UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

GeV-Scale Neutrinos: Modeling, Discovery and Characterization at Forward Physics Experiments

Permalink

https://escholarship.org/uc/item/25v386s5

ISBN

9798247968092

Author

La Rocco, Daniel

Publication Date

2026-06-04

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA,
IRVINE

GeV-Scale Neutrinos: Modeling, Discovery and Characterization at Forward Physics Experiments

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Daniel La Rocco

Dissertation Committee:
Distinguished Professor Jonathan L. Feng, Chair
Assistant Professor Shirley Li
Professor Timothy M.P. Tait
Professor Daniel Whiteson

2026

© 2026 Daniel La Rocco

DEDICATION

To my wife, Susana, with whom I share a love for science.

TABLE OF CONTENTS

Page

LIST OF FIGURES v

LIST OF TABLES ix

ACKNOWLEDGMENTS xi

VITA xii

ABSTRACT OF THE DISSERTATION xiv

1 Heavy Neutral Leptons 1
- 1.1 The Standard Model . . . 1
- 1.2 Neutrino Mass Problem . . . 2
- 1.3 HNLs, Phenomenologically . . . 4
- 1.4 Mass-scale Landscape . . . 6
- 1.5 UV Completions . . . 8

2 Searches for MeV–GeV HNLs 10
- 2.1 History of HNL Searches . . . 10
- 2.2 Forward Physics . . . 13
  - 2.2.1 Forward Kinematics at Hadron Colliders . . . 13
  - 2.2.2 The Forward-LHC Experimental Program . . . 15
- 2.3 Monte-Carlo for LLPs: FORESEE . . . 18
- 2.4 Dissertation Outline . . . 20

3 Simulating Heavy Neutral Leptons with General Couplings at Collider and Fixed Target Experiments 21
- 3.1 Introduction . . . 21
- 3.2 Model . . . 24
- 3.3 HNL Production . . . 25
- 3.4 HNL Decays . . . 27
- 3.5 HNLs at FASER and FASER2 . . . 31
  - 3.5.1 Collider and Detector Setup . . . 33
  - 3.5.2 HNL Production Rates at the LHC . . . 34
  - 3.5.3 Signal and Background Considerations . . . 34
- 3.6 Reach . . . 38
- 3.7 Conclusions . . . 41

4 Characterizing Heavy Neutral Leptons: Measuring Parameters, Discriminating Majorana versus Dirac, and Using FASER2 as a Trigger for ATLAS **43**
4.1 Introduction . . . 43
4.2 Model . . . 47
4.3 Simulating HNL Signals at FASER2 . . . 49
4.4 Statistical Methods . . . 52
4.4.1 Profiled Likelihood . . . 52
4.4.2 Estimating Model Parameters . . . 53
4.4.3 Model Discrimination . . . 54
4.5 Measuring HNL Parameters with FASER2 . . . 55
4.6 Differentiating Majorana versus Dirac HNLs at FASER2 . . . 57
4.7 FASER2 as a Trigger for ATLAS . . . 57
4.7.1 Forward Muons at ATLAS . . . 58
4.7.2 Pileup Background . . . 59
4.7.3 Results . . . 61
4.8 Conclusions . . . 63

5 Conclusions and Outlook **66**
Reflections . . . 67

Bibliography **69**

Appendix A HNL Production and Decay Modeling **81**

LIST OF FIGURES

Page

1.1 Reactor $\bar{\nu}_e$ survival probability measured by the KamLAND experiment as a function of $L_0/E_{\bar{\nu}_e}$, where $L_0 = 180\,\text{km}$ is the flux-weighted average distance to the surrounding Japanese reactors. Data points (background- and geo-neutrino-subtracted) trace the characteristic oscillatory pattern of Eq. (1.4) over two full periods; the curve is the expectation from a global fit to the KamLAND data. Figure obtained from [1]. . 3

2.1 Present 90% C.L. constraints on the heavy-neutrino mixing $|U_{eN}|^2$ as a function of the HNL mass m_N, in the single-flavor-dominance benchmark in which the HNL mixes with the electron flavor only. The colored regions are excluded by the experiments indicated and the gray region is excluded by big-bang nucleosynthesis. Figure obtained from [2]. 12

2.2 Predicted forward neutrino flux at the high-luminosity LHC ($\sqrt{s} = 14\,\text{TeV}$, $\mathcal{L} = 3\,\text{ab}^{-1}$), as a function of pseudorapidity, for electron (left), muon (center), and tau (right) flavors. The flux is shown in particles per area per bin, evaluated at $z = 620\,\text{m}$ downstream of the ATLAS interaction point; the lower horizontal axis converts pseudorapidity into the equivalent radial displacement from the line-of-sight beam axis at that distance. Red and blue lines decompose the flux into contributions from light-hadron decays ($\pi^\pm$, $K^\pm$, hyperons) and charmed-hadron decays respectively, which correspond to the two production mechanisms discussed in the text. Different line styles correspond to different forward-flux generators (`SIBYLL 2.3d`, `DPMJET 3.2017`, `EPOS-LHC`, `QGSJET II-04`, `Pythia 8.2`); the spread between them is the dominant systematic on the forward flux and is discussed further in Sec. 2.3. Pseudorapidity coverages of FASERν and SND@LHC (Run 3) and the proposed FPF detectors FASERν2 and FLArE are indicated at the bottom of each panel. Although the observable plotted is the neutrino flux, the figure is a downstream visualization of the forward-peaked parent-meson production discussed in the text. Figure obtained from [3]. 14

2.3 Locations of the on-axis far-forward sites relative to the ATLAS interaction point (IP1). The proposed Forward Physics Facility (FPF) is sited $\sim 650\,\text{m}$ west of IP1, on the French side. The currently-operating FASER detector sits $\sim 480\,\text{m}$ east of IP1 in the TI12 service tunnel adjacent to the existing UJ12 cavern, on the Swiss side. SND@LHC sits at $\sim 480\,\text{m}$ west of IP1 in the TI18 service tunnel, also on the French side. Reproduced from [3]. 16

2.4 Labeled schematic of the FASER detector as installed in the TI12 service tunnel for LHC Run 3, looking towards the ATLAS interaction point. Particles incident from the right pass first through the FASERν tungsten/emulsion neutrino detector, then a front scintillator veto system and the interface tracker (IFT) that mark the upstream face of the cylindrical decay volume. The decay volume is followed by the trigger/timing scintillator station, three tracking spectrometer stations interleaved with permanent dipole magnets, and finally the electromagnetic calorimeter with its pre-shower scintillator at the downstream end. Reproduced from [4]. 17

3.1 Representative Feynman diagrams for the HNL production processes listed in Table 3.1. The subscripts m and n are generation indices. 26
3.2 Branching fractions $B(P, \tau \to NX)$ as functions of the HNL mass m_N for the 011 (left) and 111 (right) benchmarks and $\epsilon = 10^{-3}$. P represents pseudoscalar mesons, and NX represents any final state containing the HNL N. 28
3.3 Representative Feynman diagrams for HNL decays. The subscripts m, n are generation indices. 28
3.4 HNL branching fractions in the 011 benchmark model for the dominant decays via CC (left) and NC (right) interactions as functions of m_N. For $m_N \lesssim 1.0$ GeV, the decays are primarily to pions, while for $m_N \gtrsim 1.0$ GeV, the decays are dominated by decays to 3 or more hadrons. Decays to D mesons are subdominant, with $B(N \to lD) < 0.1\%$. 31
3.5 As in Fig. 3.4, but for the 111 benchmark model. 32
3.6 Lifetimes for the 011 and 111 benchmark models with $\epsilon = 10^{-3}$. A visible kink in the lifetime can be observed at $m_N \approx m_\pi$, where the decay $N \to \nu\pi^0$ becomes kinematically accessible. 32
3.7 Production rate of HNLs in the direction of FASER with $\theta < 0.2$ mrad. The rates are grouped by parent particle and are obtained by MC integration in `FORESEE` for the 011 (left) and 111 (right) benchmarks, with a coupling of $\epsilon = 10^{-3}$. 35
3.8 Production rate of HNLs in the direction of FASER2, that is, that pass through FASER2's transverse area of 3 m $\times$ 1 m. The rates are grouped by parent particle and are obtained by MC integration in `FORESEE` for the 011 (left) and 111 (right) benchmarks, with a coupling of $\epsilon = 10^{-3}$. 35
3.9 $N = 3$ signal event contours for the single-coupling 100 (top), 010 (middle), and 001 (bottom) benchmark models. Each blue contour corresponds to a detector/collider configuration given in Table 3.3. Current bounds from other particle experiments exclude the gray shaded regions, and BBN constraints exclude the region below the dashed contour. All of the production modes in Table 3.1 are included, and all of the decay modes in Table 3.2 are included, except for the invisible $\nu\bar{\nu}\nu$ modes. For the FASER and FASER2 contours, a momentum cut of $p_N > 100$ GeV has been applied. 37
3.10 As in Fig. 3.9, but for the mixed-coupling 011 (top) and 111 (bottom) benchmark models. In contrast to the single-coupling models, constraints from other experiments are not available in the literature. The current bounds shown in the gray region are roughly estimated based on a method described in the text. 38

4.1 Schematic diagrams showing example HNL events at ATLAS and FASER2 [5]. A charged B meson decays through $B^{\mp} \to \overset{(-)}{D}{}^{*0}\mu^{\mp}N$ at the ATLAS IP. The HNL then travels approximately 650 m and decays through $N \to \mu^{\mp}\pi^{\pm}$ in FASER2. For Majorana HNLs, both the top (lepton-number conserving) and bottom (lepton-number-violating) processes are allowed. For Dirac HNLs, only the top process is allowed. 44
4.2 Schematic diagram illustrating the major steps in simulating HNL events in the ATLAS and FASER2 detectors. The parameter w is the event weight; it starts with the number of hadrons produced for a particular range of energy and angle, and then is reduced by branching fractions and detector acceptance to determine the number of HNL events detected in FASER2. 50
4.3 Kinematic distributions for HNLs that decay through $N \to \mu^{\pm}\pi^{\mp}$ in FASER2. The HNLs are produced in meson decays at the HL-LHC with $E_{\mathrm{COM}} = 14$ TeV and an integrated luminosity of 3 ab^{-1}. Results are shown for Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$, as well as for variations from this model with Dirac HNLs, and with $U_\mu = 0.0040$, as indicated. The shaded error bars illustrate the envelope of expected outcomes in each bin, given by $\sqrt{\mu_i}$, where μ_i is the number of events in bin i. 51
4.4 Two-dimensional kinematic distributions for the Model 1 Asimov dataset, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$ (left), and the best-fit Dirac model with $\hat{m}_N = 1.84$ GeV and $\hat{U}_\mu = 0.0051$, obtained by minimizing $t(\vec{d}_A \,|\, m_N, U_\mu)$ (right). The Dirac best-fit point recreates the signal shape well, but at the expense of a significantly larger event rate. 54
4.5 Expected parameter estimation contours, assuming the underlying data corresponds to Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$ (left), and Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$ (right). The HNL flux normalization and shape uncertainties are assumed to be $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively. 55
4.6 Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using FASER2 measurements under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$. Left: Assuming HNL flux normalization and shape uncertainties of $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively, the Dirac hypothesis can be rejected at 87.8% CL (1.2σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 99.8% CL (2.9σ). 56
4.7 Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using FASER2 measurements under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$. Left: Assuming HNL flux normalization and shape uncertainties of $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively, the Dirac hypothesis can be rejected at 69.2% CL (0.5σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 98.2% CL (2.1σ). 56
4.8 The number of muons produced per bunch crossing in the (η, p_T) plane, including both muons that are produced in light meson ($\pi^{\pm}$, $K^{\pm}$, K^0_L) decays and heavy meson (charm and bottom meson) decays within the ITk. 60

4.9 Two-dimensional kinematic distributions for the Model 2 Asimov dataset with a Majorana HNL with $m_N = 2$ GeV and $U_\mu = 0.002$ with SS (left) and OS (right) muons seen in ATLAS, assuming an integrated luminosity of 3 ab^{-1}. These distributions include both muons from HNL events and background muons from pileup. Muons produced in association with Majorana HNLs can be found with OS or SS muons in FASER2 with equal rates, resulting in a symmetric distribution between the two histograms. 61
4.10 Two-dimensional kinematic distributions for the Dirac HNL best-fit point to the Model 2 Asimov dataset with SS (left) and OS (right) muons seen in ATLAS, obtained by minimizing the test-statistic $t(\vec{d}_A \,|\, m_N, U_\mu)$, assuming an integrated luminosity of 3 ab^{-1}. These distributions include both muons from HNL events and background muons from pileup. Muons produced in association with Dirac HNL's can only be found with OS muons in FASER2, resulting in an asymmetric distribution between the two histograms. 62
4.11 Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using correlated FASER2 and ATLAS data, under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$. Left: Assuming HNL flux normalization and shape uncertainties of 60% and 10%, respectively, the Dirac hypothesis can be rejected at 91.3% CL (1.4σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 99.7% CL (2.7σ). 63

LIST OF TABLES

Page

3.1 HNL production processes included in `HNLCalc`. The processes are ordered by increasing parent particle mass; P, V, and H denote pseudoscalar mesons, vector mesons, and hadrons, respectively; $l = e, \mu, \tau$; and N is the HNL. Charge-conjugate processes are also implemented, but are not explicitly listed here. Representative Feynman diagrams for these processes are shown in Fig. 3.1. 25

3.2 HNL decay modes included in `HNLCalc`, where H denotes hadrons. Representative Feynman diagrams for these modes are shown in Fig. 3.3. Quark level decays, $\nu q\bar{q}$ and $lu\bar{d}$, are used when computing the total hadronic width when $m_N > 1.0$ GeV. . 29

3.3 Parameters for the three experimental configurations considered: FASER (Run 3), FASER (HL-LHC), and FASER2 (HL-LHC). L is the distance from the interaction point to the front of the detector, Δ is the length of the detector along the beam axis, Geometry specifies the cross sectional area (cylindrical for FASER, rectangular for FASER2), and $\mathcal{L}$ is the integrated luminosity. Both FASER and FASER2 are assumed to be centered on the beam collision axis. 33

4.1 Summary of the two benchmark HNL models considered in this study, including the HNL mass m_N, active-sterile mixing parameter U_μ, and decay length $c\tau_N$. $\langle E_N \rangle_{\text{FASER2}}$ is the average energy of HNLs that decay in FASER2, while $\langle E_N \rangle_{\text{FASER2+ATLAS}}$ is the average energy of HNLs that decay in FASER2 and are accompanied by a signal muon that passes through ATLAS. The dominant production modes are listed, along with the number of HNLs they produce that can be detected at FASER2 and FASER2+ATLAS, assuming the baseline HL-LHC integrated luminosity of 3 ab^{-1}. The dominant decay modes and their branching fractions are also listed. All production and decay modes include both charge-conjugate processes. The fully-reconstructible decay mode $N \to \mu\pi$ used in this study is highlighted. Model 1 represents models with a large number of decays in FASER2, while Model 2 has far fewer decays in FASER, but predicts that a sizable fraction of these decays are accompanied by signal muons that pass through ATLAS. 47

4.2 Number of muons produced with $p_T > 0.4$ GeV in the ITk per bunch crossing at ATLAS from various parent particles, shown for the full muon coverage $|\eta| < 4.9$ and the forward region $3.5 < \eta < 4.9$. "Heavy Mesons" includes charm and bottom mesons. Yields include both μ^+ and μ^-. Muons are categorized by their direct parent particle, so, for example, muons from $K^0_S \to \pi^+\pi^-$, with the pions decay to muons, are included in the $\pi^\pm$ row. The number of muons produced directly from K^0_S decays is negligible. Muons from light meson decays dominate the background, yielding around 50 per bunch crossing in the full acceptance, but only about 0.7 within the forward region after cuts. 60

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Jonathan L. Feng, whose guidance and support has shaped this dissertation and my development as a particle physicist. I could not have asked for a better mentor. Additionally, I thank the members of my dissertation committee, Shirley Li, Timothy M.P. Tait, and Daniel Whiteson, for their time, careful reading, and thoughtful feedback throughout my candidacy and dissertation. I am especially grateful to Daniel Whiteson, who joined the work of Chapter 4 as a collaborator and whose guidance on the statistical analysis and on the writing itself was immensely helpful.

I had the good fortune to work alongside excellent collaborators. Alec Hewitt was a co-author on the projects of both Chapters 3 and 4 and we worked together closely for much of our PhD experience. I am grateful to have had a steady research partner throughout it all and the projects in this dissertation and the years behind them were both better for it. Felix Kling co-authored the Chapter 3 work and shared his expertise on the `FORESEE` framework and forward-physics phenomenology with patience and generosity, well beyond what any co-authorship required. Without these collaborations, neither the tools nor the analyses in this dissertation would exist in anything like their current form.

Finally, I thank my wife, Susana, and our families whose love, support, and advice throughout my PhD I could not have done without. I am forever grateful for our loving pets who have kept me company through sleepless nights and grounded to what mattered most. Lastly, I want to specifically thank my parents, Robin and Lou, my aunt Tammy, my sister Bella, and my many close friends and family for a lifetime of love, support and belief in me across the many years that lead here.

The work in Chapters 3 and 4 was previously published in *Physical Review D* as Refs. [6, 7], and I thank the American Physical Society for permission to reproduce that material here.

I acknowledge the use of Claude Opus 4.6/4.7 as a writing assistant in preparing this dissertation, primarily to draft and revise prose. All scientific content, analysis, and final editorial decisions are my own.

This work was supported in part by U.S. National Science Foundation Grants PHY-2210283 and PHY-2514888, Simons Foundation Grant 623683, and Simons Investigator Award #376204. The work of my advisor is supported in part by NSF Grant PHY-2111427 and Heising-Simons Foundation Grants 2019-1179 and 2020-1840.

VITA

Daniel La Rocco

EDUCATION

Doctor of Philosophy in Physics — **2026**
University of California, Irvine — *Irvine, CA*

Bachelor of Arts in Physics — **2021**
University of California, Irvine — *Berkeley, CA*

Associate of Science in Physics — **2019**
Orange Coast College — *Costa Mesa, CA*

RESEARCH EXPERIENCE

Graduate Researcher — **2021–2026**
University of California, Irvine — *Irvine, California*

Undergraduate Researcher — **2020–2021**
Lawrence Berkeley National Lab — *Berkeley, California*

Undergraduate Researcher — **2020–2021**
University of California, Irvine — *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant — **2021–2026**
University of California — *Irvine, CA*

REFEREED JOURNAL PUBLICATIONS

Characterizing Heavy Neutral Leptons: Measuring Parameters, Discriminating Majorana versus Dirac, and Using FASER2 as a Trigger for ATLAS **2026**
American Physical Society, Physical Review D **113**, 035006

Simulating heavy neutral leptons with general couplings at collider and fixed target experiments **2024**
American Physical Society, Physical Review D **110**, 035029

CONFERENCE PRESENTATIONS

Simulating Heavy Neutral Leptons with General Couplings at Collider and Fixed Target Experiments **June 2026**
XXXII International Conference on Neutrino Physics and Astrophysics

Characterizing Heavy Neutral Leptons: Measuring Parameters, Discriminating Majorana versus Dirac, and Using FASER2 as a Trigger for ATLAS **Mar 2026**
9th Forward Physics Facility Meeting

Simulating Heavy Neutral Leptons with General Couplings at Collider and Fixed Target Experiments **Aug 2025**
The 32nd International Conference on Supersymmetry and Unification of Fundamental Interactions

Characterizing long-lived particles at forward-physics experiments in the presence of a signal excess **Mar 2025**
The American Physical Society Global Physics Summit

Simulating Heavy Neutral Leptons with General Couplings at Collider and Fixed Target Experiments **Jun 2024**
The 10th Workshop on Flavor Symmetries and Consequences in Accelerators and Cosmology

SOFTWARE

HNLCalc https://github.com/laroccod/HNLCalc
A comprehensive python package for calculating quantities relevant to HNL phenomenology

ABSTRACT OF THE DISSERTATION

GeV-Scale Neutrinos: Modeling, Discovery and Characterization at Forward Physics Experiments

By

Daniel La Rocco

Doctor of Philosophy in Physics

University of California, Irvine, 2026

Distinguished Professor Jonathan L. Feng, Chair

In this dissertation, we work through three aspects of MeV-to-GeV heavy neutral lepton (HNL) phenomenology relevant to the forward physics experiments at the LHC: modelling, discovery, and characterization. For *modelling*, we created `HNLCalc`, a public python library that computes HNL production and decay branching ratios for arbitrary coupling ratios $|U_e|^2 : |U_\mu|^2 : |U_\tau|^2$ across the MeV-to-GeV window. Paired with the companion HNL module we built for the `FORESEE` simulation package, the two together form an end-to-end pipeline for simulating general HNL coupling patterns at any forward-LHC experiment. For *discovery*, we use these tools to compute the sensitivity reach of FASER and FASER2 in the five Physics Beyond Colliders benchmark coupling patterns, including the mixed-flavor benchmarks (011 and 111) for which the simultaneous contribution of multiple couplings to the total HNL width reshapes event rates and kinematic distributions in ways that cannot be reconstructed by combining single-coupling results. For *characterization*, we assume that a Majorana HNL has been discovered and ask how precisely FASER2 can pin down its mass, coupling, and spinor nature through standalone FASER2 measurements of fully-visible $N \to \mu\pi$ decays. Additionally, we propose the use of FASER2 as a trigger for ATLAS, the first proposal to use a far-forward detector as a hard trigger for a central LHC experiment. This sharpens Majorana-versus-Dirac discrimination by tagging events as lepton-number-conserving or lepton-number-violating from the charge of the lepton produced in association with the HNL and subsequently detected by ATLAS.

Chapter 1

Heavy Neutral Leptons

1.1 The Standard Model

The Standard Model (SM) is a gauge theory with the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$, three generations of fermions, and a single complex scalar doublet H. The matter content of each generation i comprises a quark doublet $Q_i = (u_L, d_L)_i^T$, a lepton doublet $L_i = (\nu_L, e_L)_i^T$, and three $SU(2)_L$ singlets u_{Ri}, d_{Ri}, e_{Ri}. The Yukawa interactions among these fields, invariant under the full gauge group, are given by

$$\mathcal{L}_{\text{Yuk}} \supset -y_d^{ij}\,\overline{Q_i}\,H\,d_{Rj} - y_u^{ij}\,\overline{Q_i}\,\tilde{H}\,u_{Rj} - y_e^{ij}\,\overline{L_i}\,H\,e_{Rj} + \text{h.c.}, \tag{1.1}$$

where $\tilde{H} = i\sigma_2 H^*$. Each term pairs the left-handed doublet of a charged fermion with the corresponding right-handed singlet through H, and is gauge-invariant prior to electroweak symmetry breaking (EWSB). EWSB proceeds when the Higgs potential drives H to a vacuum expectation value $\langle H \rangle = (0, v/\sqrt{2})^T$ with $v \simeq 246\,\text{GeV}$, reducing $SU(2)_L \times U(1)_Y$ to electromagnetic $U(1)_{\text{em}}$. Substituting $\langle H \rangle$ for H in Eq. (1.1) and diagonalizing the resulting mass matrices assigns each charged fermion f a Dirac mass $m_f = y_f\,v/\sqrt{2}$.

