seen: non-zero RHN-LHN mixing worsens the fit (compared to the SM-predicted value) of the sine of the Weinberg angle, while the invisible decay width of the Z prefers non-zero mixings.

The preference for non-zero coupling by the latter observable was also observed in [253]. Additionally, it was observed in that study that the presence of RHN-LHN mixing could help ease the tension between observed and (SM-)predicted values for some of the branching ratios of W boson decays (section 4.2.1). This comes with a caveat, though, since the non-zero mixing which would help with one decay ($W \rightarrow e\nu$) worsens the tension in another ($W \rightarrow \tau\nu$). In our analysis, this tussle leads to this observable being “neutral” with respect to non-zero RHN mixing; neither a preference or aversion is shown.

It should be noted that although some papers such as [304] present bounds coming from electroweak precision observables that are stronger than the limits seen in our results, this does not represent a physical effect. The action of the observables that cause the excess likelihoods (fits to CKM matrix elements, Γ_{inv} , R_τ) leads to these constraints getting circumvented.

The triangle plots shown in section 5.5 are the first of their kind for the type-I seesaw model with three RHNs. As observed there, with an increase in the allowed values of the lightest active neutrino mass, the relations among the three flavour mixings becomes weaker, allowing larger and larger regions of the triangle to be filled, until about 10 meV, which is approximately the value beyond which the triangle will always be filled. We were also successful in reproducing the results shown in [110], for the model with two RHNs. Our

results seem to imply that if an experiment like NA62, for example, detects flavour mixings in ratios outside the boundaries set in [110], the lightest active neutrino should have masses larger than 0.01 meV.

Chapter 6

Conclusion and outlook

Right-handed neutrinos (RHNs) are theoretically well-motivated and well-studied candidates in the realm of BSM physics. Extending the Standard Model with these particles allows for the solution of many unanswered questions. Foremost among them is the mystery of the origin of neutrino masses and neutrino oscillations; explanations are offered within a range of models containing RHNs. One of these, which is among the oldest and hence most examined is the type-I seesaw model [8, 22, 23], a variant of the seesaw mechanism. They can also explain the baryon asymmetry of the universe via leptogenesis (see [40] and [41], for example) and are also postulated as dark matter candidates [7, 46].

The parameter space of these particles, especially in the mass-mixing plane, is hence very useful to study. Constraints on the properties of RHNs can be derived from a multitude of observables, ranging from efforts to directly detect them in beam dump and peak search experiments [69, 73, 252, 277] to examining the footprint they indirectly leave behind in electroweak precision observables [59, 80, 253], lepton flavour violating processes [78] and their contribution to neutrinoless double-beta decay [320].

Given the rich phenomenology of RHNs and their wide-ranging effects, a rigorous study of their parameter space must consider the combined influence of constraints. There have been multiple prior studies that have attempted this, combining different direct detection and cosmological results [57], lepton number violating processes [83], lepton universality observables [85] and combinations of these and more [59, 87, 88, 329]. These have involved models with a single, two and three right-handed neutrinos.

This thesis represents an effort to extend previously obtained results and provide a solid base for future work. It is also the first comprehensive frequentist analysis of the right-handed neutrino parameter space containing three RHNs, for which GAMBIT [93], an open-source global fitting software framework has been used. We included all of the most constraining observables that affect this space: experiments which place the most stringent upper limits on the mixing strength, constraints coming from neutrino oscillation data, electroweak precision observables, lepton flavour violating processes, lepton universality observables, cosmological limits from Big Bang nucleosynthesis and the unitarity of the CKM matrix. In order to scan the parameter space, we employed the differential evolution-based algorithm Diver and parameterized the model using the Casas-Ibarra scheme.

The profile likelihood contours that were obtained have been presented in section 5.2 of chapter 5. Constructing them was not trivial whatsoever; the parameter space of the type-I seesaw model contains fine-tuned regions and is complex as it is. Performing a likelihood-based analysis as opposed to employing hard cuts on the parameters to reflect existing constraints complicates this further. Coupled with the fact that RHNs can relieve the tension between observed and SM-predicted values in some observables, this led to us coming up with new and efficient procedures to sample the parameter space effectively, as described in section 5.1. This included the introduction of an extra, auxiliary likelihood to push selected flavour couplings to high values and dividing the full scan into subscans covering smaller mass ranges.

Nevertheless, we were successful in mapping the profile likelihood in the M_I - $U_{\alpha I}^2$ planes. Our results show that below the W boson mass, the upper limits on the RHN-LHN mixing comes from direct detection experiments, as seen in figures 5.3 and 5.4. At higher masses, for the electron and muon flavour mixings, the EWPO $\sin^2 \theta_W$ (the sine of the Weinberg angle) provides the upper limit. For the tau flavour mixing, another EWPO Γ_{inv} (the invisible decay width of the Z boson) plays this role; in fact, this observable also limits $U_{\tau I}^2$ at masses between 0.3-0.5 GeV, the region between the end of the regime of CHARM's limits (according to the re-interpretation described in section 4.3.3 of chapter 4) and that of DELPHI (section 4.3.4). A key point that helped in reaching and exploring the upper limits was also the B-L settings that were used in the scans, since being in this regime allows one to obtain the largest mixings that are allowed. Additionally, we derived upper limits on the three flavour mixings for both mass hierarchies in the high RHN

mass regime (table 5.5); for both hierarchies, U_{eI}^2 and $U_{\mu I}^2$ can reach mixing of $\mathcal{O}(10^{-4})$, while $U_{\tau I}^2$ is allowed to take even higher values of $\mathcal{O}(10^{-2})$, since there are much fewer constraints on this flavour mixing. The B-L settings, along with accidental cancellations in the active neutrino mass matrix, also play a role in overcoming the strong bounds on U_{eI}^2 that come from neutrinoless double-beta decay. Indeed, we see that our limits on U_{eI}^2 are far weaker than those seen in [139], for example, where the B-L regime was not explored.

On the other hand, a combination of BBN and oscillation data provide the lower limit on the total mixing of an RHN to all flavours (U_I^2). We also examined the behavior of the lower limit as a function of the lightest active neutrino mass m_{ν_0} in section 5.3 of chapter 5, observing that the limit gets weaker with higher m_{ν_0} .

The causes of the regions of excess likelihood were then studied in section 5.4, where it was seen that the limiting behaviour of Γ_{inv} counters the preference for high tau flavour couplings coming from the CKM unitarity constraint, hence providing the upper limits in these regions. A similar tussle was seen in the electron flavour mixing, where the preference for large couplings due to the lepton universality observable R_K is countered by the bound coming from CHARM. Although they are an intriguing feature, the fact that the level of significance of these excesses is not high leads us to conclude that they are not indicative of a detection of right-handed neutrinos. Additionally, this is not the first time they have been observed; the behavior in both cases was observed in earlier studies [80, 253].

This study also resulted in the construction of triangle plots (plots that display the allowed ratios of flavour-to-total mixing) for the model with three RHNs. These are discussed in section 5.5 of chapter 5. In contrast to the $n = 2$ case, where there are strongly constraining relations among the three flavour mixings, we observe the possibility of these relations becoming very weak. This happens when the lightest active neutrino is allowed to be “heavy”, i.e. $m_{\nu_0} > 10$ meV. By carefully distinguishing between symmetry-protected and fine-tuned parameter values, we were also able reproduce the results for the $n = 2$ case, as presented in [110].

Finally, the effect of upcoming experimental investigations on the parameter space was looked at in section 5.6. Large parts of the allowed space below ~ 100 GeV from our analysis will be probed in the future, as seen in figures 5.15 and 5.16. The potential exploration of the regions of excess at low

RHN masses by SHiP, NA62, CEPC and ILC is of particular interest.

As mentioned at the beginning of this chapter, one of the purposes of this work was to lay the ground for future efforts. By working out the difficulties in sampling the parameter space of this model in the frequentist context, we have laid the basis for robust and faster Bayesian scans over specific implementations of the type-I seesaw model.

The treatment of Big Bang nucleosynthesis offers one window for improvement. This can be done on multiple fronts: by taking into account the color factor of 3 that should enter the decay width of RHNs to jets (see appendix B), including the relevant neutral current coefficients for high RHN masses (absent in [6, 323], on which our analysis is based) and formulating the constraint as a likelihood and not a hard cut. Since BBN plays a role in determining the lower limit of U_1^2 , coupled with improved sampling of the parameter space, one could fully map this limit over the entire mass range. In the next few years, limits on the active neutrino mass(es) will likely to be revised. KATRIN is expected to place a limit on the mass of the lightest left-handed neutrino of about 0.2 eV [167], for example. Recent results from the PLANCK collaboration have already limited the sum of the masses of the left-handed neutrinos to be $\sum m_\nu < 0.12$ eV [146] using a combination of CMB, lensing and BAO data. Coupled with the lower limit on this sum for the inverted mass hierarchy, $\sum m_\nu > 0.1$ eV, if the former upper limit is pushed down to ~ 0.1 eV in the future, we may be able to answer the question of which mass hierarchy the left-handed neutrinos follow.

Our utilization of the Casas-Ibarra parameterization also placed some restrictions on the analysis: as outlined in section 2.3 of chapter 2, it is extremely hard to formulate well-motivated but transparent priors on the parameters when using this scheme. The flexible nature of GAMBIT will allow the study of a much larger range of theoretical models in a Bayesian context, which would be a very interesting avenue to pursue.

Appendix A

Decay widths and form factors for LFV observables

The decay widths of LFV processes, as described in section 4.2.2, are listed here; they are taken from [308, 334].

$$\Gamma_{l_\alpha^- \rightarrow l_\beta^- \gamma} = \frac{\alpha_{\text{em}} m_{l_\alpha}^5}{4} \left(|K_2^L|^2 + |K_2^R|^2 \right), \quad (\text{A.1})$$

$$\begin{aligned} \Gamma_{l_\alpha^- \rightarrow l_\beta^- l_\beta^+ l_\beta^+} &= \frac{m_{l_\alpha}^5}{512\pi^3} \left[e^4 |K_2^L|^2 \left(\frac{16}{3} \log \frac{m_{l_\alpha}}{m_{l_\beta}} - \frac{22}{3} \right) \right. \\ &\quad + \frac{1}{24} (|A_{LL}^S|^2 + 2|A_{LR}^S|^2) + \frac{1}{3} (2|\hat{A}_{LL}^V|^2 + |\hat{A}_{LR}^V|^2) \\ &\quad + \frac{e^2}{3} (K_2^L (A_{RL}^{S*} - 2\hat{A}_{RL}^{V*} - 4\hat{A}_{RR}^{V*}) + \text{h.c.}) \\ &\quad \left. - \frac{1}{6} (A_{LR}^S \hat{A}_{LR}^{V*} + \text{h.c.}) \right] + (L \leftrightarrow R), \quad (\text{A.2}) \end{aligned}$$

$$\begin{aligned} \Gamma_{l_\alpha^- \rightarrow l_\beta^- l_\gamma^+ l_\gamma^+} &= \frac{m_{l_\alpha}^5}{512\pi^3} \left[e^4 |K_2^L|^2 \left(\frac{16}{3} \log \frac{m_{l_\alpha}}{m_{l_\gamma}} - 8 \right) \right. \\ &\quad + \frac{1}{12} (|A_{LL}^S|^2 + |A_{LR}^S|^2) + \frac{1}{3} (|\hat{A}_{LL}^V|^2 + |\hat{A}_{LR}^V|^2) \\ &\quad \left. + \frac{2e^2}{3} (K_2^L (\hat{A}_{RL}^{V*} + \hat{A}_{RR}^{V*}) + \text{h.c.}) \right] + (L \leftrightarrow R), \quad (\text{A.3}) \end{aligned}$$

$$\Gamma_{l_\alpha^- \rightarrow l_\gamma^- l_\beta^- l_\beta^+} = \frac{m_{l_\alpha}^5}{512\pi^3} \left[\frac{1}{24} (|A_{LL}^S|^2 + 2|A_{LR}^S|^2) + \frac{1}{3} (2|\hat{A}_{LL}^V|^2 + |\hat{A}_{LR}^V|^2) - \frac{1}{6} (A_{LR}^S \hat{A}_{LR}^{V*} + \text{h.c.}) \right] + (L \leftrightarrow R), \quad (\text{A.4})$$

where $\hat{A}_{XY}^V \equiv A_{XY}^V + e^2 K_1^X$. The couplings e , g_1 & g_2 correspond to the electromagnetic, hypercharge and weak couplings of the Standard Model.

The form factors K_1^X , K_2^X , A_{XY}^S and A_{XY}^V are taken in the flavour basis where the charged lepton mass matrix is diagonal. The dipole form factors K_1^X and K_2^X are given by [308]

$$K_1^L = 0, \quad (\text{A.5})$$

$$K_1^R = \frac{G_F}{4\sqrt{2}\pi^2} \sum_{a=1}^6 \Theta_{\alpha a} \Theta_{\beta a}^* M \left(\frac{m_{\nu_a}^2}{m_W^2} \right), \quad (\text{A.6})$$

$$K_2^L = \frac{G_F}{4\sqrt{2}\pi^2} \frac{m_{l_\beta}}{m_{l_\alpha}} \sum_{a=1}^6 \Theta_{\alpha a} \Theta_{\beta a}^* G \left(\frac{m_{\nu_a}^2}{m_W^2} \right) \quad \text{and} \quad (\text{A.7})$$

$$K_2^R = \frac{G_F}{4\sqrt{2}\pi^2} \sum_{a=1}^6 \Theta_{\alpha a} \Theta_{\beta a}^* G \left(\frac{m_{\nu_a}^2}{m_W^2} \right). \quad (\text{A.8})$$

The four lepton form factors A_{XY}^V and A_{XY}^S corresponding to the process $l_\alpha^- \rightarrow l_\beta^- l_\gamma^- l_\delta^+$, with a vector or scalar mediator respectively, are [308]

$$\begin{aligned} A_{LL}^V = & \left[\frac{g_2^2}{32\pi^2} \frac{g_{-\delta\gamma\delta} \Theta_{\alpha a} \Theta_{\beta a}^*}{2m_Z^2} \left(g_+ x_a C_0(x_a, x_a) + g_2 c_W (1 - 2(B_0(1) + 2C_{00}(x_a, 1) \right. \right. \\ & \left. \left. + x_a C_0(x_a, 1)) \right) - \frac{g_+}{2} (1 - 2(B_0(x_a) - 2C_{00}(x_a, x_a) + C_0(x_a, x_a))) \right) \\ & + \frac{g_-^2 (\Theta_{\alpha a} \Theta_{\beta a}^* m_{l_\alpha}^2 - \Theta_{\beta a} \Theta_{\alpha a}^* m_{l_\beta}^2) \delta_{\gamma\delta}}{4m_Z^2 (m_{l_\alpha}^2 - m_{l_\beta}^2)} \times (1 + 2B_1(x_a)) \Big]_{\text{penguin}} \\ & + \left[\frac{g_2^4}{32\pi^2 m_W^2} \Theta_{\alpha a} \left(\Theta_{\gamma a} \Theta_{\beta c}^* \Theta_{\delta c}^* x_a x_c (D_0(x_a, x_c) + (a \leftrightarrow c)) \right. \right. \\ & + 2\Theta_{\gamma c} (\Theta_{\beta a}^* \Theta_{\delta c}^* (C_0(x_c, x_a) + D_0(x_c, x_a)) + (a \leftrightarrow c)) \\ & \left. \left. - 6\Theta_{\gamma c} (\Theta_{\beta a}^* \Theta_{\delta c}^* + \Theta_{\beta c}^* \Theta_{\delta a}^*) D_{27}(x_a, x_c) \right) \right]_{\text{box}}, \quad (\text{A.9}) \end{aligned}$$

$$A_{LR}^V = \frac{2g_1 s_W}{g_-} [A_{LL}^V]_{\text{penguin}}, \quad (\text{A.10})$$

$$A_{RL}^V = \frac{\delta_\gamma \delta g_2^2 g_- g_1 s_W}{64\pi^2 m_Z^2} \frac{\Theta_{\alpha a} \Theta_{\beta a}^* m_{l_\alpha}^2 - \Theta_{\beta a} \Theta_{\alpha a}^* m_{l_\beta}^2}{m_{l_\alpha}^2 - m_{l_\beta}^2} (1 + 2B_1(x_a)), \quad (\text{A.11})$$

$$A_{RR}^V = \frac{2g_1 s_W}{g_-} A_{RL}^V, \quad (\text{A.12})$$

$$A_{LY}^S = \frac{Y_\gamma^l \delta_\gamma \delta g_2^2}{64\pi^2 m_h^2} \frac{Y_\alpha^l \Theta_{\beta a} \Theta_{\alpha a}^* m_{l_\beta}^2 - Y_\beta^l \Theta_{\alpha a} \Theta_{\beta a}^* m_{l_\alpha}^2}{m_{l_\alpha}^2 - m_{l_\beta}^2} (1 + 2B_1(x_a)) \quad \text{and} \quad (\text{A.13})$$

$$A_{RY}^S = \frac{Y_\gamma^l \delta_\gamma \delta g_2^2}{64\pi^2 m_h^2} \frac{Y_\alpha^l \Theta_{\beta a} \Theta_{\alpha a}^* - Y_\beta^l \Theta_{\alpha a} \Theta_{\beta a}^*}{m_{l_\alpha}^2 - m_{l_\beta}^2} m_{l_\alpha} m_{l_\beta} (1 + 2B_1(x_a)), \quad (\text{A.14})$$

where a summation over a and c is assumed, $x_a = m_{\nu_a}^2 / m_W^2$, $s_W = \sin \theta_W$ and $g_\pm = g_1 \sin \theta_W \pm g_2 \cos \theta_W$.

The $\mu - e$ conversion ratio described in section 4.2.2, for a general nucleus, can be written as

$$R_{\mu-e} = \frac{\alpha_{\text{em}}^3 m_\mu^5 Z_{\text{eff}}^4 F_p^2}{4\pi^4 Z \Gamma_{\text{capt}}} \left| \sum_{q=u,d,s} \left\{ \left(e^2 Q_q (K_1^L - K_2^R) - \frac{1}{2} (B_{LL}^V + B_{LR}^V) \right) \left(ZG_V^{(q,p)} + NG_V^{(q,n)} \right) - \frac{1}{2} (B_{LL}^S + B_{LR}^S) \left(ZG_S^{(q,p)} + NG_S^{(q,n)} \right) \right\} \right|^2 + (L \leftrightarrow R),$$

where $B_{XY}^K \leftrightarrow C_{XY}^K$ for up-type quarks and the numerical factors G_K are given in [335]. The nuclear form factor F_p , the effective atomic number Z_{eff} and the capture rate Γ_{capt} of the nuclei [106] studied (^{48}Ti and ^{208}Pb) are in table A.1.

Nucleus	Z_{eff}	F_p	$\Gamma_{\text{capt}} (\times 10^6 \text{ s}^{-1})$
$^{48}_{22}\text{Ti}$	17.6	0.54	2.59
$^{208}_{82}\text{Pb}$	34	0.15	13.45

TABLE A.1: Effective atomic number, nuclear form factor and capture rate for the relevant nuclei.

The form factors K_1^X, K_2^X are defined above and B_{XY}^K and C_{XY}^K are given by

$$B_{LL}^V = -\frac{1}{3} \frac{g_d}{g_-} [A_{LL}^V]_{\text{penguin}} + \frac{g_2^4 \Theta_{\alpha\alpha} \Theta_{\beta\alpha}^* V_{\gamma c} V_{\delta c}}{16\pi^2 m_W^2} \times \\ \times (C_0(x_a, x_c^u) + D_0(x_a, x_c^u) - 3D_{27}(x_a, x_c^u)), \quad (\text{A.15})$$

$$B_{RL}^V = -\frac{1}{3} \frac{g_d}{g_-} A_{RL}^V, \quad (\text{A.16})$$

$$B_{XR}^V = \frac{1}{3} A_{XR}^V, \quad (\text{A.17})$$

$$B_{XY}^S = \frac{Y_\gamma^d}{Y_\gamma^l} A_{XY}^S, \quad (\text{A.18})$$

$$C_{LL}^V = -\frac{1}{3} \frac{g_u}{g_-} [A_{LL}^V]_{\text{penguin}} + \frac{g_2^4 \Theta_{\alpha\alpha} \Theta_{\beta\alpha}^* V_{\gamma c} V_{\delta c}}{4\pi^2 m_W^2} D_{27}(x_a, x_c^d), \quad (\text{A.19})$$

$$C_{RL}^V = -\frac{1}{3} \frac{g_u}{g_-} A_{RL}^V, \quad (\text{A.20})$$

$$C_{XR}^V = -\frac{2}{3} A_{XR}^V \quad \text{and} \quad (\text{A.21})$$

$$C_{XY}^S = \frac{Y_\gamma^u}{Y_\gamma^l} A_{XY}^S, \quad (\text{A.22})$$

with $g_d = 3g_2 \cos \theta_W + g_1 \sin \theta_W$, $g_u = -3g_2 \cos \theta_W + g_1 \sin \theta_W$, $x_c^{(u,d)} = m_{(u,d)_c}^2 / m_W^2$ and V_{ij} is the relevant CKM matrix element.

Lastly, the loop functions used in equations A.6-A.22 are defined by [308]

$$G(x) = \frac{-7 + 33x - 57x^2 + 31x^3 + 6x^2(1 - 3x)\log(x)}{12(-1 + x)^4}, \quad (\text{A.23})$$

$$M(x) = \frac{6x^2(x - 3)\log(x) - (x - 1)(5x^2 - 22x + 5)}{9(x - 1)^4}, \quad (\text{A.24})$$

$$B_0(x) = 0.252183 - \log x, \quad (\text{A.25})$$

$$B_1(x) = \frac{-1 + 4x - 3x^2 + 0.504365(x - 1)^2 + 2x^2\log(x)}{4(x - 1)^2}, \quad (\text{A.26})$$

$$C_0(x, y) = \frac{(x - y)\log(x) + (x - 1)y\log(\frac{y}{x})}{(x - 1)(x - y)(y - 1)}, \quad (\text{A.27})$$

$$C_{00}(x, y) = 0.438046 - \frac{(xy - x - y)\log(x)}{4(x - 1)(y - 1)} + \frac{y^2\log(\frac{y}{x})}{4(x - y)(y - 1)}, \quad (\text{A.28})$$

$$D_0(x, y) = \frac{(xy - 1)\log(x)}{(x - 1)^2(y - 1)^2} + \frac{y\log(\frac{y}{x})}{(x - y)(y - 1)^2} - \frac{1}{(x - 1)(y - 1)} \quad \text{and} \quad (\text{A.29})$$

$$D_{27}(x, y) = \frac{(2xy - x - y)\log(x)}{4(x - 1)^2(y - 1)^2} + \frac{y^2\log(\frac{y}{x})}{4(x - y)(y - 1)^2} - \frac{1}{4(x - 1)(y - 1)}. \quad (\text{A.30})$$

Appendix B

Decay widths relevant to Big Bang nucleosynthesis

Listed below are the formulae for the decay widths of the processes included within the Big Bang nucleosynthesis constraint [6, 323].

$$\Gamma_{N_I \rightarrow \pi^0 \nu_\alpha} = \frac{|\Theta_{\alpha I}|^2}{32\pi} G_F^2 f_\pi^2 M_I^3 \left(1 - \frac{M_\pi^2}{M_I^2}\right)^2, \quad (\text{B.1})$$

$$\begin{aligned} \Gamma_{N_I \rightarrow H^+ l_\alpha^-} &= \frac{|\Theta_{\alpha I}|^2}{16\pi} G_F^2 |V_H|^2 f_H^2 M_I^3 \times \\ &\times \left(\left(1 - \frac{M_I^2}{M_I^2}\right)^2 - \frac{M_H^2}{M_I^2} \left(1 + \frac{M_I^2}{M_I^2}\right) \right) \times \\ &\times \sqrt{\left(1 - \frac{(M_H - M_I)^2}{M_I^2}\right) \left(1 - \frac{(M_H + M_I)^2}{M_I^2}\right)}, \quad (\text{B.2}) \end{aligned}$$

$$\Gamma_{N_I \rightarrow \eta \nu_\alpha} = \frac{|\Theta_{\alpha I}|^2}{32\pi} G_F^2 f_\eta^2 M_I^3 \left(1 - \frac{M_\eta^2}{M_I^2}\right)^2, \quad (\text{B.3})$$

$$\Gamma_{N_I \rightarrow \eta' \nu_\alpha} = \frac{|\Theta_{\alpha I}|^2}{32\pi} G_F^2 f_{\eta'}^2 M_I^3 \left(1 - \frac{M_{\eta'}^2}{M_I^2}\right)^2, \quad (\text{B.4})$$

$$\begin{aligned} \Gamma_{N_I \rightarrow \rho^+ l_\alpha^-} &= \frac{|\Theta_{\alpha I}|^2}{8\pi} \frac{g_\rho^2}{M_\rho^2} G_F^2 |V_{ud}|^2 M_I^3 \times \\ &\times \left(\left(1 - \frac{M_I^2}{M_\rho^2} \right)^2 + \frac{M_\rho^2}{M_I^2} \left(1 + \frac{M_I^2 - 2M_\rho^2}{M_I^2} \right) \right) \times \\ &\times \sqrt{\left(1 - \frac{(M_\rho - M_I)^2}{M_I^2} \right) \left(1 - \frac{(M_\rho + M_I)^2}{M_I^2} \right)}, \quad (\text{B.5}) \end{aligned}$$

$$\Gamma_{N_I \rightarrow \rho^0 \nu_\alpha} = \frac{|\Theta_{\alpha I}|^2}{16\pi} \frac{g_\rho^2}{M_\rho^2} G_F^2 M_I^3 \left(1 + 2 \frac{M_\rho^2}{M_I^2} \right) \left(1 - \frac{M_\rho^2}{M_I^2} \right)^2, \quad (\text{B.6})$$

$$\Gamma_{N_I \rightarrow \sum_{\alpha,\beta} \nu_\alpha \bar{\nu}_\beta \nu_\beta} = \frac{G_F^2 M_I^5}{192\pi^3} \sum_\alpha |\Theta_{\alpha I}|^2, \quad (\text{B.7})$$

$$\begin{aligned} \Gamma_{N_I \rightarrow l_{\alpha \neq \beta}^- l_\beta^+ \nu_\beta} &= \frac{G_F^2 M_I^5}{192\pi^3} |\Theta_{\alpha I}|^2 \left(S(x_\alpha, x_\beta) g(x_\alpha, x_\beta) \right. \\ &\quad - 12x_\alpha^4 \log \left[\frac{1 - S(x_\alpha, x_\beta)(1 + x_\alpha^2 - x_\beta^2)}{2x_\alpha^2} \right. \\ &\quad \left. \left. - \frac{2x_\beta^2 + (x_\alpha^2 - x_\beta^2)^2}{2x_\alpha^2} \right] - 12x_\beta^4 \log \left[\frac{1}{2x_\beta^2} \right. \right. \\ &\quad \left. \left. - \frac{S(x_\alpha, x_\beta)(1 - x_\alpha^2 + x_\beta^2) - 2x_\alpha^2 + (x_\alpha^2 - x_\beta^2)^2}{2x_\beta^2} \right] \right. \\ &\quad \left. + 12x_\alpha^4 x_\beta^4 \log \left[\frac{1 - S(x_\alpha, x_\beta)(1 - x_\alpha^2 - x_\beta^2)}{2x_\alpha^2 x_\beta^2} \right. \right. \\ &\quad \left. \left. - \frac{2x_\alpha^2 - 2x_\beta^2 + x_\alpha^4 + x_\beta^4}{2x_\alpha^2 x_\beta^2} \right] \right), \quad (\text{B.8}) \end{aligned}$$

$$\begin{aligned}
\Gamma_{N_I \rightarrow \nu_\alpha l_\beta^+ l_\beta^-} &= \frac{G_F^2 M_I^5}{192\pi^3} |\Theta_{\alpha I}|^2 \left[(C_1(1 - \delta_{\alpha\beta}) + C_3\delta_{\alpha\beta}) \times \right. \\
&\quad \times \left((1 - 14x_l^2 - 2x_l^4 - 12x_l^6) \sqrt{1 - 4x_l^2} + 12x_l^4(x_l^4 - 1)L \right) \\
&\quad + 4(C_2(1\delta_{\alpha\beta}) + C_4\delta_{\alpha\beta}) \left(x_l^2(2 + 10x_l^2 - 12x_l^4) \sqrt{1 - 4x_l^2} \right. \\
&\quad \left. \left. + 6x_l^4(1 - 2x_l^2 + 2x_l^4)L \right) \right], \quad (\text{B.9})
\end{aligned}$$

$$\begin{aligned}
\Gamma_{N_I \rightarrow \nu_\alpha u \bar{u}} &= \frac{G_F^2 M_I^5}{192\pi^3} |\Theta_{\alpha I}|^2 \left(f^u(x_q) S(x_q, x_q) \right. \\
&\quad + x_q^4 \left(3 - \frac{16}{3}C_5x_q^2 + (3 - 8C_5)x_q^4 \right) \times \\
&\quad \left. \times \log \left[\frac{1 - 4x_q^2 + 2x_q^4 + S(x_q, x_q)(1 - 2x_q^2)}{2x_q^4} \right] \right), \quad (\text{B.10})
\end{aligned}$$

$$\begin{aligned}
\Gamma_{N_I \rightarrow \nu_\alpha d \bar{d}} &= \frac{G_F^2 M_I^5}{192\pi^3} |\Theta_{\alpha I}|^2 \left(f^d(x_q) S(x_q, x_q) \right. \\
&\quad + x_q^4 \left(3 - \frac{8}{3}C_6x_q^2 - (1 - \frac{4}{3}C_6)x_q^4 \right) \times \\
&\quad \left. \times \log \left[\frac{1 - 4x_q^2 + 2x_q^4 + S(x_q, x_q)(1 - 2x_q^2)}{2x_q^4} \right] \right), \quad (\text{B.11})
\end{aligned}$$

$$\begin{aligned}
\Gamma_{N_I \rightarrow l_\alpha u_n \bar{d}_m} &= \frac{G_F^2 M_I^5}{192\pi^3} |V_{nm}|^2 |\Theta_{\alpha I}|^2 \left(S(a, b) g(a, b) \times \right. \\
&\quad - 12a^4 \log \left[\frac{1 - S(a, b)(1 + a^2 - b^2) - 2b^2 + (a^2 - b^2)^2}{2a^2} \right] \\
&\quad - 12b^4 \log \left[\frac{1 - S(a, b)(1 - a^2 + b^2) - 2a^2 + (a^2 - b^2)^2}{2b^2} \right] \\
&\quad + 12a^4 b^4 \log \left[\frac{1 - S(a, b)(1 - a^2 - b^2)}{2a^2 b^2} \right. \\
&\quad \left. \left. - \frac{2a^2 - 2b^2 + a^4 + b^4}{2a^2 b^2} \right] \right), \quad (\text{B.12})
\end{aligned}$$

with

$$x_* = \frac{M_{l_*}}{M_I} \text{ in B.8,} \quad (\text{B.13})$$

$$x_l = \frac{M_{l_\beta}}{M_I} \text{ in B.9,} \quad (\text{B.14})$$

$$x_q = \frac{M_q}{M_I} \text{ in B.10 and B.11,} \quad (\text{B.15})$$

$$a, b = \frac{\text{Two heavier masses}}{M_I} \text{ in B.12,} \quad (\text{B.16})$$

$$S(x, y) = \sqrt{(1 - (x + y)^2)(1 - (x - y)^2)}, \quad (\text{B.17})$$

$$g(x, y) = 1 - 7x^2 - 7y^2 - 7x^4 - 7y^4 + 12x^2y^2 - 7x^2y^4 - 7x^4y^2 + x^6 + y^6, \quad (\text{B.18})$$

$$f^u(x) = \frac{1}{4} - \frac{2}{9}C_5 - \left(\frac{7}{2} - \frac{20}{9}C_5\right)x^2 \left(\frac{1}{2} + 4C_5\right)x^4 - (3 - 8C_5)x^6, \quad (\text{B.19})$$

$$f^d(x) = \frac{1}{4} - \frac{1}{9}C_6 - \left(\frac{2}{7} - \frac{10}{9}C_6\right)x^2 \left(\frac{1}{2} + 2C_6\right)x^4 - (3 - 4C_6)x^6, \quad (\text{B.20})$$

$$C_1 = \frac{1}{4}(1 - 4\sin^2\theta_W + 8\sin^4\theta_W), \quad (\text{B.21})$$

$$C_2 = \frac{1}{2}\sin^2\theta_W(2\sin^2\theta_W - 1), \quad (\text{B.22})$$

$$C_3 = \frac{1}{4}(1 + 4\sin^2\theta_W + 8\sin^4\theta_W), \quad (\text{B.23})$$

$$C_4 = \frac{1}{2}\sin^2\theta_W(2\sin^2\theta_W + 1), \quad (\text{B.24})$$

$$C_5 = \sin^2\theta_W(3 - 4\sin^2\theta_W), \quad (\text{B.25})$$

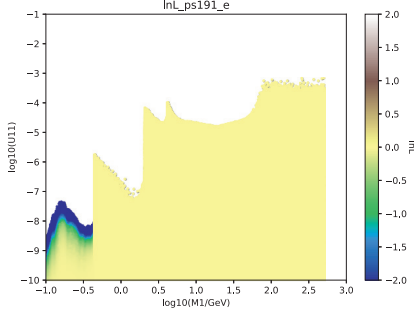
$$C_6 = \sin^2\theta_W(3 - 2\sin^2\theta_W) \quad \text{and} \quad (\text{B.26})$$

$$L = \log \left[\frac{1 - 3x_l^2 - (1 - x_l^2)\sqrt{1 - 4x_l^2}}{x_l^2(1 + \sqrt{1 - 4x_l^2})} \right]. \quad (\text{B.27})$$

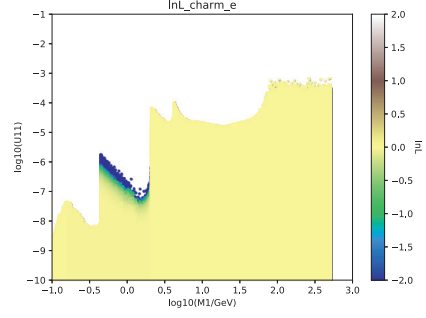
Appendix C

Additional plots for normal hierarchy

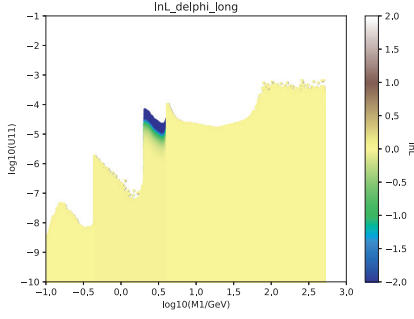
This appendix contains partial likelihood plots for every observable that affects the parameter space (similar to the plots in 5.6 in chapter 5) for the normal hierarchy. The construction of the partial likelihood plots was described in the section referred to above, but is reiterated here: the difference between the value of the partial likelihood and the value of the same partial likelihood at the global best-fit point are calculated; the plot is colour-coded according to this difference, with the parameter points having the largest difference plotted last. So, a value of zero (or the dull yellow colour seen in most of the plot) means that the partial likelihood value at that point is the same as the partial likelihood value at the best-fit point. If the colour shifts to green/blue, the partial likelihood values are less than the value at the best-fit point, implying that this likelihood begins to constrain the parameter space in these regions. A shift in colour to brown conversely implies that the partial likelihood is higher than the corresponding value at the best-fit point. If the plot for an observable is uniformly dull yellow, signifying a uniform likelihood contribution over the entire region, it is not included here.



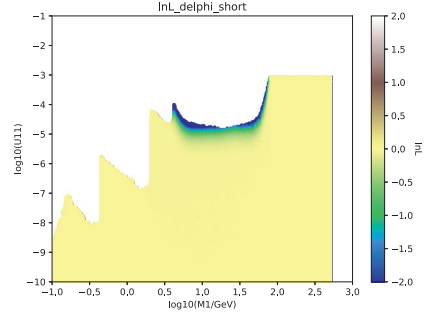
PS191's (electron channel) influence on U_{eI}^2 .



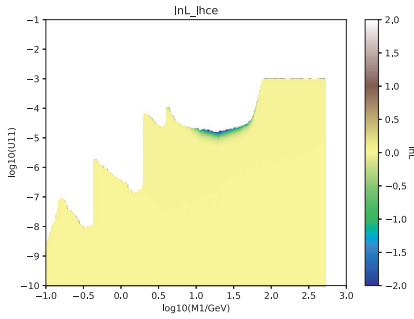
CHARM's (electron channel) influence on U_{eI}^2 .



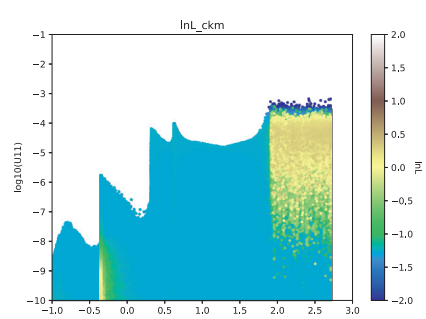
DELPHI's (long-lived RHN analysis) influence on U_{eI}^2 .



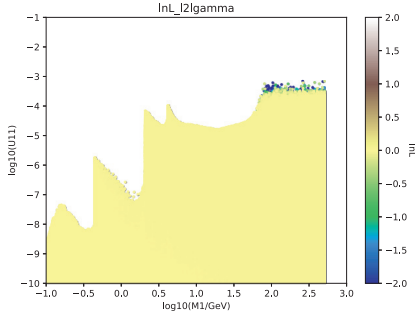
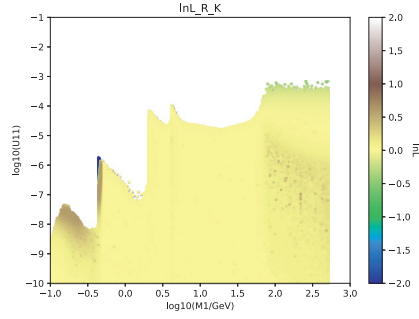
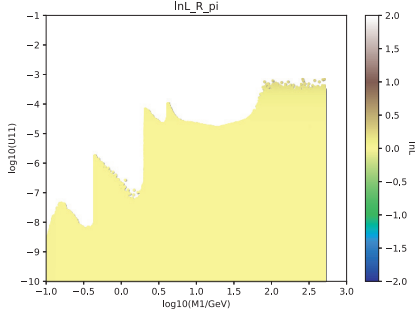
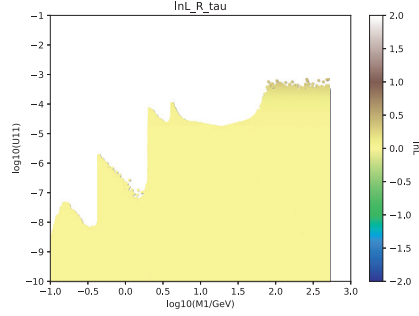
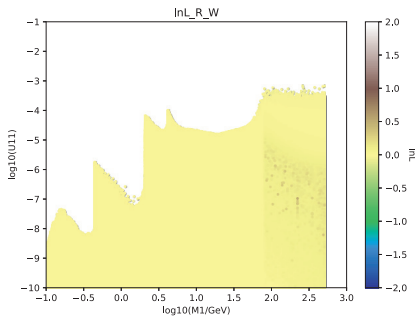
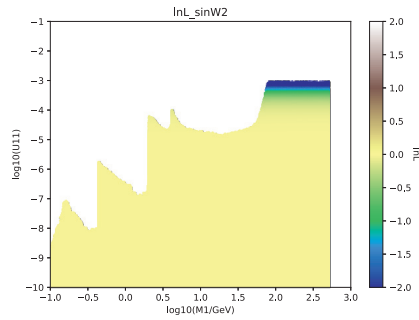
DELPHI's (short-lived RHN analysis) influence on U_{eI}^2 .

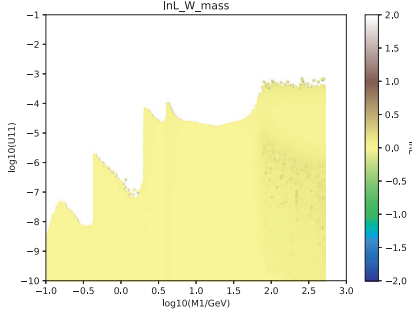


CMS' (electron channel) influence on U_{eI}^2 .

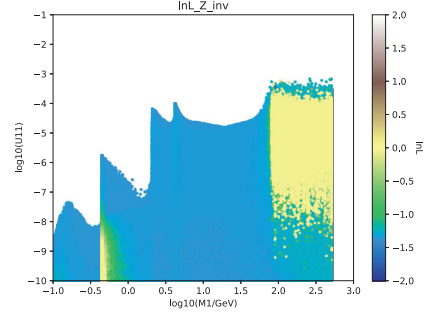


CKM unitarity's influence on U_{eI}^2 .

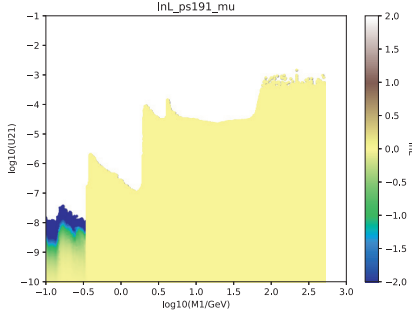

 $\ell \rightarrow \ell\gamma$ processes' influence on U_{eI}^2 .

 R_K 's influence on U_{eI}^2 .

 R_π 's influence on U_{eI}^2 .

 R_τ 's influence on U_{eI}^2 .

 R_W 's influence on U_{eI}^2 .

 $\sin^2 \theta_W$'s influence on U_{eI}^2 .



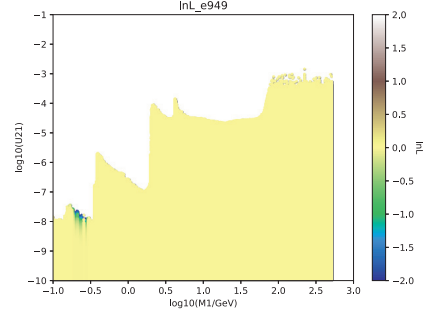
m_W 's influence on U_{eI}^2 .



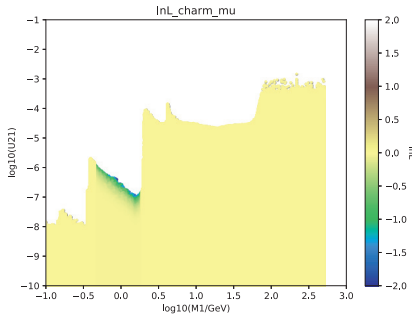
Γ_{inv} 's influence on U_{eI}^2 .



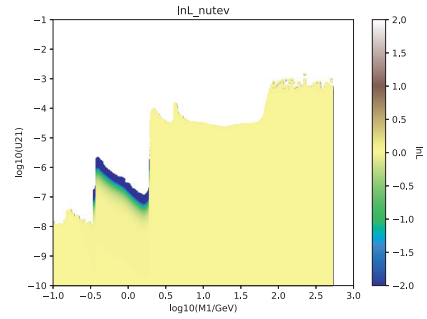
PS191's (muon channel) influence on $U_{\mu I}^2$.



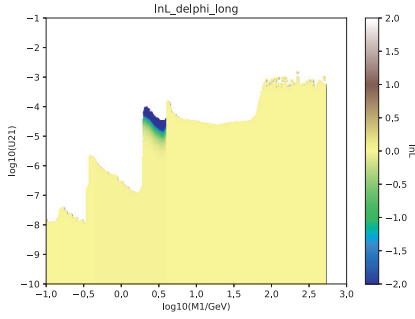
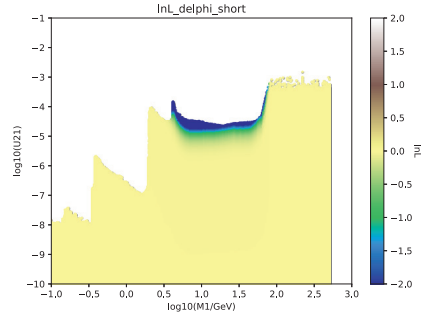
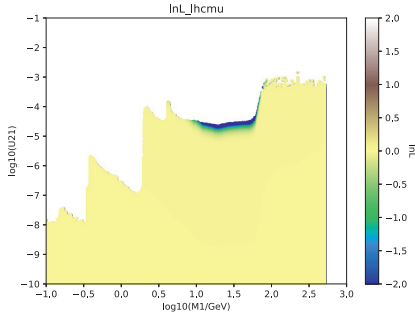
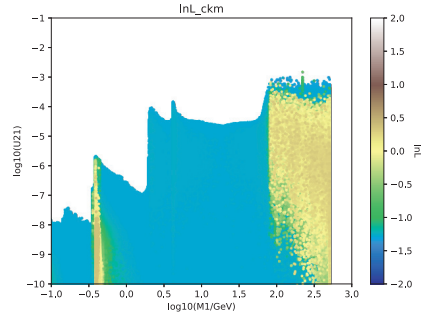
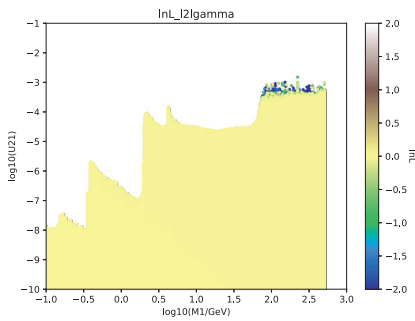
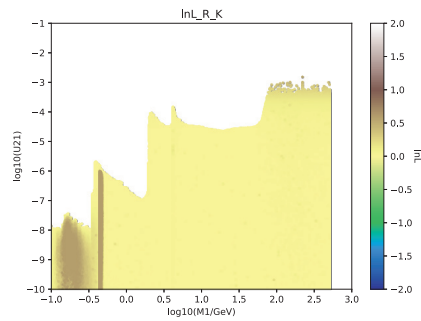
E949's influence on $U_{\mu I}^2$.

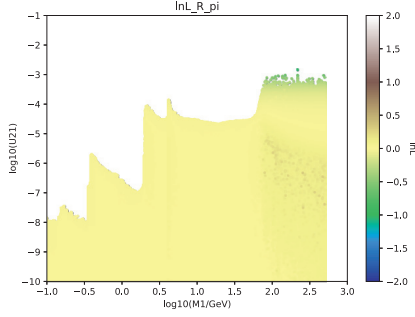
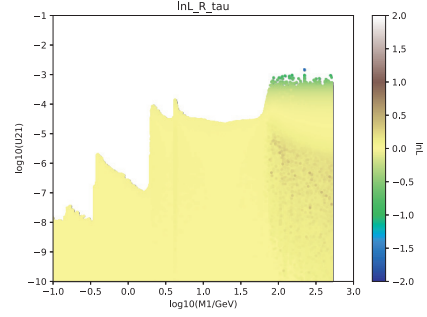
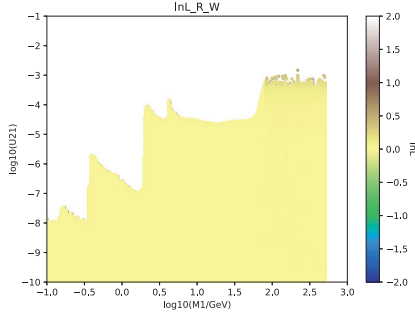
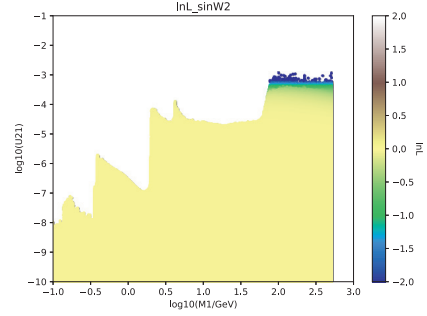
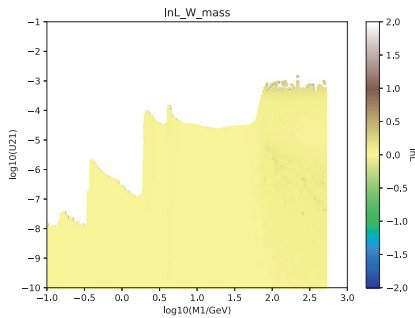
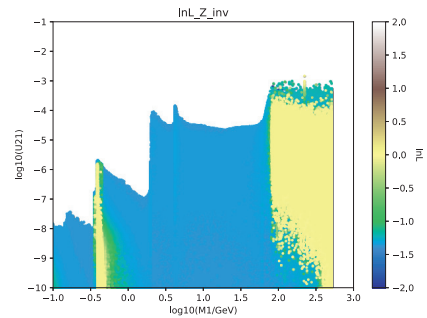


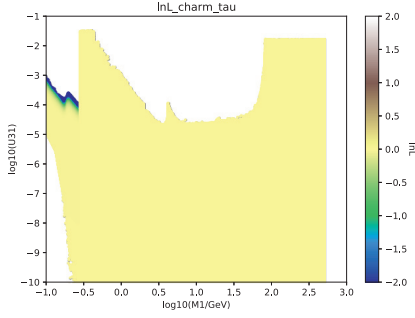
CHARM's (muon channel) influence on $U_{\mu I}^2$.



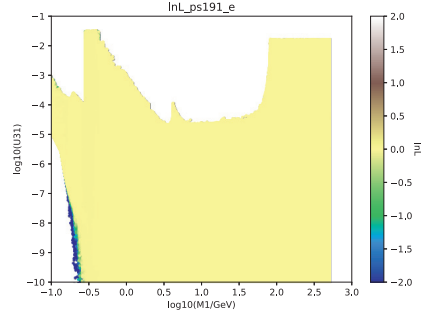
NuTeV's influence on $U_{\mu I}^2$.

DELPHI's (long-lived RHN analysis) influence on $U_{\mu I}^2$.DELPHI's (short-lived RHN analysis) influence on $U_{\mu I}^2$.CMS' (muon channel) influence on $U_{\mu I}^2$.CKM unitarity's influence on $U_{\mu I}^2$. $\ell \rightarrow \ell \gamma$ processes' influence on $U_{\mu I}^2$. R_K 's influence on $U_{\mu I}^2$.

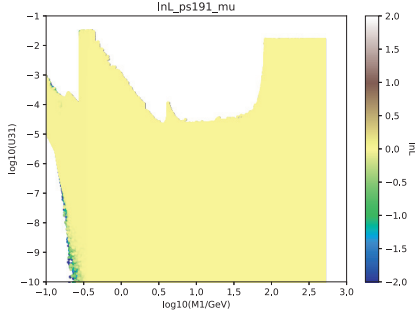
 R_π 's influence on $U_{\mu I}^2$. R_τ 's influence on $U_{\mu I}^2$. R_W 's influence on $U_{\mu I}^2$. $\sin^2 \theta_W$'s influence on $U_{\mu I}^2$. m_W 's influence on $U_{\mu I}^2$. Γ_{inv} 's influence on $U_{\mu I}^2$.



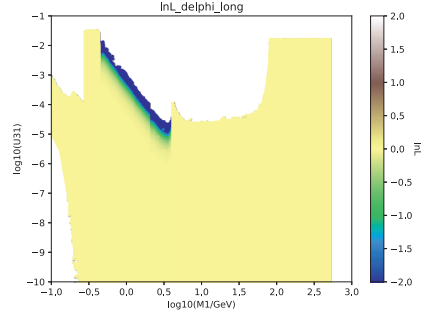
CHARM's (tau channel) influence on $U_{\tau I}^2$.



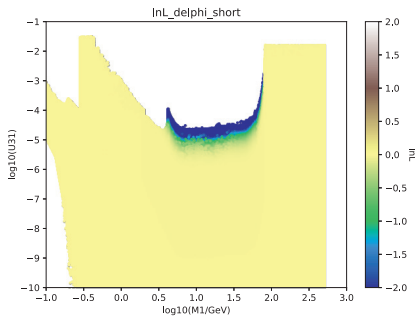
PS191's (electron channel) influence on $U_{\tau I}^2$.



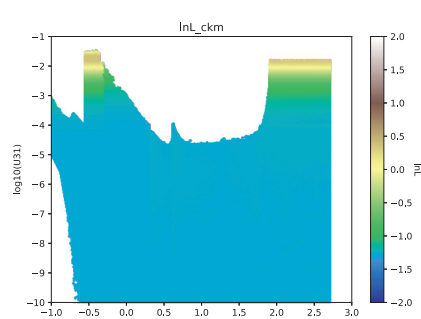
PS191's (muon channel) influence on $U_{\tau I}^2$.



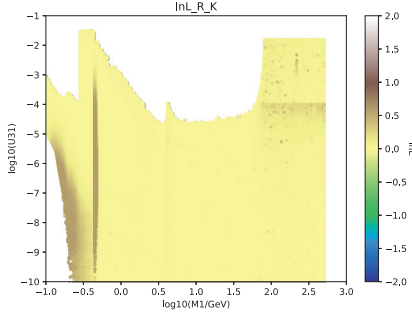
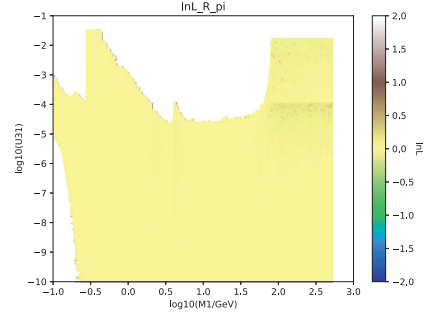
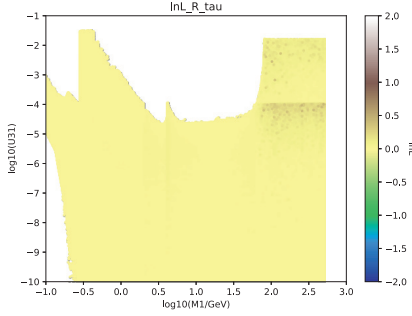
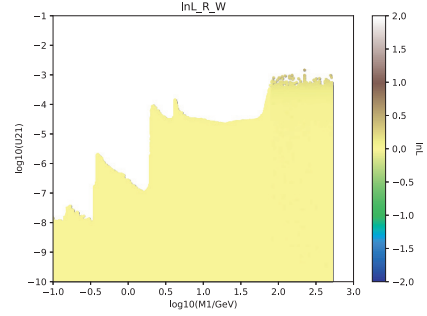
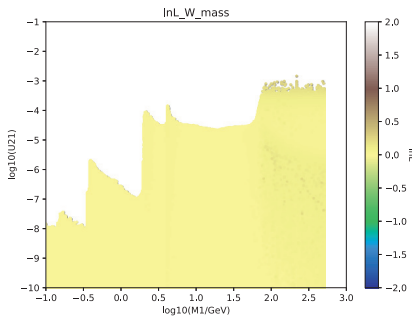
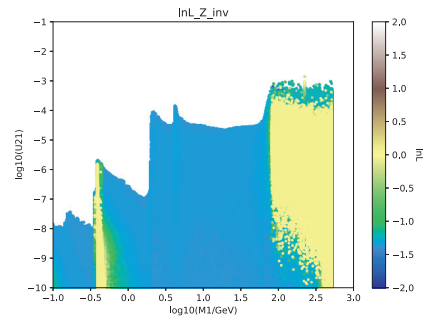
DELPHI's (long-lived RHN analysis) influence on $U_{\tau I}^2$.



DELPHI's (short-lived RHN analysis) influence on $U_{\tau I}^2$.



CKM unitarity's influence on $U_{\tau I}^2$.

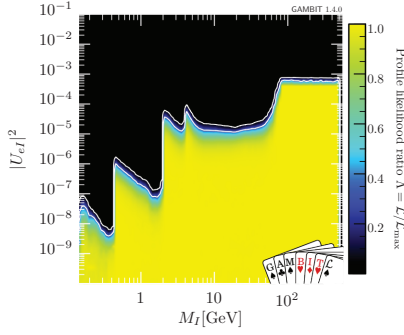
 R_K 's influence on $U_{\tau I}^2$. R_π 's influence on $U_{\tau I}^2$. R_τ 's influence on $U_{\tau I}^2$. R_W 's influence on $U_{\tau I}^2$. m_W 's influence on $U_{\tau I}^2$. Γ_{inv} 's influence on $U_{\tau I}^2$.

Appendix D

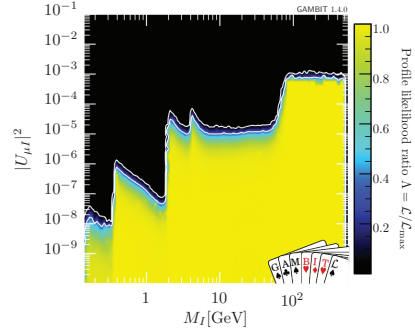
Additional plots for inverted hierarchy

This appendix contains those plots for the inverted hierarchy that are identical to their normal hierarchy counterparts in chapter 5. In addition, partial likelihood plots for every observable that affects the parameter space for this mass hierarchy are also included here.

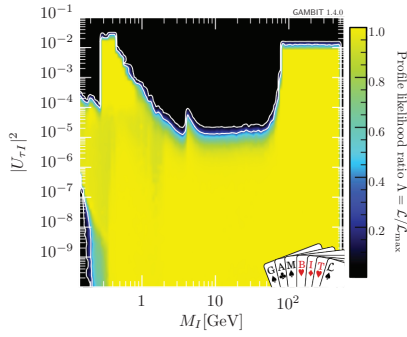
The plots corresponding to those in chapter 5 come first and begin on the following page.



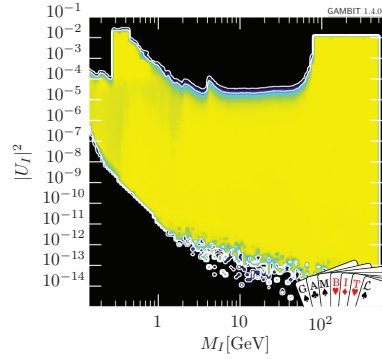
Profile likelihood in the $M_I - U_{eI}^2$ plane
with total likelihood capped.