Neutrinos are the lone exception. The minimal SM contains only the left-handed neutrino, packaged inside the lepton doublet L, and no right-handed counterpart is built in. Eq. (1.1) therefore has no

neutrino term, and the minimal SM predicts $m_\nu = 0$.

1.2 Neutrino Mass Problem

Empirically, neutrinos have mass. The smoking-gun evidence is the pheonomenon known as neutrino flavor oscillation. In this process, a neutrino beam produced in one flavor at the source is detected with a reduced fraction in that same flavor at the detector, and with non-zero probability in another flavor. This in fact requires that the active flavor states are coherent superpositions of mass eigenstates with distinct masses. Concretely, in the two-flavor case parametrized by a mixing angle θ, a flavor eigenstate at the source is the superposition

$$|\nu_\alpha\rangle \;=\; \cos\theta\;|\nu_1\rangle + \sin\theta\;|\nu_2\rangle \tag{1.2}$$

of two mass eigenstates with masses m_1 and m_2. While the neutrino propagates, each mass eigenstate evolves with its own phase, $|\nu_i(t)\rangle = e^{-iE_i t}\,|\nu_i\rangle$, with $E_i = \sqrt{p^2 + m_i^2} \approx E_\nu + m_i^2/(2E_\nu)$ in the ultrarelativistic limit $E_\nu \gg m_i$. The state at the detector, with $t \approx L$ ($c = 1$), therefore is

$$|\nu_\alpha(L)\rangle \;=\; e^{-iE_\nu L}\Big[\cos\theta\, e^{-i\,m_1^2 L/(2E_\nu)}\;|\nu_1\rangle + \sin\theta\, e^{-i\,m_2^2 L/(2E_\nu)}\;|\nu_2\rangle\Big]. \tag{1.3}$$

The overall factor $e^{-iE_\nu L}$ is a global phase common to both mass-eigenstate components and cancels in any modulus-squared amplitude and is therefore physically irrelevant. What matters is the relative phase between the two terms inside the bracket, $\Delta m^2 L/(2E_\nu)$ with $\Delta m^2 \equiv m_2^2 - m_1^2$, since this is the interference phase that drives the oscillation. Projecting onto the orthogonal flavor state $|\nu_\beta\rangle = -\sin\theta\;|\nu_1\rangle + \cos\theta\;|\nu_2\rangle$ and squaring the resulting interference between the two phase-rotated terms gives the standard two-flavor oscillation probability,

$$P(\nu_\alpha \to \nu_\beta) \;=\; \sin^2(2\theta)\,\sin^2\!\Big(\frac{\Delta m^2\, L}{4E_\nu}\Big). \tag{1.4}$$

The transition probability is non-zero only if $\Delta m^2 \neq 0$; observation of oscillation therefore implies that at least two of the three neutrino mass eigenstates have non-zero mass. Fig. 1.1 illustrates the dependence directly: the KamLAND collaboration's measurement of the reactor $\bar{\nu}_e$ survival

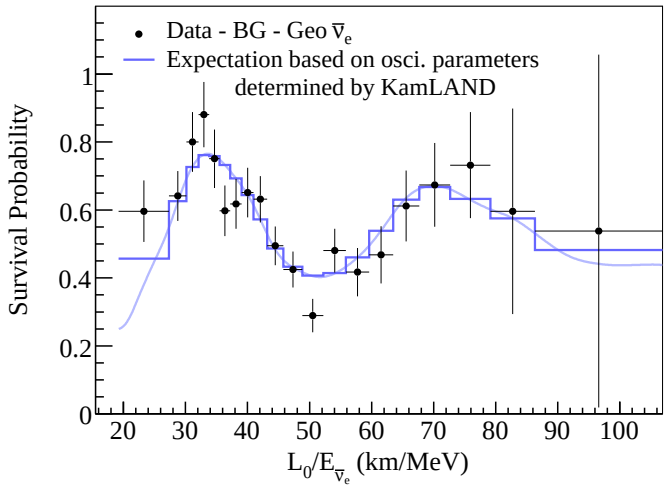


Figure 1.1: Reactor $\bar{\nu}_e$ survival probability measured by the KamLAND experiment as a function of $L_0/E_{\bar{\nu}_e}$, where $L_0 = 180\,\text{km}$ is the flux-weighted average distance to the surrounding Japanese reactors. Data points (background- and geo-neutrino-subtracted) trace the characteristic oscillatory pattern of Eq. (1.4) over two full periods; the curve is the expectation from a global fit to the KamLAND data. Figure obtained from [1].

probability traces out the sinusoid predicted by Eq. (1.4) as a function of $L_0/E_{\bar{\nu}_e}$.

Three-flavor oscillations generalize Eq. (1.4) via a unitary mixing matrix between the flavor and mass bases. In the basis where the charged-lepton mass matrix is diagonal, the active flavor states decompose as

$$\nu_{\alpha L} \;=\; \sum_{i=1}^{3} U_{\alpha i}\,\nu_{iL}, \qquad \alpha = e,\mu,\tau, \tag{1.5}$$

where the unitary matrix U is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. It is parametrized by three mixing angles $(\theta_{12},\theta_{23},\theta_{13})$, one Dirac CP-violating phase δ_{CP}, and (for Majorana neutrinos) two further phases that do not enter oscillation probabilities. Sec. 1.3 generalizes Eq. (1.5) by appending heavy mass eigenstates and identifies their active-sector mixings as the experimentally accessible heavy neutral lepton (HNL) couplings.

The phenomenon was established across three sectors during the late 1990s and early 2000s. Super-Kamiokande's observation of a deficit of atmospheric ν_μ in cosmic-ray-induced air showers [8] fixed the atmospheric mixing angle θ_{23}. The solar mixing θ_{12} was established by SNO, which resolved the long-standing solar ν_e deficit into flavor conversion through a comparison of charged-current

and neutral-current rates on deuterium [9], and confirmed by KamLAND through reactor $\bar{\nu}_e$ disappearance over baselines of order $\sim 100\,\mathrm{km}$ [10]. The smallest of the three angles, θ_{13}, was the last to be measured: it was pinned down through reactor measurements at Daya Bay [11] and through accelerator long-baseline disappearance and appearance at K2K [12], MINOS [13], T2K [14], and NOvA [15]. Global fits to the combined data yield mass-squared splittings $\Delta m^2_{21} \approx 7.4 \times 10^{-5}\,\mathrm{eV}^2$ from the solar sector and $|\Delta m^2_{31}| \approx 2.5 \times 10^{-3}\,\mathrm{eV}^2$ from the atmospheric sector [16, 17]. Oscillations therefore fix the splittings but neither the absolute mass scale nor the sign of Δm^2_{31}.

With the field content of the SM alone, both candidate renormalizable mass operators for the active neutrino are forbidden. A Dirac term $-\,y_\nu\,\overline{L}\,\tilde{H}\,\nu_R + \mathrm{h.c.}$ requires a right-handed singlet ν_R that the minimal SM does not contain. A Majorana term $-\frac{1}{2}\,m_M\,\overline{(\nu_L)^c}\,\nu_L + \mathrm{h.c.}$, written in terms of the active field alone, is not gauge-invariant under $SU(2)_L$ (it carries weak isospin $T_3 = 1$) and in any case violates lepton number by two units, an accidental symmetry of the SM. The lowest-dimension gauge-invariant Majorana operator one can write from SM fields is the dimension-five Weinberg operator [18] $(\overline{L^c}\,\tilde{H}^*)(\tilde{H}^\dagger\,L)/\Lambda$, which is non-renormalizable by inspection and therefore signals new physics at the scale Λ. Either resolution requires an extension of the SM; the natural one, adding a singlet ν_R, is taken up in Sec. 1.3. Whether the active neutrinos themselves are Dirac fermions, with a conserved lepton number, or Majorana fermions, with lepton number violated by two units, remains experimentally undetermined; with leading constriants from neutrinoless double-beta decay at KamLAND-Zen [19]. Chapter 4 develops a complementary handle on the same question via the spinor-nature measurement of a discovered right-handed neutrino, which probes the lepton-number structure of the UV completions detailed in Sec. 1.5.

1.3 HNLs, Phenomenologically

Each charged fermion of the SM appears in two pieces: a left-handed component which sits inside an electroweak doublet and a right-handed component which is an $SU(2)_L$ singlet, carrying only hypercharge. The left- and right-handed components are coupled through a Higgs Yukawa coupling, generating a Dirac mass after EWSB. Neutrinos are the lone exception: the minimal SM contains only the left-handed component of the neutrino field, with no right-handed counterpart built in. The

right-handed neutral lepton that would restore the pattern carries no SM gauge charge whatsoever, forming a *complete* singlet under $SU(3)_c \times SU(2)_L \times U(1)_Y$. We refer to this gauge eigenstate as a *sterile neutrino*. As a bonus, adding such a state to the SM supplies the most economical cure for the neutrino-mass problem of Sec. 1.2. The mechanism by which it does so depends a particular choice of UV ccompletion, which we will discuss in Sec. 1.5.

Direct experimental search probes not the sterile neutrino itself (by construction a gauge eigenstate, with no SM interactions) but rather the heavy mass eigenstates that emerge once the sterile neutrino is added to the theory and the resulting mass matrix is diagonalized. We refer to these heavy mass eigenstates as *heavy neutral leptons* (HNLs), and we will use the term HNL throughout to mean a mass eigenstate, treating the underlying mass mechanism as a separate, downstream question. Concretely, in the basis of mass eigenstates the left-handed component of each flavor neutrino field decomposes as

$$\nu_L^\alpha \;=\; \sum_{i=1}^{3} V_{\alpha i}\,\nu_i \;+\; \sum_{I} U_{\alpha I}\,N_I^c, \tag{1.6}$$

where $\alpha = e, \mu, \tau$ runs over lepton flavor, the ν_i are the three (mostly-active) light mass eigenstates, and the N_I are an unspecified number of (mostly-sterile) HNLs. The matrix V, to leading order in $|U|^2$, reduces to the PMNS matrix; the matrix U characterizes the small active–sterile admixture[1] that each HNL inherits from the active sector.

This admixture is the *only* mechanism by which the HNLs interact with SM matter. Substituting Eq. (1.6) into the SM weak currents, every charged- and neutral-current interaction involving an active neutrino acquires partner interactions in which ν_α is replaced by an HNL mass eigenstate N_I and the coupling is rescaled by $U_{\alpha I}$. The resulting charged- and neutral-current interactions read

$$\mathcal{L}_{\rm CC} \;\supset\; -\frac{g}{\sqrt{2}} \sum_{\alpha,I} U_{\alpha I}^*\,\overline{N_I^c}\,\gamma^\mu P_L\,\ell_\alpha\,W_\mu^+ + \text{h.c.}, \tag{1.7}$$

$$\mathcal{L}_{\rm NC} \;\supset\; -\frac{g}{2\cos\theta_W} \sum_{\alpha,I} U_{\alpha I}^*\,\overline{N_I^c}\,\gamma^\mu P_L\,\nu_\alpha\,Z_\mu + \text{h.c.}, \tag{1.8}$$

[1] Mass mixing is not the only portal an HNL can have to the SM. Dipole-HNLs [20], for instance, couple to the active sector through a transition magnetic-moment operator rather than through mass mixing, and other portals (light scalar mediators, hidden gauge bosons) have been studied as well. We restrict attention to the mixing portal throughout this dissertation.

where ν_α denotes the active flavor state. Every process an MeV–GeV HNL can participate in (production from heavy-flavor meson decays, decay to leptons and hadrons, scattering on detector material) follows from these two vertices alone [21, 22].

Both production and decay rates therefore carry one factor of $|U_{\alpha I}|^2$. When restricting our attention to a single HNL, neglecting interference with any others, the phenomenology is then set by four numbers, namely its mass m_N and the three mixings $|U_e|^2$, $|U_\mu|^2$, $|U_\tau|^2$. It is convenient to define the total mixing

$$\epsilon^2 \equiv |U_e|^2 + |U_\mu|^2 + |U_\tau|^2, \tag{1.9}$$

which characterizes the overall strength of the HNL's coupling to SM matter and, as we will see in Sec. 2.1, sets its proper lifetime. The pair (m_N, U_α) is what direct experimental searches actually constrain: Chapter 3 computes the FASER and FASER2 sensitivity reach as exclusion contours in the $(m_N, |U_\alpha|^2)$ plane in the benchmark coupling patterns adopted by the Physics Beyond Colliders working group [23], and Chapter 4 addresses how well a discovered HNL's (m_N, U_α) can be reconstructed. Neither calculation depends on a specific UV completion; that question is taken up next.

1.4 Mass-scale Landscape

The mass m_N of an HNL is unconstrained by the phenomenological framework of Sec. 1.3: any value is in principle allowed, and which value nature picks is fixed by the UV physics that generates the mixing in the first place. Theoretically motivated values for m_N range from the eV scale at the lower end up to the GUT scale ($\sim 10^{15}\,\text{GeV}$) at the upper end, more than twenty orders of magnitude in m_N. Direct laboratory production caps out near the few-TeV reach of the LHC, beyond which the high-scale regime can be probed only indirectly through, e.g., charged-lepton flavor violation or the matter–antimatter asymmetry from leptogenesis. Within this overall range, complementary observational handles (cosmology, astrophysics, baryogenesis, accelerator-based search) single out four qualitatively different windows. Each motivates its own search strategy, and only one is the focus of this dissertation.

$\sim$ **eV.** Motivated by short-baseline anomalies in reactor, gallium, and accelerator neutrino data [24], which point to a fourth, slightly heavier mass eigenstate mixing with the three active flavors at the percent level. Strongly constrained by big-bang nucleosynthesis and the cosmic microwave background, which together close most of the parameter space favored by the oscillation experiments [24].

$\sim$ **keV.** A keV-mass HNL is a viable warm dark matter candidate, produced either through active–sterile oscillations in the early universe (Dodelson–Widrow [25]) or resonant production in the presence of a primordial lepton asymmetry (Shi–Fuller [26]). Constrained by Lyman-α structure formation [27] and by searches for monoenergetic X-ray lines from radiative decays, a signature first proposed by Abazajian, Fuller, and Patel [28], with the most-discussed candidate the 3.5 keV line [29]; see Ref. [30] for a comprehensive review.

MeV–GeV. Accessible at high-intensity accelerator and collider experiments. The most-studied benchmark in this window is the νMSM [31, 32], which extends the SM by three HNLs: one at the keV scale to play the role of dark matter, and two nearly-degenerate at the GeV scale to generate the active neutrino masses and to drive baryogenesis through HNL oscillations. *This is the regime targeted in the rest of this dissertation.*

$\gtrsim$ **TeV.** The high-scale HNLs envisaged in the original Type-I see-saw mechanism for generating neutrino masses.(Sec. 1.5). Direct production is hopeless at any foreseeable collider, but indirect signatures (charged-lepton flavor violation, electric dipole moments, and the matter–antimatter asymmetry from high-scale leptogenesis [33]) provide complementary probes; see Refs. [22, 34] for a comprehensive review.

The choice of the MeV–GeV window in this dissertation is therefore both theoretically motivated, by its tight connection to the νMSM solution to the neutrino-mass and baryogenesis problems, and experimentally motivated, by the maturation of beam-dump and forward-LHC programs (Secs. 2.1 and 2.2) capable of probing it across the full parameter space of interest.

1.5 UV Completions

A minimal and canonical extension of the SM to include HNLs that supplies a mass for active neutrino sector is the Type-I see-saw [35, 36]. One adds a Majorana mass term for the right-handed singlet alongside a Yukawa coupling that connects it to the active sector,

$$\mathcal{L} \supset -(y_\nu)_{\alpha I}\,\overline{L_\alpha}\,\tilde{H}\,N_I \;-\; \tfrac{1}{2}\,(M_R)_{IJ}\,\overline{N_I^c}\,N_J \;+\; \text{h.c.}, \tag{1.10}$$

where $\alpha = e, \mu, \tau$ runs over lepton flavor, $I, J = 1, \ldots, n_R$ index the n_R right-handed singlets, L_α is the lepton doublet of flavor α, $\tilde{H} = i\sigma_2 H^*$, and $(M_R)_{IJ}$ is the Majorana mass matrix. After EWSB the Yukawa generates a Dirac mass matrix $(m_D)_{\alpha I} = (y_\nu)_{\alpha I}\, v/\sqrt{2}$, and diagonalizing the resulting $(3+n_R)\times(3+n_R)$ mass matrix in the limit where the eigenvalues of M_R are much larger than the entries of m_D yields three light mass eigenstates with mass matrix

$$m_\nu \;\approx\; -\,m_D\,M_R^{-1}\,m_D^T \tag{1.11}$$

and n_R heavy mass eigenstates whose masses are the eigenvalues of M_R, with active–sterile mixings $U_{\alpha I} \approx (m_D\,M_R^{-1})_{\alpha I}$. The size of the mixing is therefore tied to the light neutrino mass scale: a heavy state of mass m_N contributes a typical mixing

$$|U|^2 \;\sim\; \frac{m_\nu}{m_N} \;\approx\; 5\times 10^{-11}\left(\frac{1\,\text{GeV}}{m_N}\right), \tag{1.12}$$

where the numerical estimate uses the atmospheric mass-splitting scale $m_\nu \sim 0.05\,\text{eV}$. The locus of points satisfying Eq. (1.12) traces a diagonal line across the $(m_N, |U|^2)$ plane: the Type-I see-saw line. **This line lies one to two orders of magnitude below the reach of any current or proposed accelerator-based search.**

Direct experiments therefore probe the regime in which the Type-I relation has been deformed in some way. Several theoretically motivated mechanisms accomplish exactly this. The inverse see-saw [37] augments the singlet sector with an additional gauge-singlet S and a small Majorana mass μ_S, so that lepton number is preserved in the limit $\mu_S \to 0$; the light neutrino mass picks up a factor of $\mu_S/M_R \ll 1$ relative to the Type-I result, and at fixed m_ν the active–sterile mixing

$|U|^2$ can be much larger than its Type-I value. The linear see-saw [38] achieves a parametrically similar suppression through a different singlet structure. More generally, lepton-number-symmetric (pseudo-Dirac) arrangements pair two singlets at nearly degenerate masses, with an approximate $U(1)_L$ symmetry under which the lepton-number-violating combinations of mass parameters are small [39]; the active neutrino mass is suppressed by these small combinations rather than by the heavy scale m_N. Concretely, the splitting ΔM between the two pseudo-Dirac partners replaces m_N in the see-saw denominator, so $m_\nu \sim |U|^2 \Delta M$, and a splitting ΔM small enough to enforce approximate lepton-number symmetry but large enough to generate $m_\nu \sim 0.05\,\text{eV}$ at experimentally accessible mixings $|U|^2 \sim 10^{-9}$ requires ΔM of order tens of MeV. Two HNLs are sufficient to generate two nonzero active masses, leaving the lightest active state massless. This is consistent with the observed atmospheric and solar mass-squared splittings provided the ordering is normal ($m_1 = 0$) or inverted ($m_3 = 0$). The νMSM [31, 32] is a specific example, with two quasi-degenerate HNLs at the GeV scale (N_2, N_3) playing this role; a third HNL at the keV scale (N_1) is the dark matter candidate of Sec. 1.4 and contributes negligibly to active masses. The common feature of all of these mechanisms is that an approximate lepton-number symmetry (or a structurally similar device) decouples $|U|^2$ from m_ν and populates the region above the Type-I line.

The MeV–GeV experiments examined in this dissertation operate in this deformed-Type-I regime. A discovery of an HNL by the searches developed in Chapters 3 and 4 would therefore not vindicate the original Type-I see-saw but rather provide evidence for a low-scale extension of it, and the natural question becomes which extension. Chapter 4 returns to this question through the Dirac-versus-Majorana discrimination it develops: a true Majorana signature would point towards a Type-I-like origin, while an effectively Dirac signature would point towards a pseudo-Dirac like construction (lepton-number-symmetric inverse see-saw, νMSM, etc.). The spinor-nature measurement is therefore not merely a property of HNLs, but rather a probe of which UV completion of the neutrino mass mechanism is realized in nature.

Chapter 2

Searches for MeV–GeV HNLs

2.1 History of HNL Searches

The systematic study of HNL signals at accelerators began as early as the late 1970's [40–42], with the proposal of peak searches in pseudoscalar-meson decays. In the SM, the leptonic decays $\pi^+ \to \ell^+\nu_\ell$ and $K^+ \to \ell^+\nu_\ell$ are chirally suppressed by a factor m_ℓ^2/m_M^2. In the HNL case, the partial width for a charged pseudoscalar meson M^+ to decay into a charged lepton and N_I is given by [21, 42, 43]

$$\Gamma(M^+ \to \ell^+ N_I) \;=\; \frac{G_F^2\,|V_{qq'}|^2\,f_M^2\,m_M^3}{8\pi}\,|U_{\ell I}|^2\,\lambda^{1/2}\!\left(1,\,x_\ell^2,\,x_N^2\right)\Big[\,x_\ell^2 + x_N^2 - \left(x_\ell^2 - x_N^2\right)^2\Big], \tag{2.1}$$

where f_M is the meson decay constant, $V_{qq'}$ the CKM element associated with the parent meson's valence quarks, $x_i \equiv m_i/m_M$, and $\lambda(a,b,c) = a^2 + b^2 + c^2 - 2(ab + bc + ca)$ is the Källén function. Setting $m_N \to 0$ in Eq. (2.1) reproduces the SM result with its characteristic helicity-suppression factor $x_\ell^2(1 - x_\ell^2)^2$. For a heavy daughter neutrino with mass $m_N \lesssim m_M$, however, the x_N^2 term in the bracket comes to dominate and the helicity suppression is lifted, so any HNL admixture in the final-state neutrino mass eigenstates would manifest as a discrete peak in the charged-lepton energy spectrum at $E_\ell = (m_M^2 + m_\ell^2 - m_N^2)/(2m_M)$. The signature is purely kinematic: the HNL itself need not be detected.

Peak searches were quickly implemented at SIN [44] and KEK [45] in the early 1980s using π^+ and K^+ decays, and the sensitivity of the resulting bounds has been progressively improved through searches at TRIUMF [46] and BNL [47]. The PIENU experiment at TRIUMF [48] and NA62 at CERN [49] represent the current precision frontier, setting bounds at the $|U_{e,\mu}|^2 \sim 10^{-8}$ level over the kinematically accessible mass window.

The window itself, however, is narrow since peak searches in π and K decays cap kinematically at $m_N \lesssim m_K - m_\ell \sim 0.4\,\mathrm{GeV}$. Reaching the GeV regime requires moving to heavier parent mesons (charm and beauty), and to a different signature: rather than a peak in the production-side lepton spectrum, the HNL is detected through its *displaced decay* downstream of the production target.

The reason this strategy works is a kinematic feature of HNL decays. Because the HNL inherits its couplings from the active sector through the small mixing ϵ^2, its total width is parametrically suppressed,

$$\Gamma_N \;\propto\; \epsilon^2\, G_F^2\, m_N^5, \tag{2.2}$$

and the corresponding proper decay length $c\tau_N \propto 1/(\epsilon^2 G_F^2 m_N^5)$ ranges from millimeters to kilometers across the parameter space accessible to current and proposed experiments. An MeV–GeV HNL with the small mixings of interest therefore travels a macroscopic distansce before decaying: it is a *long-lived particle* (LLP), and the experimental strategy of choice is to instrument a region downstream of the production target where the HNL is likely to decay.