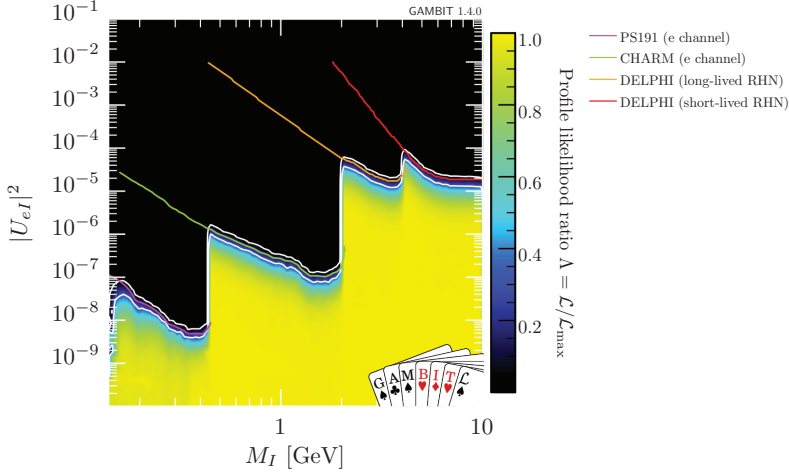
Profile likelihood in the $M_I - U_{\mu I}^2$ plane
with total likelihood capped.



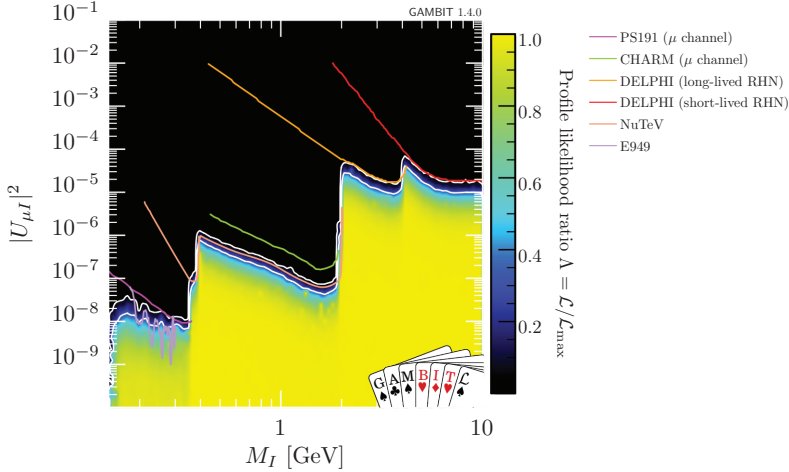
Profile likelihood in the $M_I - U_{\tau I}^2$ plane
with total likelihood capped.



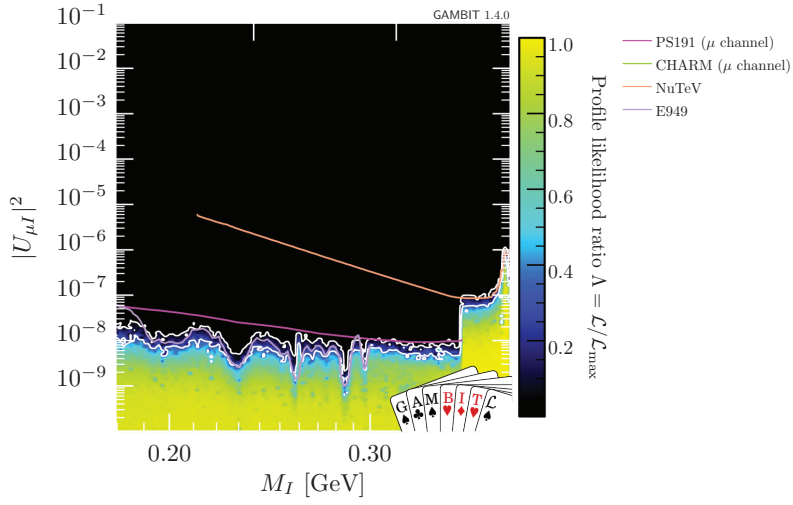
Profile likelihood in the $M_I - U_I^2$
plane with total likelihood capped.



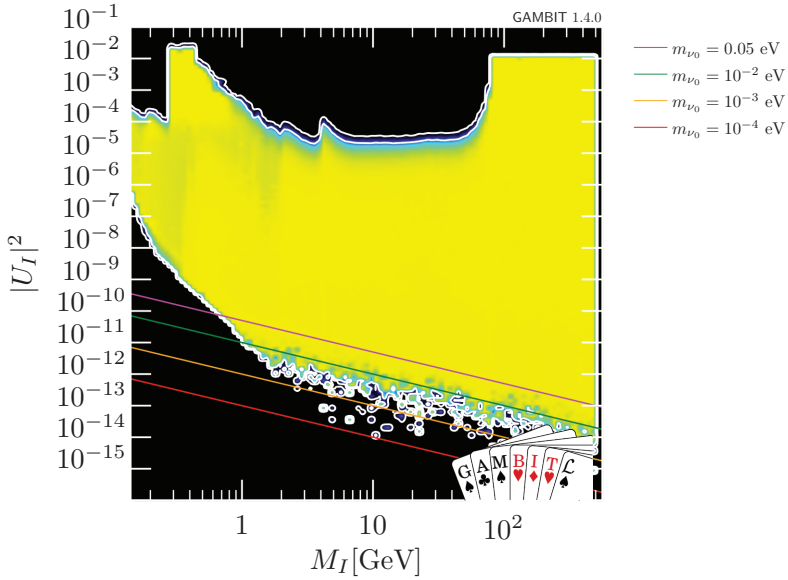
Profile likelihood in $M_I - U_{eI}^2$ plane, for $M_I < 10$ GeV, with the total likelihood capped and experimental limits overlayed.



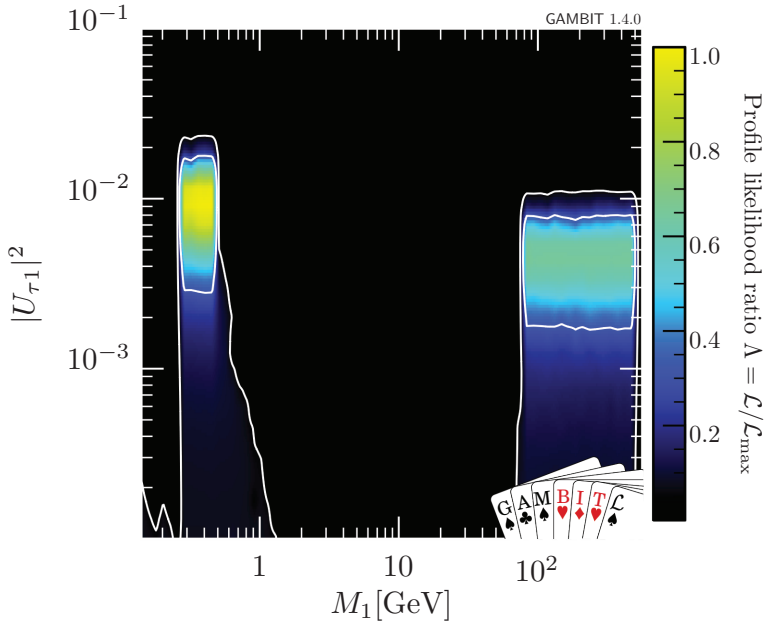
Profile likelihood in $M_I - U_{\mu I}^2$ plane, for $M_I < 10$ GeV, with the total likelihood capped and experimental limits overlayed.



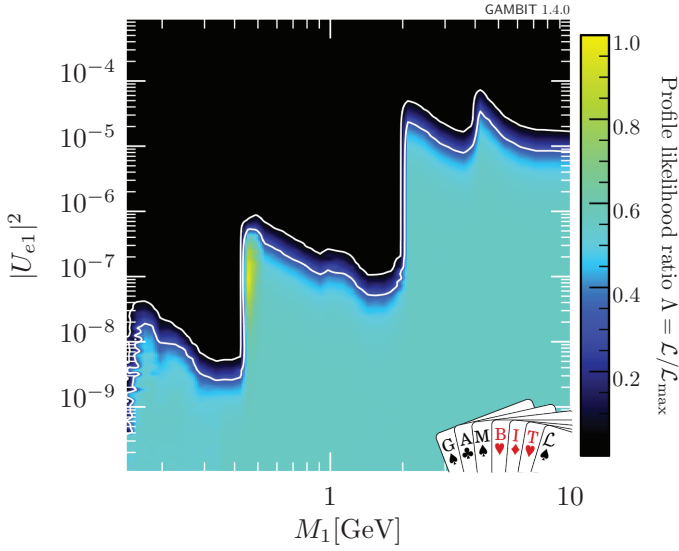
Profile likelihood in $M_I - U_{\mu I}^2$ plane, for $M_I < 0.4$ GeV, with the total likelihood capped and experimental limits overlayed.



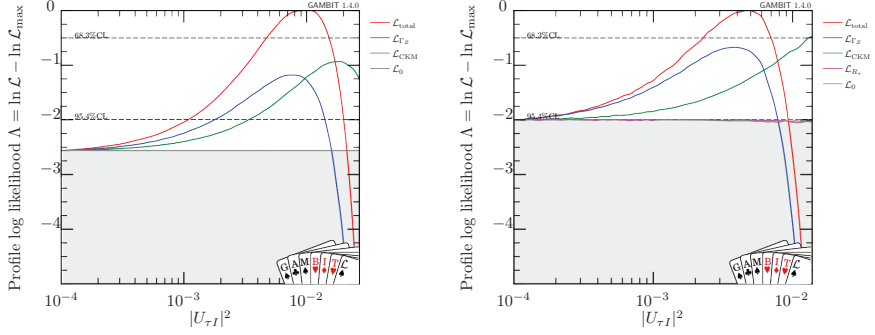
Profile likelihood in $M_I - U_I^2$ plane, with the total likelihood capped, along with the naive seesaw estimates for a few m_{ν_0} values.



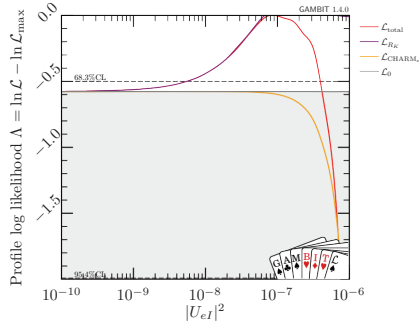
Profile likelihood in $M_I - U_{\tau I}^2$ plane, with the total likelihood uncapped.



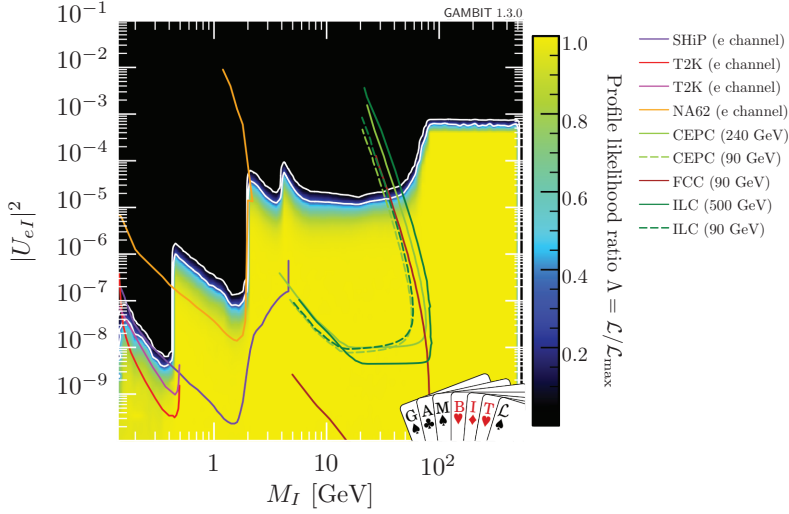
Profile likelihood in $M_I - U_{e I}^2$ plane, with the total likelihood uncapped.



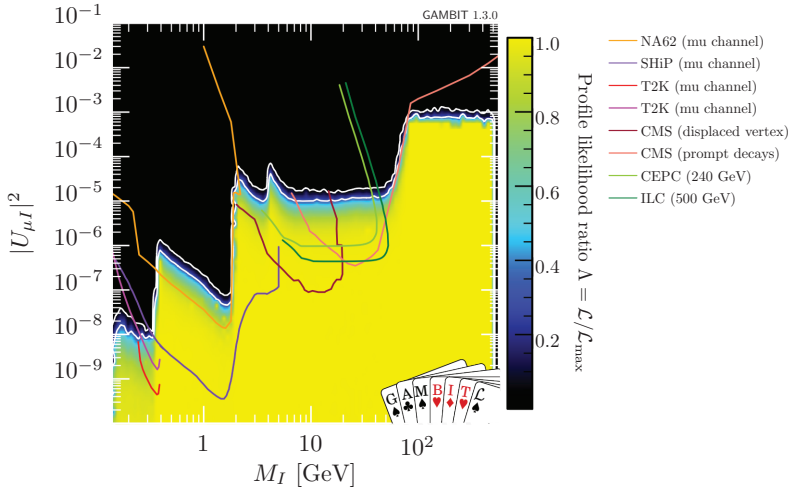
Profile likelihood of $U_{\tau 1}^2$ for $M_1 < 1$ GeV. Profile likelihood of $U_{\tau 1}^2$ for $M_1 > 60$ GeV.



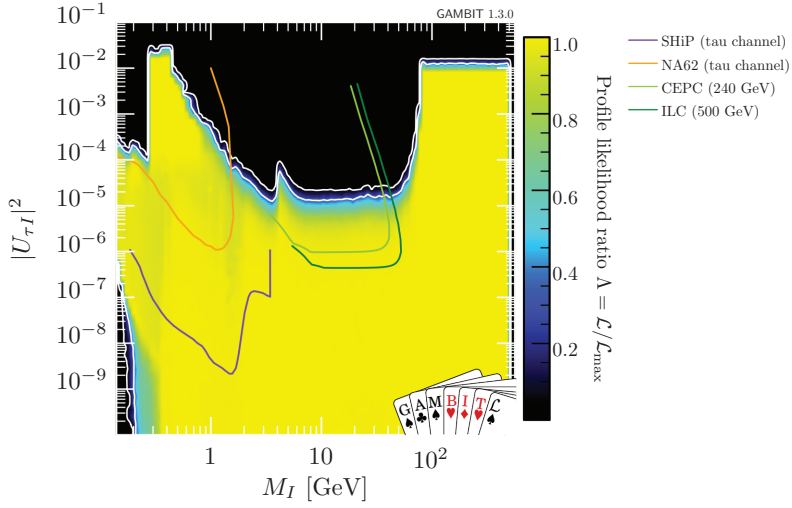
Profile likelihood of $U_{e 1}^2$ for $M_1 < 10$ GeV.



Profile likelihood in $M_I - U_{eI}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlaid.

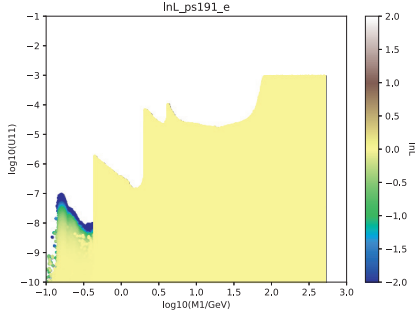


Profile likelihood in $M_I - U_{\mu I}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlaid.

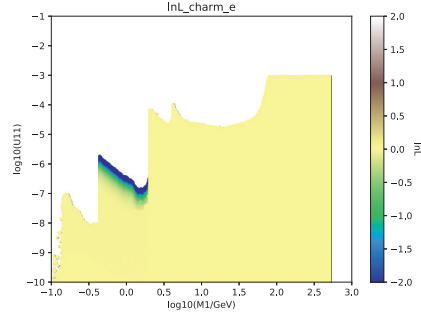


Profile likelihood in $M_I - U_{\tau I}^2$ plane, with the total likelihood capped and relevant future experiments' projected limits overlaid.

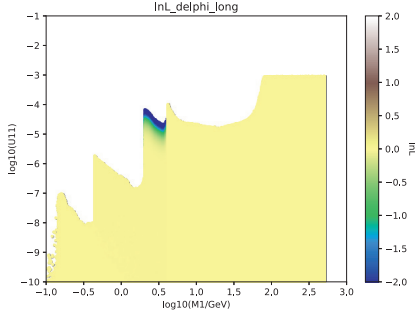
All the relevant partial likelihood plots (analogous to those in appendix C for the normal hierarchy) are shown from the next page.



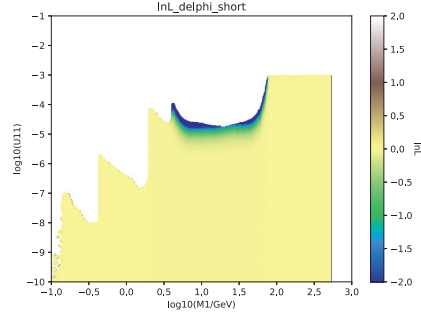
PS191's (electron channel) influence on U_{el}^2 .



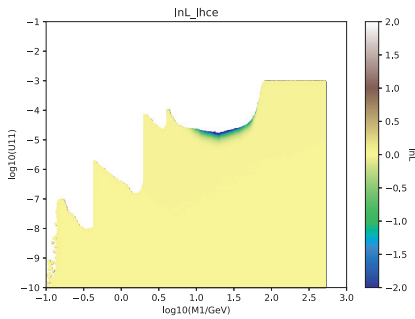
CHARM's (electron channel) influence on U_{el}^2 .



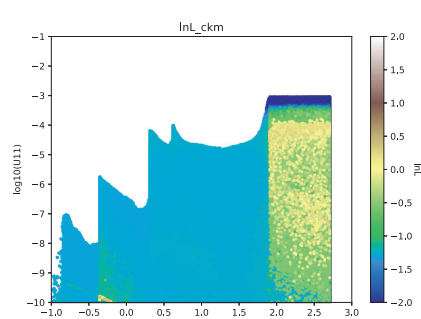
DELPHI's (long-lived RHN analysis) influence on U_{el}^2 .



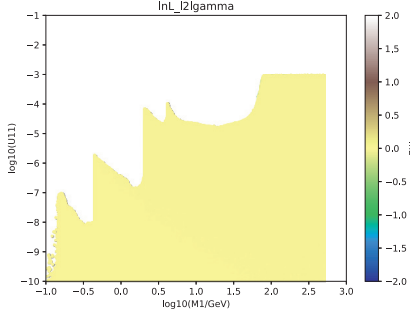
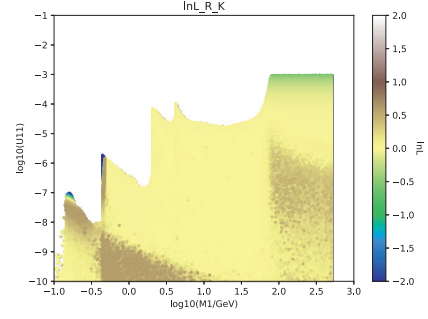
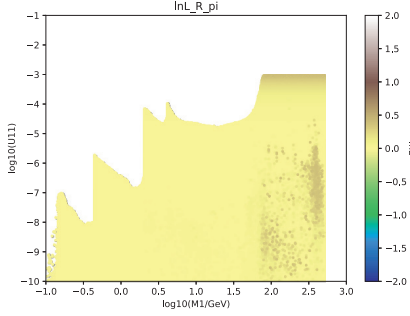
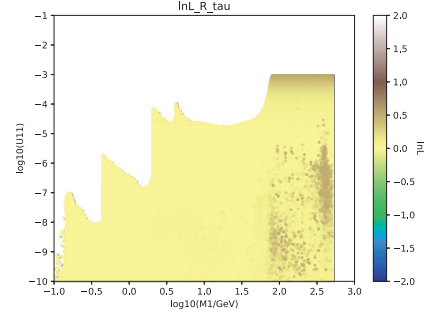
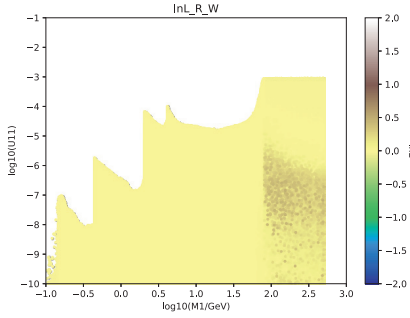
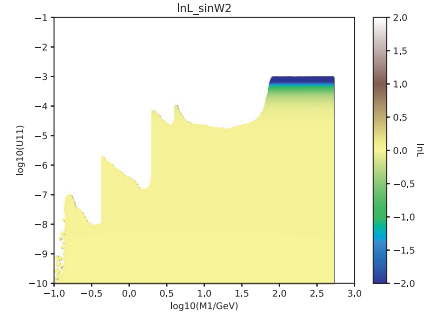
DELPHI's (short-lived RHN analysis) influence on U_{el}^2 .

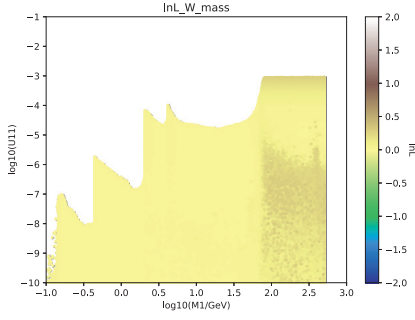
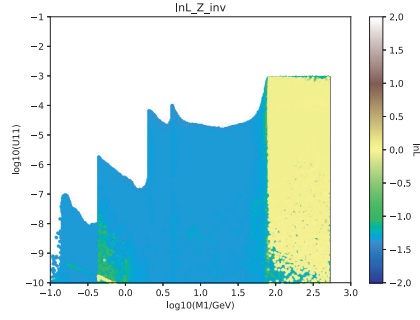
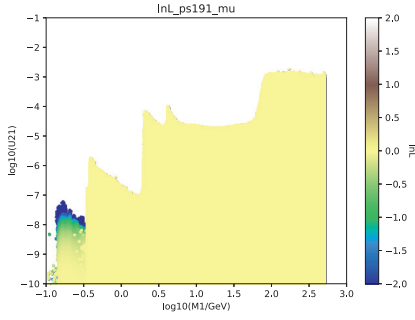
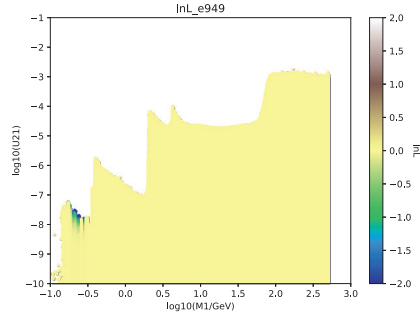
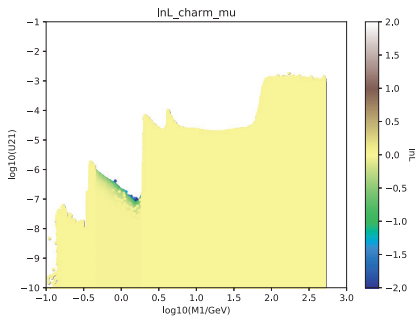
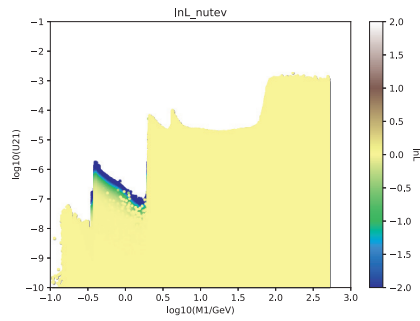


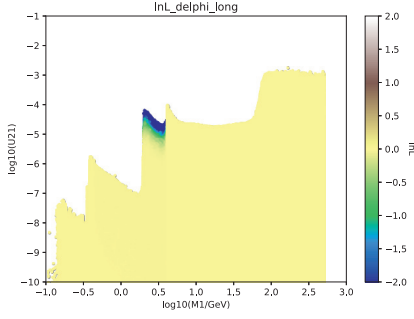
CMS' (electron channel) influence on U_{el}^2 .



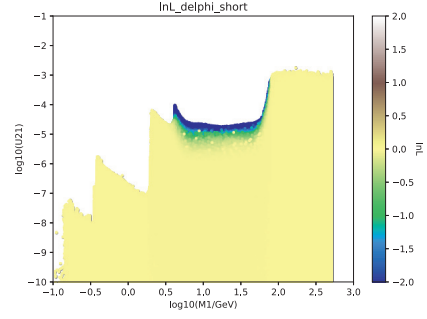
CKM unitarity's influence on U_{el}^2 .


 $\ell \rightarrow \ell\gamma$ processes' influence on U_{eI}^2 .

 R_K 's influence on U_{eI}^2 .

 R_π 's influence on U_{eI}^2 .

 R_τ 's influence on U_{eI}^2 .

 R_W 's influence on U_{eI}^2 .

 $\sin^2 \theta_W$'s influence on U_{eI}^2 .

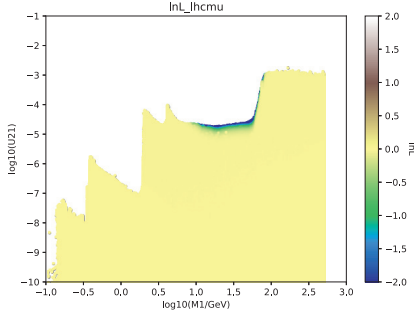
 m_W 's influence on U_{eI}^2 . Γ_{inv} 's influence on U_{eI}^2 .PS191's (muon channel) influence on $U_{\mu I}^2$.E949's influence on $U_{\mu I}^2$.CHARM's (muon channel) influence on $U_{\mu I}^2$.NuTeV's influence on $U_{\mu I}^2$.



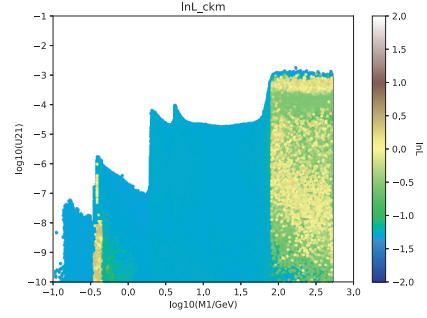
DELPHI's (long-lived RHN analysis)
influence on $U_{\mu I}^2$.



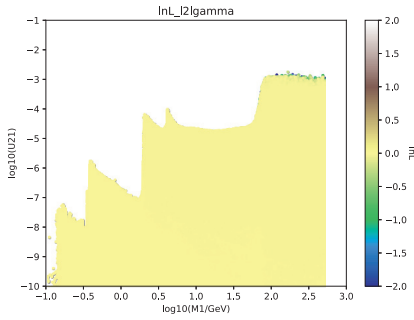
DELPHI's (short-lived RHN analysis)
influence on $U_{\mu I}^2$.



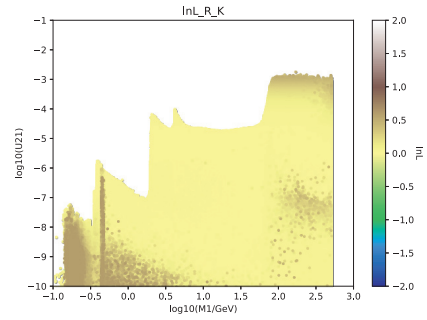
CMS' (muon channel) influence on $U_{\mu I}^2$.



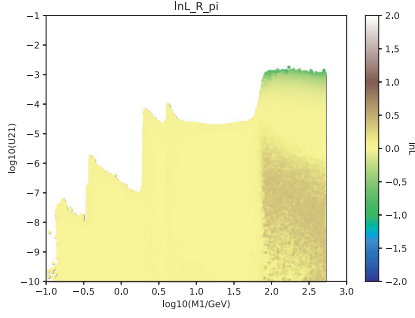
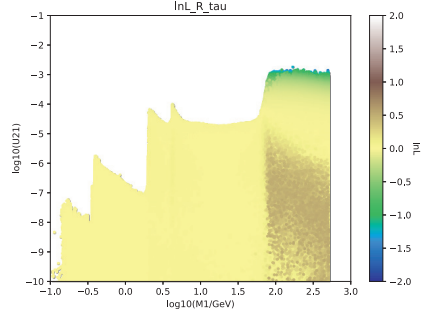
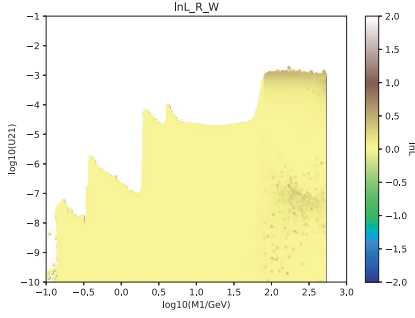
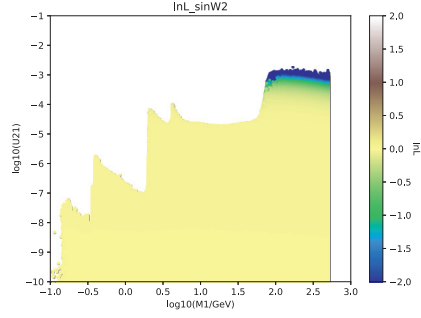
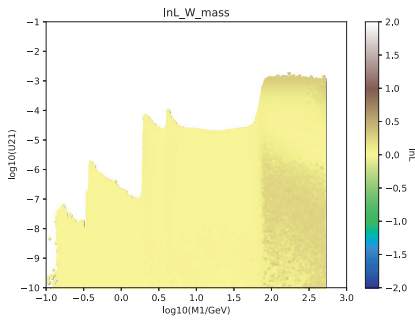
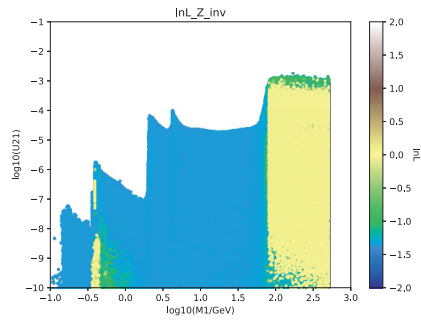
CKM unitarity's influence on $U_{\mu I}^2$.