The resulting class of *displaced-vertex searches* was developed in the 1980s by fixed-target beam-dump experiments using charm-driven HNL production: PS191 [50], CHARM [51], and NuTeV [52]. More recent realizations of the same beam-dump strategy include T2K-ND280 [53] and MicroBooNE [54], and the proposed SHiP experiment at the CERN SPS [55] is a direct descendant. The displaced-vertex strategy has since been adopted at colliding-beam machines as well, with very different production kinematics: at LEP, DELPHI [56] searched for HNL decays from Z production with both prompt and displaced topologies; at the LHC, the ATLAS [57], CMS [58], and LHCb [59] collaborations have set displaced-vertex limits using HNLs produced in W, Z, and heavy-flavor meson decays in pp collisions. A comprehensive recent survey of these constraints is given in [2, 22];

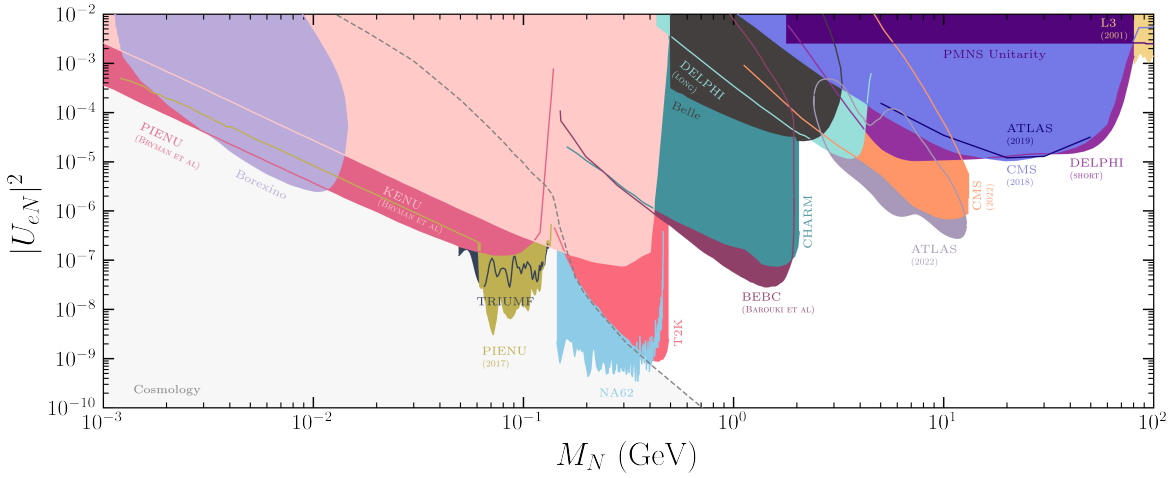


Figure 2.1: Present 90% C.L. constraints on the heavy-neutrino mixing $|U_{eN}|^2$ as a function of the HNL mass m_N, in the single-flavor-dominance benchmark in which the HNL mixes with the electron flavor only. The colored regions are excluded by the experiments indicated and the gray region is excluded by big-bang nucleosynthesis. Figure obtained from [2].

Fig. 2.1 reproduces the bounds on the electron-flavor mixing $|U_{eN}|^2$ as a representative example.

The ForwArd Search ExpeRiment (FASER) [60, 61] (currently running) is the first dedicated *forward-LLP* detectors at a hadron collider. Its distinguishing geometry is threefold: the detector sits roughly 480 m downstream of the ATLAS interaction point (IP) on the line-of-sight (LOS) beam axis; it is shielded from the IP by hundreds of meters of rock and the LHC's bending magnets, which suppress beam-induced backgrounds below the levels achievable in fixed-target beam-dump configurations; and they exploit the LHC's high-pseudorapidity flux of heavy-flavor hadrons. The combination yields sensitivity competitive with much larger fixed-target detectors. The displaced-decay signature itself is a 1980s development; the novelty of FASER and FASER2 lies in the geometry that allows a small forward detector to apply this signature at LHC luminosities. This is the experimental setting for the rest of this dissertation, and the forward-collider geometry that makes the strategy work is the subject of Sec. 2.2.

2.2 Forward Physics

2.2.1 Forward Kinematics at Hadron Colliders

The natural angular variable for hadron-collider physics is pseudorapidity. For a particle emerging from the interaction point at polar angle θ measured from the beam axis, the pseudorapidity is

$$\eta \equiv -\ln\tan(\theta/2), \tag{2.3}$$

which depends only on the angle that a detector actually measures and runs from zero in the transverse direction to $\pm\infty$ along the beam axis. The forward direction at the LHC corresponds to large $|\eta|$: the central trackers of ATLAS and CMS cover $|\eta| \lesssim 2.5$, their forward calorimeters extend to $|\eta| \approx 5$, and the on-axis far-forward region targeted by FASER reaches $\eta \gtrsim 9$.

Two distinct production mechanisms populate the forward LHC hadron flux. The first is the soft production of light hadrons. The π^0 and K mesons that drive low-mass production are emitted predominantly through soft minimum-bias inelastic processes including diffractive scattering and multi-parton interactions. The differential cross sections for these processes are computed from phenomenological string-fragmentation models implemented in event generators such as PYTHIA, EPOS-LHC, and SIBYLL.

The second is the perturbative production of heavy-flavor hadrons. The D and B mesons that drive production at the GeV scale originate in the hard process $gg \to c\bar{c}, b\bar{b}$, whose matrix element is calculable in perturbative QCD at NLO and NNLO with the heavy-quark mass providing the hard scale; the perturbative matrix element is then matched to a parton shower and convolved with the same Lund-string hadronization models that handle the light hadrons. Chapter 3 takes its forward-charm spectra from POWHEG matched with PYTHIA, following the prescription of Ref. [62].

What unites the two mechanisms phenomenologically is that both yield strongly forward-peaked differential cross sections at $\sqrt{s} \simeq 13.6\,\text{TeV}$, with substantial yields beyond $\eta \gtrsim 6$ [3, 62, 63]. Fig. 2.2 illustrates the consequence empirically: the forward neutrino flux peaks at $\eta \gtrsim 7$ across all

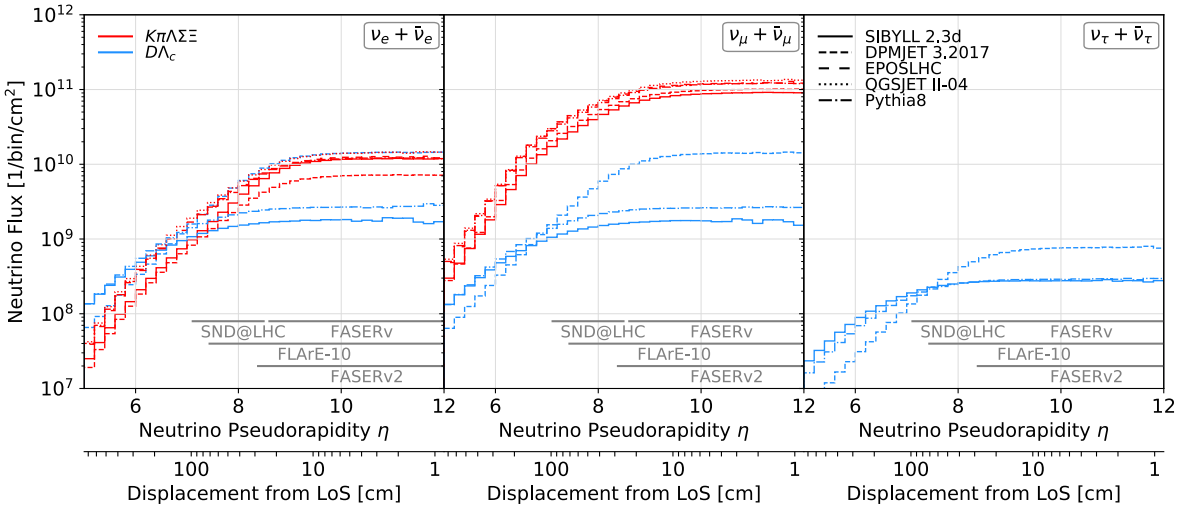


Figure 2.2: Predicted forward neutrino flux at the high-luminosity LHC ($\sqrt{s} = 14\,\text{TeV}$, $\mathcal{L} = 3\,\text{ab}^{-1}$), as a function of pseudorapidity, for electron (left), muon (center), and tau (right) flavors. The flux is shown in particles per area per bin, evaluated at $z = 620\,\text{m}$ downstream of the ATLAS interaction point; the lower horizontal axis converts pseudorapidity into the equivalent radial displacement from the line-of-sight beam axis at that distance. Red and blue lines decompose the flux into contributions from light-hadron decays ($\pi^\pm$, $K^\pm$, hyperons) and charmed-hadron decays respectively, which correspond to the two production mechanisms discussed in the text. Different line styles correspond to different forward-flux generators (`SIBYLL 2.3d`, `DPMJET 3.2017`, `EPOS-LHC`, `QGSJET II-04`, `Pythia 8.2`); the spread between them is the dominant systematic on the forward flux and is discussed further in Sec. 2.3. Pseudorapidity coverages of FASERν and SND@LHC (Run 3) and the proposed FPF detectors FASERν2 and FLArE are indicated at the bottom of each panel. Although the observable plotted is the neutrino flux, the figure is a downstream visualization of the forward-peaked parent-meson production discussed in the text. Figure obtained from [3].

three neutrino flavors, with distinct contributions from the light-hadron and charm-meson decay chains. The spread between generators visible in Fig. 2.2 sets the dominant theoretical uncertainty on the forward flux and is taken up in Sec. 2.3.

Any light particle produced in the decay of highly-boosted forward hadrons, including a hypothetical light BSM particle X, will be collimated along the beam axis. A meson of mass m_M produced with energy E_M has Lorentz factor $\gamma_M = E_M/m_M$. In the meson rest frame, a two-body decay $M \to X + \text{rec}$ emits the daughter X isotropically into the full 4π solid angle with a momentum

$$p^* \;\sim\; \tfrac{1}{2}\sqrt{m_M^2 - (m_X + m_{\text{rec}})^2} \tag{2.4}$$

fixed by two-body kinematics, where m_X is the daughter mass and m_{rec} is the mass of the recoiling

SM partner. Under the Lorentz boost to the lab frame, the transverse component of this momentum is preserved while the longitudinal component is amplified by γ_M, so the daughter emerges in the lab with $p_T^X \sim p^*$ and $p_z^X \sim \gamma_M\, m_M = E_M$. The opening angle of the daughter relative to the parent direction is the ratio of the two,

$$\theta_X \sim \frac{p_T^X}{p_z^X} \sim \frac{m_M}{E_M}, \tag{2.5}$$

The argument generalizes to any decay in which the recoiling system carries a rest-frame momentum of order m_M leads to the same parametric scaling. For typical LHC kinematics, parent-meson energies of a few hundred GeV to several TeV combined with parent-meson masses spanning $m_M \sim$ 0.1–5 GeV place θ_X in the milliradian range. This collimation argument was the original physics motivation for an on-axis far-forward detector at the LHC to search for dark photons produced in $\pi^0 \to \gamma A'$ decays [64], later expanding their proposal to include a range of LLP models like dark Higgs bosons, axion-like particles, and HNLs [60, 61].

2.2.2 The Forward-LHC Experimental Program

The forward-LHC program is hosted in a small set of locations along the line-of-sight beam axis emerging from the ATLAS interaction point. The LHC's magnets bend the proton beams off-axis and the TAN neutral-particle absorber blocks neutral hadrons, leaving a region of empty rock through which only the forward flux of weakly-interacting and long-lived particles continues to travel. Fig. 2.3 situates the relevant on-axis sites: the existing UJ12 cavern at ~ 480 m east of IP1 on the Swiss side, with the adjacent TI12 service tunnel currently housing FASER and FASERν, and the proposed Forward Physics Facility (FPF) cavern at ~ 650 m west of IP1 on the French side, the baseline site for the next generation of forward-LHC experiments at the high-luminosity LHC (HL-LHC) [3, 63].

Three experiments currently exploit the on-axis far-forward region during LHC Run 3. The detector considered in this dissertation is FASER [60, 61, 64], the dedicated forward-LLP detector with a cylindrical decay volume of radius 10 cm and length 1.5 m centered on the beam axis at $L =$ 480 m, instrumented with a magnetic spectrometer, electromagnetic calorimeter, and front and

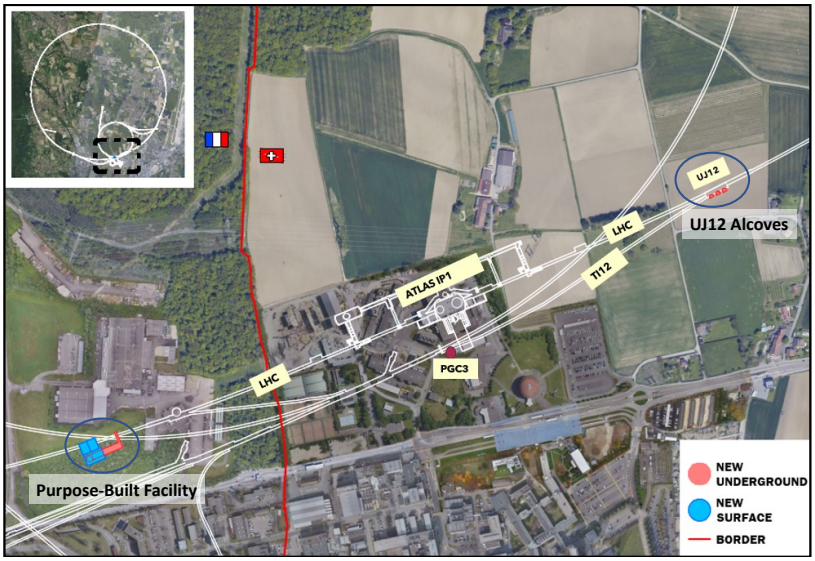


Figure 2.3: Locations of the on-axis far-forward sites relative to the ATLAS interaction point (IP1). The proposed Forward Physics Facility (FPF) is sited $\sim 650\,\text{m}$ west of IP1, on the French side. The currently-operating FASER detector sits $\sim 480\,\text{m}$ east of IP1 in the TI12 service tunnel adjacent to the existing UJ12 cavern, on the Swiss side. SND@LHC sits at $\sim 480\,\text{m}$ west of IP1 in the TI18 service tunnel, also on the French side. Reproduced from [3].

rear scintillator veto stations; see Fig. 2.4 for a labeled schematic of the as-built detector.

FASERν, an emulsion neutrino detector located immediately upstream of FASER [65], has since produced the first dedicated cross-section measurements of electron and muon neutrinos at TeV energies [66, 67]. SND@LHC, sited in the TI18 tunnel on the opposite side of IP1, covers a complementary off-axis pseudorapidity range with a hybrid emulsion-and-electronic neutrino detector [68].

The proposed FPF is a dedicated underground cavern designed to host several complementary experiments running concurrently with the HL-LHC's $\mathcal{L} = 3\,\text{ab}^{-1}$ of integrated luminosity [3, 63]. The experiment of relevance for this dissertation in the FPF is FASER2, the LLP-search successor to FASER, with a rectangular decay volume of $3\,\text{m} \times 1\,\text{m} \times 10\,\text{m}$ chosen to expand both the angular acceptance (to $\theta \approx 2\,\text{mrad}$) and the fiducial volume by roughly two orders of magnitude relative to FASER. The same cavern would also house FLArE, a liquid-argon TPC for high-statistics neutrino scattering and BSM scattering signatures, and FORMOSA, a scintillator-based search for millicharged particles. The forward-LHC physics-case documents [3, 63] cover the broader scientific motivation; for the purposes of this dissertation only FASER and FASER2 are analyzed in detail, with the other FPF experiments listed for context.

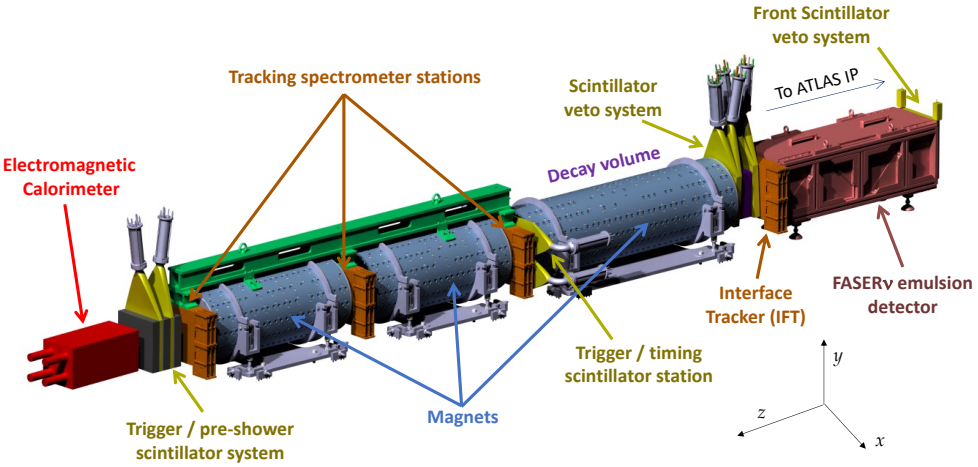


Figure 2.4: Labeled schematic of the FASER detector as installed in the TI12 service tunnel for LHC Run 3, looking towards the ATLAS interaction point. Particles incident from the right pass first through the FASERν tungsten/emulsion neutrino detector, then a front scintillator veto system and the interface tracker (IFT) that mark the upstream face of the cylindrical decay volume. The decay volume is followed by the trigger/timing scintillator station, three tracking spectrometer stations interleaved with permanent dipole magnets, and finally the electromagnetic calorimeter with its pre-shower scintillator at the downstream end. Reproduced from [4].

The single feature shared by every on-axis far-forward LLP detector is the unusually low background environment from infrastructural shielding. A hundred meters of rock between the ATLAS interaction point and the FASER and FPF locations, together with the LHC's bending magnets, suppress backgrounds from the IP to negligible levels. The remaining background components are penetrating high-energy forward muons, vetoed by dedicated front-veto scintillator stations, and the residual neutrino flux through the decay volume, which interacts at calculable rates and has now been measured directly by FASER and FASERν [66, 69].

The forward-LHC infrastructure described in this section is the experimental setting for both Chapter 3, which computes the sensitivity reach of FASER and FASER2 to MeV–GeV HNLs with generic flavour structure, and Chapter 4, which assumes a discovery and asks how well FASER2 can pin down a discovered HNL's mass, coupling, and spinor nature.

2.3 Monte-Carlo for LLPs: FORESEE

The simulation framework used in this dissertation for computing the predicted HNL signal rate at FASER is built around the publicly available `FORESEE` (FORward Experiment SEnsitivity Estimator) package [70], together with the companion `HNLCalc` library developed in Chapter 3, which supplies the model-dependent HNL production and decay branching ratios that `FORESEE` requires. The combined pipeline runs in three stages: a forward-flux Monte-Carlo (MC) generator supplies the parent-particle flux into the on-axis far-forward region; `FORESEE` convolves this flux with the HNL production branching ratios from `HNLCalc` and the parent-decay kinematics to obtain the forward HNL flux; the HNL is then propagated to the detector and the probability that it decays inside the fiducial volume is folded against the visible-decay branching fraction to give the expected event yield.

The two production mechanisms identified in Sec. 2.2.1 dictate the choice of generator. Light hadrons (π, K, hyperons) are produced through soft inelastic processes with no perturbative hard scale, modelled by hadronic-interaction generators originally developed and tuned for cosmic-ray air-shower physics: `EPOS-LHC` [71], `SIBYLL 2.3d` [72], `QGSJET II-04` [73], and `DPMJET 3.2017` [74]. Charm and beauty hadrons originate in the perturbative process $gg \to c\bar{c}, b\bar{b}$ and are simulated at NLO with `POWHEG` [75] matched to the `PYTHIA 8` [76] parton shower; Chapter 3 takes the forward-charm spectrum tabulated by Ref. [62]. τ leptons are produced primarily through $D_s \to \tau\nu$ and inherit the heavy-flavor flux modelling. The output of these generators, and the primary input of `FORESEE`, is the doubly-differential flux of parent-particles in the on-axis far-forward region,

$$\frac{d^2 N_P}{dE_P\, d\theta_P} \;=\; \mathcal{L}\, \frac{d^2 \sigma_{pp\to P\,X}}{dE_P\, d\theta_P}\,, \tag{2.6}$$

with $\mathcal{L}$ the integrated luminosity, pp the colliding proton beams, and the index P ranging over the parent species above.

`FORESEE` [70] is a publicly-available simulation package that takes the parent fluxes of Eq. (2.6) as input and produces, for any specified LLP model, the doubly-differential particle flux at a

downstream detector face,

$$\frac{d^2N_N}{dE_N\,d\theta_N} \;=\; \sum_P \int dE_P\,d\theta_P\;\frac{d^2N_P}{dE_P\,d\theta_P}\,\frac{d^2\mathrm{BR}_{P\to N}}{dE_N\,d\theta_N}(E_P,\theta_P;E_N,\theta_N)\,, \tag{2.7}$$

where the sum runs over parent species and $d^2\mathrm{BR}_{P\to N}/(dE_N\,d\theta_N)$ is the doubly-differential branching fraction for a parent of kinematics (E_P,θ_P) to produce an HNL of kinematics (E_N,θ_N). The integral is evaluated with the use of Monte-Carlo sampling. The model-dependent input is the differential branching fraction itself, which for an HNL of mass m_N and mixings U_α is computed in Chapter 3. `FORESEE` ships with modules for several other LLP candidates (dark photons, dark Higgs bosons, axion-like particles); Chapter 3 adds an HNL module to that set.

After calculating the forward-HNL flux, the second half of the calculation determines the fraction of HNLs that actually decay inside the fiducial volume. An HNL with energy E_N and proper lifetime $\tau_N = 1/\Gamma_N$ obtained from the total HNL width of Eq. (2.2) travels a lab-frame distance

$$\ell_N \;=\; \beta\gamma\,c\tau_N \;\approx\; \frac{E_N}{m_N}\,c\tau_N\,, \qquad c\tau_N \;\propto\; \frac{1}{\epsilon^2\,G_F^2\,m_N^5}\,, \tag{2.8}$$

before decaying. For an MeV–GeV HNL with TeV-scale energies and the mixings of experimental interest, $\epsilon^2 \sim 10^{-7}$ to 10^{-9}, ℓ_N ranges from meters to many kilometers. The on-axis far-forward sites at $L \sim 480$ to 650 m sit squarely inside this window, which is what makes a search at this baseline sensitive to HNL couplings well below those probed by closer-baseline experiments. The probability that an HNL produced at the IP survives to the detector at a distance L along the beam line and then decays inside its fiducial length Δ is given by

$$P_{\mathrm{dec}}(\ell_N;L,\Delta) \;=\; e^{-L/\ell_N}\left(1-e^{-\Delta/\ell_N}\right). \tag{2.9}$$

Putting everything together, the expected number of reconstructed HNL decays in a detector of fiducial length Δ at distance L is the integral of the doubly-differential HNL flux from the `FORESEE` convolution above against the per-event weight,

$$N_{\mathrm{sig}} \;=\; \int dE_N\,d\theta_N\;\frac{d^2N_N}{dE_N\,d\theta_N}\;\varepsilon_{\mathrm{geom}}(\theta_N)\;P_{\mathrm{dec}}\big(\ell_N(E_N);L,\Delta\big)\;\mathrm{BR}(N\to\mathrm{vis})\,, \tag{2.10}$$

where $\varepsilon_{\text{geom}}$ is the geometric acceptance and $\text{BR}(N \to \text{vis})$ is the fraction of HNL decays into visible final states.

2.4 Dissertation Outline

In this dissertation, we work through all aspects of MeV-to-GeV HNL phenomenology relevant to the forward physics experiments: modelling, discovery, and characterization. In Chapter 3 we introduce `HNLCalc` and a companion `FORESEE` HNL module, simulation tools that compute HNL production and decay branching ratios for arbitrary coupling ratios $|U_e|^2 : |U_\mu|^2 : |U_\tau|^2$ in the MeV-to-GeV window, and uses them to determine the sensitivity reach of FASER and FASER2 in the five benchmark coupling patterns of [23]. In Chapter 4 we assume that a Majorana HNL has been discovered and ask how precisely FASER2 can pin down its mass, coupling, and spinor nature from either the kinematics of its visible decays or by using FASER2 as a Level-0 trigger for ATLAS, the first proposal to use a far-forward detector as an ATLAS hard trigger and the first to combine far-forward and central-detector measurements to probe HNL spinor nature. Lastly, we collect our conclusions and outlook in Chapter 5. Detailed expressions for the HNL production and decay branching fractions implemented in `HNLCalc` are collected in Secs. A.1 and A.2.

Chapter 3

Simulating Heavy Neutral Leptons with General Couplings at Collider and Fixed Target Experiments

The following chapter is largely based on Ref. [6] in collaboration with Jonathan L. Feng, Alec Hewitt and Felix Kling.

3.1 Introduction

In Chapter 1, we established the model-independent framework for HNLs in the MeV–GeV mass window, parametrized by the HNL mass m_N and the active–sterile mixings $U_{\alpha I}$ (Eqs. (1.6) to (1.9)). Such HNLs have attracted interest since at least the early 1970's [77, 78]; for a review that includes references to the early literature, see Ref. [79] or Sec. 2.1. In recent years, however, there has been a resurgence of interest in HNLs with masses in the MeV to GeV range. HNLs with such masses may be used to generate neutrino masses and mixings consistent with experimental measurements and simultaneously address many of the cosmological problems of the SM; see, for example, Ref. [32] or Sec. 1.5. More generally, from a purely phenomenological perspective, such HNLs are amenable to a wide variety of experimental probes, since they are light enough to be copiously produced

at many collider and accelerator facilities, but also heavy enough to decay to visible particles on length scales that can be observed in particle detectors [21, 40, 43, 80–85]; see also Sec. 1.3. For these reasons, HNLs have become a leading example of LLPs, and they have helped motivate the growing world-wide research program in search of LLPs.

To evaluate the prospects for HNL discovery, it is essential to have a user-friendly, fast, and flexible tool that can model HNL production and decay and also be used to estimate event rates in current and proposed experiments. The number of production channels and decay modes that may be important for HNLs dwarfs the corresponding number for dark photons and other well-known LLPs. Thankfully, a great deal of work has been done to identify and quantify the leading production and decay processes [41–43, 80, 81, 83, 84, 86, 87]. Building on this work, in this study, we present the python library `HNLCalc`, which computes the properties of general HNL models.[1] It includes all of the potentially dominant processes and accurately describes the production and decay of HNLs with $\mathcal{O}$(GeV) masses and arbitrary couplings to the e, μ, and τ leptons.

HNLs may be produced in both fixed target and particle collider experiments. Fixed target experiments may produce many HNLs in the MeV to GeV range, and it is certainly of interest to predict the sensitivity of current and proposed experiments for general HNL searches. At the same time, particle colliders, like the LHC, are also of interest, particularly in the forward direction, where event rates for such HNLs are greatly enhanced. Our results may be used to model general HNL models and determine discovery prospects at both accelerator and collider experiments.

As examples, in this study, we consider FASER [4, 64], a current experiment that is purpose-built to search for LLPs in the far-forward region at the LHC, and FASER2, a future experiment to be housed in the proposed FPF [3, 63] at the HL-LHC, which will have an even greater discovery potential for LLPs. For forward physics experiments at colliders, the `FORESEE` [70] simulation package has become a very useful tool. For a fixed proton-proton center-of-mass (COM) energy, `FORESEE` takes as input the forward hadron production rates given by MC generators, determines the resulting energy and angular distribution for various LLPs, and then calculates the signal rate in particle detectors. `FORESEE` modules have been written to simulate a variety of LLPs, including

[1] `HNLCalc` is publicly available on GitHub at https://github.com/laroccod/HNLCalc.

dark photons, other gauge bosons, and also scalars, but, before this work, not HNLs. In this chapter, we extend `FORESEE` to simulate HNLs, using the HNL properties provided by the `HNLCalc` package. We will focus on models with Majorana-like HNLs in which total lepton number L is violated.

To illustrate the flexibility of both `HNLCalc` and the `FORESEE` HNL module, we analyze the sensitivity reach for five benchmark models, where the ratios of HNL couplings are

$$\begin{array}{lccccc} & |U_e|^2 & : & |U_\mu|^2 & : & |U_\tau|^2 \\ \text{Benchmark 1}: & 1 & : & 0 & : & 0 \\ \text{Benchmark 2}: & 0 & : & 1 & : & 0 \\ \text{Benchmark 3}: & 0 & : & 0 & : & 1 \\ \text{Benchmark 4}: & 0 & : & 1 & : & 1 \\ \text{Benchmark 5}: & 1 & : & 1 & : & 1\,. \end{array} \tag{3.1}$$

We will refer to these benchmarks using the convenient shorthand 100, 010, 100, 011, and 111. The 100, 010, and 001 models are simple cases, each with only one nonzero coupling, and they are among the benchmarks typically considered, for example, by the Physics Beyond Colliders study group [23]. The reach for FASER in these scenarios has been analyzed previously in Refs. [60, 88, 89], and we reproduce these earlier results. The 011 and 111 models are less minimal, but have been proposed as more representative of models that explain the observed neutrino masses and mixings [90]. The reach for FASER and FASER2 in the 011 and 111 models has not been determined previously, but is analyzed here rather easily, given the flexibility of the work described here. We note that sensitivities for models with more than one non-zero coupling cannot be estimated simply by adding together event rates from the 100, 010, and 001 models, since all couplings enter the decay width, and so turning on a second or third coupling impacts both the event rates and the kinematic distributions of events mediated by the first coupling in highly non-trivial ways. This interplay can only be included through the detailed simulation work described here.

This chapter is organized as follows: In Sec. 3.2, we discuss the HNL model and establish the

notation we will use. In Secs. 3.3 and 3.4 we discuss the many HNL production and decay processes, respectively, and we describe the detector and collider configuations we will consider in Sec. 3.5. In Sec. 3.6 we use all of these results to determine the sensitivity and discovery prospects for FASER and FASER2 in the five benchmark models discussed above. We collect our conclusions in Sec. 3.7. Detailed expressions for the production and decay branching fractions are contained in Secs. A.1 and A.2, respectively. In the course of this work, typos and errors in the existing literature were identified, and these are noted in the Appendices.

3.2 Model

The Lagrangian, mixing parametrization, and electroweak interactions of HNLs were established in Sec. 1.3: the Type-I extension of the SM by gauge-singlet sterile neutrinos (Eq. (1.10)), the active–sterile mixing matrix $U_{\alpha I}$ that defines the HNL mass eigenstates N_I (Eq. (1.6)), and the resulting charged- and neutral-current interactions through which the HNLs couple to SM matter (Eqs. (1.7) and (1.8)). We adopt those conventions throughout this chapter and focus on Majorana-like HNLs, in which total lepton number L is violated.

In this study, for simplicity, we consider models in which the phenomenology is dominated by a single HNL, which we denote N, while others either contribute subdominantly or are outside the mass range of interest. Note, however, that we do not assume that the N's mixing is dominated by mixing with only one active neutrino. We will consider the N field to have a mass in the MeV to GeV range, and so a wealth of existing observables, for example, the width of the Z boson, constrain their mixings with the active neutrinos to be small. Barring rather special scenarios, for example, where the HNLs are almost mass-degenerate and can oscillate into each other on relevant length scales [91, 92], the signal rates for HNLs in models with two or more HNLs with significant mixings can therefore be determined simply by adding together the signal rates for each HNL considered separately.

The production and decay rates for the HNL are completely determined by the mass m_N and the couplings U_α. As noted in Sec. 3.1, in this study, we consider benchmark models with fixed coupling ratios $|U_e|^2 : |U_\mu|^2 : |U_\tau|^2$ defined in Eq. (3.1); the overall size of the coupling is parametrized

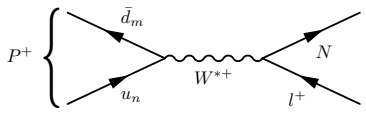
HNL Production Processes					
$P \to lN$	Fig. 3.1 (a)	$\pi^+ \to l^+N$ $B^+ \to l^+N$	$K^+ \to l^+N$ $B_c^+ \to l^+N$	$D^+ \to l^+N$	$D_s^+ \to l^+N$
$P \to P'lN$	Fig. 3.1 (b)	$K^+ \to \pi^0 l^+N$ $\bar{D}^0 \to \pi^+ l^- N$ $D^+ \to \bar{K}^0 l^+N$ $B^+ \to \pi^0 l^+N$ $B^0 \to \pi^- l^+N$ $B_c^+ \to D^0 l^+N$	$K_S \to \pi^+ l^- N$ $D^+ \to \pi^0 l^+N$ $D_s^+ \to \bar{K}^0 l^+N$ $B^+ \to \eta l^+N$ $B^0 \to D^- l^+N$ $B_c^+ \to \eta_c l^+N$	$K_L \to \pi^+ l^- N$ $D^+ \to \eta l^+N$ $D_s^+ \to \eta l^+N$ $B^+ \to \eta' l^+N$ $B_s^0 \to K^- l^+N$ $B_c^+ \to B^0 l^+N$	$D^0 \to K^- l^+N$ $D^+ \to \eta' l^+N$ $D_s^+ \to \eta' l^+N$ $B^+ \to \bar{D}^0 l^+N$ $B_s^0 \to D_s^- l^+N$ $B_c^+ \to B_s^0 l^+N$
$P \to Vl N$	Fig. 3.1 (b)	$D^0 \to \rho^- l^+N$ $D^+ \to \bar{K}^{*0} l^+N$ $B^+ \to \omega l^+N$ $B_s^0 \to K^{*-} l^+N$ $B_c^+ \to B^{*0} l^+N$	$D^0 \to K^{*-} l^+N$ $D_s^+ \to K^{*0} l^+N$ $B^+ \to \bar{D}^{*0} l^+N$ $B_s^0 \to D_s^{*-} l^+N$ $B_c^+ \to B_s^{*0} l^+N$	$D^+ \to \rho^0 l^+N$ $D_s^+ \to \phi l^+N$ $B^0 \to \rho^- l^+N$ $B_c^+ \to D^{*0} l^+N$	$D^+ \to \omega l^+N$ $B^+ \to \rho^0 l^+N$ $B^0 \to D^{*-} l^+N$ $B_c^+ \to J/\psi\, l^+N$
$\tau \to HN$	Fig. 3.1 (c)	$\tau^+ \to \pi^+ N$	$\tau^+ \to K^+ N$	$\tau^+ \to \rho^+ N$	$\tau^+ \to K^{*+} N$
$\tau^+ \to l^+ \nu N$	Fig. 3.1 (d,e)	$\tau^+ \to l^+ \bar{\nu}_\tau N$	$\tau^+ \to l^+ \nu_l N$		

Table 3.1: HNL production processes included in `HNLCalc`. The processes are ordered by increasing parent particle mass; P, V, and H denote pseudoscalar mesons, vector mesons, and hadrons, respectively; $l = e, \mu, \tau$; and N is the HNL. Charge-conjugate processes are also implemented, but are not explicitly listed here. Representative Feynman diagrams for these processes are shown in Fig. 3.1.

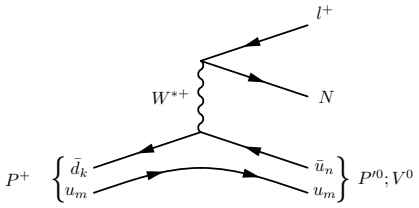
by $\epsilon^2 \equiv |U_e|^2 + |U_\mu|^2 + |U_\tau|^2$ (Eq. (1.9)). With these conventions, given a particular benchmark model, the HNL phenomenology is completely determined by (m_N, ϵ), and we present all of our results below as functions of one of these two parameters or in the (m_N, ϵ) plane.

3.3 HNL Production

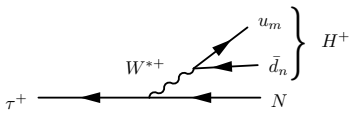
At the LHC, high-energy proton-proton collisions produce hadrons and leptons that can decay to HNLs through the CC and NC interactions of Eq. (1.7) (see Sec. 1.3). In the mass range of interest here, HNLs are primarily produced in the decays of mesons and tau leptons. The production of these parent particles at the LHC is modeled using a variety of generators, which will be discussed below in Sec. 3.5.2.



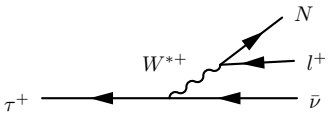
(a) 2-body decay $P \to lN$



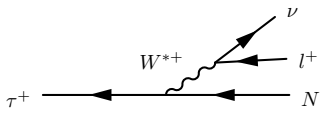
(b) 3-body decays $P \to P'lN$ and $P \to VlN$



(c) 2-body decay $\tau \to HN$



(d) 3-body decay $\tau^+ \to l^+\bar{\nu}N$



(e) 3-body decay $\tau^+ \to l^+\nu N$

Figure 3.1: Representative Feynman diagrams for the HNL production processes listed in Table 3.1. The subscripts m and n are generation indices.

The complete list of all HNL production processes included in this work is given in Table 3.1. They include both 2-body and 3-body decays, and are divided into six categories: $P \to lN$, $P \to P'lN$, $P \to VlN$, $\tau \to HN$, $\tau^+ \to l^+\bar{\nu}_\tau N$, and $\tau^+ \to l^+\nu_\ell N$, where P, V, and H represent pseudoscalar mesons, vector mesons, and hadrons, respectively. Representative quark-level Feynman diagrams for these processes are given in Fig. 3.1, and expressions for the corresponding branching fractions are given in Sec. A.1.

The HNL production processes of Table 3.1 include all of the possibly leading contributions. Note that all of the parent hadrons are pseudoscalar mesons. Vector meson decays do not typically produce HNLs with significant branching fractions, because the decays to HNLs compete with decays mediated by the strong interactions and so are highly suppressed. HNL production in baryon decays is also not included, since they are subdominant.

The branching fractions for meson and tau lepton decays into HNL's are presented in Fig. 3.2, as functions of HNL mass. For illustrative purposes, we show results for the 011 and 111 benchmark models with coupling parameter $\epsilon = 10^{-3}$. For each benchmark model there are as many as 150 HNL production processes. Rather than show branching fractions for each of these modes, in Fig. 3.2 we show the total BSM branching fraction for each parent meson, which includes final states with all possible hadrons and lepton flavors.

Because the total branching fractions include many modes, distinctive aspects of individual modes are not always apparent. However, we note a few general features:

- The branching fractions are typically larger for the longer-lived mesons, where the competing SM decay modes having smaller widths. This implies that the BSM branching fractions are typically larger for the lighter parent mesons.

- Of course, the branching fractions vanish when the decay modes become kinematically inaccessible. For example, $B(\pi^+ \to NX)$ vanishes at $m_N = m_{\pi^+} - m_\mu \simeq 34$ MeV for the 011 benchmark and at $m_N = m_{\pi^+} - m_e \simeq 139$ MeV for the 111 benchmark.

- At the same time, although for many modes the branching fraction drops as m_N approaches the kinematic threshold, this is not always the case. For parent mesons where the dominant production mode is chirality suppressed, for example, for K^+ mesons, where the dominant HNL decay mode is $K^+ \to l^+ N$, the branching fraction vanishes for $m_N = 0$, and grows as m_N increases. In these cases, the branching fraction can be large and even maximal very near threshold.

These facts, coupled with the fact that more light mesons than heavy mesons are produced at the LHC, imply that typically the dominant production mechanism for HNLs is decays of the lightest parent mesons for which the decay is kinematically allowed.

In Fig. 3.2, we have fixed $\epsilon = 10^{-3}$, a value that is very roughly at the limit of current constraints. We see that branching fractions of 10^{-7} to 10^{-6} are allowed. Given the enormous flux of far-forward mesons at the LHC, this implies that the flux of highly collimated far-forward HNLs can be significant. Of course, to be detected, the HNLs must decay in the detector to a visible final state.

3.4 HNL Decays

Once produced, HNLs decay to SM final states through the CC and NC interactions of Eq. (1.7). The complete list of HNL decay modes included in this work is given in Table 3.2. We consider

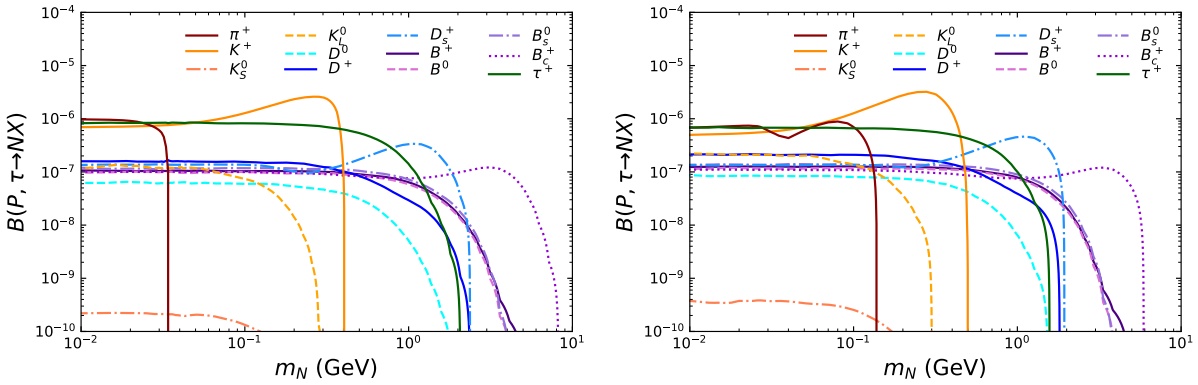
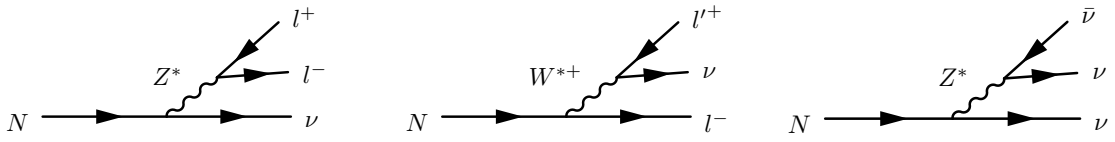


Figure 3.2: Branching fractions $B(P, \tau \to NX)$ as functions of the HNL mass m_N for the 011 (left) and 111 (right) benchmarks and $\epsilon = 10^{-3}$. P represents pseudoscalar mesons, and NX represents any final state containing the HNL N.



(a) NC-mediated $N \to \nu l^+ l^-$ (b) CC-mediated $N \to \nu l'^+ l^-$ (c) NC-mediated $N \to \nu\bar{\nu}\nu$

(d) NC-mediated $N \to \nu H^0$ (e) CC-mediated $N \to l^- H^+$

Figure 3.3: Representative Feynman diagrams for HNL decays. The subscripts m, n are generation indices.

purely leptonic 3-body decays, semi-leptonic 2-body decays, and semi-leptonic 3-body decays. Representative Feynman diagrams for each category are shown in Fig. 3.3. The HNL decay branching fractions are computed using formulae derived in Refs. [21, 84], and the expressions for the branching fractions are given in Sec. A.2.

To properly simulate the response of the detector to an HNL decay, it is important to know and accurately represent its final states. For example, it can make a big difference experimentally whether an HNL decays to states with charged tracks, e.g., $N \to \nu\rho \to \nu\pi^+\pi^-$, or to states with only final state photons, e.g., $N \to \nu\pi^0 \to \nu\gamma\gamma$. For this reason, it is generally insufficient to specify HNL decays into quarks, as it would rely on hadronization tools to obtain hadronic final

HNL Decay Modes							
$\nu\, l^+ l^-$	Fig. 3.3 (a)	$\nu_l e^+ e^-$	$\nu_l \mu^+ \mu^-$	$\nu_l \tau^+ \tau^-$			
$l^\pm\, \nu_{l'} l'^\mp$	Fig. 3.3 (b)	$l^\pm \nu_e e^\mp$	$l^\pm \nu_\mu \mu^\mp$	$l^\pm \nu_\tau \tau^\mp$			
$\nu_l\, \overline{\nu} \nu$	Fig. 3.3 (c)	$\nu_l \bar{\nu}_e \nu_e$	$\nu_l \bar{\nu}_\mu \nu_\mu$	$\nu_l \bar{\nu}_\tau \nu_\tau$			
$\nu_l\, H^0$	Fig. 3.3 (d)	$\nu_l \pi^0$	$\nu_l \eta$	$\nu_l \eta'$	$\nu_l \rho^0$	$\nu_l \omega$	$\nu_l \phi$
$l^\pm\, H^\mp$	Fig. 3.3 (e)	$l^\pm \pi^\mp$	$l^\pm K^\mp$	$l^\pm D^\mp$	$l^\pm D_s^\mp$	$l^\pm \rho^\mp$	$l^\pm K^{*\mp}$
$\nu_l\, q\bar{q}$	Fig. 3.3 (d)	$\nu_l u\bar{u}$	$\nu_l d\bar{d}$	$\nu_l s\bar{s}$	$\nu_l c\bar{c}$	$\nu_l b\bar{b}$	
$l^\pm\, u\overline{d}'$	Fig. 3.3 (e)	$l^- u\bar{d}$ $l^- c\bar{d}$	$l^- u\bar{s}$ $l^- c\bar{s}$	$l^- u\bar{b}$ $l^- c\bar{b}$	$l^+ \bar{u} d$ $l^+ \bar{c} d$	$l^+ \bar{u} s$ $l^+ \bar{c} s$	$l^+ \bar{u} b$ $l^+ \bar{c} b$

Table 3.2: HNL decay modes included in `HNLCalc`, where H denotes hadrons. Representative Feynman diagrams for these modes are shown in Fig. 3.3. Quark level decays, $\nu q\bar{q}$ and $lu\bar{d}$, are used when computing the total hadronic width when $m_N > 1.0$ GeV.

states. These tools are known not to work well when the invariant mass of the hadronic final state is close to or below the QCD confinement scale, because, for example, this treatment fails to model hadronic resonances and kinematic thresholds. More generally, in this regime, the factorization theorem loses validity, meaning that one cannot factorize the HNL decay into quarks and use quark hadronization into hadrons anymore. On the other hand, for $m_N \gtrsim 1.0$ GeV, decays into single mesons are insufficient for computing the total HNL decay width, since in this region multi-meson decays become important.