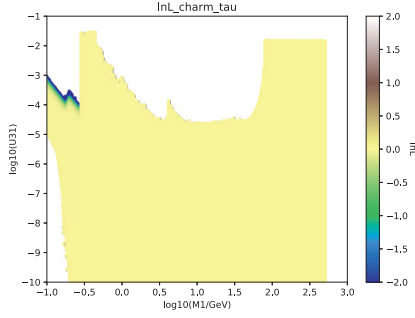


$\ell \rightarrow \ell \gamma$ processes' influence on $U_{\mu I}^2$.

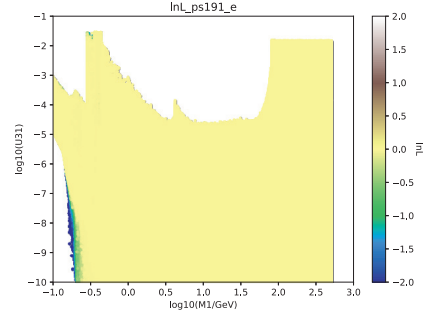


R_K 's influence on $U_{\mu I}^2$.

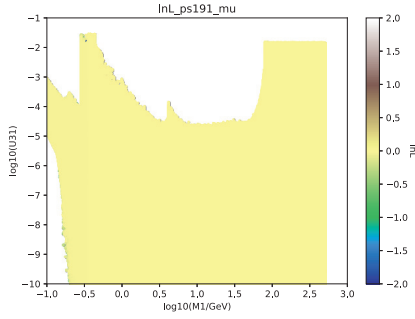
 R_π 's influence on $U_{\mu I}^2$. R_τ 's influence on $U_{\mu I}^2$. R_W 's influence on $U_{\mu I}^2$. $\sin^2 \theta_W$'s influence on $U_{\mu I}^2$. m_W 's influence on $U_{\mu I}^2$. Γ_{inv} 's influence on $U_{\mu I}^2$.



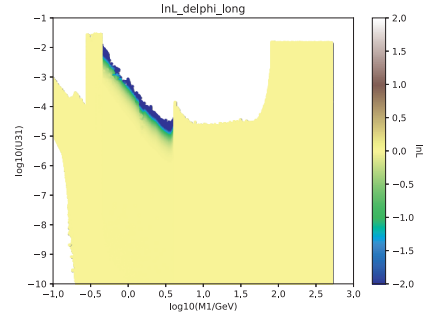
CHARM's (tau channel) influence on $U_{\tau I}^2$.



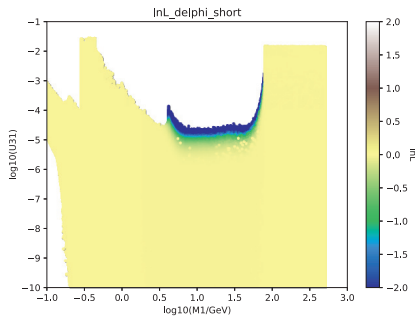
PS191's (electron channel) influence on $U_{\tau I}^2$.



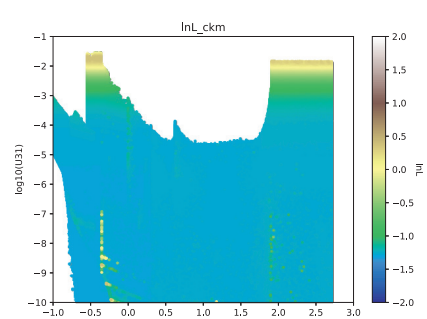
PS191's (muon channel) influence on $U_{\tau I}^2$.



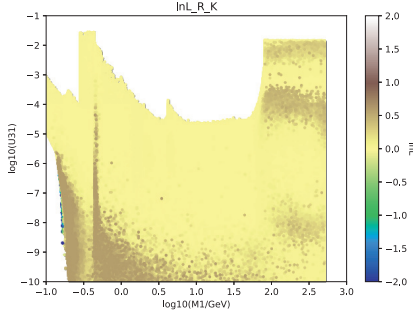
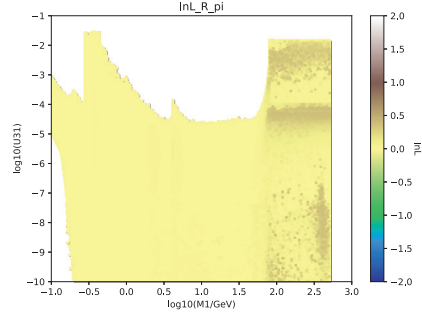
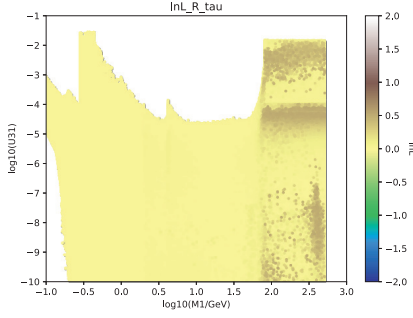
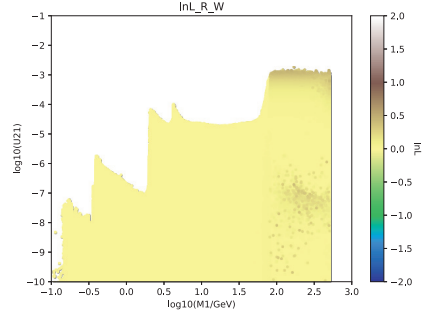
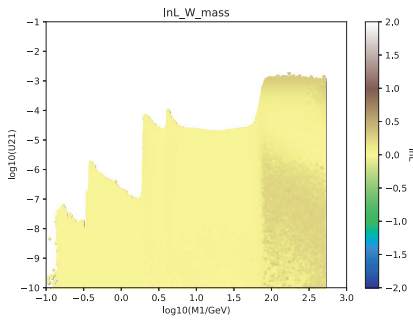
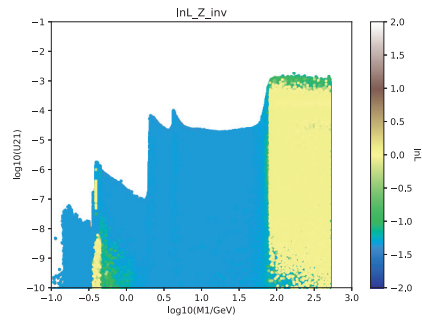
DELPHI's (long-lived RHN analysis) influence on $U_{\tau I}^2$.



DELPHI's (short-lived RHN analysis) influence on $U_{\tau I}^2$.



CKM unitarity's influence on $U_{\tau I}^2$.

 R_K' 's influence on $U_{\tau I}^2$. R_{π}' 's influence on $U_{\tau I}^2$. R_{τ}' 's influence on $U_{\tau I}^2$. R_W' 's influence on $U_{\tau I}^2$. m_W' 's influence on $U_{\tau I}^2$. Γ_{inv}' 's influence on $U_{\tau I}^2$.

Appendix E

GAMBIT capabilities & module functions

E.1 Neutrino models

Within `NeutrinoBit.cpp`, the left-handed neutrino mass matrix `m_nu`, mixing matrix `UPMNS`, as well as the type of hierarchy `ordering`, the squared mass splittings `md21`, `md31` and `md32` and the minimal neutrino mass `min_mass` are defined or calculated. It is possible to fix the hierarchy of a scan by providing the option `ordering` to the capability `m_nu` in the YAML file.

The right-handed neutrino sector contains the capabilities `SeesawI_Vnu`, which is the active neutrino mixing matrix in the type-I seesaw model (`UPMNS` corrected by the presence of the right-handed neutrinos) and `SeesawI_Theta`, the active-sterile mixing matrix, currently implemented using the C-I parameterization.

Another useful capability defined here is `Unitarity`, which is fulfilled by two module functions according to whether the model scanned is the SM or an RHN model, and checks whether the full mixing matrix is unitary. All these capabilities relating to neutrino masses and mixings and the module functions that fulfill them, along with their dependencies and options can be seen in table E.1.

`NeutrinoBit.cpp` also contains likelihoods for the active neutrinos, which are implemented following the results from the NuFit collaboration. The capabilities associated with these are `md21_lnL` and `md31_lnL` for the mass splittings, `deltaCP_lnL`, `theta12_lnL`, `theta23_lnL` and `theta13_lnL` for the phases and mixing angles, and `sum_mnu_lnL` for the cosmological limit on the sum of neutrino masses. Refer to table E.2 for a listing.

E.2 Right-handed neutrino likelihood functions

Every observable and likelihood described in chapter 3 has an assigned capability in GAMBIT. Most of these have been implemented in the new GAMBIT module `NeutrinoBit`, since they mostly concern neutrino physics. Their module functions are coded in `RightHandedNeutrinos.cpp`, to keep them separate from the likelihoods and observables concerning only active neutrinos in `NeutrinoBit.cpp`. The exception to this is the LFV observables and semileptonic lepton universality tests, which live in `FlavBit` [336] (implemented in `FlavBit.cpp`) and the electroweak precision observables, which live in `PrecisionBit` [337] and are coded in `PrecisionBit.cpp`; implementation details follow.

Electroweak precision observables

The EWPO capabilities `sinW2`, `mW` and `W_to_l_decays` can be seen in table E.3, along with their module functions and dependencies. The log-likelihoods, provided by the capabilities `lnL_sinW2`, `lnL_W_mass` and `lnL_W_decays`, can also be seen in the same table. Additionally, the module `DecayBit` [337] contains the capabilities for the invisible width of the Z boson: `Z_gamma_nu` and `lnL_Z_inv`. A list is in table E.3.

Lepton flavour violation

The capabilities related to lepton flavour violation can be found in `FlavBit`. They are `muegamma`, `tauegamma`, `taumugamma`, `mueeee`, `taueeee`, `taumumumu`, `taumuee`, `taueemu`, `tauemumu`, `taumumue`, `mueTi` and `muePb`. Table E.4 shows these capabilities, the module functions that provide them and their dependencies. The likelihoods, shown in table E.5, are collated into three capabilities; `l2lgamma_lnL` for $l \rightarrow l\gamma$, `l2lll_lnL` for $l^- \rightarrow l^- l^- l^+$ and `mu2e_lnL` for $\mu - e$ conversion in nuclei.

Lepton universality

The observables and likelihoods associated with lepton universality constraints are spread between the `NeutrinoBit` and `FlavBit` modules. Those involving fully leptonic decays are implemented in `RightHandedNeutrinos.cpp`; those

for semi-leptonic decays of B mesons are in `FlavBit.cpp`. The capabilities for leptonic decays are `R_pi`, `R_K`, `R_tau` and `R_W`, and for semi-leptonic ones, `RK`, `RKstar_0045_11` and `RKstar_11_60`. They can be seen in table E.6 together with the module functions that provide them and their dependencies. The capability `LUV_LL` collates all semi-leptonic universality observables into the FlavBit-defined class `FlavBit::predictions_measurements_covariances`¹. The capabilities that compute the related likelihoods are `lnLlepuniv` for leptonic and `LUV_LL` for semi-leptonic decays, and they, the module functions and dependencies can be seen in table E.7.

CKM unitarity

The capability `calc_Vus`, implemented in `RightHandedNeutrinos.cpp`, calculates the value of V_{us} that maximizes the likelihood for a given Θ , and `lnLckm_Vusmin` utilizes this value to compute the log-likelihood. The capabilities, module functions and dependencies defined in GAMBIT for the calculation of these quantities are listed in table E.8.

Neutrinoless double beta decay

There are two calculations for the observable of neutrinoless double-beta decay, according to the experiment that measures it. The capability `mbb_0nubb_Xe` calculates $m_{\beta\beta}$ for the process as measured by KamLAND-Zen. The log-likelihood is computed by `lnL_mbb_0nubb_KamLAND_Zen`. For the GERDA measurement the capability `mbb_0nubb_Ge` provides $m_{\beta\beta}$ and `lnL_mbb_0nubb_GERDA` its likelihood. Lastly, the total log-likelihood is given by the capability `lnL_mbb_0nubb`. Table E.9 shows the defined capabilities and associated module functions; the module functions' dependencies are also listed.

Big Bang nucleosynthesis

A number of processes contribute to the decay width of the right-handed neutrinos, and each of them is computed by a capability. These are `Gamma_RHN2pi0nu`, `Gamma_RHN2pipusl`, `Gamma_RHN2Kplusl`, `Gamma_RHN2Dplusl`, `Gamma_RHN2Dsl`, `Gamma_RHN2Bplusl`, `Gamma_RHN2Bsl`, `Gamma_RHN2Bcl`, `Gamma_RHN2etanu`, `Gamma_RHN2etaprimenu`, `Gamma_RHN2rhoplusl`, `Gamma_RHN2rho0nu`, `Gamma_RHN23nu`,

¹For more details about FlavBit types, see [336].

`Gamma_RHN2l1nu`, `Gamma_RHN2null`, `Gamma_RHN2nuubar`, `Gamma_RHN2nuddbar` and `Gamma_RHN2ludbar`. The total decay width of each of the right-handed neutrinos is given by `Gamma_BBN` and the log-likelihood for BBN by `lnL_bbn`. Table E.10 shows these capabilities, and the module functions that satisfy them, along with dependencies that other module functions fulfill. The decay process considered in each function is mentioned below its name.

Direct searches

As mentioned in chapter 3, the likelihoods for direct searches depend on the active-sterile matrix elements $U_{\alpha I}$. Hence, for simplicity the capabilities `Ue1`, `Ue2`, `Ue3`, `Um1`, `Um2`, `Um3`, `Ut1`, `Ut2` and `Ut3` are implemented in `NeutrinoBit`, as well as the phases of each of the matrix elements: `Ue1_phase`, `Ue2_phase`, `Ue3_phase`, `Um1_phase`, `Um2_phase`, `Um3_phase`, `Ut1_phase`, `Ut2_phase` and `Ut3_phase`. These can be seen in table E.12. All the capabilities `UaI` can take a pair of options: `upper_limit` and `lower_limit`, to force the values within a given range. Using these quantities, the likelihoods for the different direct search experiments are calculated, and their capabilities are `lnLpienu`, `lnLps191e`, `lnLps191mu`, `lnLcharme`, `lnLcharmmu`, `lnLcharmtau`, `lnLdelphi_shortlived`, `lnLdelphi_longlived`, `lnLatlase`, `lnLatlasmu`, `lnLlhce`, `lnLlhcmu`, `lnLe949` and `lnLnutev`. The capabilities, module functions and their dependencies for all relevant direct search experiments are tabulated in table E.13.

Other capabilities

Finally, the theoretical constraint for perturbativity of the Yukawa couplings has been implemented as well. The capability `perturbativity_lnL` calculates a step function likelihood for this constraint. Table E.11 shows the module function that provides this capability and its dependencies.

Capability	Function (Return Type): Brief description	Dependencies
ordering	<code>ordering(bool)</code> : Specifies the hierarchy type.	None
m_nu	<code>M_nu(Eigen::Matrix3cd)</code> : Calculates the diagonalized LHN mass matrix.	ordering
md21	<code>md21(double)</code> : Calculates the squared mass splitting Δm_{21}^2 .	m_nu
md31	<code>md31(double)</code> : Calculates the squared mass splitting Δm_{31}^2 .	m_nu
md32	<code>md32(double)</code> : Calculates the squared mass splitting Δm_{32}^2 .	m_nu
min_mass	<code>min_mass(double)</code> : Calculates the minimal neutrino mass.	ordering, m_nu
UPMNS	<code>UPMNS(Eigen::Matrix3cd)</code> : Calculates the PMNS matrix.	None
SeesawI_Theta	<code>CI_Theta(Eigen::Matrix3cd)</code> : Calculates the LHN-RHN mixing matrix in the type-I seesaw model using the C-I parameterization.	m_nu, UPMNS, SMINPUTS
SeesawI_Vnu	<code>Vnu(Eigen::Matrix3cd)</code> : Calculates the LHN mixing matrix.	UPMNS, SeesawI_Theta
Unitarity	<code>Unitarity_UPMNS(bool)</code> : Checks for unitarity of the SM neutrino mixing matrix. <code>Unitarity_SeesawI(bool)</code> : Checks for unitarity of the full neutrino mixing matrix in the type-I seesaw model.	m_nu, UPMNS m_nu, SeesawI_Theta, SeesawI_Vnu

TABLE E.1: Capabilities and module functions implemented for left-handed and right-handed neutrino masses and mixings.

Capability	Function (Return Type): Brief description	Dependencies
md21_lnl	md21_lnl(double): Computes the log-likelihood for Δm_{21}^2 .	md21
md31_lnl	md31_lnl(double): Computes the log-likelihood for Δm_{31}^2 for normal hierarchy or Δm_{32}^2 for inverted.	ordering, md31, md32
deltaCP_lnl	deltaCP_lnl(double): Computes the log-likelihood for δ_{CP} .	ordering, deltaCP
theta12_lnl	theta12_lnl(double): Computes the log-likelihood for θ_{12} .	theta12
theta23_lnl	theta23_lnl(double): Computes the log-likelihood for θ_{23} .	ordering, theta23
theta13_lnl	theta13_lnl(double): Computes the log-likelihood for θ_{13} .	ordering, theta13
sum_mnu_lnl	sum_mnu_lnl(double): Computes the log-likelihood for $\sum m_\nu$	None

TABLE E.2: Capabilities and module functions implemented that calculate log-likelihoods for the left-handed neutrino parameters.

Capability	Function (Return Type): Brief description	Dependencies
<code>sinW2</code>	<code>RHN_sinW2(triplet<double>):</code> Calculates $\sin^2 \theta_{W,\text{eff}}$.	<code>SeesawI_Theta</code>
<code>mW</code>	<code>RHN_mW(triplet<double>):</code> Calculates m_W .	<code>sinW2</code> , <code>SeesawI_Theta</code>
<code>Z_gamma_nu</code>	<code>Z_gamma_nu_2l(triplet<double>):</code> Calculates the decay width of Z to neutrinos.	<code>SM_spectrum</code> , <code>SeesawI_Theta</code> , <code>SeesawI_Vnu</code>
<code>W_to_l_decays</code>	<code>RHN_W_to_l_decays(vector<double>):</code> Calculates the decay width of the processes $W \rightarrow l\nu$.	<code>SMINPUTS</code> , <code>mw</code> , <code>SeesawI_Theta</code>
<code>lnL_sinW2</code>	<code>lnL_sinW2_chi2(double):</code> Computes the log-likelihood for $\sin^2 \theta_{W,\text{eff}}$.	<code>sinW2</code>
<code>lnL_W_mass</code>	<code>lnL_W_mass_chi2(double):</code> Computes the log-likelihood for m_W .	<code>mW</code>
<code>lnL_Z_inv</code>	<code>lnL_Z_inv(double):</code> Computes the log-likelihood for Γ_{inv} .	<code>Z_gamma_nu</code>
<code>lnL_W_decays</code>	<code>lnL_W_decays_chi2(double):</code> Computes the log-likelihood for $\Gamma_{W \rightarrow l\nu}$.	<code>W_to_l_decays</code> , <code>W_plus_decay_rates</code>

TABLE E.3: Capabilities and module functions that calculate electroweak precision observables and their likelihoods.

Capability	Function (Return Type): Brief description	Dependencies
muegamma	<code>RHN_muegamma(double)</code> : Calculates $BR(\mu^- \rightarrow e^- \gamma)$.	SMINPUTS, m_nu, SeesawI_Vnu, SeesawI_Theta, mu_minus_decay_rates
mueee	<code>RHN_mueee(double)</code> : Calculates $BR(\mu^- \rightarrow e^- e^- e^+)$.	Ditto.
tauegamma	<code>RHN_taugamma(double)</code> : Calculates $BR(\tau^- \rightarrow e^- \gamma)$.	SMINPUTS, m_nu, SeesawI_Vnu, SeesawI_Theta, tau_minus_decay_rates
taumugamma	<code>RHN_taumugamma(double)</code> : Calculates $BR(\tau^- \rightarrow \mu^- \gamma)$.	Ditto.
taueee	<code>RHN_taujee(double)</code> : Calculates $BR(\tau^- \rightarrow e^- e^- e^+)$.	Ditto.
taumumumu	<code>RHN_taumumumu(double)</code> : Calculates $BR(\tau^- \rightarrow \mu^- \mu^- \mu^+)$.	Ditto.
taumuee	<code>RHN_taumuee(double)</code> : Calculates $BR(\tau^- \rightarrow \mu^- e^- e^+)$.	Ditto.
taueemu	<code>RHN_tauemu(double)</code> : Calculates $BR(\tau^- \rightarrow e^- e^- \mu^+)$.	Ditto.
taumumu	<code>RHN_tauemumu(double)</code> : Calculates $BR(\tau^- \rightarrow e^- \mu^- \mu^+)$.	Ditto.
taumumue	<code>RHN_taumumue(double)</code> : Calculates $BR(\tau^- \rightarrow \mu^- \mu^- e^+)$.	Ditto.
mueTi	<code>RHN_mueTi(double)</code> : Calculates $R(\mu - e)$ in a Ti nucleus.	SMINPUTS, m_nu, SeesawI_Vnu, SeesawI_Theta
muePb	<code>RHN_muePb(double)</code> : Calculates $R(\mu - e)$ in a Pb nucleus.	Ditto.

TABLE E.4: Capabilities and module functions to calculate LFV observables.

Capability	Function (Return Type): Brief description	Dependencies
<code>l2lgamma_1nL</code>	<code>l2lgamma_likelihood(double)</code> : Computes the log-likelihood for $l^- \rightarrow l^- \gamma$ processes.	<code>muegamma</code> , <code>tauegamma</code> , <code>taumugamma</code>
<code>l2l1l_1nL</code>	<code>l2l1l_likelihood(double)</code> : Computes the log-likelihood for $l^- \rightarrow l^- l^- l^+$ processes.	<code>mueee</code> , <code>taueee</code> , <code>taumumu</code> , <code>taumumue</code> , <code>tauemumu</code>
<code>mu2e_1nL</code>	<code>mu2e_likelihood(double)</code> : Computes the log-likelihood associated with $\mu - e$ conversion.	<code>mueTi</code> , <code>muePb</code>

TABLE E.5: Capabilities and module functions for the likelihoods related to LFV observables.

Capability	Function (Return Type): Brief description	Dependencies
<code>R_pi</code>	<code>RHN_R_pi(double)</code> : Calculates the test of lepton universality $R_{e\mu}^\pi$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>R_K</code>	<code>RHN_R_K(double)</code> : Calculates the test of lepton universality $R_{e\mu}^K$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>R_tau</code>	<code>RHN_R_tau(double)</code> : Calculates the test of lepton universality $R_{\mu e}^\tau$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>R_W</code>	<code>RHN_R_W(vector<double>)</code> : Calculates the test of lepton universality $R_{\alpha\beta}^W$.	<code>SeesawI_Theta</code> , <code>mw</code>
<code>RK</code>	<code>RHN_RK(double)</code> : Calculates the test of lepton universality R_K .	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>RKstar_0045_11</code>	<code>RHN_RKstar_0045_11(double)</code> : Calculates the test of lepton universality R_{K^*} for the range $0.045 < q^2 < 1.1 \text{ GeV}^2$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>RKstar_11_60</code>	<code>RHN_RKstar_11_60(double)</code> : Calculates the test of lepton universality R_{K^*} for the range $1.1 < q^2 < 6.0 \text{ GeV}^2$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>

TABLE E.6: Capabilities and module functions that calculate lepton universality observables.

Capability	Function (Return Type): Brief description	Dependencies
<code>lnLlepuniv</code>	<code>lnL_lepuniv(double)</code> : Calculates the total log-likelihood for lepton universality tests of leptonic decays using π and K mesons, τ , and W bosons.	<code>R_pi</code> , <code>R_K</code> , <code>R_tau</code> , <code>R_W</code>
<code>LUV_M</code>	<code>LUV_measurements()</code> : Collates measurements of semi-leptonic tests of lepton universality in B meson decays.	<code>RK</code> , <code>RKstar_0045_11</code> , <code>RKstar_11_60</code>
<code>LUV_LL</code>	<code>lnL_lepuniv(double)</code> : Calculates the total log-likelihood for semi-leptonic tests of lepton universality in B meson decays.	<code>LUV_M</code>

TABLE E.7: Capabilities and module functions for the likelihoods related to lepton universality tests.

Capability	Function (Return Type): Brief description	Dependencies
<code>calc_Vus</code>	<code>calc_Vus(double)</code> : Calculates V_{us} for a particular Θ .	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>lnLckm_Vusmin</code>	<code>lnL_ckm_Vusmin(double)</code> : Calculates the total log-likelihood.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code> , <code>calc_Vus</code>

TABLE E.8: Capabilities and module functions implemented to calculate CKM unitarity-related observables and the likelihood.

Capability	Function (Return Type): Brief description	Dependencies
<code>mbb_Onubb_Xe</code>	<code>RHN_mbb_Onubb_Xe(double):</code> Calculates $m_{\beta\beta}$ for Xe, used in KamLAND-Zen.	<code>m_nu</code> , <code>UPMNS</code> , <code>SeesawI_Theta</code>
<code>mbb_Onubb_Ge</code>	<code>RHN_mbb_Onubb_Ge(double):</code> Calculates $m_{\beta\beta}$ for Ge, used in GERDA.	<code>m_nu</code> , <code>UPMNS</code> , <code>SeesawI_Theta</code>
<code>lnL_mbb_Onubb_KamLAND_Zen</code>	<code>lnL_mbb_Onubb_KamLAND_Zen(double):</code> Calculates the log-likelihood for KamLAND-Zen.	<code>mbb_Onubb_Xe</code>
<code>lnL_mbb_Onubb_GERDA</code>	<code>lnL_mbb_Onubb_GERDA(double):</code> Calculates the log-likelihood for GERDA.	<code>mbb_Onubb_Ge</code>
<code>lnL_mbb_Onubb</code>	<code>lnL_mbb_Onubb(double):</code> Calculates the total log-likelihood.	<code>lnL_mbb_Onubb_*</code> , as listed above

TABLE E.9: Capabilities and module functions implemented to calculate neutrinoless double-beta decay observables and likelihoods.

Capability	Function (Return Type): Brief description	Dependencies
<code>Gamma_RHN2pi0nu</code>	<code>Gamma_RHN2pi0nu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \pi^0 \nu_\alpha)$.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>
<code>Gamma_RHN2piplus1</code>	<code>Gamma_RHN2piplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \pi^+ l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Kplus1</code>	<code>Gamma_RHN2Kplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow K^+ l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Dplus1</code>	<code>Gamma_RHN2Dplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow D^+ l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Ds1</code>	<code>Gamma_RHN2Ds1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow D_s l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Bplus1</code>	<code>Gamma_RHN2Bplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow B^+ l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Bs1</code>	<code>Gamma_RHN2Bs1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow B_s l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2Bc1</code>	<code>Gamma_RHN2Bc1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow B_c l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2etanu</code>	<code>Gamma_RHN2etanu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \eta \nu_\alpha)$.	Ditto.
<code>Gamma_RHN2etaprimenu</code>	<code>Gamma_RHN2etaprimenu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \eta' \nu_\alpha)$.	Ditto.
<code>Gamma_RHN2rhoplus1</code>	<code>Gamma_RHN2rhoplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \rho^+ l_\alpha^-)$.	Ditto.
<code>Gamma_RHN2rho0nu</code>	<code>Gamma_RHN2rho0nu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \rho^0 \nu_\alpha)$.	Ditto.
<code>Gamma_RHN23nu</code>	<code>Gamma_RHN23nu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \nu_\alpha \bar{\nu}_\beta \nu_\beta)$.	Ditto.
<code>Gamma_RHN211nu</code>	<code>Gamma_RHN211nu(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow l_{\alpha \neq \beta}^- l_\beta^+ \nu_\beta)$.	Ditto.
<code>Gamma_RHN2null1</code>	<code>Gamma_RHN2Kplus1(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \nu_\alpha l_\beta^+ l_\beta^-)$.	Ditto.
<code>Gamma_RHN2nuuubar</code>	<code>Gamma_RHN2nuuubar(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \nu_\alpha q^u \bar{q}^u)$.	Ditto.
<code>Gamma_RHN2nuddbar</code>	<code>Gamma_RHN2nuddbar(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow \nu_\alpha q^d \bar{q}^d)$.	Ditto.
<code>Gamma_RHN2ludbar</code>	<code>Gamma_RHN2ludbar(std::vector<double>):</code> Calculates $\Gamma(N_I \rightarrow l_\alpha q_\beta^u \bar{q}_\gamma^d)$.	Ditto.
<code>Gamma_BBN</code>	<code>Gamma_BBN(std::vector<double>):</code> Calculates the total decay width for each RHN.	<code>Gamma_*</code> , as listed above
<code>lnL_bbn</code>	<code>lnL_bbn(double):</code> Calculates the log-likelihood.	<code>Gamma_BBN</code>

TABLE E.10: Capabilities and module functions implemented in NeutrinoBit to calculate BBN observables and likelihood.

Capability	Function (Return Type): Brief description	Dependencies
<code>perturbativity_lnl</code>	<code>perturbativity_likelihood(double):</code> Calculates the log-likelihood related to the enforcing of perturbativity of Yukawa couplings.	<code>SMINPUTS</code> , <code>SeesawI_Theta</code>

TABLE E.11: Capability and module function in NeutrinoBit related to the perturbativity constraint.