To best model hadronic decays at all scales, while also accurately computing the HNL lifetime, we follow the approach taken in Refs. [21, 84]. For $m_N < 1.0$ GeV, decays into single mesons, $H = \pi^{\pm,0}, K^{\pm}, \rho^{\pm,0}, \omega, K^{*\pm}, \eta, \eta'$, are calculated using their respective decay constants, f_H, and these are used to obtain the total hadronic decay width. For $m_N \geq 1.0$ GeV, the total hadronic width is instead calculated by summing up HNL decay widths into the quark-level final states $lq\bar{q}'$ and $\nu q\bar{q}$. A QCD loop correction to account for hadronization is applied. This estimated correction is obtained from the known corrections up to $\mathcal{O}(\alpha_s^3)$ in hadronic τ decay; see Sec. A.2 for its explicit form.

Included in our single meson decays are decays into ρ mesons. It has been shown in Ref. [21] that the dominant source of two pion final states are decays into ρ mesons followed by $\rho \to \pi\pi$. We therefore take the difference between the quark-level decay width and the single-meson total width as an estimate of the decay width to final states with ≥ 3 hadrons. Additionally, to avoid overestimating the contribution of νss to the hadronic width below the $2m_K$ kinematic threshold, a phase space suppression factor $\sqrt{1 - 4m_K^2/m_N^2}$ is applied to this decay. This same approach is also applied to the decays τud and τus, with kinematic thresholds at $m_\tau + 2m_\pi$ and $m_\tau + m_\pi + m_K$, respectively.

In Figs. 3.4 and 3.5, we plot the branching fractions for all relevant modes in the 011 and 111 benchmarks. Additionally, the HNL lifetimes are given in Fig. 3.6. For $m_N < m_\pi$, HNL decays are dominated by NC decays into the 3-body final states $\nu\bar{\nu}\nu$ and $\nu e^+ e^-$. The invisible decay mode dominates, and the visible branching fraction is only approximately 10% and 20% in the 011 and 111 benchmark models, respectively. However, for $m_N > m_\pi$, the 2-body decays $N \to l^{\pm}\pi^{\mp}$ and

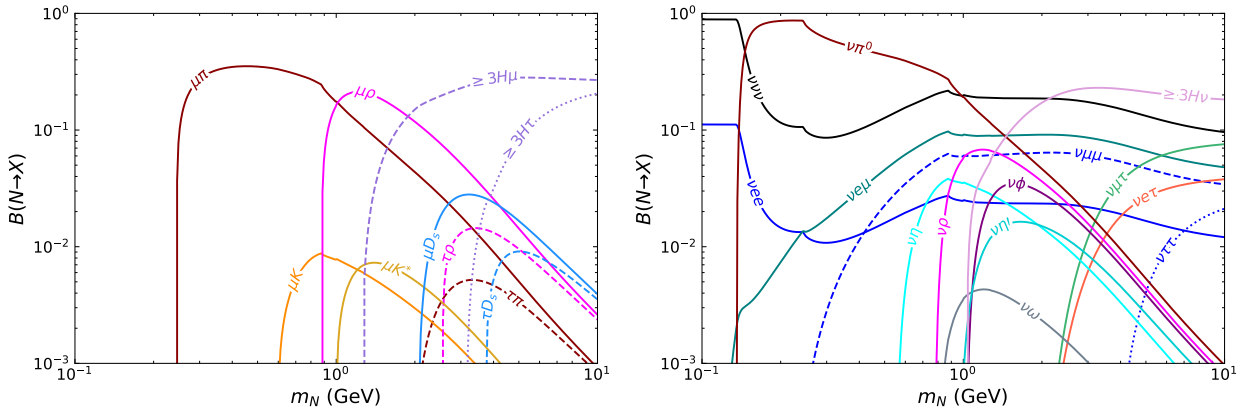


Figure 3.4: HNL branching fractions in the 011 benchmark model for the dominant decays via CC (left) and NC (right) interactions as functions of m_N. For $m_N \lesssim 1.0$ GeV, the decays are primarily to pions, while for $m_N \gtrsim 1.0$ GeV, the decays are dominated by decays to 3 or more hadrons. Decays to D mesons are subdominant, with $B(N \to lD) < 0.1\%$.

$N \to \nu\pi^0$ become kinematically accessible, resulting in a sharp drop in the lifetime, as seen in Fig. 3.6. For masses $m_N > m_\pi$, the hadronic decay modes become dominant, and the invisible decay branching fraction is below 20% in both scenarios.

It is important to note that in this work we have neglected the effects of spin correlations between production and decay. These do not change the HNL lifetimes and event rates, but they can impact the kinematic distributions of the final state particles. In the case of Majorana HNLs, where lepton-number-conserving and lepton-number-violating processes, for example, $B_c \to \mu^+\mu^+\tau^-\nu$ and $B_c \to \mu^+\mu^+\tau^-\bar{\nu}$, are both possible, interesting effects in kinematic distributions have the potential to differentiate between the two final states, which are otherwise experimentally identical; see, for example, Ref. [93].

3.5 HNLs at FASER and FASER2

In this section, we describe the modeling of meson and τ production at the LHC and the resulting flux of HNLs at FASER [64] and FASER2 [3] (see also Sec. 2.2.2) during Run 3 and the HL-LHC era. Given the modeling of production and decay of HNLs in Secs. 3.3 and 3.4, respectively, with the use of `HNLCalc`, the `FORESEE` simulation framework can determine the sensitivity for HNL searches with forward detectors at the LHC.

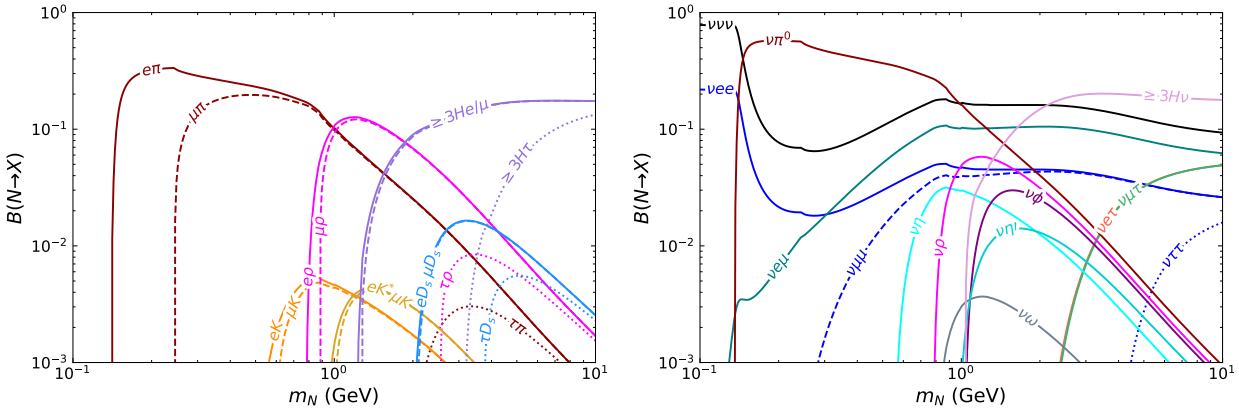


Figure 3.5: As in Fig. 3.4, but for the 111 benchmark model.

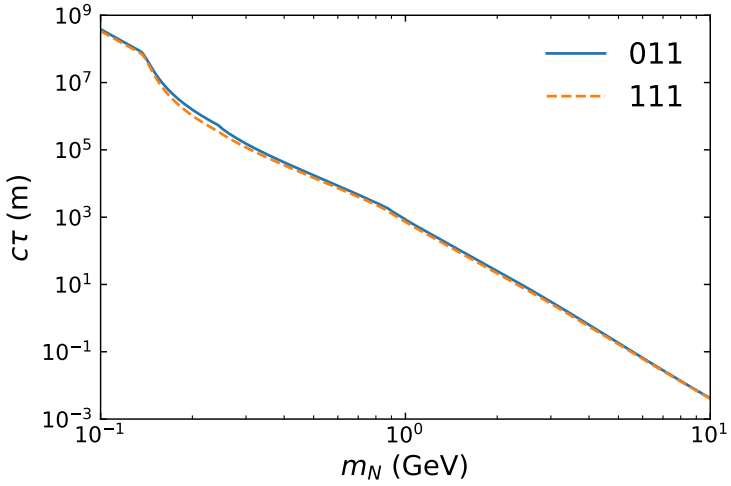


Figure 3.6: Lifetimes for the 011 and 111 benchmark models with $\epsilon = 10^{-3}$. A visible kink in the lifetime can be observed at $m_N \approx m_\pi$, where the decay $N \to \nu\pi^0$ becomes kinematically accessible.

	L	Δ	Geometry	$\mathcal{L}$
FASER (Run 3)	480 m	1.5 m	Cyl. $R = 10$ cm	250 fb^{-1}
FASER (HL-LHC)	480 m	1.5 m	Cyl. $R = 10$ cm	3 ab^{-1}
FASER2 (HL-LHC)	650 m	10 m	Rect. 3 m $\times$ 1 m	3 ab^{-1}

Table 3.3: Parameters for the three experimental configurations considered: FASER (Run 3), FASER (HL-LHC), and FASER2 (HL-LHC). L is the distance from the interaction point to the front of the detector, Δ is the length of the detector along the beam axis, Geometry specifies the cross sectional area (cylindrical for FASER, rectangular for FASER2), and $\mathcal{L}$ is the integrated luminosity. Both FASER and FASER2 are assumed to be centered on the beam collision axis.

3.5.1 Collider and Detector Setup

FASER is located in a tunnel $L = 480$ m downstream from the ATLAS interaction point (IP). The FASER decay volume is a cylinder with a radius of $R = 10$ cm and a length of $\Delta = 1.5$ m along the beam collision axis. The decay volume extends to angles $\theta \simeq 0.2$ mrad from the beamline and pseudorapidities $\eta \gtrsim 9.2$. We will consider the reach of FASER during Run 3 of the LHC with an integrated luminosity of 250 fb^{-1}. We will also consider the scenario in which the FASER detector runs throughout the HL-LHC era with an integrated luminosity of 3 ab^{-1}.

In addition, we consider FASER2, one of the experiments proposed for the future FPF [3, 63]. FASER2 will be positioned roughly $L = 650$ m away from the ATLAS IP. Its decay volume has length $\Delta = 10$ m and a rectangular cross-section with dimensions 3 m $\times$ 1 m [94]. The FASER2 decay volume extends to angles $\theta \simeq 2.4$ mrad from the beamline and pseudorapidities $\eta \gtrsim 6.7$.

During Run 3, the LHC has operated with a pp COM energy of $\sqrt{s} = 13.6$ TeV. For the HL-LHC era, the COM energy is expected to be increased to 14 TeV. We have found negligible differences in sensitivities from this change in COM energy, and for simplicity, we assume $\sqrt{s} = 14$ TeV for all results derived here. The detector configurations are summarized in Table 3.3.

3.5.2 HNL Production Rates at the LHC

HNL production rates at the LHC are obtained in `FORESEE` through the MC sampling of the decays of parent particles [70]; see also Sec. 2.3. The spectra for parent pions and kaons are obtained using the dedicated hadronic interaction model `EPOS LHC` [71]. For charm and bottom hadrons, we use the spectra obtained from `POWHEG` [75] matched with `Pythia` [76], as presented in Ref. [62]. Tau leptons are produced primarily in charm and bottom hadron decay, so we follow the same prescription there. The primary source of uncertainty on the HNL flux originates from the modeling of hadron production. To estimate this uncertainty, we follow the prescription of Ref. [67] and consider the range of predictions from `EPOS-LHC`, `QGSJET 2.04` [73], and `Sibyll 2.3d` [72] to model the uncertainty for light hadron production, and we take account of scale variations [62] to model the flux uncertainty for heavy hadron production.

The production rates for the 011 and 111 benchmarks at the LHC are shown in Figs. 3.7 and 3.8. These figures show the total flux of HNLs produced in the direction of FASER and FASER2 respectively. For FASER, Fig. 3.7 shows the rate for production with an angle $\theta \leq 0.1\,\text{m}/480\,\text{m} = 0.2\,\text{mrad}$ relative to the beam collision axis, and for FASER2, Fig. 3.8 shows the rate for the production of parent particles that pass through FASER2's 3 m$\times$1 m transverse area. As anticipated in Sec. 3.3, the dominant production rate is typically from the lightest meson for which the decay is kinematically accessible. For the 111 scenario, the decay $\pi \to eN$ is possible, and so this is the dominant production process for very light N with masses $m_N \sim 100$ MeV. For the 011 scenario, this decay is not allowed, and pion decays are never dominant for $m_N \gtrsim 100$ MeV. For larger m_N, the results for the 011 and 111 scenarios are similar, with the production rates cutting off as the kaon, tau lepton, D meson, and B meson kinematic thresholds are passed.

3.5.3 Signal and Background Considerations

Depending on their mass, HNLs can decay into a variety of final states shown in Figs. 3.4 and 3.5. The dominant decay channels either consist of two charged particles, such as the decays $N \to \nu l^+ l^-$ and $N \to l^+ \pi^-$, or photons, especially from $N \to \nu\pi^0 \to \nu\gamma\gamma$. The corresponding experimental signatures in the FASER detector would consist of highly energetic charged particles or photons that emerge from the decay volume. The charged particle signal would leave high-momentum tracks

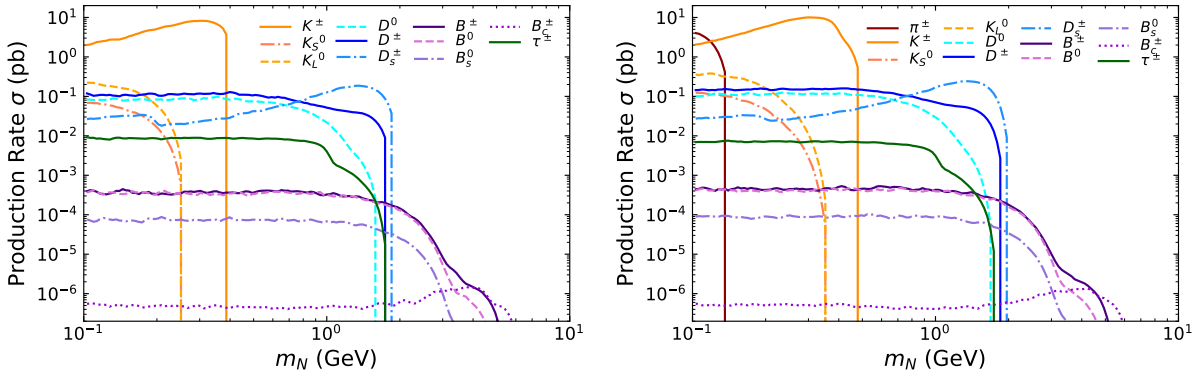


Figure 3.7: Production rate of HNLs in the direction of FASER with $\theta < 0.2$ mrad. The rates are grouped by parent particle and are obtained by MC integration in `FORESEE` for the 011 (left) and 111 (right) benchmarks, with a coupling of $\epsilon = 10^{-3}$.

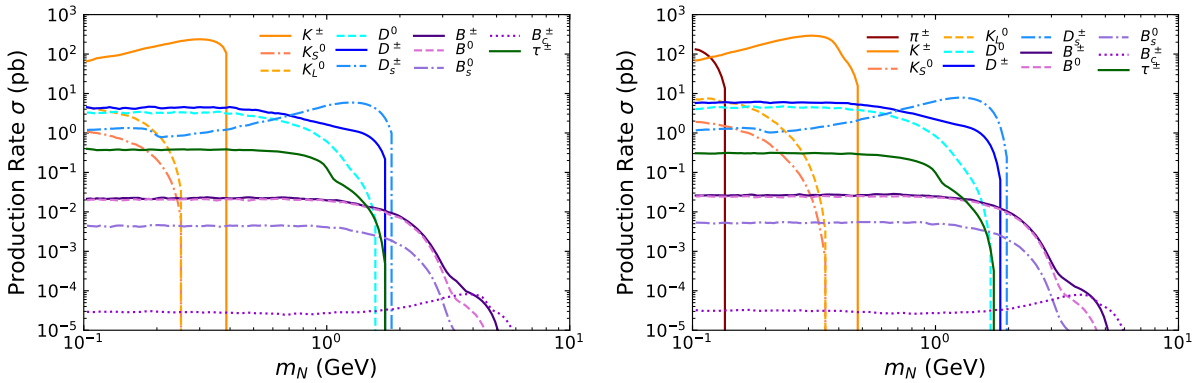


Figure 3.8: Production rate of HNLs in the direction of FASER2, that is, that pass through FASER2's transverse area of 3 m $\times$ 1 m. The rates are grouped by parent particle and are obtained by MC integration in `FORESEE` for the 011 (left) and 111 (right) benchmarks, with a coupling of $\epsilon = 10^{-3}$.

in the FASER spectrometer, that are consistent with a single vertex inside the decay volume and whose combined momentum points back to the IP. If the final state contains electrons, they would additionally leave a sizbale energy deposit in the calorimeter. The multi-photon signal would leave a characteristic signal in the preshower and deposit a large amount of energy in the calorimeter. In both cases, the signal would not trigger the front veto.

The potential backgrounds for LLP searches at FASER are induced by either high-energy muons or neutrinos coming from the direction of the IP. However, muons are efficiently detected by FASER's front veto, while neutrino interactions are relatively rare and typically have different kinematics. The FASER Collaboration presented their first analysis on dark photons, providing a two-track signal, in which they accounted for a variety of possible background sources associated with veto inefficiencies, neutral hadrons, muons missing the veto, neutrinos, and non-collision background [95]. These backgrounds were determined to be either very small or negligible, and no events with two reconstructed tracks passing the veto requirement have been observed. We therefore assume that backgrounds for multi-track signatures can be considered negligible.

More recently, the FASER Collaboration also presented results for a search for axion-like particles, considering a multi-photon signal [96]. A potentially sizable background from neutrinos interacting at the end of the detector was identified, particularly at energies below 1 TeV. To address this issue, a high-precision preshower upgrade is planned to be installed at the end of 2024, which will allow the experiment to identify multi-photon signatures and differentiate them from neutrino backgrounds [97]. We therefore assume that backgrounds for multi-photon signatures are also negligible.

The design for FASER2 is currently under development [3]. It will have a similar conceptual architecture as the currently-operating FASER detector, and background rejection is one of the key considerations for the detector design. Although the design has not yet been finalized, it is envisioned that the same searches will be background free.

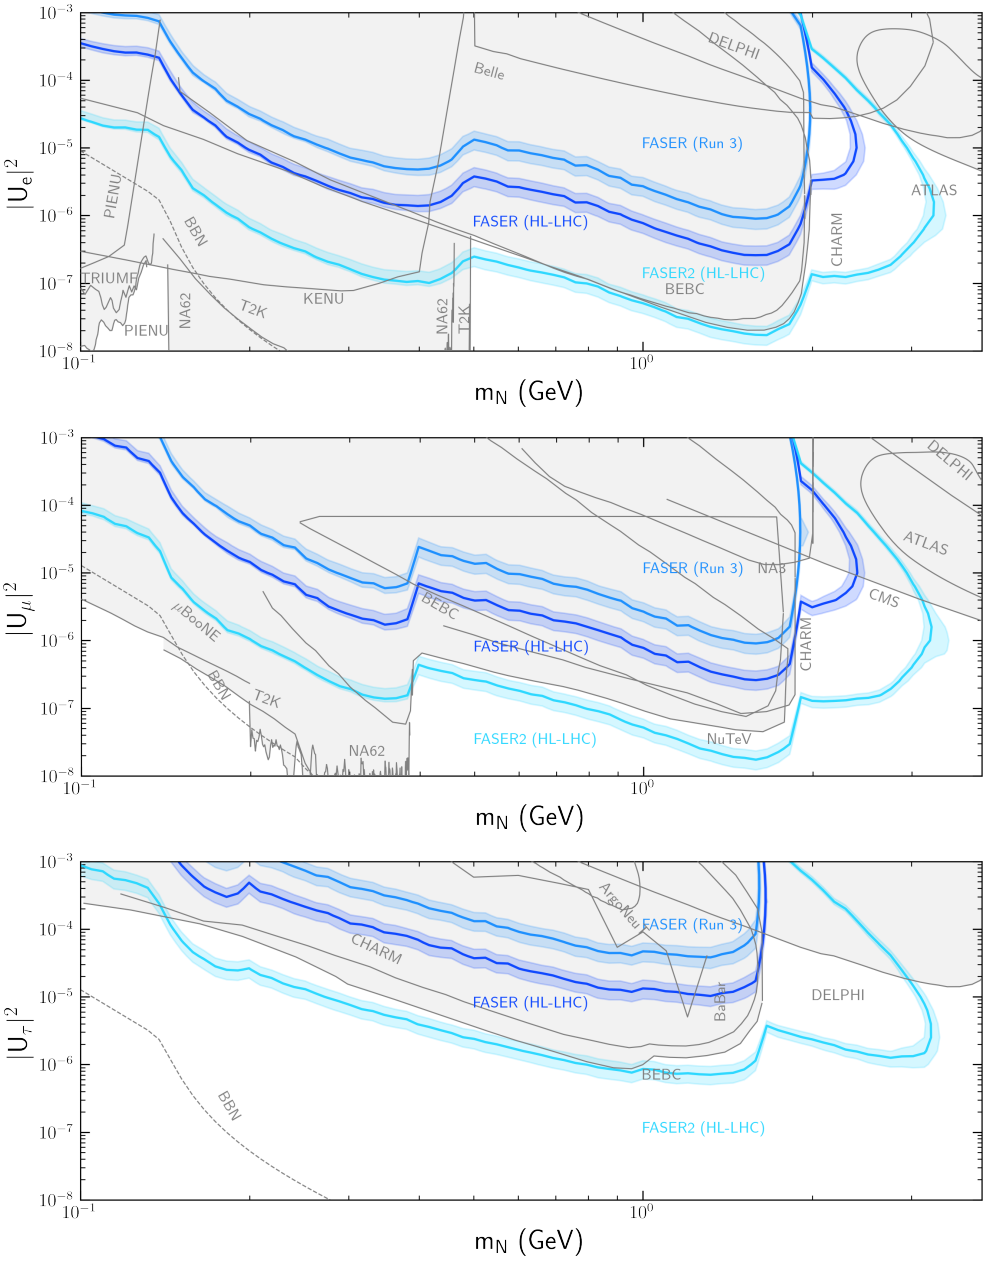


Figure 3.9: $N = 3$ signal event contours for the single-coupling 100 (top), 010 (middle), and 001 (bottom) benchmark models. Each blue contour corresponds to a detector/collider configuration given in Table 3.3. Current bounds from other particle experiments exclude the gray shaded regions, and BBN constraints exclude the region below the dashed contour. All of the production modes in Table 3.1 are included, and all of the decay modes in Table 3.2 are included, except for the invisible $\nu\bar{\nu}\nu$ modes. For the FASER and FASER2 contours, a momentum cut of $p_N > 100$ GeV has been applied.