Capability	Function (Return Type): Brief description	Dependencies	Options
<code>UeI</code>	<code>UeI(double):</code> Magnitude of the matrix element $U_{eI} = \Theta_{eI} $.	<code>SeesawI_Theta</code>	<code>upper_limit(double)</code> , <code>lower_limit(double)</code>
<code>UmuI</code>	<code>UmuI(double):</code> Magnitude of the matrix element $U_{\mu I} = \Theta_{\mu I} $.	<code>SeesawI_Theta</code>	Ditto.
<code>UtauI</code>	<code>UtauI(double):</code> Magnitude of the matrix element $U_{\tau I} = \Theta_{\tau I} $.	<code>SeesawI_Theta</code>	Ditto.
<code>UeI_phase</code>	<code>UeI_phase(double):</code> Argument of the matrix element Θ_{eI} .	<code>SeesawI_Theta</code>	None.
<code>UmuI_phase</code>	<code>UmuI_phase(double):</code> Argument of the matrix element $\Theta_{\mu I}$.	<code>SeesawI_Theta</code>	None.
<code>UtauI_phase</code>	<code>UtauI_phase(double):</code> Argument of the matrix element $\Theta_{\tau I}$.	<code>SeesawI_Theta</code>	None.

TABLE E.12: Capabilities and module functions that calculate the magnitudes and arguments of the matrix elements of Θ ($I = 1, 2, 3$).

Capability	Function (Return Type): Brief description	Dependencies
<code>lnLpienu</code>	<code>lnL_pienu(double)</code> : Calculates the log-likelihood for PIENU.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code>
<code>lnLps191e</code>	<code>lnL_ps191_e(double)</code> : Calculates the log-likelihood for PS-191 in the electron sector.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code> , <code>Um1</code> , <code>Um2</code> , <code>Um3</code> , <code>Ut1</code> , <code>Ut2</code> , <code>Ut3</code>
<code>lnLps191mu</code>	<code>lnL_ps191_mu(double)</code> : Calculates the log-likelihood for PS-191 in the muon sector.	Ditto.
<code>lnLcharm_e</code>	<code>lnL_charm_e(double)</code> : Calculates the log-likelihood for CHARM in the electron sector.	Ditto.
<code>lnLcharmmu</code>	<code>lnL_charm_mu(double)</code> : Calculates the log-likelihood for CHARM in the muon sector.	Ditto.
<code>lnLcharmtau</code>	<code>lnL_charm_tau(double)</code> : Calculates the log-likelihood for CHARM in the tau sector.	<code>Ut1</code> , <code>Ut2</code> , <code>Ut3</code>
<code>lnLdelphi_shortlived</code>	<code>lnL_delphi_short_lived(double)</code> : Calculates the log-likelihood for DELPHI's short-lived RHN analyses.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code> , <code>Um1</code> , <code>Um2</code> , <code>Um3</code> , <code>Ut1</code> , <code>Ut2</code> , <code>Ut3</code>
<code>lnLdelphi_longlived</code>	<code>lnL_delphi_long_lived(double)</code> : Calculates the log-likelihood for DELPHI's long-lived RHN analyses.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code> , <code>Um1</code> , <code>Um2</code> , <code>Um3</code> , <code>Ut1</code> , <code>Ut2</code> , <code>Ut3</code>
<code>lnLatlase</code>	<code>lnL_atlas_e(double)</code> : Calculates the log-likelihood for ATLAS in the electron sector.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code>
<code>lnLatlasmu</code>	<code>lnL_atlas_mu(double)</code> : Calculates the log-likelihood for ATLAS in the muon sector.	<code>Um1</code> , <code>Um2</code> , <code>Um3</code>
<code>lnLlhce</code>	<code>lnL_lhc_e(double)</code> : Calculates the log-likelihood for CMS in the electron sector.	<code>Ue1</code> , <code>Ue2</code> , <code>Ue3</code>
<code>lnLlhcmu</code>	<code>lnL_lhc_mu(double)</code> : Calculates the log-likelihood for CMS in the muon sector.	<code>Um1</code> , <code>Um2</code> , <code>Um3</code>
<code>lnLe949</code>	<code>lnL_e949(double)</code> : Calculates the log-likelihood for E949.	<code>Um1</code> , <code>Um2</code> , <code>Um3</code>
<code>lnLnutev</code>	<code>lnL_nutev(double)</code> : Calculates the log-likelihood for NuTeV.	<code>Um1</code> , <code>Um2</code> , <code>Um3</code>

TABLE E.13: Capabilities and module functions implemented to calculate direct search likelihoods.

Bibliography

- [1] M. Chruszcz et al. “A frequentist analysis of three right-handed neutrinos with GAMBIT”. In: (2019). arXiv: 1908.02302 [hep-ph].
- [2] R. Bartels, S. Krishnamurthy, and C. Weniger. “Strong support for the millisecond pulsar origin of the Galactic center GeV excess”. In: *Phys. Rev. Lett.* 116 (2016), p. 051102. arXiv: 1506.05104 [astro-ph.HE].
- [3] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>.
- [4] Steven Weinberg. *The Quantum theory of fields. Vol. 1: Foundations*. Cambridge University Press, 2005. ISBN: 9780511252044.
- [5] R. Oerter. *The theory of almost everything: The standard model, the unsung triumph of modern physics*. 2006.
- [6] Laurent Canetti et al. “Dark Matter, Baryogenesis and Neutrino Oscillations from Right Handed Neutrinos”. In: *Phys. Rev. D* 87 (2013), p. 093006. DOI: 10.1103/PhysRevD.87.093006. arXiv: 1208.4607 [hep-ph].
- [7] M. Drewes et al. “A White Paper on keV Sterile Neutrino Dark Matter”. In: *JCAP* 1701.01 (2017), p. 025. DOI: 10.1088/1475-7516/2017/01/025. arXiv: 1602.04816 [hep-ph].
- [8] Peter Minkowski. “ $\mu \rightarrow e\gamma$ 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.
- [9] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. “Matter and Antimatter in the Universe”. In: *New J. Phys.* 14 (2012), p. 095012. DOI: 10.1088/1367-2630/14/9/095012. arXiv: 1204.4186 [hep-ph].
- [10] J. N. Abdurashitov et al. “Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period”. In: *Phys. Rev. C* 80 (2009), p. 015807. DOI: 10.1103/PhysRevC.80.015807. arXiv: 0901.2200 [nucl-ex].

- [11] T. A. Kirsten. "GALLEX solar neutrino results and status of GNO". In: *Nucl. Phys. Proc. Suppl.* 77 (1999). [26(1999)], pp. 26–34. DOI: 10.1016/S0920-5632(99)00389-8.
- [12] B. Aharmim et al. "Searches for High Frequency Variations in the ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory". In: *Astrophys. J.* 710 (2010), pp. 540–548. DOI: 10.1088/0004-637X/710/1/540. arXiv: 0910.2433 [astro-ph.SR].
- [13] S. Abe et al. "Precision Measurement of Neutrino Oscillation Parameters with KamLAND". In: *Phys. Rev. Lett.* 100 (2008), p. 221803. DOI: 10.1103/PhysRevLett.100.221803. arXiv: 0801.4589 [hep-ex].
- [14] T. Kajita, E. Kearns, and M. Shiozawa. "Establishing atmospheric neutrino oscillations with Super-Kamiokande". In: *Nucl. Phys.* B908 (2016), pp. 14–29. DOI: 10.1016/j.nuclphysb.2016.04.017.
- [15] J. Zalipska. "Neutrino oscillation in the K2K experiment". In: *Acta Phys. Polon.* B37 (2006), pp. 1935–1945.
- [16] Keigo Nakamura. "The T2K cross-section results and prospects from the oscillation perspective". In: *PoS NuFact2017* (2018), p. 046. DOI: 10.22323/1.295.0046.
- [17] Justin Evans. "New results from MINOS and MINOS+". In: *J. Phys. Conf. Ser.* 888.1 (2017), p. 012017. DOI: 10.1088/1742-6596/888/1/012017.
- [18] N. Agafonova et al. "Final Results of the OPERA Experiment on ν_τ Appearance in the CNGS Neutrino Beam". In: *Phys. Rev. Lett.* 120.21 (2018). [Erratum: *Phys. Rev. Lett.* 121, no. 13, 139901 (2018)], p. 211801. DOI: 10.1103/PhysRevLett.121.139901, 10.1103/PhysRevLett.120.211801. arXiv: 1804.04912 [hep-ex].
- [19] Y. Abe et al. "Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector". In: *JHEP* 10 (2014). [Erratum: *JHEP* 02, 074 (2015)], p. 086. DOI: 10.1007/JHEP02(2015)074, 10.1007/JHEP10(2014)086. arXiv: 1406.7763 [hep-ex].
- [20] S. H. Seo et al. "Spectral Measurement of the Electron Antineutrino Oscillation Amplitude and Frequency using 500 Live Days of RENO Data". In: *Phys. Rev.* D98.1 (2018), p. 012002. DOI: 10.1103/PhysRevD.98.012002. arXiv: 1610.04326 [hep-ex].
- [21] J. P. Ochoa-Ricoux. "Latest results from the Daya Bay experiment". In: *PoS NuFact2017* (2018), p. 039. DOI: 10.22323/1.295.0039.

- [22] Rabindra N. Mohapatra and Goran Senjanovic. "Neutrino Mass and Spontaneous Parity Violation". In: *Phys. Rev. Lett.* 44 (1980). [231(1979)], p. 912. DOI: 10.1103/PhysRevLett.44.912.
- [23] Rabindra N. Mohapatra and Goran Senjanovic. "Neutrino Masses and Mixings in Gauge Models with Spontaneous Parity Violation". In: *Phys. Rev. D.* 23 (1981), p. 165. DOI: 10.1103/PhysRevD.23.165.
- [24] Patrick Huet and Eric Sather. "Electroweak baryogenesis and standard model CP violation". In: *Phys. Rev. D* 51 (1995), pp. 379–394. DOI: 10.1103/PhysRevD.51.379. arXiv: hep-ph/9404302 [hep-ph].
- [25] Renata Jora. "Baryon asymmetry in the standard model revisited". In: (2018). arXiv: 1806.00597 [hep-ph].
- [26] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov. "On the Anomalous Electroweak Baryon Number Nonconservation in the Early Universe". In: *Phys. Lett.* 155B (1985), p. 36. DOI: 10.1016/0370-2693(85)91028-7.
- [27] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov. "Anomalous Electroweak Baryon Number Nonconservation and GUT Mechanism for Baryogenesis". In: *Phys. Lett.* B191 (1987), pp. 171–173. DOI: 10.1016/0370-2693(87)91340-2.
- [28] M. D'Onofrio, K. Rummukainen, and A. Tranberg. "Sphaleron Rate in the Minimal Standard Model". In: *Phys. Rev. Lett.* 113.14 (2014), p. 141602. DOI: 10.1103/PhysRevLett.113.141602. arXiv: 1404.3565 [hep-ph].
- [29] Takehiko Asaka and Mikhail Shaposhnikov. "The nuMSM, dark matter and baryon asymmetry of the universe". In: *Phys. Lett. B.* 620 (2005), pp. 17–26. DOI: 10.1016/j.physletb.2005.06.020. arXiv: hep-ph/0505013 [hep-ph].
- [30] Thomas Hambye and Daniele Teresi. "Higgs doublet decay as the origin of the baryon asymmetry". In: *Phys. Rev. Lett.* 117.9 (2016), p. 091801. DOI: 10.1103/PhysRevLett.117.091801. arXiv: 1606.00017 [hep-ph].
- [31] Thomas Hambye and Daniele Teresi. "Baryogenesis from L-violating Higgs-doublet decay in the density-matrix formalism". In: *Phys. Rev. D* 96.1 (2017), p. 015031. DOI: 10.1103/PhysRevD.96.015031. arXiv: 1705.00016 [hep-ph].
- [32] M. Fukugita and T. Yanagida. "Baryogenesis Without Grand Unification". In: *Phys. Lett. B.* 174 (1986), pp. 45–47. DOI: 10.1016/0370-2693(86)91126-3.

- [33] M. Fukugita and T. Yanagida. “Sphaleron Induced Baryon Number Nonconservation and a Constraint on Majorana Neutrino Masses”. In: *Phys. Rev. D* 42 (1990), pp. 1285–1286. DOI: 10.1103/PhysRevD.42.1285.
- [34] Masataka Fukugita et al. “CP violating phase from minimal texture neutrino mass matrix: Test of the phase relevant to leptogenesis”. In: *Phys. Lett. B* 764 (2017), pp. 163–166. DOI: 10.1016/j.physletb.2016.11.024. arXiv: 1609.01864 [hep-ph].
- [35] J. Racker, Manuel Pena, and Nuria Rius. “Leptogenesis with small violation of B-L”. In: *JCAP* 1207 (2012), p. 030. DOI: 10.1088/1475-7516/2012/07/030. arXiv: 1205.1948 [hep-ph].
- [36] Sacha Davidson and Alejandro Ibarra. “A Lower bound on the right-handed neutrino mass from leptogenesis”. In: *Phys. Lett. B* 535 (2002), pp. 25–32. DOI: 10.1016/S0370-2693(02)01735-5. arXiv: hep-ph/0202239 [hep-ph].
- [37] Thomas Hambye et al. “Constraints on neutrino masses from leptogenesis models”. In: *Nucl. Phys. B* 695 (2004), pp. 169–191. DOI: 10.1016/j.nuclphysb.2004.06.027. arXiv: hep-ph/0312203 [hep-ph].
- [38] Stefan Antusch et al. “Non-unitary Leptonic Mixing and Leptogenesis”. In: *JHEP* 01 (2010), p. 017. DOI: 10.1007/JHEP01(2010)017. arXiv: 0910.5957 [hep-ph].
- [39] Evgeny K. Akhmedov, V. A. Rubakov, and A. Yu. Smirnov. “Baryogenesis via neutrino oscillations”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1359–1362. DOI: 10.1103/PhysRevLett.81.1359. arXiv: hep-ph/9803255 [hep-ph].
- [40] Laurent Canetti, Marco Drewes, and Björn Garbrecht. “Probing leptogenesis with GeV-scale sterile neutrinos at LHCb and Belle II”. In: *Phys. Rev. D* 90.12 (2014), p. 125005. DOI: 10.1103/PhysRevD.90.125005. arXiv: 1404.7114 [hep-ph].
- [41] Marco Drewes et al. “Testing the low scale seesaw and leptogenesis”. In: *JHEP* 08 (2017), p. 018. DOI: 10.1007/JHEP08(2017)018. arXiv: 1609.09069 [hep-ph].
- [42] Valentin V. Khoze and Gunnar Ro. “Leptogenesis and Neutrino Oscillations in the Classically Conformal Standard Model with the Higgs Portal”. In: *JHEP* 10 (2013), p. 075. DOI: 10.1007/JHEP10(2013)075. arXiv: 1307.3764 [hep-ph].

- [43] Asmaa Abada et al. "Lepton number violation as a key to low-scale leptogenesis". In: *JCAP* 1511.11 (2015), p. 041. DOI: 10.1088/1475-7516/2015/11/041. arXiv: 1507.06215 [hep-ph].
- [44] Alexander Kartavtsev, Peter Millington, and Hendrik Vogel. "Lepton asymmetry from mixing and oscillations". In: *JHEP* 06 (2016), p. 066. DOI: 10.1007/JHEP06(2016)066. arXiv: 1601.03086 [hep-ph].
- [45] P. Hernández et al. "Testable Baryogenesis in Seesaw Models". In: *JHEP* 08 (2016), p. 157. DOI: 10.1007/JHEP08(2016)157. arXiv: 1606.06719 [hep-ph].
- [46] Scott Dodelson and Lawrence M. Widrow. "Sterile-neutrinos as dark matter". In: *Phys. Rev. Lett.* 72 (1994), pp. 17–20. arXiv: hep-ph/9303287 [hep-ph].
- [47] Xiang-Dong Shi and George M. Fuller. "A New DM candidate: Non-thermal sterile neutrinos". In: *Phys. Rev. Lett.* 82 (1999), pp. 2832–2835. DOI: 10.1103/PhysRevLett.82.2832. arXiv: astro-ph/9810076 [astro-ph].
- [48] S. Tremaine and J. E. Gunn. "Dynamical Role of Light Neutral Leptons in Cosmology". In: *Phys. Rev. Lett.* 42 (1979). [66(1979)], pp. 407–410. DOI: 10.1103/PhysRevLett.42.407.
- [49] Esra Bulbul et al. "Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters". In: *Astrophys. J.* 789 (2014), p. 13. DOI: 10.1088/0004-637X/789/1/13. arXiv: 1402.2301 [astro-ph.CO].
- [50] Alexey Boyarsky et al. "Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster". In: *Phys. Rev. Lett.* 113 (2014), p. 251301. DOI: 10.1103/PhysRevLett.113.251301. arXiv: 1402.4119 [astro-ph.CO].
- [51] Christopher Dessert, Nicholas L. Rodd, and Benjamin R. Safdi. "Evidence against the decaying dark matter interpretation of the 3.5 keV line from blank sky observations". In: (2018). arXiv: 1812.06976.
- [52] Michael E. Anderson, Eugene Churazov, and Joel N. Bregman. "Non-Detection of X-Ray Emission From Sterile Neutrinos in Stacked Galaxy Spectra". In: *Mon. Not. Roy. Astron. Soc.* 452.4 (2015), pp. 3905–3923. DOI: 10.1093/mnras/stv1559. arXiv: 1408.4115 [astro-ph.HE].
- [53] F. A. Aharonian et al. "Hitomi constraints on the 3.5 keV line in the Perseus galaxy cluster". In: *Astrophys. J.* 837.1 (2017), p. L15. DOI: 10.3847/2041-8213/aa61fa. arXiv: 1607.07420 [astro-ph.HE].

- [54] Marco Drewes. "The Phenomenology of Right Handed Neutrinos". In: *Int. J. Mod. Phys. E* 22 (2013), p. 1330019. arXiv: 1303.6912 [hep-ph].
- [55] Andre de Gouvea. "See-saw energy scale and the LSND anomaly". In: *Phys. Rev. D*. 72 (2005), p. 033005. DOI: 10.1103/PhysRevD.72.033005. arXiv: hep-ph/0501039 [hep-ph].
- [56] Oleg Ruchayskiy and Artem Ivashko. "Restrictions on the lifetime of sterile neutrinos from primordial nucleosynthesis". In: *JCAP* 1210 (2012), p. 014. DOI: 10.1088/1475-7516/2012/10/014. arXiv: 1202.2841 [hep-ph].
- [57] Oleg Ruchayskiy and Artem Ivashko. "Experimental bounds on sterile neutrino mixing angles". In: *JHEP* 06 (2012), p. 100. DOI: 10.1007/JHEP06(2012)100. arXiv: 1112.3319 [hep-ph].
- [58] P. Hernandez, M. Kekic, and J. Lopez-Pavon. " N_{eff} in low-scale seesaw models versus the lightest neutrino mass". In: *Phys. Rev. D* 90.6 (2014), p. 065033. DOI: 10.1103/PhysRevD.90.065033. arXiv: 1406.2961 [hep-ph].
- [59] Marco Drewes and Björn Garbrecht. "Combining experimental and cosmological constraints on heavy neutrinos". In: *Nucl. Phys.* B921 (2017), pp. 250–315. DOI: 10.1016/j.nuclphysb.2017.05.001. arXiv: 1502.00477 [hep-ph].
- [60] Alexander Dueck and Werner Rodejohann. "Fits to SO(10) Grand Unified Models". In: *JHEP* 09 (2013), p. 024. DOI: 10.1007/JHEP09(2013)024. arXiv: 1306.4468 [hep-ph].
- [61] Francesco Vissani. "Do experiments suggest a hierarchy problem?" In: *Phys. Rev. D*. 57 (1998), pp. 7027–7030. DOI: 10.1103/PhysRevD.57.7027. arXiv: hep-ph/9709409 [hep-ph].
- [62] T. Asaka, S. Blanchet, and M. Shaposhnikov. "The nuMSM, dark matter and neutrino masses". In: *Phys. Lett. B*. 631 (2005), pp. 151–156. DOI: 10.1016/j.physletb.2005.09.070. arXiv: hep-ph/0503065 [hep-ph].
- [63] 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.
- [64] D. Wyler and L. Wolfenstein. "Massless Neutrinos in Left-Right Symmetric Models". In: *N. Phys. B*. 218 (1983), pp. 205–214. DOI: 10.1016/0550-3213(83)90482-0.
- [65] Jörn Kersten and Alexei Yu. Smirnov. "Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation". In:

- Phys. Rev. D* 76 (2007), p. 073005. DOI: 10.1103/PhysRevD.76.073005. arXiv: 0705.3221 [hep-ph].
- [66] Mikhail Shaposhnikov. "A Possible symmetry of the nuMSM". In: *N. Phys. B.* 763 (2007), pp. 49–59. DOI: 10.1016/j.nuclphysb.2006.11.003. arXiv: hep-ph/0605047 [hep-ph].
- [67] Valentin V. Khoze and Alexis D. Plascencia. "Dark Matter and Leptogenesis Linked by Classical Scale Invariance". In: *JHEP* 11 (2016), p. 025. DOI: 10.1007/JHEP11(2016)025. arXiv: 1605.06834 [hep-ph].
- [68] G. Bernardi et al. "Further Limits on Heavy Neutrino Couplings". In: *Phys. Lett. B* 203 (1988), pp. 332–334. DOI: 10.1016/0370-2693(88)90563-1.
- [69] F. Bergsma et al. "A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-GeV to 2.8-GeV". In: *Phys. Lett.* 166B (1986), pp. 473–478. DOI: 10.1016/0370-2693(86)91601-1.
- [70] A. Vaitaitis et al. "Search for neutral heavy leptons in a high-energy neutrino beam". In: *Phys. Rev. Lett.* 83 (1999), pp. 4943–4946. DOI: 10.1103/PhysRevLett.83.4943. arXiv: hep-ex/9908011 [hep-ex].
- [71] Amanda M. Cooper-Sarkar et al. "Search for Heavy Neutrino Decays in the BEBC Beam Dump Experiment". In: *Phys. Lett.* 160B (1985), pp. 207–211. DOI: 10.1016/0370-2693(85)91493-5.
- [72] P. Astier et al. "Search for heavy neutrinos mixing with tau neutrinos". In: *Phys. Lett. B* 506 (2001), pp. 27–38. DOI: 10.1016/S0370-2693(01)00362-8. arXiv: hep-ex/0101041 [hep-ex].
- [73] P. Abreu et al. "A Study of the reaction $e^+e^- \rightarrow \mu^+\mu^-\gamma$ (ISR) at LEP and search for new physics at annihilation energies near 80 GeV". In: *Z. Phys. C* 75 (1997), pp. 581–592. DOI: 10.1007/s002880050504.
- [74] Arindam Das and Nobuchika Okada. "Improved bounds on the heavy neutrino productions at the LHC". In: *Phys. Rev. D* 93.3 (2016), p. 033003. DOI: 10.1103/PhysRevD.93.033003. arXiv: 1510.04790 [hep-ph].
- [75] Arindam Das and Nobuchika Okada. "Bounds on heavy Majorana neutrinos in type-I seesaw and implications for collider searches". In: *Phys. Lett. B* 774 (2017), pp. 32–40. DOI: 10.1016/j.physletb.2017.09.042. arXiv: 1702.04668 [hep-ph].
- [76] Arindam Das, Nobuchika Okada, and Digesh Raut. "Enhanced pair production of heavy Majorana neutrinos at the LHC". In: *Phys. Rev. D* 97.11 (2018), p. 115023. DOI: 10.1103/PhysRevD.97.115023. arXiv: 1710.03377 [hep-ph].

- [77] S. M. Bilenky and C. Giunti. “Neutrinoless Double-Beta Decay: a Probe of Physics Beyond the Standard Model”. In: *Int. J. Mod. Phys. A* 30.04n05 (2015), p. 1530001. DOI: 10.1142/S0217751X1530001X. arXiv: 1411.4791 [hep-ph].
- [78] J. A. Casas and A. Ibarra. “Oscillating neutrinos and $\mu \rightarrow e\gamma$ ”. In: *Nucl. Phys. B* 618 (2001), pp. 171–204. DOI: 10.1016/S0550-3213(01)00475-8. arXiv: hep-ph/0103065 [hep-ph].
- [79] A. Abada et al. “Tree-level lepton universality violation in the presence of sterile neutrinos: impact for R_K and R_π ”. In: *JHEP* 02 (2013), p. 048. DOI: 10.1007/JHEP02(2013)048. arXiv: 1211.3052 [hep-ph].
- [80] Stefan Antusch and Oliver Fischer. “Non-unitarity of the leptonic mixing matrix: Present bounds and future sensitivities”. In: *JHEP* 10 (2014), p. 094. DOI: 10.1007/JHEP10(2014)094. arXiv: 1407.6607 [hep-ph].
- [81] F. del Aguila, J. de Blas, and M. Perez-Victoria. “Effects of new leptons in Electroweak Precision Data”. In: *Phys. Rev. D* 78 (2008), p. 013010. DOI: 10.1103/PhysRevD.78.013010. arXiv: 0803.4008 [hep-ph].
- [82] 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].
- [83] Dmitry Gorbunov and Inar Timiryasov. “Testing ν MSM with indirect searches”. In: *Phys. Lett. B* 745 (2015), pp. 29–34. DOI: 10.1016/j.physletb.2015.02.060. arXiv: 1412.7751 [hep-ph].
- [84] Takehiko Asaka, Shintaro Eijima, and Hiroyuki Ishida. “Mixing of Active and Sterile Neutrinos”. In: *JHEP* 04 (2011), p. 011. DOI: 10.1007/JHEP04(2011)011. arXiv: 1101.1382 [hep-ph].
- [85] Takehiko Asaka, Shintaro Eijima, and Kazuhiro Takeda. “Lepton Universality in the ν MSM”. In: *Phys. Lett. B* 742 (2015), pp. 303–309. DOI: 10.1016/j.physletb.2015.01.049. arXiv: 1410.0432 [hep-ph].
- [86] Enrique Fernandez-Martinez et al. “Loop level constraints on See-saw neutrino mixing”. In: *JHEP* 10 (2015), p. 130. arXiv: 1508.03051 [hep-ph].
- [87] E. Fernandez-Martinez, J. Hernandez-Garcia, and J. Lopez-Pavon. “Global constraints on heavy neutrino mixing”. In: *JHEP* 08 (2016), p. 033. DOI: 10.1007/JHEP08(2016)033. arXiv: 1605.08774 [hep-ph].
- [88] A. de Gouvea and A. Kobach. “Global Constraints on a Heavy Neutrino”. In: *Phys. Rev. D* 93.3 (2016), p. 033005. DOI: 10.1103/PhysRevD.93.033005. arXiv: 1511.00683 [hep-ph].

- [89] Alexey Boyarsky et al. “The Masses of active neutrinos in the nuMSM from X-ray astronomy”. In: *JETP Lett.* 83 (2006), pp. 133–135. DOI: 10.1134/S0021364006040011. arXiv: hep-ph/0601098 [hep-ph].
- [90] Apostolos Pilaftsis and Thomas E. J. Underwood. “Resonant leptogenesis”. In: *N. Phys. B.* 692 (2004), pp. 303–345. DOI: 10.1016/j.nuclphysb.2004.05.029. arXiv: hep-ph/0309342 [hep-ph].
- [91] Marco Drewes and Björn Garbrecht. “Leptogenesis from a GeV Seesaw without Mass Degeneracy”. In: *JHEP* 03 (2013), p. 096. DOI: 10.1007/JHEP03(2013)096. arXiv: 1206.5537 [hep-ph].
- [92] Bjorn Garbrecht. “More Viable Parameter Space for Leptogenesis”. In: *Phys. Rev. D.* 90.6 (2014), p. 063522. DOI: 10.1103/PhysRevD.90.063522. arXiv: 1401.3278 [hep-ph].
- [93] P. Athron et al. “GAMBIT: The Global and Modular Beyond the Standard Model Inference Tool”. In: *EPJC* 77 (May 2017), p. 784. arXiv: 1705.07908 [hep-ph].
- [94] Gregory D. Martinez et al. “Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module”. In: *Eur. Phys. J. C* 77.11 (2017), p. 761. DOI: 10.1140/epjc/s10052-017-5274-y. arXiv: 1705.07959 [hep-ph].
- [95] Yashar Akrami et al. “A Profile Likelihood Analysis of the Constrained MSSM with Genetic Algorithms”. In: *JHEP* 04 (2010), p. 057. DOI: 10.1007/JHEP04(2010)057. arXiv: 0910.3950 [hep-ph].
- [96] Ghulam Mustafa, Faisal Akram, and Bilal Masud. “Optimization of Neutrino Oscillation Parameters using Differential Evolution”. In: *Commun. Theor. Phys.* 59 (2013), pp. 324–330. DOI: 10.1088/0253-6102/59/3/14. arXiv: 1109.2431 [physics.comp-ph].
- [97] Peter Athron et al. “Status of the scalar singlet dark matter model”. In: *Eur. Phys. J. C* 77.8 (2017), p. 568. DOI: 10.1140/epjc/s10052-017-5113-1. arXiv: 1705.07931 [hep-ph].
- [98] Peter Athron et al. “Global analyses of Higgs portal singlet dark matter models using GAMBIT”. In: *Eur. Phys. J. C* 79.1 (2019), p. 38. DOI: 10.1140/epjc/s10052-018-6513-6. arXiv: 1808.10465 [hep-ph].
- [99] Peter Athron et al. “A global fit of the MSSM with GAMBIT”. In: *Eur. Phys. J. C* 77.12 (2017), p. 879. DOI: 10.1140/epjc/s10052-017-5196-8. arXiv: 1705.07917 [hep-ph].
- [100] Peter Athron et al. “Global fits of GUT-scale SUSY models with GAMBIT”. In: *Eur. Phys. J. C* 77.12 (2017), p. 824. DOI: 10.1140/epjc/s10052-017-5167-0. arXiv: 1705.07935 [hep-ph].

- [101] Sebastian Hoof et al. "Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT". In: (2018). arXiv: 1810.07192 [hep-ph].
- [102] Ivan Esteban et al. "Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity". In: *JHEP* 01 (2017), p. 087. DOI: 10.1007/JHEP01(2017)087. arXiv: 1611.01514 [hep-ph].
- [103] Ivan Esteban et al. *NuFit 3.2 (2018)*. URL: <http://www.nu-fit.org> (visited on 05/08/2018).
- [104] Andrea Ferroglia and Alberto Sirlin. "Comparison of the standard theory predictions of M_W and $\sin^2\theta_{eff}^{lept}$ with their experimental values". In: *Phys. Rev. D* 87.3 (2013), p. 037501. DOI: 10.1103/PhysRevD.87.037501. arXiv: 1211.1864 [hep-ph].
- [105] Y. Amhis et al. "Averages of b -hadron, c -hadron, and τ -lepton properties as of summer 2016". In: (2016). arXiv: 1612.07233 [hep-ex].
- [106] R. Alonso et al. "Muon conversion to electron in nuclei in type-I seesaw models". In: *JHEP* 01 (2013), p. 118. DOI: 10.1007/JHEP01(2013)118. arXiv: 1209.2679 [hep-ph].
- [107] R. Aaij et al. "Test of lepton universality with $B^0 \rightarrow K^{*0}\ell^+\ell^-$ decays". In: *JHEP* 08 (2017), p. 055. DOI: 10.1007/JHEP08(2017)055. arXiv: 1705.05802 [physics.data-an].
- [108] Roel Aaij et al. "Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \rightarrow D^{*+}\tau^-\bar{\nu}_\tau)/\mathcal{B}(\bar{B}^0 \rightarrow D^{*+}\mu^-\bar{\nu}_\mu)$ ". In: *Phys. Rev. Lett.* 115.11 (2015). DOI: 10.1103/PhysRevLett.115.159901, 10.1103/PhysRevLett.115.111803. arXiv: 1506.08614 [hep-ex].
- [109] Roel Aaij et al. "Test of lepton universality using $B^+ \rightarrow K^+\ell^+\ell^-$ decays". In: *Phys. Rev. Lett.* 113 (2014), p. 151601. arXiv: 1406.6482 [hep-ex].
- [110] Marco Drewes et al. "NA62 sensitivity to heavy neutral leptons in the low scale seesaw model". In: *JHEP* 07 (2018), p. 105. DOI: 10.1007/JHEP07(2018)105. arXiv: 1801.04207 [hep-ph].
- [111] Asmaa Abada et al. "Low-scale leptogenesis with three heavy neutrinos". In: *JHEP* 01 (2019). [JHEP19,164(2020)], p. 164. DOI: 10.1007/JHEP01(2019)164. arXiv: 1810.12463 [hep-ph].
- [112] A. C. Hayes and Petr Vogel. "Reactor Neutrino Spectra". In: *Ann. Rev. Nucl. Part. Sci.* 66 (2016), pp. 219–244. DOI: 10.1146/annurev-nucl-102115-044826. arXiv: 1605.02047 [hep-ph].

- [113] E. Kearns, T. Kajita, and Y. Totsuka. "Detecting massive neutrinos". In: *Sci. Am.* 281N2 (1999). [Spektrum Wiss. Dossier2003N1,64(2003)], pp. 48–55.
- [114] R. Schaeffer, Y. Declais, and S. Jullian. "The Neutrino Emission of SN1987A". In: *Nature* 330 (1987), pp. 142–144. DOI: 10.1038/330142a0.
- [115] Charles D. Dermer. "Best-Bet Astrophysical Neutrino Sources". In: *J. Phys. Conf. Ser.* 60 (2007), pp. 8–13. DOI: 10.1088/1742-6596/60/1/002. arXiv: astro-ph/0611191 [astro-ph].
- [116] Teresa Montaruli. "First results of the IceCube observatory on high energy neutrino astronomy". In: *J. Phys. Conf. Ser.* 120 (2008), p. 062009. DOI: 10.1088/1742-6596/120/6/062009.
- [117] Tommaso Chiarusi. "Neutrino astrophysics with the ANTARES Cherenkov Detector". In: *J. Phys. Conf. Ser.* 409.1 (2013), p. 012133. DOI: 10.1088/1742-6596/409/1/012133.
- [118] M. G. Aartsen et al. "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector". In: *Science* 342 (2013), p. 1242856. DOI: 10.1126/science.1242856. arXiv: 1311.5238 [astro-ph.HE].
- [119] U. F. Katz. "Status and progress of the KM3NeT project". In: *Italian Phys. Soc. Proc.* 103 (2011), pp. 563–574.
- [120] David Waters. "Latest Results from NEMO-3 & Status of the SuperNEMO Experiment". In: *J. Phys. Conf. Ser.* 888.1 (2017), p. 012033. DOI: 10.1088/1742-6596/888/1/012033.
- [121] Petros A. Rapidis. "The NESTOR underwater neutrino telescope project". In: *Nucl. Instrum. Meth. A* 602 (2009), pp. 54–57. DOI: 10.1016/j.nima.2008.12.216.
- [122] S. Bilenky. "Neutrino oscillations: From a historical perspective to the present status". In: *Nucl. Phys. B* 908 (2016), pp. 2–13. DOI: 10.1016/j.nuclphysb.2016.01.025. arXiv: 1602.00170 [hep-ph].
- [123] H. A. Bethe. "Energy production in stars". In: *Phys. Rev.* 55 (1939), pp. 434–456. DOI: 10.1103/PhysRev.55.434.
- [124] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. "Search for neutrinos from the sun". In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [125] John N. Bahcall, Neta A. Bahcall, and G. Shaviv. "Present status of the theoretical predictions for the Cl-36 solar neutrino experiment". In: *Phys. Rev. Lett.* 20 (1968). [45(1968)], pp. 1209–1212. DOI: 10.1103/PhysRevLett.20.1209.

- [126] B. Pontecorvo. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge". In: *Sov. Phys. JETP* 26 (1968). [Zh. Eksp. Teor. Fiz.53,1717(1967)], pp. 984–988.
- [127] K. A. Olive et al. "Neutrino mass, mixing and oscillations". In: *Chin. Phys. C* 38 (2014).
- [128] Y. Fukuda et al. "Evidence for oscillation of atmospheric neutrinos". In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003 [hep-ex].
- [129] S. N. Ahmed et al. "Measurement of the total active B-8 solar neutrino flux at the Sudbury Neutrino Observatory with enhanced neutral current sensitivity". In: *Phys. Rev. Lett.* 92 (2004), p. 181301. DOI: 10.1103/PhysRevLett.92.181301. arXiv: nucl-ex/0309004 [nucl-ex].
- [130] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the unified model of elementary particles". In: *Prog. Theor. Phys.* 28 (1962). [34(1962)], pp. 870–880. DOI: 10.1143/PTP.28.870.
- [131] M. Tanabashi et al. "Review of Particle Physics". In: *Phys. Rev. D* 98.3 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001.
- [132] Res Jost. "Eine Bemerkung zum PCT Theorem". In: *Helv. Phys. Acta* 30 (1957), p. 409.
- [133] Andrei D Sakharov. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe". In: *Soviet Physics Uspekhi* 34.5 (May 1991), pp. 392–393. DOI: 10.1070/pu1991v034n05abeh002497. URL: <https://doi.org/10.1070%2Fpu1991v034n05abeh002497>.
- [134] M. Agostini et al. "First Simultaneous Precision Spectroscopy of pp , ${}^7\text{Be}$, and pep Solar Neutrinos with Borexino Phase-II". In: (2017). arXiv: 1707.09279 [hep-ex].
- [135] M. Agostini et al. "Improved measurement of 8B solar neutrinos with 1.5 kt y of Borexino exposure". In: (2017). arXiv: 1709.00756 [hep-ex].
- [136] M Ambrosio et al. "Atmospheric neutrino oscillations from upward through going muon multiple scattering in MACRO". In: *Phys. Lett. B* 566 (2003), pp. 35–44. DOI: 10.1016/S0370-2693(03)00806-2. arXiv: hep-ex/0304037 [hep-ex].
- [137] R. Abbasi et al. "The Design and Performance of IceCube DeepCore". In: *Astropart. Phys.* 35 (2012), pp. 615–624. arXiv: 1109.6096.
- [138] J. A. Formaggio. "Direct neutrino mass measurements after PLANCK". In: *Phys. Dark Univ.* 4 (2014), pp. 75–80. DOI: 10.1016/j.dark.2014.10.004.

- [139] Amand Faessler et al. "Arbitrary mass Majorana neutrinos in neutrinoless double beta decay". In: *Phys. Rev. D* 90.9 (2014), p. 096010. DOI: 10.1103/PhysRevD.90.096010. arXiv: 1408.6077 [hep-ph].
- [140] M. Agostini et al. "Improved Limit on Neutrinoless Double- β Decay of ^{76}Ge from GERDA Phase II". In: *Phys. Rev. Lett.* 120.13 (2018), p. 132503. DOI: 10.1103/PhysRevLett.120.132503. arXiv: 1803.11100 [nucl-ex].
- [141] A. Gando et al. "Limit on Neutrinoless $\beta\beta$ Decay of ^{136}Xe from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in ^{76}Ge ". In: *Phys. Rev. Lett.* 110.6 (2013), p. 062502. DOI: 10.1103/PhysRevLett.110.062502. arXiv: 1211.3863 [hep-ex].
- [142] J. B. Albert et al. "Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector". In: *Phys. Rev. Lett.* 120.7 (2018), p. 072701. DOI: 10.1103/PhysRevLett.120.072701. arXiv: 1707.08707 [hep-ex].
- [143] R. K. Sachs and A. M. Wolfe. "Perturbations of a cosmological model and angular variations of the microwave background". In: *Astrophys. J.* 147 (1967). [Gen. Rel. Grav.39,1929(2007)], pp. 73–90. DOI: 10.1007/s10714-007-0448-9.
- [144] Z. Hou et al. "Constraints on Cosmology from the Cosmic Microwave Background Power Spectrum of the 2500 deg² SPT-SZ Survey". In: *Astrophys. J.* 782 (2014), p. 74. DOI: 10.1088/0004-637X/782/2/74. arXiv: 1212.6267 [astro-ph.CO].
- [145] Sai Wang, Yi-Fan Wang, and Dong-Mei Xia. "Constraints on the sum of neutrino masses using cosmological data including the latest extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample". In: *Chin. Phys.* C42.6 (2018), p. 065103. DOI: 10.1088/1674-1137/42/6/065103. arXiv: 1707.00588 [astro-ph.CO].
- [146] N. Aghanim et al. "Planck 2018 results. VI. Cosmological parameters". In: (2018). arXiv: 1807.06209 [astro-ph.CO].
- [147] Ettore Majorana. "Teoria simmetrica dell'elettrone e del positrone". In: *Nuovo Cim.* 14 (1937), pp. 171–184. DOI: 10.1007/BF02961314.
- [148] Palash B. Pal. "Dirac, Majorana and Weyl fermions". In: *Am. J. Phys.* 79 (2011), pp. 485–498. DOI: 10.1119/1.3549729. arXiv: 1006.1718 [hep-ph].
- [149] R. N. Mohapatra and P. B. Pal. "Massive neutrinos in physics and astrophysics. Second edition". In: *World Sci. Lect. Notes Phys.* 60 (1998). [World Sci. Lect. Notes Phys.72,1(2004)], pp. 1–397.

- [150] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 9780198508717.
- [151] Apostolos Pilaftsis. “Radiatively induced neutrino masses and large Higgs neutrino couplings in the standard model with Majorana fields”. In: *Z. Phys.* C55 (1992), pp. 275–282. DOI: 10.1007/BF01482590. arXiv: hep-ph/9901206 [hep-ph].
- [152] M. Magg and C. Wetterich. “Neutrino Mass Problem and Gauge Hierarchy”. In: *Phys. Lett.* 94B (1980), pp. 61–64. DOI: 10.1016/0370-2693(80)90825-4.
- [153] George Lazarides, Q. Shafi, and C. Wetterich. “Proton Lifetime and Fermion Masses in an SO(10) Model”. In: *Nucl. Phys.* B181 (1981), pp. 287–300. DOI: 10.1016/0550-3213(81)90354-0.
- [154] Pei-Hong Gu, He Zhang, and Shun Zhou. “A Minimal Type II Seesaw Model”. In: *Phys. Rev.* D74 (2006), p. 076002. DOI: 10.1103/PhysRevD.74.076002. arXiv: hep-ph/0606302 [hep-ph].
- [155] Alejandra Melfo et al. “Type II Seesaw at LHC: The Roadmap”. In: *Phys. Rev.* D85 (2012), p. 055018. DOI: 10.1103/PhysRevD.85.055018. arXiv: 1108.4416 [hep-ph].
- [156] Zhi-Long Han, Ran Ding, and Yi Liao. “LHC Phenomenology of Type II Seesaw: Nondegenerate Case”. In: *Phys. Rev.* D91 (2015), p. 093006. DOI: 10.1103/PhysRevD.91.093006. arXiv: 1502.05242 [hep-ph].
- [157] Robert Foot et al. “Seesaw Neutrino Masses Induced by a Triplet of Leptons”. In: *Z. Phys.* C44 (1989), p. 441. DOI: 10.1007/BF01415558.
- [158] A. Abada et al. “ $\mu \rightarrow e\gamma$ and $\tau \rightarrow l\gamma$ decays in the fermion triplet seesaw model”. In: *Phys. Rev.* D78 (2008), p. 033007. DOI: 10.1103/PhysRevD.78.033007. arXiv: 0803.0481 [hep-ph].
- [159] Pavel Fileviez Perez. “Type III Seesaw and Left-Right Symmetry”. In: *JHEP* 03 (2009), p. 142. DOI: 10.1088/1126-6708/2009/03/142. arXiv: 0809.1202 [hep-ph].
- [160] Carl H. Albright and S. M. Barr. “Leptogenesis in the type III seesaw mechanism”. In: *Phys. Rev.* D69 (2004), p. 073010. DOI: 10.1103/PhysRevD.69.073010. arXiv: hep-ph/0312224 [hep-ph].
- [161] R. N. Mohapatra and J. W. F. Valle. “Neutrino Mass and Baryon Number Nonconservation in Superstring Models”. In: *Phys. Rev.* D34 (1986). [235(1986)], p. 1642. DOI: 10.1103/PhysRevD.34.1642.
- [162] M. C. Gonzalez-Garcia and J. W. F. Valle. “Fast Decaying Neutrinos and Observable Flavor Violation in a New Class of Majoron Models”.

- In: *Phys. Lett.* B216 (1989), pp. 360–366. DOI: 10.1016/0370-2693(89)91131-3.
- [163] F. Deppisch and J. W. F. Valle. “Enhanced lepton flavor violation in the supersymmetric inverse seesaw model”. In: *Phys. Rev. D* 72 (2005), p. 036001. DOI: 10.1103/PhysRevD.72.036001. arXiv: hep-ph/0406040 [hep-ph].
- [164] Asmaa Abada and Michele Lucente. “Looking for the minimal inverse seesaw realisation”. In: *Nucl. Phys.* B885 (2014), pp. 651–678. DOI: 10.1016/j.nuclphysb.2014.06.003. arXiv: 1401.1507 [hep-ph].
- [165] J. Lopez-Pavon, E. Molinaro, and S. T. Petcov. “Radiative Corrections to Light Neutrino Masses in Low Scale Type I Seesaw Scenarios and Neutrinoless Double Beta Decay”. In: *JHEP* 11 (2015), p. 030. DOI: 10.1007/JHEP11(2015)030. arXiv: 1506.05296 [hep-ph].
- [166] M. Fukugita and T. Yanagida. *Physics of neutrinos and applications to astrophysics*. 2003.
- [167] R. G. Hamish Robertson. “KATRIN: an experiment to determine the neutrino mass from the beta decay of tritium”. In: *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. arXiv: 1307.5486 [physics.ins-det].
- [168] C. Ahdida et al. “Sensitivity of the SHiP experiment to Heavy Neutral Leptons”. In: (2018). arXiv: 1811.00930 [hep-ph].
- [169] K. Abe et al. “Search for heavy neutrinos with the T2K near detector ND280”. In: (2019). arXiv: 1902.07598 [hep-ex].
- [170] K. Moffat, S. Pascoli, and C. Weiland. “Equivalence between massless neutrinos and lepton number conservation in fermionic singlet extensions of the Standard Model”. In: (2017). arXiv: 1712.07611 [hep-ph].
- [171] Marco Drewes, Juraj Klarić, and Philipp Klose. “On Lepton Number Violation in Heavy Neutrino Decays at Colliders”. In: (2019). arXiv: 1907.13034 [hep-ph].
- [172] G. Anamjati, M. Hirsch, and E. Nardi. “Quasi-Dirac neutrinos at the LHC”. In: *JHEP* 10 (2016), p. 010. DOI: 10.1007/JHEP10(2016)010. arXiv: 1607.05641 [hep-ph].
- [173] Stefan Antusch, Eros Cazzato, and Oliver Fischer. “Heavy neutrino-antineutrino oscillations at colliders”. In: (2017). arXiv: 1709.03797 [hep-ph].

- [174] Stefan Antusch et al. “Probing Leptogenesis at Future Colliders”. In: *JHEP* 09 (2018), p. 124. DOI: 10.1007/JHEP09(2018)124. arXiv: 1710.03744 [hep-ph].
- [175] Daniel Boyanovsky. “Nearly degenerate heavy sterile neutrinos in cascade decay: mixing and oscillations”. In: *Phys. Rev. D* 90.10 (2014). DOI: 10.1103/PhysRevD.90.105024. arXiv: 1409.4265 [hep-ph].
- [176] Gorazd Cvetič et al. “Oscillation of heavy sterile neutrino in decay of $B \rightarrow \mu e \pi$ ”. In: *Phys. Rev. D* 92 (2015), p. 013015. DOI: 10.1103/PhysRevD.92.013015. arXiv: 1505.04749 [hep-ph].
- [177] Claudio O. Dib et al. “Distinguishing Dirac/Majorana Sterile Neutrinos at the LHC”. In: *Phys. Rev. D* 94.1 (2016), p. 013005. DOI: 10.1103/PhysRevD.94.013005. arXiv: 1605.01123 [hep-ph].
- [178] Arindam Das, P. S. Bhupal Dev, and Rabindra N. Mohapatra. “Same Sign versus Opposite Sign Dileptons as a Probe of Low Scale Seesaw Mechanisms”. In: *Phys. Rev. D* 97.1 (2018), p. 015018. DOI: 10.1103/PhysRevD.97.015018. arXiv: 1709.06553 [hep-ph].
- [179] Gorazd Cvetič, Arindam Das, and Jilberto Zamora-Saá. “Probing heavy neutrino oscillations in rare W boson decays”. In: *J. Phys. G* 46 (2019), p. 075002. DOI: 10.1088/1361-6471/ab1212. arXiv: 1805.00070 [hep-ph].
- [180] P. Hernández, J. Jones-Pérez, and O. Suarez-Navarro. “Majorana vs Pseudo-Dirac Neutrinos at the ILC”. In: *Eur. Phys. J. C* 79.3 (2019), p. 220. DOI: 10.1140/epjc/s10052-019-6728-1. arXiv: 1810.07210 [hep-ph].
- [181] Gorazd Cvetič et al. “Measuring the heavy neutrino oscillations in rare W boson decays at the Large Hadron Collider”. In: (2019). arXiv: 1905.03097 [hep-ph].
- [182] Takehiko Asaka and Takanao Tsuyuki. “Perturbativity in the seesaw mechanism”. In: *Phys. Lett. B* 753 (2016), pp. 147–149. DOI: 10.1016/j.physletb.2015.12.013. arXiv: 1509.02678 [hep-ph].
- [183] R. N. Mohapatra. “Mechanism for Understanding Small Neutrino Mass in Superstring Theories”. In: *Phys. Rev. Lett.* 56 (1986), pp. 561–563. DOI: 10.1103/PhysRevLett.56.561.
- [184] J. Bernabeu et al. “Lepton Flavor Nonconservation at High-Energies in a Superstring Inspired Standard Model”. In: *Phys. Lett. B* 187 (1987), pp. 303–308. DOI: 10.1016/0370-2693(87)91100-2.

- [185] Evgeny K. Akhmedov et al. "Left-right symmetry breaking in NJL approach". In: *Phys. Lett.* B368 (1996), pp. 270–280. DOI: 10.1016/0370-2693(95)01504-3. arXiv: hep-ph/9507275 [hep-ph].
- [186] Evgeny K. Akhmedov et al. "Dynamical left-right symmetry breaking". In: *Phys. Rev. D* 53 (1996), pp. 2752–2780. DOI: 10.1103/PhysRevD.53.2752. arXiv: hep-ph/9509255 [hep-ph].
- [187] Thomas Appelquist and Robert Shrock. "Neutrino masses in theories with dynamical electroweak symmetry breaking". In: *Phys. Lett.* B548 (2002), pp. 204–214. DOI: 10.1016/S0370-2693(02)02854-X. arXiv: hep-ph/0204141 [hep-ph].
- [188] Thomas Appelquist and Robert Shrock. "Dynamical symmetry breaking of extended gauge symmetries". In: *Phys. Rev. Lett.* 90 (2003). DOI: 10.1103/PhysRevLett.90.201801. arXiv: hep-ph/0301108 [hep-ph].
- [189] Louis Lyons. "Bayes and Frequentism: a Particle Physicist's perspective". In: *Contemp. Phys.* 54 (2013), p. 1. DOI: 10.1080/00107514.2012.756312. arXiv: 1301.1273 [physics.data-an].
- [190] T. Bayes. "An essay towards solving a problem in the doctrine of chances". In: *Phil. Trans. of the Royal Soc. of London* 53 (1763), pp. 370–418.
- [191] Pierre Simon Laplace. "Memoir on the Probability of the Causes of Events". In: *Statist. Sci.* 1.3 (Aug. 1986), pp. 364–378. DOI: 10.1214/ss/1177013621. URL: <https://doi.org/10.1214/ss/1177013621>.
- [192] H. Jeffreys. *Theory of Probability*. Third. Oxford, 1961.
- [193] J. Neyman. "Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability". In: *Phil. Trans. Roy. Soc. Lond.* A236.767 (1937), pp. 333–380. DOI: 10.1098/rsta.1937.0005.
- [194] Phil Gregory. *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica® Support*. Cambridge University Press, 2005. DOI: 10.1017/CB09780511791277.
- [195] S. S. Wilks. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". In: *Ann. Math. Statist.* 9.1 (Mar. 1938), pp. 60–62. DOI: 10.1214/aoms/1177732360. URL: <https://doi.org/10.1214/aoms/1177732360>.
- [196] C. L. Cowan et al. "Detection of the free neutrino: A Confirmation". In: *Science* 124 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103.
- [197] B de Finetti. *Theory of Probability: A critical introductory treatment*. Dec. 2016, pp. 1–582. DOI: 10.1002/9781119286387.

- [198] R. T. Cox. “Readings in Uncertain Reasoning”. In: ed. by Glenn Shafer and Judea Pearl. 1990. Chap. Probability, Frequency and Reasonable Expectation, pp. 353–365.
- [199] E. T. Jaynes. *Probability Theory: The Logic of Science*. Ed. by G. Larry Bretthorst. Cambridge University Press, 2003.
- [200] A. B. Arbuzov et al. “ZFITTER: A Semi-analytical program for fermion pair production in $e^+ e^-$ annihilation, from version 6.21 to version 6.42”. In: *Comput. Phys. Commun.* 174 (2006), pp. 728–758. DOI: 10 . 1016/j . cpc . 2005 . 12 . 009. arXiv: hep-ph/0507146 [hep-ph].
- [201] Arif Akhundov et al. “The ZFITTER project”. In: *Phys. Part. Nucl.* 45.3 (2014), pp. 529–549. DOI: 10 . 1134/S1063779614030022. arXiv: 1302 . 1395 [hep-ph].
- [202] M. Baak et al. “Updated Status of the Global Electroweak Fit and Constraints on New Physics”. In: *Eur. Phys. J. C* 72 (2012), p. 2003. DOI: 10 . 1140/epjc/s10052-012-2003-4. arXiv: 1107 . 0975 [hep-ph].
- [203] J. Charles et al. “CP violation and the CKM matrix: Assessing the impact of the asymmetric B factories”. In: *Eur. Phys. J. C* 41.1 (2005), pp. 1–131. DOI: 10 . 1140 / epjc / s2005 - 02169 - 1. arXiv: hep - ph / 0406184 [hep-ph].
- [204] D. V. Forero, M. Tortola, and J. W. F. Valle. “Neutrino oscillations refitted”. In: *Phys. Rev. D* 90.9 (2014), p. 093006. DOI: 10 . 1103/PhysRevD . 90 . 093006. arXiv: 1405 . 7540 [hep-ph].
- [205] P. F. de Salas et al. “Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved CP sensitivity”. In: *Phys. Lett. B* 782 (2018), pp. 633–640. DOI: 10 . 1016/j . physletb . 2018 . 06 . 019. arXiv: 1708 . 01186 [hep-ph].
- [206] Johannes Bergstrom et al. “Bayesian global analysis of neutrino oscillation data”. In: *JHEP* 09 (2015), p. 200. DOI: 10 . 1007/JHEP09(2015) 200. arXiv: 1507 . 04366 [hep-ph].
- [207] M. C. Gonzalez-Garcia, Michele Maltoni, and Thomas Schwetz. “Global Analyses of Neutrino Oscillation Experiments”. In: *Nucl. Phys. B* 908 (2016), pp. 199–217. DOI: 10 . 1016/j . nucphysb . 2016 . 02 . 033. arXiv: 1512 . 06856 [hep-ph].
- [208] Edward A. Baltz and Paolo Gondolo. “Markov chain Monte Carlo exploration of minimal supergravity with implications for dark matter”. In: *JHEP* 10 (2004), p. 052. DOI: 10 . 1088/1126 - 6708/2004/10/052. arXiv: hep-ph/0407039 [hep-ph].

- [209] B. C. Allanach and C. G. Lester. “Multi-dimensional mSUGRA likelihood maps”. In: *Phys. Rev. D* 73 (2006), p. 015013. DOI: 10.1103/PhysRevD.73.015013. arXiv: hep-ph/0507283 [hep-ph].
- [210] Remi Lafaye et al. “SFitter: Determining supersymmetric parameters”. In: *Statistical issues for LHC physics. Proceedings, Workshop, PHYSTAT-LHC, Geneva, Switzerland, June 27-29, 2007*. 2007, pp. 159–162. URL: <http://cds.cern.ch/record/1099985/files/p159.pdf>.
- [211] M. Rauch et al. “SFitter: Reconstructing the MSSM Lagrangian from LHC data”. In: *SUSY 2007 Proceedings, 15th International Conference on Supersymmetry and Unification of Fundamental Interactions, July 26 - August 1, 2007, Karlsruhe, Germany*. 2007, pp. 228–231. arXiv: 0710.2822 [hep-ph]. URL: <http://www.susy07.uni-karlsruhe.de/Proceedings/proceedings/susy07.pdf>.
- [212] Roberto Ruiz de Austri, Roberto Trotta, and Leszek Roszkowski. “A Markov chain Monte Carlo analysis of the CMSSM”. In: *JHEP* 05 (2006), p. 002. DOI: 10.1088/1126-6708/2006/05/002. arXiv: hep-ph/0602028 [hep-ph].
- [213] Leszek Roszkowski, Roberto Ruiz de Austri, and Roberto Trotta. “A Bayesian approach to the constrained MSSM”. In: *Statistical issues for LHC physics. Proceedings, Workshop, PHYSTAT-LHC, Geneva, Switzerland, June 27-29, 2007*. 2007, pp. 163–166. URL: <http://cds.cern.ch/record/1099986/files/p163.pdf>.
- [214] Roberto Trotta et al. “The Impact of priors and observables on parameter inferences in the Constrained MSSM”. In: *JHEP* 12 (2008), p. 024. DOI: 10.1088/1126-6708/2008/12/024. arXiv: 0809.3792 [hep-ph].
- [215] Farhan Feroz et al. “Challenges of Profile Likelihood Evaluation in Multi-Dimensional SUSY Scans”. In: *JHEP* 06 (2011), p. 042. DOI: 10.1007/JHEP06(2011)042. arXiv: 1101.3296 [hep-ph].
- [216] C. Strege et al. “Updated global fits of the cMSSM including the latest LHC SUSY and Higgs searches and XENON100 data”. In: *JCAP* 1203 (2012), p. 030. DOI: 10.1088/1475-7516/2012/03/030. arXiv: 1112.4192 [hep-ph].
- [217] C. Strege et al. “Profile likelihood maps of a 15-dimensional MSSM”. In: *JHEP* 09 (2014), p. 081. DOI: 10.1007/JHEP09(2014)081. arXiv: 1405.0622 [hep-ph].
- [218] P. Scott et al. “Use of event-level neutrino telescope data in global fits for theories of new physics”. In: *JCAP* 1211 (2012), p. 057. DOI: 10.1088/1475-7516/2012/11/057. arXiv: 1207.0810 [hep-ph].