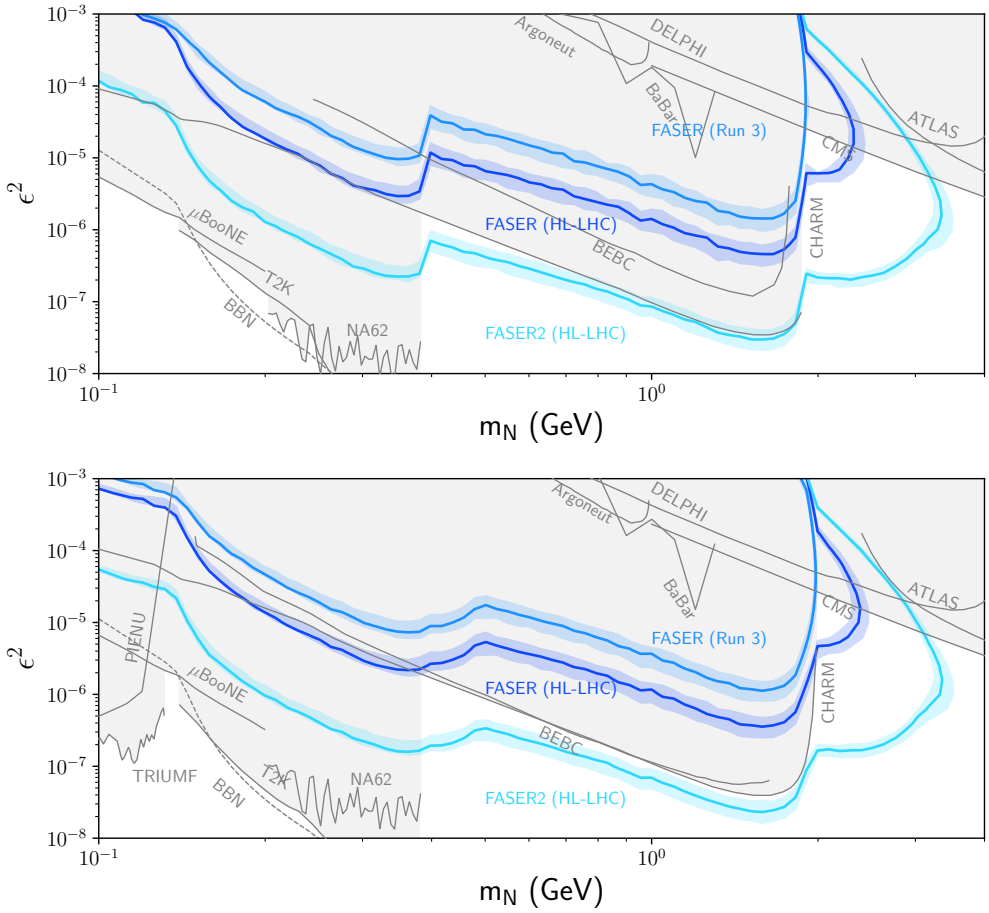


Figure 3.10: As in Fig. 3.9, but for the mixed-coupling 011 (top) and 111 (bottom) benchmark models. In contrast to the single-coupling models, constraints from other experiments are not available in the literature. The current bounds shown in the gray region are roughly estimated based on a method described in the text.

3.6 Reach

We now present the $N = 3$ signal event contours in the (m_N, ϵ) plane for the various benchmark scenarios. In the absence of background, these may be considered to be discovery contours. The results for the 100, 010, and 011 benchmark models are given in Fig. 3.9, and the results for the 011 and 111 benchmark models are given in Fig. 3.10. The signal event rates include all of the production modes listed in Table 3.1 and all of the decay modes shown in Table 3.2, except for the invisible decay modes $\nu\bar{\nu}\nu$. Additionally, a momentum cutoff of $p_N > 100$ GeV is applied. The event rates are determined by setting `nsample` $= 25$ in the `FORESEE` MC simulation [70], and the (m_N, ϵ) parameter space is scanned using 100 equally log-spaced masses m_N from 100 MeV to 4

GeV, and 50 equally log-spaced couplings ϵ from 10^{-5} to 1.

Results are presented for the three FASER/FASER2 configurations shown in Table 3.3. The shaded band corresponds to the flux uncertainty, following the prescription in Sec. 3.5.2. Sensitivity contours for the 100, 010, and 001 models have been investigated previously [60, 88]. Our obtained sensitivities are consistent with these results over most of the mass range. They are slightly improved for $m_N \lesssim m_K$, due to an issue that was found in the kaon decay modes of these previous results. Sensitivity reaches for FASER and FASER2 for the 011 and 111 models have not been determined before.

Various threshold effects can be observed in the sensitivity contours, depending on which production channels dominate HNL production. In the 100 benchmark, the production of HNLs is predominantly governed by pions in the region 100 MeV $\leq m_N \leq m_\pi$, transitions to being dominated by kaons in the range $m_\pi \leq m_N \leq m_K$, shifts to being influenced by D and τ particles within the interval $m_K \leq m_N \leq m_D$, and eventually becomes dominated by B mesons for $m_N \geq m_D$. Additionally, one can infer from the sensitivity contours that the 011 benchmark sensitivity is primarily driven by the U_μ coupling and the 111 benchmark sensitivity is primarily driven by the U_e coupling.

It is, of course, interesting to compare the sensitivity of FASER and FASER2 to current constraints. For the 100, 010, and 001 models, the exclusion bounds have been determined for many past experiments, and are available using the Python package `HNL-Limits` [2]. For the 100 and 001 benchmarks, the BEBC and CHARM bounds in `HNL-Limits` have been refined in Ref. [98], and we have adopted those limits. For the 010 benchmark, we use the lower bound provided by the CHARM Collaboration [51]. To establish the upper limit for this benchmark, we use the bounds derived for the 100 benchmark model from the CHARM recasting detailed in Ref. [98] as an approximation. We further apply a cutoff at $m_{D_s} - m_\mu$ in the 010 and 011 CHARM bounds to account for the allowed phase space. In addition, for small masses, there are constraints from Big Bang nucleosynthesis (BBN), which constrains the lifetime of the HNL to be less than 0.1 s [99].

For the 011 and 111 models, the sensitivities of previous experiments are not often found in the

literature. To compare FASER and FASER2 to existing constraints, an analysis similar to the one we have done for FASER and FASER2 must be done, requiring a dedicated analysis for each model and experiment. This is beyond the scope of this work. However, we may roughly estimate the sensitivities of previous experiments in the 011 and 111 models by extrapolating their limits in the 100, 010, and 001 models, as we now describe.

In peak search experiments, one looks for evidence of HNL production by searching for peaks in the energy spectrum of co-produced particles. One such example are the bounds placed by the PIENU Collaboration [48], where they performed a search for peaks in the positron energy spectrum in the process $\pi^+ \to N e^+$. Bounds in these types of experiments place limits on U_α that are independent of the HNL lifetime and decay products and thus will not be effected by the introduction of mixed couplings. The bounds on U_α can therefore be re-scaled directly to obtain bounds for ϵ.

In prompt decay or decay-in-flight searches, such as those at ATLAS [57] or CHARM [100], one searches for evidence of HNLs by looking for HNL decay products within the detector. In these experiments, the introduction of mixed couplings can effect the kinematic distributions of final states, as well as the HNL lifetime, causing a change in the overall expected event rate. As a very rough approximation, we can neglect the effect on the kinematics by making the simplifying assumptions that the detector acceptance and efficiency are uniform across all benchmarks. In the long lifetime limit, then, the expected event rate in these experiments scales as

$$N_\alpha \propto \frac{|U_\alpha|^2}{\tau} \ , \tag{3.2}$$

where N_α is the contribution to the total event rate occurring through ν_α mixing. This allows one to relate the expected event rates between the mixed coupling and single coupling scenario by

$$\frac{N_{\alpha,m}}{N_{\alpha,s}} = \frac{|U_{\alpha,m}|^2}{|U_{\alpha,s}|^2} \frac{\tau_s}{\tau_m} \ , \tag{3.3}$$

where the subscript m (s) denotes the mixed (single) coupling scenario. Approximate exclusion bounds derived in this way for the 011 and 111 benchmarks are plotted in Fig. 3.10.

For the sensitivity reaches in the single coupling scenarios shown in Fig. 3.9, the results may be

summarized as follows. For FASER (Run 3), the reach beyond current bounds is rather modest and is limited to the 100 model. For FASER (HL-LHC), the reach is extended, with discovery prospects in the 100 and 010 scenarios for $m_N \sim 2$ GeV. Finally, for FASER2 (HL-LHC), there are significant regions of parameter space in all 3 models with sensitivity beyond current bounds. This includes regions of parameter space with $10^{-7} \lesssim \epsilon^2 \lesssim 10^{-5}$ and 2 GeV $\lesssim m_N \lesssim 4$ GeV, where the HNLs are produced in B meson decays, and also regions of low ϵ at lower m_N. This is especially notable in the 001 scenario, where the new parameter space probed extends all the way down to HNL masses of approximately 150 MeV.

For the sensitivity contours in mixed coupling scenarios shown in Fig. 3.10, we find that the comparison of the reach of FASER and FASER2 relative to current bounds for the 011 and 111 scenarios approximately mirrors those for the 010 and 100 models, respectively. In particular, for FASER2 (HL-LHC) there is again new parameter space probed with $10^{-7} \lesssim \epsilon^2 \lesssim 10^{-5}$ and 2 GeV $\lesssim m_N \lesssim 4$ GeV, and also improved sensitivity over current bounds at low ϵ and masses 400 MeV $\lesssim m_N \lesssim 2$ GeV. We reiterate, however, that in this comparison, the FASER and FASER2 bounds derived here are compared to current bounds that have been derived in the very rough way we have outlined above. Detailed comparisons require dedicated analyses of the reach of other experiments in these less minimal, mixed-coupling scenarios.

Finally, let us comment on the flux uncertainties, which are shown as shaded bands in Figs. 3.9 and 3.10. Despite substantial flux uncertainties, which can be as large as a factor of two for charm production, their overall impact on the sensitivity is relatively small. This is due to a strong coupling dependence of the event rate at both small and large couplings. The flux uncertainties mainly affect the reach at the higher mass edge of the sensitivity contour. We note that further measurement, especially the measurement of collider neutrino fluxes, will help to further decrease those flux uncertainties.

3.7 Conclusions

HNLs are well motivated in extensions of the SM designed to address the outstanding puzzles of neutrino masses, dark matter, and baryogengesis. In this chapter, we have comprehensively

studied the possibility of probing HNLs with mass in the 100 MeV to 10 GeV range and completely arbitrary mixings with the three active SM neutrinos. We have produced a comprehensive package, `HNLCalc`, that computes the properties of the HNL model, including hundreds of production and decay modes. Using `HNLCalc`, we created an HNL module in the `FORESEE` simulation package to evaluate the discovery potential of HNLs in various experimental setups.

In particular, we have estimated the sensitivity to HNLs for FASER in Run 3, FASER at the HL-LHC, and FASER2 at the HL-LHC. As an illustration of the flexibility of `HNLCalc`, we have considered five benchmark models. For the well-studied 100, 010, and 001 benchmarks, we find that FASER in Run 3 has rather limited discovery prospects. However, FASER at the HL-LHC can probe new parameter space with $m_N \sim 2$ GeV, and FASER2 at the HL-LHC can probe new parameter space for a wide range of HNL masses from 150 MeV to 4 GeV.

Additionally, we have considered two new mixed-coupling benchmark models, 011 and 111, and determined the sensitivity of FASER and FASER2. These new benchmarks are not well studied and there are no comprehensive studies of the sensitivity of other current and proposed experiments in these models. We have roughly estimated the sensitivity of past experiments, and we find that FASER2 at the HL-LHC can likely probe new parameter space for a wide range of HNL masses from 400 MeV to 4 GeV.

Overall, this study contributes significantly to the understanding of HNLs and their potential implications in particle physics. We have produced `HNLCalc`, a flexible, fast, comprehensive, and publicly-available tool for computing HNL decay and production rates with arbitrary couplings. In this study, we have used this tool to extend the `FORESEE` simulation package to incorporate HNLs. Models with general HNL couplings are more complicated than models with a single coupling, but they are also more well-motivated, and studies of the discovery prospects of other experiments in these models are encouraged.

Chapter 4

Characterizing Heavy Neutral Leptons: Measuring Parameters, Discriminating Majorana versus Dirac, and Using FASER2 as a Trigger for ATLAS

The following chapter is largely based on Ref. [7] in collaboration with Jonathan L. Feng, Alec Hewitt and Daniel Whiteson.

4.1 Introduction

Once an HNL is discovered, attention will immediately turn to measuring its mass and couplings and to determining whether it is a Majorana or Dirac fermion. These measurements will be essential, not only to evaluate whether HNLs have the right properties to solve the open puzzles of the SM, but also because their properties may point toward other particles, providing a target for future colliders, facilities, and experiments.

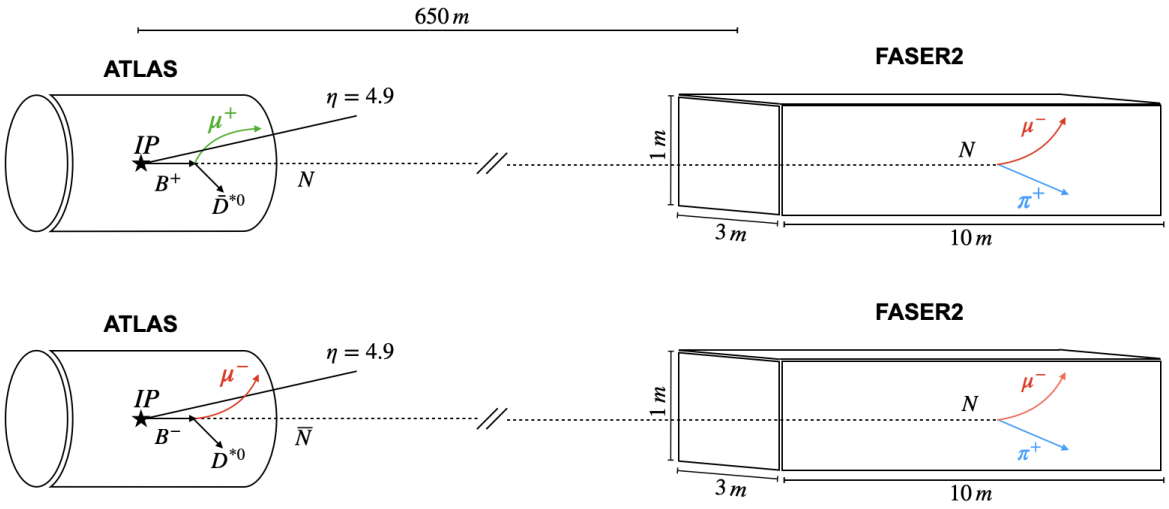


Figure 4.1: Schematic diagrams showing example HNL events at ATLAS and FASER2 [5]. A charged B meson decays through $B^{\mp} \to \overset{(-)}{D}{}^{*0}\mu^{\mp}N$ at the ATLAS IP. The HNL then travels approximately 650 m and decays through $N \to \mu^{\mp}\pi^{\pm}$ in FASER2. For Majorana HNLs, both the top (lepton-number conserving) and bottom (lepton-number-violating) processes are allowed. For Dirac HNLs, only the top process is allowed.

In this chapter, we consider the possibility that HNLs are discovered in forward detectors at the LHC. Specifically, we consider a successor to the currently-running FASER experiment [4, 64], FASER2 [5], which will be located 650 m downstream of the ATLAS IP, along the beam collision axis or LOS, in the proposed FPF [3, 101, 102] (see also Sec. 2.2.2). At the LHC, HNLs with weak couplings and masses in the MeV to few GeV range can be produced in very large numbers in a highly collimated beam along the LOS, travel hundreds of meters, and then decay to visible SM particles. FASER2 is therefore ideally positioned to discover HNLs in regions of parameter space that are currently unconstrained by particle experiments and cosmology, and the possibility of detecting HNLs in far-forward experiments at the LHC has been the subject of numerous studies [6, 60, 88, 89, 103–108].

In this paper, we assume that an HNL has been discovered, and we explore the potential of FASER2 to constrain its mass and couplings and determine its spinor nature, that is, if it is a Majorana or Dirac fermion. We begin by considering a model where a Majorana HNL has mass 1.84 GeV and couples to muons. In this model, the HNL is dominantly produced in decays $D_s^{\pm} \to \mu^{\pm}N$ at the ATLAS IP and then travels a macroscopic distance before decaying; see Fig. 4.1. This

model is currently allowed by all constraints, but it predicts remarkable event rates for FASER2: at the HL-LHC with its expected integrated luminosity of 3 ab^{-1}, $\sim 10^9$ HNLs are produced by ATLAS, $\sim 10^5$ of these decay in FASER2, and approximately 8600 of these are $N \to \mu\pi$ decays. By measuring the energy and momentum of the HNL decay products at FASER2 using only this fully-visible decay mode, we show that the HNL mass and coupling can be constrained to fractional uncertainties of 0.1% and 3% at 95% confidence level (CL), respectively, and the possibility that the HNL is a Dirac fermion can be disfavored at 99.8% CL (2.9σ), provided current uncertainties in forward hadron fluxes can be significantly reduced by future far-forward neutrino measurements, as is expected [109] and will be discussed below.

Of course, HNLs may be discovered with significantly fewer events. We therefore then consider a Majorana HNL with mass 2 GeV, for which only roughly 80 $N \to \mu\pi$ decays in FASER2 are expected at the HL-LHC. For this model, standalone FASER2 measurements are able to constrain the HNL mass and coupling to fractional uncertainties of 1% and 25% at 95% CL, respectively, and the possibility that the HNL is a Dirac fermion can be disfavored at 98.2% CL (2.1σ).

In this second model, however, we also note that a sizable fraction of FASER2 HNL events are accompanied by muons that can be observed in ATLAS. We, therefore, also consider the possibility of using FASER2 as a trigger for ATLAS and combining the resulting data from FASER2 and ATLAS. An HNL travels from the ATLAS IP to FASER2 in 2.2 μs, and a signal may be sent back to ATLAS through an optical cable in 3.2 μs. Including the time it takes to identify a signal at FASER2 as being interesting, the resulting roundtrip time is (just) within the ATLAS Level-0 (L0) trigger latency window of 6 μs.

As shown in Fig. 4.1, if the muon produced in association with the HNL can be recorded in ATLAS, its properties may be used to sharpen FASER2's measurements. In particular, by measuring the muon's charge, one can determine if the complete HNL event is lepton-number-conserving (LNC) (Fig. 4.1, top) or lepton-number-violating (LNV) (Fig. 4.1, bottom), and distinguish between Majorana and Dirac HNLs. In the absence of background, events may be classified as LNC or LNV on an event-by-event basis, and a single unambiguous LNV event may provide conclusive evidence that the HNL is a Majorana fermion. Unfortunately, the muons produced in association with HNLs

at ATLAS are very far-forward, and there are daunting backgrounds in this region, especially given pileup at the HL-LHC. Despite this, we find that, assuming detector upgrades to ensure muon detection up to pseudorapidities of $|\eta| < 4.9$, suitable cuts allow FASER2 + ATLAS to differentiate between Majorana and Dirac HNLs and reject the Dirac fermion hypothesis at 99.7% CL.

Several strategies have been proposed in the literature to distinguish Majorana from Dirac HNLs by exploiting differences between LNC and LNV processes [91, 110–112]. Previous work has also explored the use of the far-forward detector SND@LHC to trigger ATLAS [113], and a recent study has shown that FASER can distinguish Majorana and Dirac HNLs through angular measurements in the e^+e^- final states [114]. In contrast, our work uses energy and invariant mass distributions, and is the first to utilize a far-forward detector as a trigger for ATLAS to probe the spinor nature of HNLs.

We note also that, very recently, the first candidate "tagged" neutrino event [115] has been reported by the NA62 Collaboration [116], where the neutrino's production in the decay $K^+ \to \mu^+\nu_\mu$ is observed, and then its subsequent charged-current (CC) interaction in a liquid krypton calorimeter is also detected. In this study, we extend the detection of tagged neutrinos in the SM to the detection of tagged HNLs in BSM models. This work demonstrates that, with "cradle to grave" observations that track HNLs throughout their entire lifetime, one can precisely determine the HNL's properties, including if it is a Majorana fermion. Such a determination will have far-reaching consequences for our understanding of neutrino masses, baryogenesis, and the fundamental symmetries of nature.

This chapter is organized as follows. In Sec. 4.2, we describe the HNL model, including the relevant Lagrangian, mixing parameters, and the resulting interactions. In Sec. 4.3, we explain our experimental setup and simulation, and in Sec. 4.4 we detail the statistical methods we will use. In Sec. 4.5, we assess how measurements at FASER2 alone can measure HNL masses and couplings. In Sec. 4.6, we show that these same measurements may be used to determine if an HNL is a Majorana or Dirac fermion. In Sec. 4.7, we consider the more challenging HNL model, and investigate the potential of using FASER2 as a trigger for ATLAS, showing how the charge of the lepton produced in association with the HNL can be used to improve model discrimination. Finally, in Sec. 4.8 we summarize our findings.

4.2 Model

The HNL framework was established in Sec. 1.3, and we adopt those conventions here: the Type-I extension of the SM by gauge-singlet sterile neutrinos (Eq. (1.10)), the active–sterile mixing matrix $U_{\alpha I}$ that defines the HNL mass eigenstates N_I (Eq. (1.6)), and the resulting charged- and neutral-current interactions through which the HNLs couple to SM matter (Eqs. (1.7) and (1.8)). For simplicity, we focus on a minimal scenario involving a single HNL, denoted N, which represents either a Dirac spinor, composed of 2 independent Weyl spinors, or a Majorana spinor, composed of a single Weyl spinor. These interactions allow hadrons to decay to HNLs through CC interactions, as well as HNLs to decay back into SM particles through NC and CC interactions.

		Model 1		**Model 2**		
HNL Model Parameters		$m_N = 1.84$ GeV, $U_\mu = 0.0036$ Majorana		$m_N = 2.00$ GeV, $U_\mu = 0.002$ Majorana		
$c\tau_N$		2.03 m		4.33 m		
$\langle E_N \rangle_{\text{FASER2}}$		820 GeV		828 GeV		
$\langle E_N \rangle_{\text{FASER2+ATLAS}}$		458 GeV		452 GeV		
			FASER2		FASER2	FASER2+ATLAS
Production Modes	N_{events} at HL-LHC (in 3 ab^{-1})	$D_s \to \mu N$	119,000	$B \to D^{*0}\mu N$	447	56.1
		$B \to D^{*0}\mu N$	1670	$B^0 \to D^*\mu N$	405	50.7
		$B^0 \to D^*\mu N$	1570	$B \to D^0\mu N$	199	26.4
		$B_s^0 \to D_s^*\mu N$	333	$B^0 \to D\mu N$	183	22.4
Decay Modes	Branching Fraction	$\mu\rho$	0.187	$\mu\rho$	0.162	
		$\nu\nu\nu$	0.129	$\nu\nu\nu$	0.128	
		$\nu e\mu$	0.125	$\nu e\mu$	0.125	
		$\nu\mu\mu$	0.072	$\nu\mu\mu$	0.073	
		$\boldsymbol{\mu\pi}$	**0.071**	$\boldsymbol{\mu\pi}$	**0.060**	
		$\nu\pi^0$	0.038	$\nu\pi^0$	0.032	
		Other	0.378	Other	0.419	
Total FASER2 $N \to \mu\pi$ Events			8610		81.8	10.6

Table 4.1: Summary of the two benchmark HNL models considered in this study, including the HNL mass m_N, active-sterile mixing parameter U_μ, and decay length $c\tau_N$. $\langle E_N \rangle_{\text{FASER2}}$ is the average energy of HNLs that decay in FASER2, while $\langle E_N \rangle_{\text{FASER2+ATLAS}}$ is the average energy of HNLs that decay in FASER2 and are accompanied by a signal muon that passes through ATLAS. The dominant production modes are listed, along with the number of HNLs they produce that can be detected at FASER2 and FASER2+ATLAS, assuming the baseline HL-LHC integrated luminosity of 3 ab^{-1}. The dominant decay modes and their branching fractions are also listed. All production and decay modes include both charge-conjugate processes. The fully-reconstructible decay mode $N \to \mu\pi$ used in this study is highlighted. Model 1 represents models with a large number of decays in FASER2, while Model 2 has far fewer decays in FASER, but predicts that a sizable fraction of these decays are accompanied by signal muons that pass through ATLAS.