- [219] Gianfranco Bertone et al. “Global analysis of the pMSSM in light of the Fermi GeV excess: prospects for the LHC Run-II and astroparticle experiments”. In: *JCAP* 1604.04 (2016), p. 037. DOI: 10 . 1088 / 1475 - 7516 / 2016 / 04 / 037. arXiv: 1507 . 07008 [hep-ph].
- [220] Philip Bechtle, Klaus Desch, and Peter Wienemann. “Fittino, a program for determining MSSM parameters from collider observables using an iterative method”. In: *Comput. Phys. Commun.* 174 (2006), pp. 47–70. DOI: 10 . 1016 / j . cpc . 2005 . 09 . 002. arXiv: hep-ph / 0412012 [hep-ph].
- [221] P. Bechtle, K. Desch, and P. Wienemann. “SUSY parameter measurements with Fittino”. In: *Proceedings, 2005 International Linear Collider Physics and Detector Workshop and 2nd ILC Accelerator Workshop (Snowmass 2005)*. 2005. arXiv: hep-ph / 0511137 [hep-ph]. URL: <http://www.slac.stanford.edu/econf/C0508141>.
- [222] O. Buchmueller et al. “Likelihood Functions for Supersymmetric Observables in Frequentist Analyses of the CMSSM and NUHM1”. In: *Eur. Phys. J. C* 64 (2009), pp. 391–415. DOI: 10 . 1140 / epjc / s10052 - 009 - 1159 - z. arXiv: 0907 . 5568 [hep-ph].
- [223] O. Buchmueller et al. “Frequentist Analysis of the Parameter Space of Minimal Supergravity”. In: *Eur. Phys. J. C* 71 (2011), p. 1583. DOI: 10 . 1140 / epjc / s10052 - 011 - 1583 - 8. arXiv: 1011 . 6118 [hep-ph].
- [224] O. Buchmueller et al. “The CMSSM and NUHM1 in Light of 7 TeV LHC, $B_s \rightarrow \mu^+ \mu^-$ and XENON100 Data”. In: *Eur. Phys. J. C* 72 (2012), p. 2243. DOI: 10 . 1140 / epjc / s10052 - 012 - 2243 - 3. arXiv: 1207 . 7315 [hep-ph].
- [225] O. Buchmueller et al. “The CMSSM and NUHM1 after LHC Run 1”. In: *Eur. Phys. J. C* 74.6 (2014), p. 2922. DOI: 10 . 1140 / epjc / s10052 - 014 - 2922 - 3. arXiv: 1312 . 5250 [hep-ph].
- [226] K. J. de Vries et al. “The pMSSM10 after LHC Run 1”. In: *Eur. Phys. J. C* 75.9 (2015), p. 422. DOI: 10 . 1140 / epjc / s10052 - 015 - 3599 - y. arXiv: 1504 . 03260 [hep-ph].
- [227] E. A. Bagnaschi et al. “Supersymmetric Dark Matter after LHC Run 1”. In: *Eur. Phys. J. C* 75 (2015), p. 500. DOI: 10 . 1140 / epjc / s10052 - 015 - 3718 - 9. arXiv: 1508 . 01173 [hep-ph].
- [228] Andrew Fowlie et al. “Bayesian Implications of Current LHC and XENON100 Search Limits for the Constrained MSSM”. In: *Phys. Rev. D* 85 (2012), p. 075012. DOI: 10 . 1103 / PhysRevD . 85 . 075012. arXiv: 1111 . 6098 [hep-ph].

- [229] Leszek Roszkowski, Enrico Maria Sessolo, and Yue-Lin Sming Tsai. “Bayesian Implications of Current LHC Supersymmetry and Dark Matter Detection Searches for the Constrained MSSM”. In: *Phys. Rev. D* 86 (2012), p. 095005. DOI: 10.1103/PhysRevD.86.095005. arXiv: 1202.1503 [hep-ph].
- [230] Kamila Kowalska et al. “Constrained next-to-minimal supersymmetric standard model with a 126 GeV Higgs boson: A global analysis”. In: *Phys. Rev. D* 87 (2013), p. 115010. DOI: 10.1103/PhysRevD.87.115010. arXiv: 1211.1693 [hep-ph].
- [231] Csaba Balazs et al. “Should we still believe in constrained supersymmetry?” In: *Eur. Phys. J. C* 73 (2013), p. 2563. DOI: 10.1140/epjc/s10052-013-2563-y. arXiv: 1205.1568 [hep-ph].
- [232] Doyoun Kim et al. “Bayesian naturalness of the CMSSM and CN-MSSM”. In: *Phys. Rev. D* 90.5 (2014), p. 055008. DOI: 10.1103/PhysRevD.90.055008. arXiv: 1312.4150 [hep-ph].
- [233] Maria Eugenia Cabrera-Catalan et al. “Indirect and direct detection prospect for TeV dark matter in the nine parameter MSSM”. In: *Phys. Rev. D* 92.3 (2015), p. 035018. DOI: 10.1103/PhysRevD.92.035018. arXiv: 1503.00599 [hep-ph].
- [234] Daniel E. Lopez-Fogliani et al. “A Bayesian Analysis of the Constrained NMSSM”. In: *Phys. Rev. D* 80 (2009), p. 095013. DOI: 10.1103/PhysRevD.80.095013. arXiv: 0906.4911 [hep-ph].
- [235] Pat Scott et al. “Direct Constraints on Minimal Supersymmetry from Fermi-LAT Observations of the Dwarf Galaxy Segue 1”. In: *JCAP* 1001 (2010), p. 031. DOI: 10.1088/1475-7516/2010/01/031. arXiv: 0909.3300 [astro-ph.CO].
- [236] Charlotte Strege et al. “Fundamental statistical limitations of future dark matter direct detection experiments”. In: *Phys. Rev. D* 86 (2012), p. 023507. DOI: 10.1103/PhysRevD.86.023507. arXiv: 1201.3631 [hep-ph].
- [237] Csaba Balázs et al. “ColliderBit: a GAMBIT module for the calculation of high-energy collider observables and likelihoods”. In: *Eur. Phys. J. C* 77.11 (2017), p. 795. DOI: 10.1140/epjc/s10052-017-5285-8. arXiv: 1705.07919 [hep-ph].
- [238] Florian U. Bernlochner et al. “FlavBit: A GAMBIT module for computing flavour observables and likelihoods”. In: *Eur. Phys. J. C* 77.11 (2017), p. 786. DOI: 10.1140/epjc/s10052-017-5157-2. arXiv: 1705.07933 [hep-ph].

- [239] Torsten Bringmann et al. "DarkBit: A GAMBIT module for computing dark matter observables and likelihoods". In: *Eur. Phys. J. C* 77.12 (2017), p. 831. DOI: 10.1140/epjc/s10052-017-5155-4. arXiv: 1705.07920 [hep-ph].
- [240] Peter Athron et al. "SpecBit, DecayBit and PrecisionBit: GAMBIT modules for computing mass spectra, particle decay rates and precision observables". In: *Eur. Phys. J. C* 78.1 (2018), p. 22. DOI: 10.1140/epjc/s10052-017-5390-8. arXiv: 1705.07936 [hep-ph].
- [241] Message P Forum. *MPI: A Message-Passing Interface Standard*. Tech. rep. Knoxville, TN, USA, 1994.
- [242] L. Dagum and R. Menon. "OpenMP: An Industry-Standard API for Shared-Memory Programming". In: *IEEE Comput. Sci. Eng.* 5.1 (Jan. 1998), pp. 46–55. ISSN: 1070-9924. DOI: 10.1109/99.660313. URL: <https://doi.org/10.1109/99.660313>.
- [243] B. C. Allanach et al. "SUSY Les Houches Accord 2". In: *Comp. Phys. Comm.* 180 (2009), pp. 8–25. DOI: 10.1016/j.cpc.2008.08.004. arXiv: 0801.0045 [hep-ph].
- [244] Nelson Christensen et al. "II. Bayesian methods for cosmological parameter estimation from cosmic microwave background measurements". In: *Class. Quant. Grav.* 18 (2001), p. 2677. DOI: 10.1088/0264-9381/18/14/306. arXiv: astro-ph/0103134 [astro-ph].
- [245] Joanna Dunkley et al. "Fast and reliable mcmc for cosmological parameter estimation". In: *Mon. Not. Roy. Astron. Soc.* 356 (2005), pp. 925–936. DOI: 10.1111/j.1365-2966.2004.08464.x. arXiv: astro-ph/0405462 [astro-ph].
- [246] Antony Lewis and Sarah Bridle. "Cosmological parameters from CMB and other data: A Monte Carlo approach". In: *Phys. Rev. D* 66 (2002), p. 103511. DOI: 10.1103/PhysRevD.66.103511. arXiv: astro-ph/0205436 [astro-ph].
- [247] John Skilling. "Nested sampling for general Bayesian computation". In: *Bayesian Analysis* 1.4 (2006), pp. 833–859. DOI: 10.1214/06-BA127.
- [248] F. Feroz, M. P. Hobson, and M. Bridges. "MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics". In: *Mon. Not. Roy. Astron. Soc.* 398 (2009), pp. 1601–1614. DOI: 10.1111/j.1365-2966.2009.14548.x. arXiv: 0809.3437 [astro-ph].
- [249] Rainer Storn and Kenneth Price. "Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces". In: *Journal of Global Optimization* 11.4 (Dec. 1997), pp. 341–

359. ISSN: 1573-2916. DOI: 10.1023/A:1008202821328. URL: <https://doi.org/10.1023/A:1008202821328>.
- [250] Daniela Zaharie. “Statistical Properties of Differential Evolution and Related Random Search Algorithms”. In: *COMPSTAT 2008*. Ed. by Paula Brito. Heidelberg: Physica-Verlag HD, 2008, pp. 473–485. ISBN: 978-3-7908-2084-3.
- [251] J. Brest et al. “Self-Adapting Control Parameters in Differential Evolution: A Comparative Study on Numerical Benchmark Problems”. In: *IEEE Transactions on Evolutionary Computation* 10.6 (Dec. 2006), pp. 646–657. ISSN: 1089-778X. DOI: 10.1109/TEVC.2006.872133.
- [252] M. Aoki et al. “Search for Massive Neutrinos in the Decay $\pi \rightarrow e\nu$ ”. In: *Phys. Rev. D* 84 (2011), p. 052002. DOI: 10.1103/PhysRevD.84.052002. arXiv: 1106.4055 [hep-ex].
- [253] A. Abada et al. “Sterile neutrinos in leptonic and semileptonic decays”. In: *JHEP* 02 (2014), p. 091. DOI: 10.1007/JHEP02(2014)091. arXiv: 1311.2830 [hep-ph].
- [254] K. A. Olive et al. “Review of Particle Physics”. In: update to Ref. [272] (2017).
- [255] A. M. Baldini et al. “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ \gamma$ with the full dataset of the MEG experiment”. In: *Eur. Phys. J. C* 76.8 (2016), p. 434. DOI: 10.1140/epjc/s10052-016-4271-x. arXiv: 1605.05081 [hep-ex].
- [256] Bernard Aubert et al. “Searches for Lepton Flavor Violation in the Decays $\tau^\pm \rightarrow e^\pm \gamma$ and $\tau^\pm \rightarrow \mu^\pm \gamma$ ”. In: *Phys. Rev. Lett.* 104 (2010), p. 021802. DOI: 10.1103/PhysRevLett.104.021802. arXiv: 0908.2381 [hep-ex].
- [257] K. Hayasaka et al. “New Search for $\tau \rightarrow \mu \gamma$ and $\tau \rightarrow e \gamma$ Decays at Belle”. In: *Phys. Lett. B* 666 (2008), pp. 16–22. DOI: 10.1016/j.physletb.2008.06.056. arXiv: 0705.0650 [hep-ex].
- [258] Georges Aad et al. “Probing lepton flavour violation via neutrinoless $\tau \rightarrow 3\mu$ decays with the ATLAS detector”. In: *Eur. Phys. J. C* 76.5 (2016), p. 232. DOI: 10.1140/epjc/s10052-016-4041-9. arXiv: 1601.03567 [hep-ex].
- [259] Roel Aaij et al. “Search for the lepton flavour violating decay $\tau^- \rightarrow \mu^- \mu^+ \mu^-$ ”. In: *JHEP* 02 (2015), p. 121. DOI: 10.1007/JHEP02(2015)121. arXiv: 1409.8548 [hep-ex].
- [260] U. Bellgardt et al. “Search for the Decay $\mu^+ \rightarrow e^+ e^+ e^-$ ”. In: *Nucl. Phys. B* 299 (1988), pp. 1–6. DOI: 10.1016/0550-3213(88)90462-2.

- [261] K. Hayasaka et al. "Search for Lepton Flavor Violating Tau Decays into Three Leptons with 719 Million Produced Tau+Tau- Pairs". In: *Phys. Lett. B* 687 (2010), pp. 139–143. DOI: 10.1016/j.physletb.2010.03.037. arXiv: 1001.3221 [hep-ex].
- [262] J. P. Lees et al. "Limits on tau Lepton-Flavor Violating Decays in three charged leptons". In: *Phys. Rev. D* 81 (2010), p. 111101. DOI: 10.1103/PhysRevD.81.111101. arXiv: 1002.4550 [hep-ex].
- [263] J. Kaulard et al. "Improved limit on the branching ratio of $\mu^- \rightarrow e^+$ conversion on titanium". In: *Phys. Lett. B* 422 (1998), pp. 334–338. DOI: 10.1016/S0370-2693(97)01423-8.
- [264] W. Honecker et al. "Improved limit on the branching ratio of $\mu \rightarrow e$ conversion on lead". In: *Phys. Rev. Lett.* 76 (1996), pp. 200–203. DOI: 10.1103/PhysRevLett.76.200.
- [265] G. Czapek et al. "Branching ratio for the rare pion decay into positron and neutrino". In: *Phys. Rev. Lett.* 70 (1993), pp. 17–20. DOI: 10.1103/PhysRevLett.70.17.
- [266] C. Lazzeroni et al. "Precision Measurement of the Ratio of the Charged Kaon Leptonic Decay Rates". In: *Phys. Lett. B* 719 (2013), pp. 326–336. DOI: 10.1016/j.physletb.2013.01.037. arXiv: 1212.4012 [hep-ex].
- [267] Roel Aaij et al. "Measurement of forward $W \rightarrow e\nu$ production in pp collisions at $\sqrt{s} = 8\text{ TeV}$ ". In: *JHEP* 10 (2016), p. 030. DOI: 10.1007/JHEP10(2016)030. arXiv: 1608.01484 [hep-ex].
- [268] S. Schael et al. "Electroweak Measurements in Electron-Positron Collisions at W-Boson-Pair Energies at LEP". In: *Phys. Rept.* 532 (2013), pp. 119–244. DOI: 10.1016/j.physrep.2013.07.004. arXiv: 1302.3415 [hep-ex].
- [269] M. Antonelli et al. "An Evaluation of $|V_{us}|$ and precise tests of the Standard Model from world data on leptonic and semileptonic kaon decays". In: *Eur. Phys. J. C* 69 (2010), pp. 399–424. DOI: 10.1140/epjc/s10052-010-1406-3. arXiv: 1005.2323 [hep-ph].
- [270] E. Follana et al. "High Precision determination of the π , K , D and $D(s)$ decay constants from lattice QCD". In: *Phys. Rev. Lett.* 100 (2008), p. 062002. DOI: 10.1103/PhysRevLett.100.062002. arXiv: 0706.1726 [hep-lat].
- [271] Y. Amhis et al. "Averages of B-Hadron, C-Hadron, and tau-lepton properties as of early 2012". In: (2012). arXiv: 1207.1158 [hep-ex].
- [272] C. Patrignani et al. "Review of Particle Physics". In: *Chin. Phys. C* 40.10 (2016), p. 100001. DOI: 10.1088/1674-1137/40/10/100001.

- [273] S. Aoki et al. "Review of lattice results concerning low-energy particle physics". In: *Eur. Phys. J. C* 77.2 (2017), p. 112. DOI: 10.1140/epjc/s10052-016-4509-7. arXiv: 1607.00299 [hep-lat].
- [274] M. Agostini et al. "Background-free search for neutrinoless double- β decay of ^{76}Ge with GERDA". In: (2017). [Nature544,47(2017)]. DOI: 10.1038/nature21717. arXiv: 1703.00570 [nucl-ex].
- [275] A. Gando et al. "Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen". In: *Phys. Rev. Lett.* 117 (2016). [Addendum: *Phys. Rev. Lett.* 117, no. 10, 109903 (2016)], p. 082503. DOI: 10.1103/PhysRevLett.117.082503. arXiv: 1605.02889 [hep-ex].
- [276] Georges Aad et al. "Search for heavy Majorana neutrinos with the ATLAS detector in pp collisions at $\sqrt{s} = 8 \text{ TeV}$ ". In: *JHEP* 07 (2015), p. 162. DOI: 10.1007/JHEP07(2015)162. arXiv: 1506.06020 [hep-ex].
- [277] Albert M Sirunyan et al. "Search for heavy neutral leptons in events with three charged leptons in proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$ ". In: *Phys. Rev. Lett.* 120 (2018), p. 221801. arXiv: 1802.02965 [hep-ex].
- [278] A. V. Artamonov et al. "Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_H$ decays". In: *Phys. Rev. D* 91.5 (2015), p. 052001. DOI: 10.1103/PhysRevD.91.059903, 10.1103/PhysRevD.91.052001. arXiv: 1411.3963 [hep-ex].
- [279] J. Orloff, Alexandre N. Rozanov, and C. Santoni. "Limits on the mixing of tau neutrino to heavy neutrinos". In: *Phys. Lett. B* 550 (2002), pp. 8–15. DOI: 10.1016/S0370-2693(02)02769-7. arXiv: hep-ph/0208075 [hep-ph].
- [280] M. H. Ahn et al. "Measurement of Neutrino Oscillation by the K2K Experiment". In: *Phys. Rev. D* 74 (2006), p. 072003. arXiv: hep-ex/0606032 [hep-ex].
- [281] K. Abe et al. "Measurement of Neutrino Oscillation Parameters from Muon Neutrino Disappearance with an Off-axis Beam". In: *Phys. Rev. Lett.* 111.21 (2013), p. 211803. DOI: 10.1103/PhysRevLett.111.211803. arXiv: 1308.0465 [hep-ex].
- [282] P. Adamson et al. "Measurement of the Neutrino Mass Splitting and Flavor Mixing by MINOS". In: *Phys. Rev. Lett.* 106 (2011), p. 181801. DOI: 10.1103/PhysRevLett.106.181801. arXiv: 1103.0340 [hep-ex].
- [283] K. Abe et al. "Evidence for the Appearance of Atmospheric Tau Neutrinos in Super-Kamiokande". In: *Phys. Rev. Lett.* 110.18 (2013). DOI: 10.1103/PhysRevLett.110.181802. arXiv: 1206.0328 [hep-ex].

- [284] P. A. R. Ade et al. "Planck 2015 results. XIII. Cosmological parameters". In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10 . 1051 / 0004 - 6361/201525830. arXiv: 1502.01589 [astro-ph.CO].
- [285] B. T. Cleveland et al. "Measurement of the solar electron neutrino flux with the Homestake chlorine detector". In: *Astrophys. J.* 496 (1998), pp. 505–526. DOI: 10.1086/305343.
- [286] F. Kaether et al. "Reanalysis of the GALLEX solar neutrino flux and source experiments". In: *Phys. Lett. B* 685 (2010), pp. 47–54. DOI: 10 . 1016/j.physletb.2010.01.030. arXiv: 1001.2731 [hep-ex].
- [287] B. Aharmim et al. "Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory". In: *Phys. Rev. C* 88 (2013), p. 025501. DOI: 10.1103/PhysRevC.88.025501. arXiv: 1109.0763 [nucl-ex].
- [288] J. Hosaka et al. "Solar neutrino measurements in super-Kamiokande-I". In: *Phys. Rev. D* 73 (2006), p. 112001. DOI: 10 . 1103/PhysRevD . 73 . 112001. arXiv: hep-ex/0508053 [hep-ex].
- [289] J. P. Cravens et al. "Solar neutrino measurements in Super-Kamiokande-II". In: *Phys. Rev. D* 78 (2008), p. 032002. DOI: 10 . 1103/PhysRevD . 78 . 032002. arXiv: 0803.4312 [hep-ex].
- [290] K. Abe et al. "Solar neutrino results in Super-Kamiokande-III". In: *Phys. Rev. D* 83 (2011), p. 052010. DOI: 10.1103/PhysRevD.83.052010. arXiv: 1010.0118 [hep-ex].
- [291] G. Bellini et al. "Precision measurement of the ^7Be solar neutrino interaction rate in Borexino". In: *Phys. Rev. Lett.* 107 (2011), p. 141302. DOI: 10.1103/PhysRevLett.107.141302. arXiv: 1104.1816 [hep-ex].
- [292] G. Bellini et al. "Measurement of the solar ^8B neutrino rate with a liquid scintillator target and 3 MeV energy threshold in the Borexino detector". In: *Phys. Rev. D* 82 (2010), p. 033006. DOI: 10.1103/PhysRevD.82.033006. arXiv: 0808.2868 [astro-ph].
- [293] G. Bellini et al. "Neutrinos from the primary proton–proton fusion process in the Sun". In: *Nature* 512.7515 (2014), pp. 383–386. DOI: 10 . 1038/nature13702.
- [294] M. G. Aartsen et al. "Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of Ice-Cube DeepCore data". In: *Phys. Rev. D* 91.7 (2015), p. 072004. DOI: 10 . 1103/PhysRevD.91.072004. arXiv: 1410.7227 [hep-ex].

- [295] A. Gando et al. "Reactor On-Off Antineutrino Measurement with KamLAND". In: *Phys. Rev. D* 88.3 (2013), p. 033001. DOI: 10.1103/PhysRevD.88.033001. arXiv: 1303.4667 [hep-ex].
- [296] Feng Peng An et al. "Improved Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay". In: *Chin. Phys. C* 41.1 (2017), p. 013002. DOI: 10.1088/1674-1137/41/1/013002. arXiv: 1607.05378 [hep-ex].
- [297] Feng Peng An et al. "Measurement of electron antineutrino oscillation based on 1230 days of operation of the Daya Bay experiment". In: *Phys. Rev. D* 95.7 (2017), p. 072006. DOI: 10.1103/PhysRevD.95.072006. arXiv: 1610.04802 [hep-ex].
- [298] H. Seo. *New Results from RENO. - 2017. Talk given at the EPS Conference on High Energy Physics, Venice, Italy, July 5-12.*
- [299] P. Adamson et al. "Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS". In: *Phys. Rev. Lett.* 110.25 (2013), p. 251801. DOI: 10.1103/PhysRevLett.110.251801. arXiv: 1304.6335 [hep-ex].
- [300] P. Adamson et al. "Electron neutrino and antineutrino appearance in the full MINOS data sample". In: *Phys. Rev. Lett.* 110.17 (2013), p. 171801. DOI: 10.1103/PhysRevLett.110.171801. arXiv: 1301.4581 [hep-ex].
- [301] A. Izmaylov. *T2K Neutrino Experiment. Recent Results and Plans. - 2017. Talk given at the Flavour Physics Conference, Quy Nhon, Vietnam, August 13-19.*
- [302] A. Radovic. *Latest oscillation results from NOvA. - 2018. Fermilab. Joint Experimental-Theoretical Physics Seminar USA, January 12.*
- [303] R. P. Feynman and Murray Gell-Mann. "Theory of Fermi interaction". In: *Phys. Rev.* 109 (1958). [417(1958)], pp. 193–198. DOI: 10.1103/PhysRev.109.193.
- [304] Stefan Antusch and Oliver Fischer. "Testing sterile neutrino extensions of the Standard Model at future lepton colliders". In: *JHEP* 05 (2015), p. 053. DOI: 10.1007/JHEP05(2015)053. arXiv: 1502.05915 [hep-ph].
- [305] Ievgen Dubovyk et al. "Complete electroweak two-loop corrections to Z boson production and decay". In: (2018). arXiv: 1804.10236 [hep-ph].
- [306] S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry". In: *Phys. Rev. D* 2 (1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.

- [307] A. Ilakovac and A. Pilaftsis. "Flavor violating charged lepton decays in seesaw-type models". In: *Nucl. Phys.* B437 (1995), p. 491. DOI: 10.1016/0550-3213(94)00567-X. arXiv: hep-ph/9403398 [hep-ph].
- [308] A. Abada et al. "Lepton flavor violation in low-scale seesaw models: SUSY and non-SUSY contributions". In: *JHEP* 11 (2014), p. 048. DOI: 10.1007/JHEP11(2014)048. arXiv: 1408.0138 [hep-ph].
- [309] J. Beringer et al. "Review of Particle Physics (RPP)". In: *Phys. Rev.* D86 (2012), p. 010001. DOI: 10.1103/PhysRevD.86.010001.
- [310] D. Decamp et al. "Measurement of tau branching ratios". In: *Z. Phys.* C54 (1992), pp. 211–228. DOI: 10.1007/BF01566649.
- [311] "A Combination of preliminary electroweak measurements and constraints on the standard model". In: (2001). arXiv: hep-ex/0103048 [hep-ex].
- [312] Sofiane M. Boucenna, Jose W. F. Valle, and Avelino Vicente. "Are the B decay anomalies related to neutrino oscillations?" In: *Phys. Lett.* B750 (2015), pp. 367–371. DOI: 10.1016/j.physletb.2015.09.040. arXiv: 1503.07099 [hep-ph].
- [313] Robert E. Shrock. "General Theory of Weak Leptonic and Semileptonic Decays. 1. Leptonic Pseudoscalar Meson Decays, with Associated Tests For, and Bounds on, Neutrino Masses and Lepton Mixing". In: *Phys. Rev.* D24 (1981), p. 1232. DOI: 10.1103/PhysRevD.24.1232.
- [314] Vincenzo Cirigliano and Ignasi Rosell. "Two-loop effective theory analysis of $\pi(K) \rightarrow e\bar{\nu}_e\gamma$ branching ratios". In: *Phys. Rev. Lett.* 99 (2007), p. 231801. DOI: 10.1103/PhysRevLett.99.231801. arXiv: 0707.3439 [hep-ph].
- [315] Diptimoy Ghosh, Marco Nardecchia, and S. A. Renner. "Hint of Lepton Flavour Non-Universality in B Meson Decays". In: *JHEP* 12 (2014), p. 131. DOI: 10.1007/JHEP12(2014)131. arXiv: 1408.4097 [hep-ph].
- [316] Xiao-Gang He and German Valencia. "Are the B-anomalies evidence for heavy neutrinos?" In: (2017). arXiv: 1706.07570 [hep-ph].
- [317] Wolfgang Altmannshofer et al. "Symmetries and Asymmetries of $B \rightarrow K^*\mu^+\mu^-$ Decays in the Standard Model and Beyond". In: *JHEP* 01 (2009), p. 019. DOI: 10.1088/1126-6708/2009/01/019. arXiv: 0811.1214 [hep-ph].
- [318] Debajyoti Choudhury et al. " $R_{K^{(*)}}$ and $R(D^{(*)})$ anomalies resolved with lepton mixing". In: (2017). arXiv: 1712.01593 [hep-ph].
- [319] Antonio Pich et al. "Tau decays". In: *Int. J. Mod. Phys.* A24S1 (2009), pp. 715–737. DOI: 10.1142/S0217751X09046746.

- [320] Marco Drewes and Shintaro Eijima. “Neutrinoless double β decay and low scale leptogenesis”. In: *Phys. Lett.* B763 (2016), pp. 72–79. DOI: 10 . 1016/j.physletb.2016.09.054. arXiv: 1606.06221 [hep-ph].
- [321] J. Lopez-Pavon, S. Pascoli, and Chan-fai Wong. “Can heavy neutrinos dominate neutrinoless double beta decay?” In: *Phys. Rev.* D87.9 (2013), p. 093007. DOI: 10 . 1103/PhysRevD . 87 . 093007. arXiv: 1209 . 5342 [hep-ph].
- [322] A. D. Dolgov and F. L. Villante. “BBN bounds on active sterile neutrino mixing”. In: *Nucl. Phys.* B679 (2004), pp. 261–298. DOI: 10 . 1016/j.nuclphysb.2003.11.031. arXiv: hep-ph/0308083 [hep-ph].
- [323] Dmitry Gorbunov and Mikhail Shaposhnikov. “How to find neutral leptons of the ν MSM?” In: *JHEP* 10 (2007). [Erratum: JHEP11,101(2013)], p. 015. DOI: 10 . 1007/JHEP11(2013)101 , 10 . 1088/1126-6708/2007/10/015. arXiv: 0705.1729 [hep-ph].
- [324] Gary J. Feldman and Robert D. Cousins. “A Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev.* D57 (1998), pp. 3873–3889. DOI: 10 . 1103/PhysRevD . 57 . 3873. arXiv: physics / 9711021 [physics.data-an].
- [325] D. I. Britton et al. “Improved search for massive neutrinos in $\pi^+ \rightarrow e^+ \nu$ decay”. In: *Phys. Rev.* D46 (1992), R885–R887.
- [326] N. Lurkin et al. “Heavy neutrino searches and NA62 status”. In: *Proceedings, 52nd Rencontres de Moriond on Electroweak Interactions and Unified Theories: La Thuile, Italy, March 18-25, 2017*. 2017, pp. 171–178. arXiv: 1808.00827 [hep-ex].
- [327] Georges Aad et al. “Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector”. In: (2019). arXiv: 1905 . 09787 [hep-ex].
- [328] Albert M Sirunyan et al. “Search for heavy Majorana neutrinos in same-sign dilepton channels in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 01 (2019), p. 122. DOI: 10 . 1007/JHEP01(2019)122. arXiv: 1806.10905 [hep-ex].
- [329] Frank F. Deppisch, P. S. Bhupal Dev, and Apostolos Pilaftsis. “Neutrinos and Collider Physics”. In: *New J. Phys.* 17.7 (2015), p. 075019. DOI: 10 . 1088/1367-2630/17/7/075019. arXiv: 1502.06541 [hep-ph].
- [330] Wei Liao and Xiao-Hong Wu. “Signature of heavy sterile neutrinos at CEPC”. In: *Phys. Rev.* D97.5 (2018), p. 055005. DOI: 10 . 1103/PhysRevD . 97 . 055005. arXiv: 1710.09266 [hep-ph].