For this work, we consider HNLs that only couple to muons, so $U_\mu \neq 0$, but $U_e = U_\tau = 0$. The

analysis could be extended to HNLs with electron or tau couplings, with the typical caveat that the experimental signals may be more challenging to detect, particularly for tau couplings. For the muon case we consider here, the dominant production modes include $P \to \mu N$, $P \to P'\mu N$, and $P \to V\mu N$, where P and P' denote pseudoscalar mesons, and V denotes a vector meson, as well as production from τ decays $\tau \to \mu\nu N$. Possible decay channels of the HNL include $N \to H^0\nu_\mu$, $N \to H\mu$, $N \to \ell\ell\nu_\mu$, $N \to \ell\nu_\ell\mu$, and $N \to \bar{\nu}_\ell\nu_\ell\nu_\mu$, where H denotes either a pseudoscalar or a vector meson, and $\ell = e, \mu, \tau$. For a full discussion of HNL production and decay modes, see Ref. [6].

In this study, we examine a Majorana HNL with mass m_N and coupling U_μ, and test how well the signal can be described by a Dirac HNL with parameters m'_N and U'_μ. Because LNV decay modes are allowed for Majorana HNLs, but forbidden for Dirac HNLs, the lifetimes of Majorana HNLs are related to the lifetimes of their Dirac counterparts by

$$\tau_{\text{Majorana}} = \frac{1}{\Gamma_{\text{LNC}} + \Gamma_{\text{LNV}}} = \frac{1}{2}\tau_{\text{Dirac}} \ . \tag{4.1}$$

As discussed in Sec. 4.1, we will consider two representative HNL models. A summary of their relevant properties is given in Table 4.1. In Model 1 the HNL mass lies just below the D_s threshold, allowing prolific production of HNLs from forward $D_s \to \mu N$ decays at the LHC. In contrast, in Model 2, the HNL mass lies above the D_s threshold, and HNLs can only be produced in B decays, resulting in lower overall yields, but the muons produced in association with the HNLs are produced at lower η, allowing them to be detected in ATLAS. Both benchmark points lie in currently unconstrained regions of parameter space and are within reach of FASER2, offering promising opportunities to probe the Majorana vs. Dirac nature of HNLs through event rates, kinematics, and decay signatures. In this work, we focus solely on the $\mu\pi$ decay mode, where the decay products are fully visible and the HNL four-momentum is, in principle, reconstructible. Other decay modes, notably $\mu\rho$, may also be useful and can be expected to improve the results derived here.

4.3 Simulating HNL Signals at FASER2

FASER2 [3, 5, 101] is a proposed upgrade of the FASER experiment. As illustrated in Fig. 4.1, we assume that FASER2 is centered on the LOS, and the front of the FASER2 detector is $L = 650$ m from the ATLAS IP. The FASER2 decay volume is 10 m long, with a rectangular cross section that is 1 m high and 3 m wide.

The detector includes several key components critical to reconstructing HNL decay events. An electromagnetic calorimeter stops electrons and photons and measures their energies. A hadronic calorimeter measures the energy of charged pions and distinguishes them from muons. A tracking system measures muon momenta and distinguishes μ^- from μ^+ through their curvature in a magnetic field. Estimates of the uncertainties of these measurements for the current FASER2 design [117] are

$$\frac{\sigma_{p_\mu}}{p_\mu} = 3.3 \times 10^{-5}\ [p_\mu/\text{GeV}] \tag{4.2}$$

$$\frac{\sigma_{E_{\pi^\pm}}}{E_{\pi^\pm}} = \frac{35\%}{\sqrt{E_{\pi^\pm}/\text{GeV}}} \oplus 0.5\% \tag{4.3}$$

$$\sigma_{\theta_{\mu\pi}} = 250\ \mu\text{rad}\ . \tag{4.4}$$

In addition to these uncertainties, there is a systematic uncertainty on the HNL flux from uncertainties in the flux of forward hadron production at the LHC. We estimate this uncertainty by using the event generators in `FORESEE` [70] that yield the maximal and minimal number of HNL events. For the minimum, we used `QGSJET 2.04` [73] for light hadrons and `NLO-P8-Min` for heavy hadrons. For the maximum, we used `SIBYLL 2.3d` [72] for light hadrons and `NLO-P8-Max` for heavy hadrons. The heavy hadron generators are based on `Pythia8` [76], where the underlying assumptions are detailed in Ref. [62]. The range of event rates from these generators corresponds to an overall flux uncertainty of approximately $\pm 60\%$.

To perform realistic event simulations, we use the `FORESEE` framework [70] for event generation and `HNLCalc` [6, 118] to extend `FORESEE` to simulate HNL events; see also Sec. 2.3. A diagram illustrating the simulation workflow is shown in Fig. 4.2. In `FORESEE`, one specifies the pp center-of-mass energy

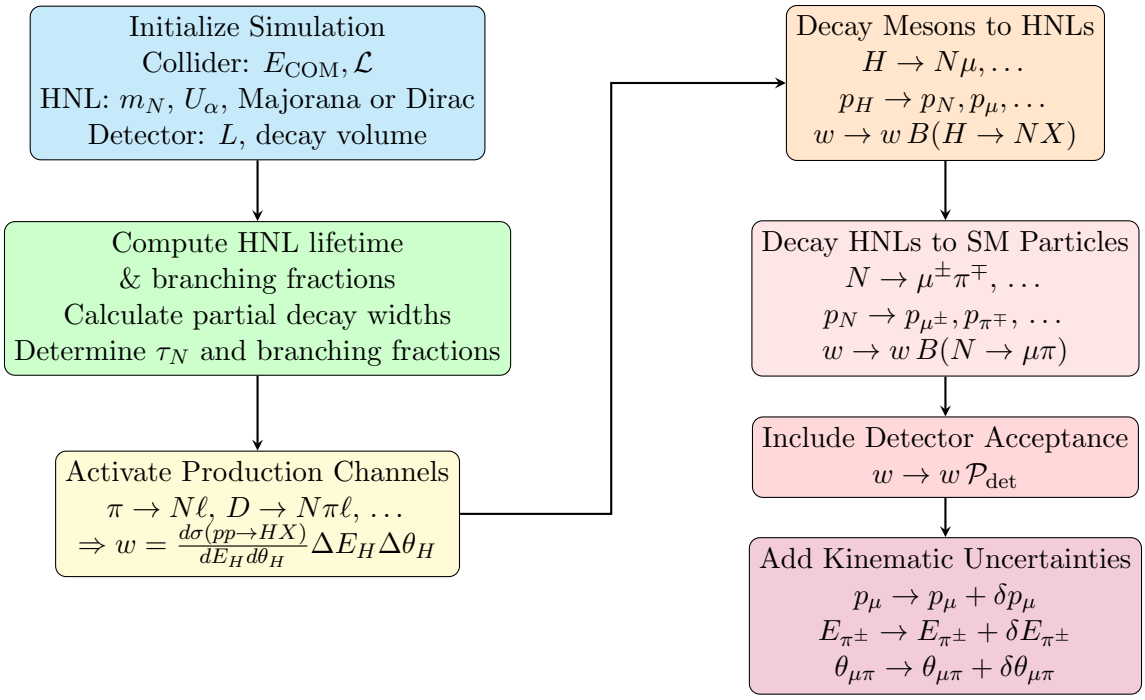


Figure 4.2: Schematic diagram illustrating the major steps in simulating HNL events in the ATLAS and FASER2 detectors. The parameter w is the event weight; it starts with the number of hadrons produced for a particular range of energy and angle, and then is reduced by branching fractions and detector acceptance to determine the number of HNL events detected in FASER2.

$E_{\rm COM} = 14$ TeV and integrated luminosity $\mathcal{L} = 3$ ab^{-1} for the HL-LHC, and chooses one of several state-of-the-art MC generators to generate the forward meson production cross sections and spectra, as discussed above. One then further specifies the HNLs mass m_N, mixing angles U_α, and spinor nature (Majorana or Dirac), and `HNLCalc` then computes the production rate for HNLs in meson decays, and the HNL decay branching fractions and lifetime. Finally, one must specify the detector's distance from the ATLAS IP L and its decay volume geometry.

With all of this information, `FORESEE` then simulates the production of HNLs in meson decays, propagates them to the FASER2 detector, and decays them in the specified decay volume. As a last step, we crudely simulate the detector's response by adding a Gaussian smearing of the relevant kinematic quantities with the resolutions given in Eq. (4.4). The output is an estimate of the number of HNL events detected in FASER2, along with all relevant kinematic distributions.

In Fig. 4.3 we show results for Model 1 of the simulation for one-dimensional kinematic distributions reconstructed from HNLs that decay via $N \to \mu^\pm\pi^\mp$ in FASER2. The invariant mass is

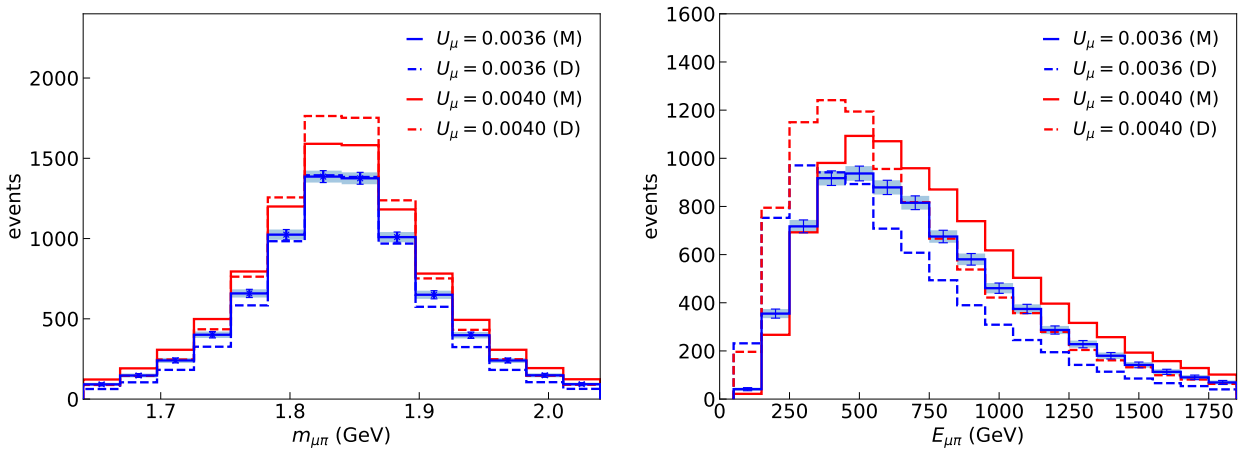


Figure 4.3: Kinematic distributions for HNLs that decay through $N \to \mu^\pm\pi^\mp$ in FASER2. The HNLs are produced in meson decays at the HL-LHC with $E_{\rm COM} = 14$ TeV and an integrated luminosity of 3 ab^{-1}. Results are shown for Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$, as well as for variations from this model with Dirac HNLs, and with $U_\mu = 0.0040$, as indicated. The shaded error bars illustrate the envelope of expected outcomes in each bin, given by $\sqrt{\mu_i}$, where μ_i is the number of events in bin i.

$m_{\mu\pi} = \left[(p_\mu + p_\pi)^2\right]^{1/2}$, and the total energy is $E_{\mu\pi} = E_\mu + E_\pi$. Although the full results are two-dimensional histograms, we display the one-dimensional projections for clarity. Kinematic smearing has been included, as described above. All available meson production channels are included in the simulation. For Model 1, approximately 8600 $N \to \mu^\pm\pi^\mp$ events are expected at the HL-LHC.

Results are shown for Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$, as well as for variations from this model. The overall event rate scales roughly as U_μ^4, and so there are significantly more events for $U_\mu = 0.004$ than for $U_\mu = 0.0036$. The $m_{\mu\pi}$ distribution is centered at m_N in all cases, as expected. However, the peak of the energy distribution shifts. A larger U_μ implies a shorter HNL lifetime, requiring the HNLs to have higher energies to reach FASER2 before decaying. Similarly, for fixed U_μ, switching from a Majorana to a Dirac HNL eliminates half of the decay modes, which increases the lifetime, leading to softer energy distributions.

As can be seen in Fig. 4.3, changing the HNL from a Majorana to a Dirac fermion or changing the coupling by a factor of 10% leads to statistically significant changes in the number of events in the bins of these one-dimensional distributions, suggesting that these properties of HNLs can be precisely determined by FASER2 measurements alone. In the following, we investigate whether

these conclusions hold up under a more rigorous statistical analysis.

4.4 Statistical Methods

4.4.1 Profiled Likelihood

To investigate the characterization capabilities of FASER2 in the event of a signal detection, we perform a binned maximum likelihood (ML) fit of the HNL parameters to an artificial dataset, referred to as an *Asimov dataset*, $\vec{d}_A = (\mu_1, \ldots, \mu_N)$, where μ_i is the mean signal yield in bin i [119], allowing us to obtain the expected measurement sensitivity to the model parameters. We will separate our data into 2D bins over the invariant mass $m_{\mu\pi}$ and total energy $E_{\mu\pi}$, as we will discuss in more detail below in Sec. 4.5.

We model the likelihood of observing a given set of data, which we denote $\vec{d}$, as a Poisson distribution in each bin with Gaussian constraints from flux and MC uncertainties:

$$L(\vec{d}|m_N, U_\mu\,;\,\eta,\vec{\gamma}) = \prod_i f_P(d_i\,|\,\eta\gamma_i\mu_i) f_G(1|\eta,\sigma_\eta) f_G(1|\gamma_i,\sigma_i)\ , \tag{4.5}$$

where $\mu_i = \mu_i(m_N, U_\mu)$ is the expected bin count for a given m_N and U_μ; η and $\vec{\gamma}$ are nuisance parameters that characterize the flux normalization and bin-by-bin uncertainties, respectively; and f_P and f_G are Poisson and Gaussian distributions, respectively.

The parameter η represents the systematic uncertainty in the HNL flux normalization. As a simplifying assumption, we have chosen to model this as a Gaussian distribution with a mean value of unity and a variance equal to $\sigma_\eta = 60\%$, in accordance with the size of the overall spread of fluxes predicted by various forward hadron generators, as discussed in Sec. 4.3.

The parameters γ_i, on the other hand, represent the relative uncertainties in each bin. We incorporate the statistical uncertainty in estimating μ_i that results from limited MC sampling. Additionally, to account for systematic uncertainties associated with the shape of the forward hadron flux, we add in quadrature a uniform fractional uncertainty, denoted σ_γ, which we estimate is currently roughly 10%. In the large MC sample limit, we can model the bin-by-bin uncertainty as a Gaussian

distribution with a mean value of unity and a variance of $\sigma_i = \frac{1}{\sqrt{n_i}} \oplus \sigma_\gamma$, where n_i is the number of MC samples in that bin.

4.4.2 Estimating Model Parameters

Using the Asimov dataset and the likelihood of Eq. (4.5), we perform a maximum likelihood estimation of the model parameters by minimizing the profiled likelihood

$$t(\vec{d}\,|\,m_N, U_\mu) := -2\ln L(\vec{d}\,|\,m_N, U_\mu\,;\,\hat{\eta}, \hat{\vec{\gamma}}) \ , \tag{4.6}$$

where $\{\hat{\eta}\,(\vec{d}), \hat{\vec{\gamma}}\,(\vec{d})\}$ are the ML estimates for the nuisance parameters that minimize t for a given set of data $\vec{d}$. One can then assess the expected goodness-of-fit of the model parameters to the data by treating t as a test-statistic and computing the median p value, where

$$p = \int_{t_{\rm exp}}^{\infty} f_t(t'\,|\,\hat{m}_{N,A}, \hat{U}_{\mu,A})\,dt' \ , \tag{4.7}$$

and f_t is the distribution of t, $t_{\rm exp} = t(\vec{d}_A|\hat{m}_{N,A}, \hat{U}_{\mu,A})$, and $\hat{m}_{N,A}$ and $\hat{U}_{\mu,A}$ are the ML estimates for the HNL mass and coupling obtained using $\vec{d}_A$. The probability distribution, f_t, is obtained by generating $\mathcal{O}(10^3)$ pseudo-experiments, where the nuisance parameters have been fixed to their best-fit values.

Additionally, we can obtain parameter estimation contours by considering relative likelihoods using the test statistic

$$\tilde{t}(\vec{d}\,|\,m_N, U_\mu) := -2\ln \frac{L(\vec{d}\,|\,m_N, U_\mu;\,\hat{\eta}, \hat{\vec{\gamma}})}{L(\vec{d}\,|\,\hat{m}_N, \hat{U}_\mu;\,\hat{\eta}, \hat{\vec{\gamma}})} \ . \tag{4.8}$$

CL contours can then be determined by computing the median p-value of this test statistic for each point in the model parameter space. Under the assumption that the underlying model is correct, Wilks' theorem states that the profiled likelihood ratio, Eq. (4.8), should follow a χ^2 distribution with degrees of freedom equal to the number of free parameters, which allows us to identify 1σ, 2σ, 3σ, and 4σ CL thresholds at $\tilde{t}(\vec{d}_A) = 2.30$, 6.18, 11.83, and 23.94.

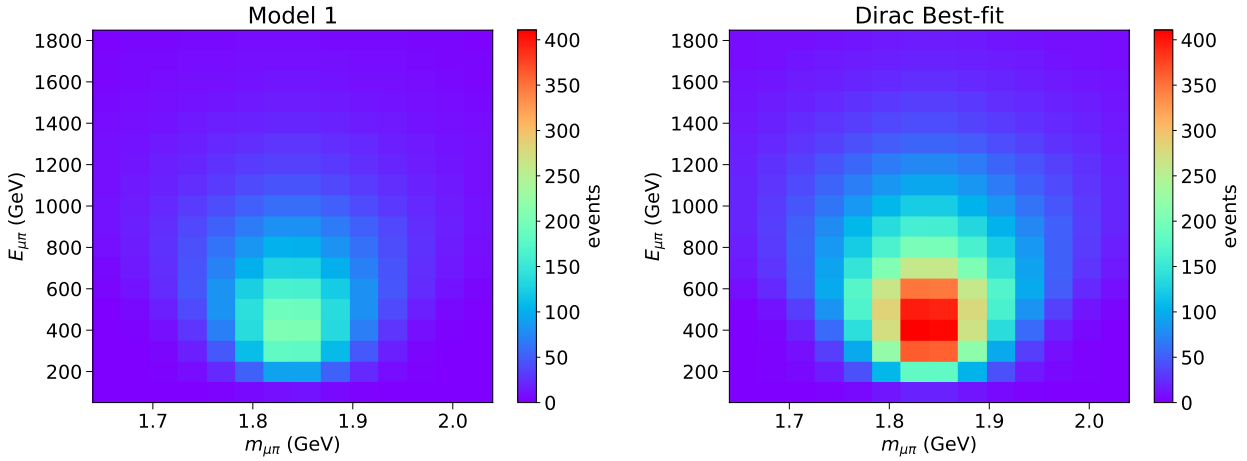


Figure 4.4: Two-dimensional kinematic distributions for the Model 1 Asimov dataset, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$ (left), and the best-fit Dirac model with $\hat{m}_N = 1.84$ GeV and $\hat{U}_\mu = 0.0051$, obtained by minimizing $t(\vec{d}_A \,|\, m_N, U_\mu)$ (right). The Dirac best-fit point recreates the signal shape well, but at the expense of a significantly larger event rate.

4.4.3 Model Discrimination

To differentiate between the Majorana and Dirac HNL models, we define the test statistic

$$\lambda(\vec{d}) := -2 \ln \frac{L_D(\vec{d} \,|\, \hat{m}_N^D, \hat{U}_\mu^D;\, \hat{\eta}^D, \hat{\vec{\gamma}}^D)}{L_M(\vec{d} \,|\, \hat{m}_N^M, \hat{U}_\mu^M;\, \hat{\eta}^M, \hat{\vec{\gamma}}^M)} \tag{4.9}$$

$$= t_D(\vec{d} \,|\, \hat{m}_N^D, \hat{U}_\mu^D) - t_M(\vec{d} \,|\, \hat{m}_N^M, \hat{U}_\mu^M) \;, \tag{4.10}$$

which is the likelihood ratio of the Majorana and Dirac models, which have different expected bin counts, for a given set of data $\vec{d}$. If $\lambda > 0$, the signal appears more "Majorana-like," and if $\lambda < 0$, the signal appears more "Dirac-like." We find the significance of this test by computing its median p-value,

$$p = \int_{\lambda_{\text{exp}}}^{\infty} f_\lambda^D(\lambda')\, d\lambda' \;, \tag{4.11}$$

where f_λ^D is the λ distribution assuming the Dirac hypothesis is true, and $\lambda_{\text{exp}} = \lambda(\vec{d}_A)$. This represents the median probability that a Dirac HNL would fluctuate to produce a signal that looks just as (or more) "Majorana-like" than the current observation. Similar to t, we calculate the p-value by generating the probability distribution, f_λ, from $\mathcal{O}(10^3)$ pseudo-experiments, where the nuisance parameters for both models have been fixed to their best-fit values.