- [331] Cosmas K. Zachos. "Ternary Plots for Neutrino Mixing Visualization". In: *EPL* 99.1 (2012), p. 11001. DOI: 10.1209/0295-5075/99/11001. arXiv: 1205.4772 [hep-ph].
- [332] Marco Drewes and Jan Hajer. "Heavy Neutrinos in displaced vertex searches at the LHC and HL-LHC". In: (2019). arXiv: 1903.06100 [hep-ph].
- [333] Alain Blondel et al. "Search for Heavy Right Handed Neutrinos at the FCC-ee". In: *Nucl. Part. Phys. Proc.* 273-275 (2016), pp. 1883–1890. DOI: 10.1016/j.nuclphysbps.2015.09.304. arXiv: 1411.5230 [hep-ex].
- [334] Yoshitaka Kuno and Yasuhiro Okada. "Muon decay and physics beyond the standard model". In: *Rev. Mod. Phys.* 73 (2001), pp. 151–202. DOI: 10.1103/RevModPhys.73.151. arXiv: hep-ph/9909265 [hep-ph].
- [335] T. S. Kosmas, Sergey Kovalenko, and Ivan Schmidt. "Nuclear muon-e- conversion in strange quark sea". In: *Phys. Lett.* B511 (2001), p. 203. DOI: 10.1016/S0370-2693(01)00657-8. arXiv: hep-ph/0102101 [hep-ph].
- [336] F. U. Bernlochner et al. "FlavBit: A GAMBIT module for computing flavour observables and likelihoods". In: *EPJC* 77 (May 2017), p. 786. arXiv: 1705.07933 [hep-ph].
- [337] P. Athron et al. "SpecBit, DecayBit and PrecisionBit: GAMBIT modules for computing mass spectra, particle decay rates and precision observables". In: *EPJC* 78, 22 (Jan. 2018), p. 22. DOI: 10.1140/epjc/s10052-017-5390-8. arXiv: 1705.07936 [hep-ph].
- [338] Takehiko Asaka and Shintaro Eijima. "Direct Search for Right-handed Neutrinos and Neutrinoless Double Beta Decay". In: *PTEP* 2013.11 (2013), 113B02. DOI: 10.1093/ptep/ptt094. arXiv: 1308.3550 [hep-ph].
- [339] C. D. Froggatt and Holger Bech Nielsen. "Hierarchy of Quark Masses, Cabibbo Angles and CP Violation". In: *Nucl. Phys.* B147 (1979), pp. 277–298. DOI: 10.1016/0550-3213(79)90316-X.
- [340] T. Bringmann et al. "DarkBit: A GAMBIT module for computing dark matter observables and likelihoods". In: *EPJC* 77 (May 2017), p. 831. arXiv: 1705.07920 [hep-ph].
- [341] G. D. Martinez et al. "Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module". In: *EPJC* 77.11 (May 2017), p. 761. DOI: 10.1140/epjc/s10052-017-5274-y. arXiv: 1705.07959 [hep-ph].
- [342] Robert D. Cousins. "Why isn't every physicist a Bayesian?" In: *Am. J. Phys.* 63 (1995), p. 398. DOI: 10.1119/1.17901.

- [343] Pat Scott. "Pippi - painless parsing, post-processing and plotting of posterior and likelihood samples". In: *Eur. Phys. J. Plus* 127 (2012), p. 138. DOI: 10.1140/epjp/i2012-12138-3. arXiv: 1206.2245.
- [344] Alexey Boyarsky, Oleg Ruchayskiy, and Mikhail Shaposhnikov. "The Role of sterile neutrinos in cosmology and astrophysics". In: *Ann. Rev. Nucl. Part. Sci.* 59 (2009), pp. 191–214. DOI: 10.1146/annurev.nucl.010909.083654. arXiv: 0901.0011 [hep-ph].
- [345] N. Bohr. "Conservation Laws in Quantum Theory". In: *Nature* 138 (1936), pp. 25–26. DOI: 10.1038/138025b0.
- [346] W. Pauli. "Aufsätze und Vorträge über Physik und Erkenntnistheorie". In: *Proceedings of the VII Solvay Congress, Brussels*. 1933, p. 324.
- [347] C. M. G. Lattes et al. "PROCESSES INVOLVING CHARGED MESONS". In: *Nature* 159 (1947). [42(1947)], pp. 694–697. DOI: 10.1038/159694a0.
- [348] F. Reines and C. L. Cowan. "Detection of the free neutrino". In: *Phys. Rev.* 92 (1953), pp. 830–831. DOI: 10.1103/PhysRev.92.830.
- [349] B. Pontecorvo. "Electron and Muon Neutrinos". In: *Sov. Phys. JETP* 10 (1960). [Zh. Eksp. Teor. Fiz.37,1751(1959)], pp. 1236–1240.
- [350] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36.
- [351] Martin L. Perl et al. "Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation". In: *Phys. Rev. Lett.* 35 (1975). [193(1975)], pp. 1489–1492. DOI: 10.1103/PhysRevLett.35.1489.
- [352] K. Kodama et al. "Observation of tau neutrino interactions". In: *Phys. Lett.* B504 (2001), pp. 218–224. DOI: 10.1016/S0370-2693(01)00307-0. arXiv: hep-ex/0012035 [hep-ex].
- [353] Paul Langacker. *The standard model and beyond*. Series in High Energy Physics, Cosmology and Gravitation. Boca Raton, FL: Taylor and Francis, 2010. URL: <https://cds.cern.ch/record/1226768>.
- [354] Mattias Blennow et al. "Neutrinoless double beta decay in seesaw models". In: *JHEP* 07 (2010), p. 096. DOI: 10.1007/JHEP07(2010)096. arXiv: 1005.3240 [hep-ph].
- [355] Luca Stanco. "The neutrino mass hierarchy from oscillation". In: *Proceedings, Prospects in Neutrino Physics (NuPhys2017): London, UK, December 20-22, 2017*. 2018, pp. 97–106. arXiv: 1803.08722 [hep-ph].
- [356] Takeshi Amemiya. *Advanced Econometrics*. Harvard University Press, 1985.

- [357] Russell Davidson and James MacKinnon. *Estimation and Inference in Econometrics*. Oxford University Press, 1993.
- [358] Christian Gourieroux and Alain Monfort. *Statistics and Econometric Models*. Vol. 2. Themes in Modern Econometrics. Cambridge University Press, 1995. DOI: 10.1017/CB09780511751950.
- [359] Torsten Bringmann et al. “DarkSUSY 6 : An Advanced Tool to Compute Dark Matter Properties Numerically”. In: *JCAP* 1807.07 (2018), p. 033. DOI: 10.1088/1475-7516/2018/07/033. arXiv: 1802.03399 [hep-ph].
- [360] Fengpeng An et al. “Neutrino Physics with JUNO”. In: *J. Phys. G* 43.3 (2016), p. 030401. DOI: 10.1088/0954-3899/43/3/030401. arXiv: 1507.05613 [physics.ins-det].

Samenvatting

Neutrino's zijn altijd al mysterieuze deeltjes geweest. Het baanbrekende experiment uit 1932 dat vond dat de tijdens het bètaverval van tritium vrijgekomen elektronen geen constante energie hebben deed Niels Bohr de wet van behoud van energie in twijfel trekken. In de daaropvolgende jaren ontwikkelde Wolfgang Pauli en Enrico Fermi de theoretische onderbouwing voor het bestaan van neutrino's. Echter, in 1950 was er nog steeds geen experimentele verificatie van hun werk. De ongrijpbare aard van de neutrino was hiervoor verantwoordelijk: niemand had neutrino's ook maar iets zien doen. Vandaag de dag weten wij dat neutrino's een extreem zwakke interactie aangaan met andere materie. Ze kunnen lichtjaren aan lood penetreren en het zijn de enige deeltjes die kunnen ontsnappen uit de kern van een imploderende ster. Zodoende was het opstellen van een experiment dat hun bestaan onomstotelijk kon bewijzen een grote uitdaging.

In de jaren '50 waren Frederick Reines en Clyde Cowan hierin wel succesvol en de rest is geschiedenis. Vandaag de dag is neutrino (astro)fysica een snel evoluerend onderzoeksgebied. Eén van de bekendste recente voorbeelden is het bewijs dat neutrino's massa bezitten, terwijl het standaardmodel van de deeltjesfysica (SM) zegt dat de massa van het neutrino nul is. Deze ontdekking werd gedaan in 1988 en resulteerde in de Nobelprijs voor Arthur McDonald en Takaaki Kajita in 2015.

Gegeven deze twee feiten, dat neutrino's massa hebben maar als massaloos worden omschreven in het standaardmodel, rest er de vraag welk mechanisme verantwoordelijk is voor de massa van neutrino's. Zodoende vereist het bestaan van neutrino massa een uitbreiding van het standaardmodel, oftewel Beyond the Standard Model (BSM) fysica. In het SM zijn alle fermionen - behalve neutrino's - aanwezig in links en rechts chirale varianten. Het idee om de neutrinosector uit te breiden met een rechts-chirale variant bestond al voor de ontdekking neutrinomassa, maar de laatstgenoemde ontdekking gaf een extra versnelling aan het onderzoek naar de rechts-chirale variant van deze deeltjes. Naast een mogelijke uitleg op de vraag waarom neutrino's

massa hebben, bieden zij ook mogelijke inzichten in baryogenese en zouden ze ook de donkere materie in het heelal kunnen vormen.

Waar de (links-chirale) neutrino's in het standaardmodel al zeker als mysterieus bestempeld kunnen worden, zijn de rechts-chirale neutrino's ware fantomen. Terwijl de links-chirale neutrino's reageren via de zwakke wisselwerking en de zwaartekracht, geldt voor de rechts-chirale neutrino's (of rechtshandige) neutrino's (RHNs) dat zij in de meeste modellen waarin zij voorkomen alleen via de zwaartekracht en Yukawa koppelingen interacties aangaan. Veel experimenten hebben naar deze deeltjes gezocht, en blijven zoeken tot op de dag van vandaag. Hieronder bevinden zich LSND, het IceCube Neutrino Observatory en MiniBooNE op het Fermilab. Daarbij komen ook nog de deeltjesdetectorexperimenten zoals LHCb en ATLAS op het CERN. Helaas is enig bewijs voor het bestaan van de rechtshandige neutrino's tot op heden uitgebleven. Desalniettemin, de resultaten van de eerdergenoemde experimenten hebben natuurkundigen in staat gesteld om de verschillende eigenschappen van RHNs te begrenzen, in het bijzonder hun massa en koppelingen.

Did proefschrift is getiteld "*A Frequentist Analysis of the $n = 3$ Type-I Seesaw Model*". Hierin wordt beschreven hoe wij op een statistisch rigoureuze manier de sterkste limieten op RHNs combineren, met het uiteindelijke doel om de profielwaarschijnlijkheid voor koppeling en massa in kaart te brengen. Dit wordt gedaan in de context van het type-I seesaw model met drie rechtshandige neutrino's.

In Hoofdstuk 1, welke dient als introductie, geef ik een motivatie voor de introductie van de rechtshandige neutrino's (zoals hierboven in het kort beschreven) en een omschrijving van de bijbehorende fenomenologie. Men heeft geprobeerd hen direct waar te nemen doormiddel van meson verval, ijkboson verval, s-kanaal uitwisseling van W-bosonen en vector-boson fusie. Daarnaast kunnen ze indirect worden waargenomen, bijvoorbeeld in neutrinoloos dubbel- β verval, lepton smaak-schendende processen, afwijkingen van lepton universaliteit en in precisie metingen van de elektrozwakke wisselwerking, om er een paar te noemen. Did hoofdstuk beschrijft ook hoe wij enkele verbeteringen hebben aangebracht ten opzichte van eerdere studies, welke benodigd zijn om het doel van dit proefschrift te bereiken.

Hoofdstuk 2 bevat de theoretische onderbouwing die relevant is voor dit proefschrift, beginnende bij het model, het type-I seesaw model. In dit model

krijgen de links-chirale (of linkshandige) neutrino's (LHNs) van het SM geen Majorana massa, de Higgssector van het SM blijft onveranderd, en de Majorana massa's van de RHNs zijn veel groter dan hun Dirac massa's. Onder deze voorwaarden kunnen de massaverschillen tussen de neutrino's in het SM verklaard worden. Dit is het meest minimale en best bestudeerde model, en ontwijkt complexiteit die wel aanwezig is in andere modellen, zoals het herzien van de Higgssector en de toevoeging van nog meer nieuwe deeltjes. Om dit model te beschrijven gebruiken wij de Casas-Ibarra (C-I) parameterizatie, welke door zijn constructie in overeenkomst is met de neutrino oscillatie data. Dit geeft een numeriek voordeel wanneer wij een aselechte steekproef doen op de parameter ruimte. Het is triviaal om met deze parameterizatie grote waarden voor de koppelingen van de RHNs te krijgen (door het kiezen van bepaalde parameter intervallen), maar dit brengt wel enkele subtiele details met zich mee, welke in ogenschouw genomen dienen te worden. De parameterizatie zoals hierboven beschreven leidt tot sterke koppelingen, maar de meeste van deze punten zijn het resultaat van ongefundeerde fine-tuning. Aan de andere kant is het ook mogelijk om deze sterke koppelingen te krijgen op een meer natuurlijke manier, namelijk als de B-L symmetrie van het SM bij benadering behouden blijft. Wij maken onderscheidt tussen deze twee types punten in de scans.

Doordat de Casas-Ibarra parameterizatie het erg lastig maakt om goede priors op te stellen is het moeilijk om een Bayesiaanse studie uit te voeren op de parameter ruimte. Daarom gaan wij voor de frequentistische aanpak. Om de scans over de parameter ruimte uit te voeren maken wij gebruik van het Global And Modular BSM Inference Tool (GAMBIT). Dit is een global fitting software raamwerk welke uitermate geschikt is voor onze studie, aangezien nieuwe variabelen relatief eenvoudig kunnen worden geïncorporeerd. De code voor dit project bevindt zich in een nieuwe module binnen het raamwerk. We gebruiken een differential evolutionary-based algorithm, Diver, om de ruimte te scannen, aangezien deze bijzonder efficient is voor optimalisatie problemen, zoals de berekening van profielwaarschijnlijkheden. Het statistisch raamwerk en een gedetailleerde omschrijving van GAMBIT en Diver worden gegeven in Hoofdstuk 3.

In Hoofdstuk 4 worden de observabelen beschreven die in de studie worden meegenomen. Ten eerste, de neutrino massaverschillen en de mixhoek zoals gemeten in de oscillatie data, omdat elke uitbreiding van het SM aan deze waarnemingen moet voldoen. Daarnaast incorporeren we de limiet op

de som van de LHN massa's zoals gemeten door PLANCK. Zoals eerdergenoemd, zijn er veel elektrozwakke precisie observabelen (electroweak precision observables, EWPOs) die worden beïnvloed door de aanwezigheid van RHNs; dit komt doordat ze de interactiesterkte van de LHNs via de zwakke wisselwerking aanpassen. We gebruiken de Fermi constante, Weinberg hoek, massa van de W boson en de onzichtbare component in het verval van het Z boson. In het SM worden smaak-veranderende neutrale processen, zoals lepton smaak schending (lepton flavour violation, LFV) onderdrukt door het Glashow-Iliopoulos-Maiani (GIM) mechanisme. Als een BSM theorie wel bijdraagt aan deze processen zouden zij een grotere bijdrage leveren dan het SM, dit maakt het een smoking gun voor de aanwezigheid van nog onbekende natuurkundige fenomenen. Ondanks dat zulke signalen tot op heden nog niet zijn waargenomen bestaat er wel een gemeten bovengrens op de verval patronen, en de belangrijkste hiervan nemen wij mee in deze studie. RHNs hebben daarnaast ook invloed op testen van lepton universaliteit en wij nemen de meest belangrijke semileptonische en leptonisch vervallen mee. Ook de aanwezige koppeling tussen RHNs en LHNs is van invloed op metingen van de CKM matrix elementen, de leptonische elektrozwakke vervallen en pion vervallen (welke het meest beïnvloedt worden door deze veranderingen) worden meegenomen. De Majorana massa van de RHNs schendt het behoud van leptongetal en maakt neutrino dubbel- β verval mogelijk. De belangrijkste limiet op de effectieve Majorana massa, welke komt van GERDA en KamLAND-Zen, wordt ook in acht genomen. Tot slot nemen we ook de limieten uit de oerknal-nucleosynthese in ogenschouw, aangezien deze vereisen dat de RHNs een koppeling moeten hebben welke groot genoeg is zodat ze vervallen in een tijdsframe dat de juiste hoeveelheden van bepaalde elementen oplevert. Verschillende directe detectie experimenten worden ook meegenomen, van de volgende zijn de metingen het meest uitsluitend voor het massa interval waar wij naar kijken: PIENU, PS-191, CHARM, DELPHI, ATLAS, CMS, E949 and NuTeV.

De resultaten van deze studie worden beschreven en bediscussieerd in Hoofdstuk 5. Het is ons gelukt de profielwaarschijnlijkheid in het koppeling-massa vlak in kaart te brengen. De grootste koppelingen worden bereikt boven de W en Z boson massa's, waar de directe detectie experimenten weinig parameter ruimte uitsluiten. In dit gebied leggen EWPOs, LFV en CKM de grootste beperkingen op aan de maximale grote van de koppeling. Tot de W boson massa zijn directe detectie experimenten het meest uitsluitend, waarbij de

combinatie van experimenten soms zorgt voor de sterkste limieten. Onder een massa van ongeveer 0.3 GeV geeft een combinatie van de ondergrens op de mixhoek vanuit BBN, bovengrenzen vanuit directe detectie zoektochten en neutrino oscillatie data een ondergrens op de koppeling die beter is dan elk van de metingen apart. Voor het eerst onderzoeken wij ook het effect van de lichtste LHN massa op het smaak mix fracties. We vinden ook bepaalde fenomenen zoals eerder waargenomen in de situatie $n = 2$. Zo zien wij regio's met een overschot aan likelihood, wat eerder is waargenomen. Dit komt doordat het type-I seesaw model een betere fit oplevert voor de observabelen (onzichtbaar verval van het Z boson, CKM en een lepton universaliteit waardes) dan het SM.

Dit proefschrift eindigt met een samenvatting van de resultaten en suggesties voor vervolg onderzoek in Hoofdstuk 6. Grote gedeeltes van de toegestane parameter ruimte onder de 100 GeV kunnen in toekomstige experimenten worden bestudeerd. Daarnaast is het mogelijk om in te zoomen op de regio's met een overschot aan likelihood. Verbeteringen in de limieten op de massa van LHNs, samen met de ondergrens op de som aan massa's, kan eindelijk een antwoord geven op de vraag wat de neutrino hiërarchie is. Een ander interessant pad zou een uitbreiding van deze studie zijn naar meerdere theoretische modellen in een Bayesiaanse context.

Summary

Neutrinos have always been elusive particles. When the hallmark experiment that studied the beta decay of tritium in 1932 found the emitted electrons having variable energy, Niels Bohr controversially (although not for the first time) was prepared to abandon the law of conservation of energy. Wolfgang Pauli's and Enrico Fermi's work over the next few years put the existence of neutrinos on firm theoretical ground, but, by 1950, there was still no experimental verification available. This was due to their elusive nature; no one had seen neutrinos *do* anything. As we now know, they are extraordinarily weakly interacting. They can penetrate light years of lead, and are the only known particles that can penetrate the extremely dense material in collapsing stars. Devising experiments that could unambiguously confirm their existence was no trivial challenge.

Fortunately, in the 1950s, Frederick Reines and Clyde Cowan were successful in doing so, and the rest is history. The disciplines of neutrino physics and astrophysics are rich and rapidly evolving. One of the most famous pieces of knowledge that has come out these fields in recent years is the confirmation of neutrino having mass; they were taken to be massless in the Standard Model (SM). This discovery was made in 1998 and resulted in the Nobel Prize being awarded to Arthur McDonald and Takaaki Kajita in 2015 for their work that led to it.

Given these two facts, that neutrinos are massive, but considered massless within the framework of the Standard Model, considerable thought has been put into figuring out how the masses can be generated. This necessitates extending the framework Beyond the Standard Model (BSM). In the SM, all fermions *except neutrinos* are present in both left- and right-chiral variations. The idea to extend the neutrino sector to include right-chiral counterparts existed before the discovery of their massive nature, but the discovery gave an impetus to research in these particles. Aside from offering an explanation for why neutrinos have mass, they can also offer plausible avenues for baryogenesis and can even be dark matter candidates.

If the (left-chiral) neutrinos of the Standard Model are elusive, their right-chiral cousins are downright phantoms. While the left-chiral variants feel the effects of the weak force and gravity, a majority of models involving right-chiral (or right-handed) neutrinos (RHNs) allow them to interact only via gravity and Yukawa couplings. Numerous experiments have searched, and continue to search for them, among them are LSND, the IceCube Neutrino Observatory and MiniBooNE at Fermilab, not to mention collider experiments such as LHCb and ATLAS at CERN. Unfortunately, evidence for their existence is not yet forthcoming. Nevertheless, the results from these experiments has enabled physicists to put limits on the possible ranges of the properties of RHNs, particularly their masses and couplings.

This thesis is titled “A Frequentist Analysis of the $n = 3$ Type-I Seesaw Model” and in it, we perform a statistically rigorous combination of the most stringent constraints on RHNs, with a primary goal of mapping the profile likelihood in the coupling-mass plane. This is done in the context of the type-I seesaw model with three right-handed neutrinos.

In Chapter 1, which serves as the introduction, I briefly summarize the motivation behind the introduction of right-handed neutrinos (which I have touched upon in the preceding paragraphs) and their phenomenology. On the direct detection side, they have been looked for in meson decays and gauge boson decays, s-channel exchange of W boson and vector boson fusion. They may also leave numerous indirect signatures of their presence: in neutrinoless double-beta decay, lepton flavour-violating processes, deviations from lepton universality and in electroweak precision observables, to list a few. The chapter also discusses the improvements we have made over previous studies in order to achieve the goal of the thesis, such as the inclusion of a comprehensive number of observables, the treatment of constraints on the observed neutrino mass differences as likelihoods and the construction of flavour mixing plots for the $n = 3$ case (for the first time).

Chapter 2 contains theoretical details relevant to the thesis, beginning with the model, the type-I seesaw. In this model, the left-chiral (or left-handed) neutrinos (LHNs) of the SM are not given a Majorana mass, the Higgs sector of the SM is untouched, and the Majorana masses of the RHNs are much larger than their Dirac masses. These considerations can lead to the observed mass differences among the neutrinos in the SM. This variant is the most minimal and well-studied, avoiding complexities inherent in other models, e.g. the reformulation of the Higgs sector and the introduction of other particles.

We parameterize the model using the Casas-Ibarra (C-I) scheme, a bottom-up scheme which, by construction, satisfies neutrino oscillation data, offering a numerical advantage in sampling the parameter space. Obtaining large couplings for the RHNs using this parameterization is trivial (by choosing certain parameter ranges), but brings many subtle details along with it, which need to be considered. The parameterization route mentioned above results in large couplings, but most of these points will be a result of an unmotivated “fine tuning”. On the other hand, one can also obtain large couplings if the B-L symmetry of the SM is approximately respected, which is a natural route of achieving this end. We make careful cuts to distinguish the two kinds of points in the scans.

The Casas-Ibarra parameterization, while offering many advantages, makes it very difficult to formulate well-motivated priors on the parameters, which makes a Bayesian study of the space hard to perform. We thus opt for the frequentist approach. To carry out the scans of the parameter space, we use the Global And Modular BSM Inference Tool (GAMBIT). It is a flexible global fitting software framework which is especially suited for the task at hand, since it allows new observables to be included relatively easily and is statistically secular. The code related to this project is contained in a new module within the framework. We use a differential evolution-based algorithm, Diver, for the purpose of scanning, since this class has been shown to be particularly efficient at “optimization problems”, like the calculation of profile likelihoods. The statistical framework and details of GAMBIT and Diver are in Chapter 3.

In Chapter 4, the observables included in the study are discussed in detail. First off, the observed neutrino mass differences and mixing angles from oscillation data are included, since any extension of the SM must respect these facts. In addition, the constraint on the sum of the LHN masses from PLANCK is also enforced. As mentioned earlier, many electroweak precision observables (EWPOs) are affected by the presence of RHNs; this is because they modify the weak current and the interaction strength of the LHNs. We use the Fermi constant, Weinberg angle, mass of the W boson and its decay width to leptons, and the invisible decay width of the Z boson. In the SM, flavour-changing neutral processes, like lepton flavour violation (LFV), are suppressed due to the Glashow-Iliopoulos-Maiani (GIM) mechanism. If a BSM theory leads to contributions to these processes, they would dominate over SM contributions, making them “smoking gun” signatures of new physics. Although no positive measurements have been made,

upper limits on the branching ratios involved have been placed, and we include the most significant of these in this study. Since RHNs modify the leptonic currents, they will also contribute to tests of lepton universality; we include the most constraining semileptonic and leptonic decays. The non-zero coupling of RHNs with LHNs also affects the measurements of the CKM matrix elements, leptonic electroweak decays and pion decays (which are most sensitive to such changes) are considered, with the contribution from RHNs required to abide with the observed values of the matrix elements. The Majorana mass of the RHNs violates lepton number and allows neutrinoless double-beta decay to take place. The most constraining measurements of the effective Majorana mass, coming from GERDA and KamLAND-Zen, along with a conservative treatment of the nuclear matrix elements is encoded. Last among the indirect signatures included here is Big Bang nucleosynthesis, which forces the RHNs to conform to certain limits, the restriction coming from the requirement that they have to have a coupling strong enough to decay within a timeframe that gives the observed abundances of certain elements. Numerous direct detection experiments are also included; these are the experiments whose exclusion limits are the strongest in the mass range we consider here. The results from PIENU, PS-191, CHARM, DELPHI, ATLAS, CMS, E949 and NuTeV.

The results of this study are presented and discussed in Chapter 5. We were successful in mapping the profile likelihood in the coupling-mass plane. The largest couplings are reached above the W and Z boson masses, where the constraints from direct experiments are weak; in this region, EWPOs, LFV and CKM constraints are responsible for limiting the coupling from above. Up until the W boson mass, direct searches are the strongest, with experiments combining to sometimes give stronger limits than a simple overlap of individual results. Below masses of about 0.3 GeV, a lower bound on the mixings from BBN, upper bounds from direct searches and neutrino oscillation data also combine to give limits stronger than the ingredients. The effect of the lightest LHN mass on the pattern of flavour mixing ratios is also examined for the first time. Utilizing the cuts mentioned earlier, we also recover the behaviour previously observed in the $n = 2$ case. We also observed the presence of regions in which excess likelihood is present; this has been previously observed, and is due to the type-I seesaw model offering a better fit to observables (the invisible decay width of the Z boson, CKM and a lepton universality quantity) than the SM.

The thesis ends with a summary of the results and an outlook for future work in Chapter 6. Large parts of the allowed parameter space below 100 GeV from this study will be probed by experiments in the future; the exploration of the regions of excess may also be possible. Improvements on the LHN mass limits, along with the lower limit on their sum, may finally provide an answer to the question of which mass hierarchy the LHNs adhere to. The extension of this study to a larger range of theoretical models in a Bayesian context is also an exciting avenue to be pursued.

Acknowledgements

I had wanted to pursue research in theoretical physics since my undergraduate years. The scarcity of good opportunities to do so in India was a prime motivation for me to travel overseas and study the subject. This step would not have been possible without my family's support, both financially and emotionally. I am very grateful for the love my parents and brother gave me and for their steadfastness. I could not have accomplished this without it.

I have fewer words for what the company of my partner, Vrinda, has meant to me over this period of more than half a decade. You have kept me grounded and focused when it most mattered. I trust my failure to be eloquent here is forgivable; you know what you mean to me.

Christoph, thank you for putting up with me and enduring me for these four (I am sure long) years. Thank you for your advice and patience and for guiding me to the end of this road.

The project which forms the basis for this thesis and a majority of what I did in my PhD would not have been possible without the enthusiastic participation of my collaborators. Marco, Julia, Marcin, Tomas and everyone in GAMBIT, you have my gratitude for that.

To all my colleagues and friends at GRAPPA, I apologize for not being the most sociable person. While I may not have expressed it, the determination and drive I saw in all of you was invaluable in keeping me motivated as well.

For their love and friendship (and all the vacations and welcome breaks from work), Shri, Sushu, Swaroop, I owe you.