4.5 Measuring HNL Parameters with FASER2

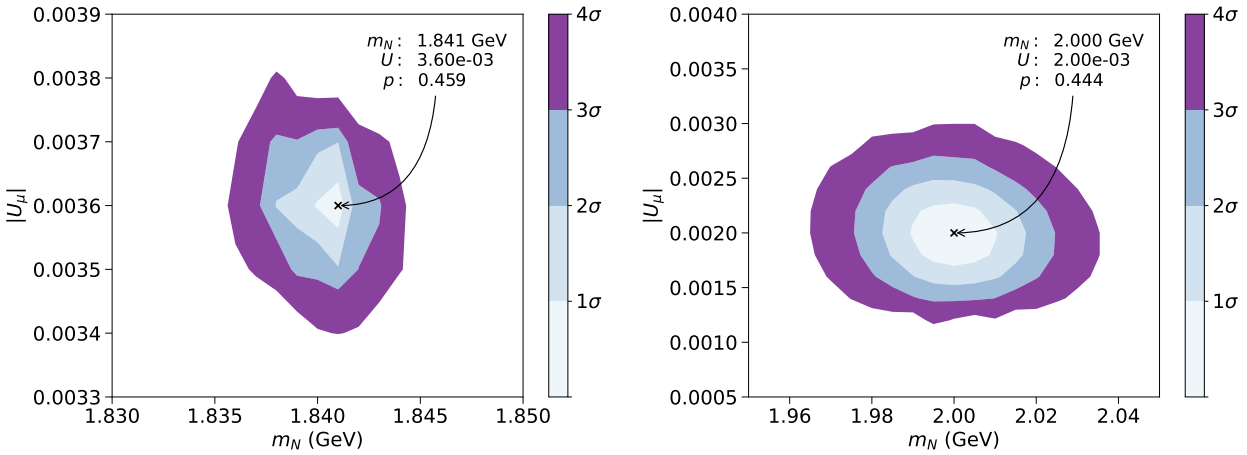


Figure 4.5: Expected parameter estimation contours, assuming the underlying data corresponds to Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$ (left), and Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$ (right). The HNL flux normalization and shape uncertainties are assumed to be $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively.

With the aim of constraining (m_N, U_μ), we separate simulated $N \to \mu\pi$ events into kinematic bins based on the invariant mass, $m_{\mu\pi}$, and total energy, $E_{\mu\pi}$, of the final state. The invariant mass of the final state is reconstructed via

$$m_{\mu\pi}^2 = m_\mu^2 + m_\pi^2 + 2(E_\mu E_\pi - p_\mu p_\pi \cos\theta_{\mu\pi}) \ , \tag{4.12}$$

where $E_{\mu(\pi)}$ is the energy of the muon (pion), $p_{\mu(\pi)}$ is the 3-momentum of the muon (pion), and $\theta_{\mu\pi}$ is the opening angle between the muon and pion, all of which are smeared according to the resolutions given in Sec. 4.3.

To understand the measurement capabilities of FASER2, we conduct our analysis on an Asimov dataset representing Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$, which predicts a large number of $N \to \mu\pi$ events in the FASER2 detector. The resulting 2D kinematic distribution of the Model 1 dataset is shown in the left panel of Fig. 4.4.

In the left panel of Fig. 4.5, we present results showing the expected FASER2 constraints on (m_N, U_μ) in the Majorana parameter space under the assumption that Model 1 is true. We find that FASER2 has the ability to measure m_N to about 2 MeV (a fractional uncertainty of 0.1%)

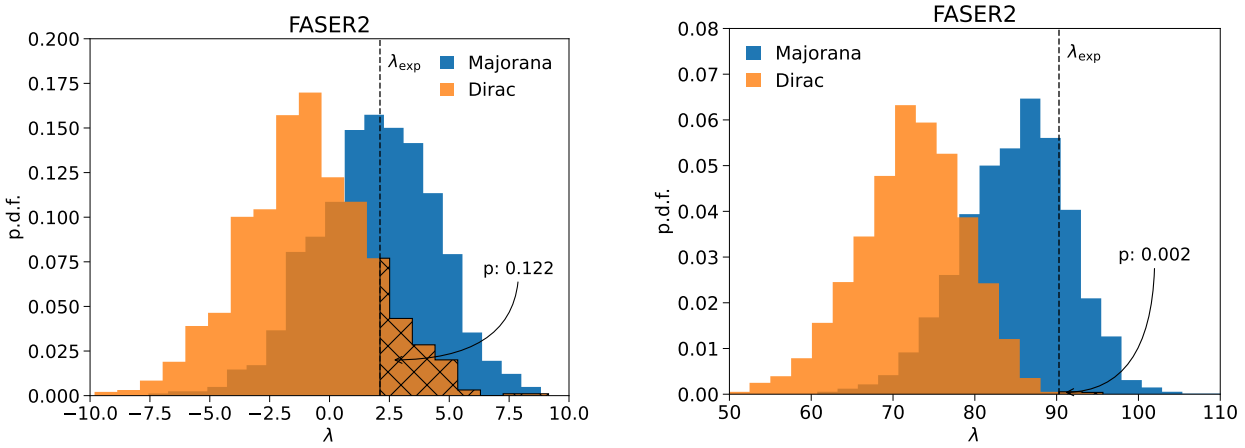


Figure 4.6: Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using FASER2 measurements under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$. Left: Assuming HNL flux normalization and shape uncertainties of $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively, the Dirac hypothesis can be rejected at 87.8% CL (1.2σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 99.8% CL (2.9σ).

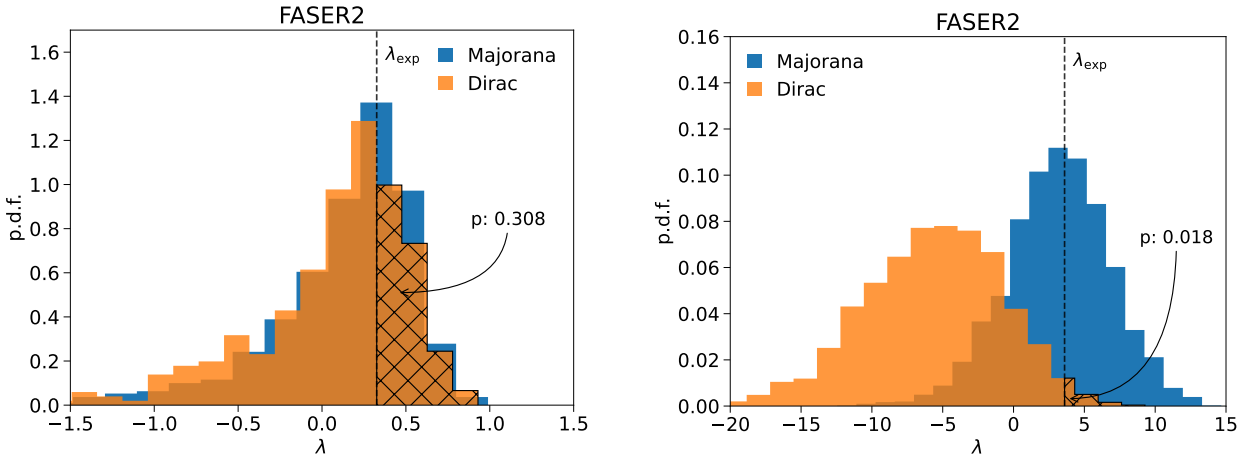


Figure 4.7: Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using FASER2 measurements under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$. Left: Assuming HNL flux normalization and shape uncertainties of $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively, the Dirac hypothesis can be rejected at 69.2% CL (0.5σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 98.2% CL (2.1σ).

and U_μ with a fractional error of approximately 3% at 95% CL. Additionally using Eq. (4.7), we find the expected goodness-of-fit p-value to be 0.46 (0.1σ), suggesting consistency with the Asimov dataset, as expected.

4.6 Differentiating Majorana versus Dirac HNLs at FASER2

Using the test statistic λ defined in Eq. (4.9), we differentiate between the Majorana and Dirac models by generating the distribution of λ under both assumptions and comparing their distributions to the observed value λ_{exp} using Eq. (4.11). We present the resulting distributions after 1000 pseudo-experiments in Fig. 4.6. We find that with HNL flux normalization and shape uncertainties of $\sigma_\eta = 60\%$ and $\sigma_\gamma = 10\%$, respectively, we expect that FASER2 would only favor the Majorana model over the Dirac model at 87.8% CL (1.2σ), as shown in the left panel of Fig. 4.6.

To understand this surprising result, in the right panel of Fig. 4.4, we show that the Dirac model best-fit recreates the shape of the mass and energy distributions of the data very well, but overestimates the total event rate by a factor of 2. However, this falls within the large normalization uncertainty in the HNL flux, resulting in limited discrimination capability.

However, as FASERν continues to constrain hadron fluxes from neutrino observations [120], the uncertainty in hadron production from the MC generators will be reduced significantly by data. These constraints may be sharpened further by additional data from FASERν2 and FLArE at the FPF [101, 102]. If the normalization uncertainty is reduced to $\sigma_\eta = 5\%$ and the shape uncertainty σ_γ becomes negligible, this would exacerbate the difference in event rates and increase our expected discrimination capability to 99.8% CL (2.9σ), as shown in the right panel of Fig. 4.6.

4.7 FASER2 as a Trigger for ATLAS

So far, we have considered what can be learned with FASER2 measurements alone. In optimistic scenarios where there are many events, as in Model 1 of Table 4.1, which we have considered so far, we have seen that FASER2 may precisely measure the HNL parameters, and help determine if it is a Majorana or Dirac fermion if hadronic fluxes are better constrained.

As noted in Sec. 4.1, however, it is also possible to discover HNLs with far fewer events, and it is interesting to explore what can be done in this case. In this section, we consider Model 2 of Table 4.1, where only approximately 80 HNL decays $N \to \mu\pi$ can be seen at FASER2 at the HL-LHC.

Applying the same method and statistical analysis as done for Model 1, we find that for Model 2, FASER2 alone can still constrain the HNL mass and coupling to fractional uncertainties of 1% and 25% at 95% CL, respectively, as shown in the right panel of Fig. 4.5. The mass measurement is still quite precise, but the coupling measurement is significantly degraded relatively to Model 1. Additionally, we find that the Dirac hypothesis can be disfavored at 98.2% CL (2.1σ) with improved hadronic fluxes, as shown in Fig. 4.7. The ability to distinguish Majorana from Dirac is therefore slightly degraded relative to the case of Model 1.

In the rest of this section, we discuss whether FASER2 may be used as a trigger for ATLAS, allowing ATLAS to record muons produced in association with the HNL, and we discuss what improvements to model discrimination can be made by the combining measurements at the two detectors.

4.7.1 Forward Muons at ATLAS

The HL-LHC is scheduled to begin operating in 2030 and continue through to the planned end of the LHC in the 2040s. During this period, the peak instantaneous luminosity is expected to rise from its current value of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, producing on average 27 pileup events per bunch crossing, to $7.18 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, producing around 200 pileup events per bunch crossing. The goal is to accumulate an integrated luminosity of approximately 3 ab^{-1} by the end of the HL-LHC era [121, 122].

In conjunction with the HL-LHC upgrade, the ATLAS detector will undergo a significant Phase-II upgrade. This includes the installation of a completely new Inner Tracker (ITk), which will replace the existing Pixel Detector, Semiconductor Tracker (SCT), and Transition Radiation Tracker (TRT). Additional upgrades include major improvements to the Trigger and Data Acquisition (TDAQ) system and discussion of an additional high η muon tagger. The muon tagger will extend muon coverage to regions up to $|\eta| < 4$, significantly increasing overall acceptance [123–125].

Calorimeter systems at ATLAS also play a critical role in particle identification and energy measurement. The electromagnetic calorimeter provides coverage up to $|\eta| < 3.2$ [124], while the Tile Hadronic Calorimeter covers the central region up to $|\eta| < 1.7$ [125, 126]. The Hadronic Endcap Calorimeter (HEC) extends coverage into the $1.5 < |\eta| < 3.2$ range [127]. Complementing these, the Forward Calorimeter (FCal) comprises three modules, FCal1 (EM) and FCal2/FCal3 (hadronic) which covers the very forward region from $3.1 < |\eta| < 4.9$ [128, 129].

Taken together, the tracking and calorimeter systems provide full coverage for energy measurements and charged particle reconstruction up to $|\eta| < 4.9$, with muon identification capabilities extending to $|\eta| < 4$. For this analysis, we consider an optimistic scenario in which both the inner tracker and muon tagger are extended to cover $|\eta| < 4.9$. Such an extension would enhance the prospects for identifying muons and their charge in the forward region.

An intriguing possibility is that FASER2 could be used as a trigger for ATLAS. FASER2 is capable of issuing a trigger signal within approximately 1 ns [4]. As noted in Sec. 4.1, the total time for a round trip signal, accounting for both the time it takes the HNL to propagate to FASER2 at the speed of light and the return signal through an optical fiber back to ATLAS, is approximately 5 μs. The ATLAS L0 trigger latency is 6 μs [123]. Therefore, it is feasible to instruct ATLAS to retain data from a bunch crossing that has produced an HNL signal in FASER2. This coordination would allow for the simultaneous use of FASER2 and ATLAS data to perform a correlated analysis [130–132].

4.7.2 Pileup Background

The muons produced in association with HNLs at ATLAS are produced in the forward region, where there is a lot of background, especially given the effects of pileup. To estimate the pileup background, we simulate proton-proton collisions at a center-of-mass energy of 14 TeV using `Pythia8`. We conservatively assume the peak instantaneous luminosity, corresponding to roughly 200 pileup events per bunch crossing, which gives the maximal expected background. The simulation includes all muons originating from either light or heavy mesons that decay within the ITk, which is modeled as a cylindrical volume with a radius of 1 m and a length of 6 m. The SoftQCD option in `Pythia8`

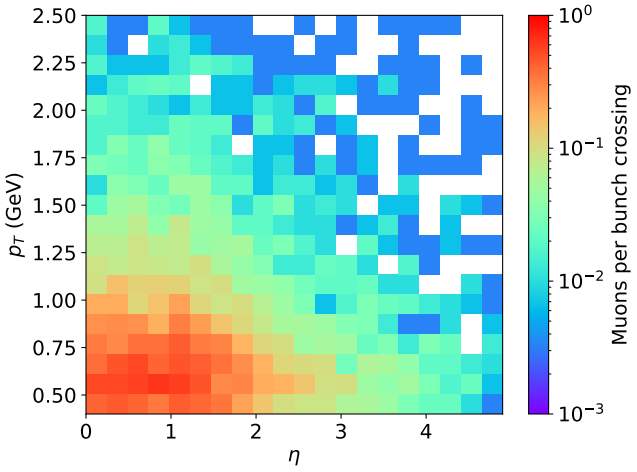


Figure 4.8: The number of muons produced per bunch crossing in the (η, p_T) plane, including both muons that are produced in light meson ($\pi^{\pm}$, $K^{\pm}$, K_L^0) decays and heavy meson (charm and bottom meson) decays within the ITk.

Parent	Muons ($\lvert\eta\rvert < 4.9$)	Muons ($3.5 < \eta < 4.9$)
$\pi^{\pm}$	31.41	0.43
$K^{\pm}$	16.09	0.26
K_L^0	0.37	0.0033
Heavy Mesons	3.34	0.33

Table 4.2: Number of muons produced with $p_T > 0.4$ GeV in the ITk per bunch crossing at ATLAS from various parent particles, shown for the full muon coverage $|\eta| < 4.9$ and the forward region $3.5 < \eta < 4.9$. "Heavy Mesons" includes charm and bottom mesons. Yields include both μ^+ and μ^-. Muons are categorized by their direct parent particle, so, for example, muons from $K_S^0 \to \pi^+\pi^-$, with the pions decay to muons, are included in the $\pi^{\pm}$ row. The number of muons produced directly from K_S^0 decays is negligible. Muons from light meson decays dominate the background, yielding around 50 per bunch crossing in the full acceptance, but only about 0.7 within the forward region after cuts.

is turned on to simulate low-energy QCD interactions.

The simulation assumes a muon detector with coverage up to $|\eta| < 4.9$, as well as a shield that allows one to distinguish muons from charged pions. The resulting background-muon distribution in (η, p_T) is shown in Fig. 4.8, peaking at lower η values. In contrast, signal muons from HNL decays in FASER2 populate higher η, motivating the cut $3.5 < \eta < 4.9$ to retain most of the signal while rejecting the bulk of the background. An additional requirement of $p_T > 0.4$ GeV avoids the region of poor reconstruction efficiency. The surviving background rates are summarized in Table 4.2.

Combining muon contributions from both light and heavy meson decays, we estimate that roughly 1 background muon per bunch crossing satisfies the final selection criteria. This sets a baseline for understanding and mitigating pileup backgrounds in the forward region of the detector, which is essential for assessing the feasibility of using FASER2 as a trigger for ATLAS.

4.7.3 Results

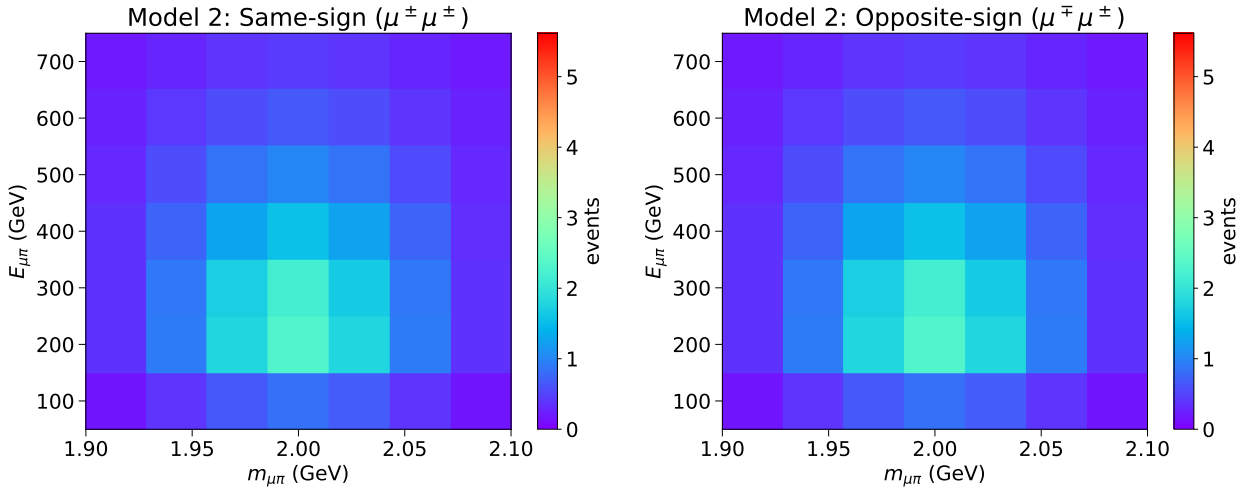


Figure 4.9: Two-dimensional kinematic distributions for the Model 2 Asimov dataset with a Majorana HNL with $m_N = 2$ GeV and $U_\mu = 0.002$ with SS (left) and OS (right) muons seen in ATLAS, assuming an integrated luminosity of 3 ab^{-1}. These distributions include both muons from HNL events and background muons from pileup. Muons produced in association with Majorana HNLs can be found with OS or SS muons in FASER2 with equal rates, resulting in a symmetric distribution between the two histograms.

We separate each triggered HNL event into the number of muons detected in ATLAS that have either the same sign (SS) or opposite sign (OS) as the muon detected in FASER2. Majorana HNLs decay to $\mu^+\pi^-$ and $\mu^-\pi^+$ and produce SS and OS signals with equal rates, whereas Dirac HNLs only produce OS signals. In the absence of background, a single SS event would be evidence of LNV and a Majorana HNL. However, uncorrelated background muons produced in ATLAS may give "fake" SS signals, and so LNV and Majorana-ness can only be established on a statistical basis. Therefore, the expected number of SS and OS events in each bin is given by

$$\mu_i^{\mathrm{SS}} = \frac{b}{2}\,\mu_i^{\mathrm{F}} + \frac{\eta_M}{2}\,\mu_i^{\mathrm{A}} \tag{4.13}$$

$$\mu_i^{\mathrm{OS}} = \frac{b}{2}\,\mu_i^{\mathrm{F}} + \left(1 - \frac{\eta_M}{2}\right)\,\mu_i^{\mathrm{A}}\ , \tag{4.14}$$

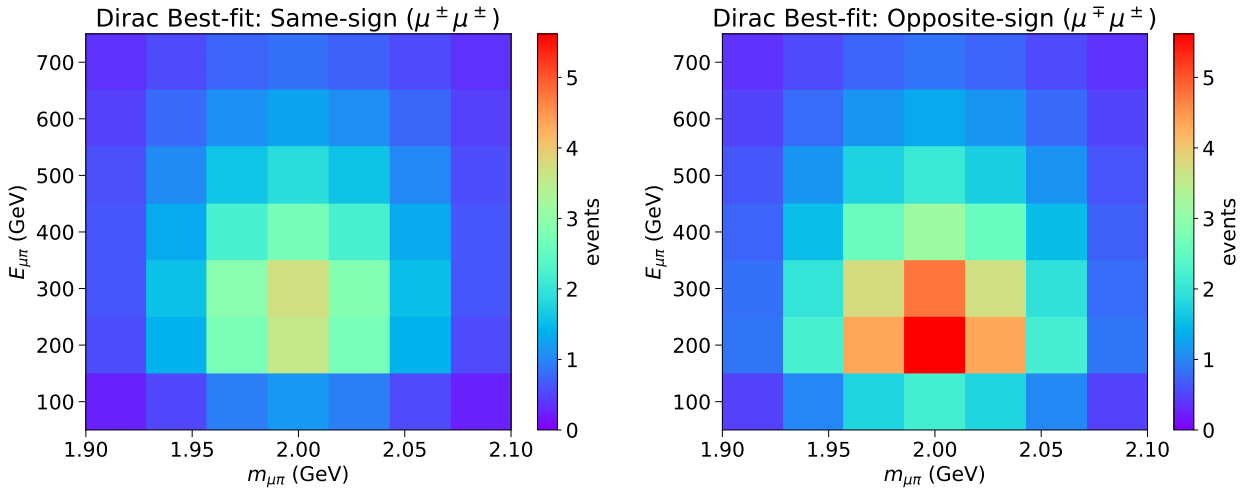


Figure 4.10: Two-dimensional kinematic distributions for the Dirac HNL best-fit point to the Model 2 Asimov dataset with SS (left) and OS (right) muons seen in ATLAS, obtained by minimizing the test-statistic $t(\vec{d}_A \,|\, m_N, U_\mu)$, assuming an integrated luminosity of 3 ab^{-1}. These distributions include both muons from HNL events and background muons from pileup. Muons produced in association with Dirac HNL's can only be found with OS muons in FASER2, resulting in an asymmetric distribution between the two histograms.

where $\mu_i^{\rm F}$ is the total expected number of HNL events in bin i as measured in FASER2, $b \approx 1$ is the expected number of uncorrelated background muons per bunch crossing produced by pileup, μ_i^A is the expected number of truly correlated muons detected in ATLAS in bin i, and $\eta_M = 1$ for the Majorana model and $\eta_M = 0$ for the Dirac. The uncertainty in each bin is estimated by adding the errors of $\mu_i^{\rm F}$ and $\mu_i^{\rm A}$ in quadrature. In Figs. 4.9 and 4.10, we plot the two-dimensional kinematic distribution of the Model 2 Asimov dataset, and the Dirac best-fit point, respectively, separated into SS and OS bins where a clear asymmetry can be observed between the distribution of events that appear LNC and LNV in the Dirac case. In Fig. 4.11, we present results for the distribution of λ utilizing the correlated events from ATLAS. We expect that the Dirac model would be rejected at 91.3% CL (1.4σ), assuming a normalization uncertainty of $\sigma_\eta = 60\%$ and a shape uncertainty of $\sigma_\gamma = 10\%$. If the normalization uncertainty is reduced to $\sigma_\eta = 5\%$ and the shape uncertainty σ_γ becomes negligible by improved constraints on forward hadron production, we find a major improvement in discrimination capability, and we expect to be able to reject the Dirac model at 99.7% CL (2.7σ).

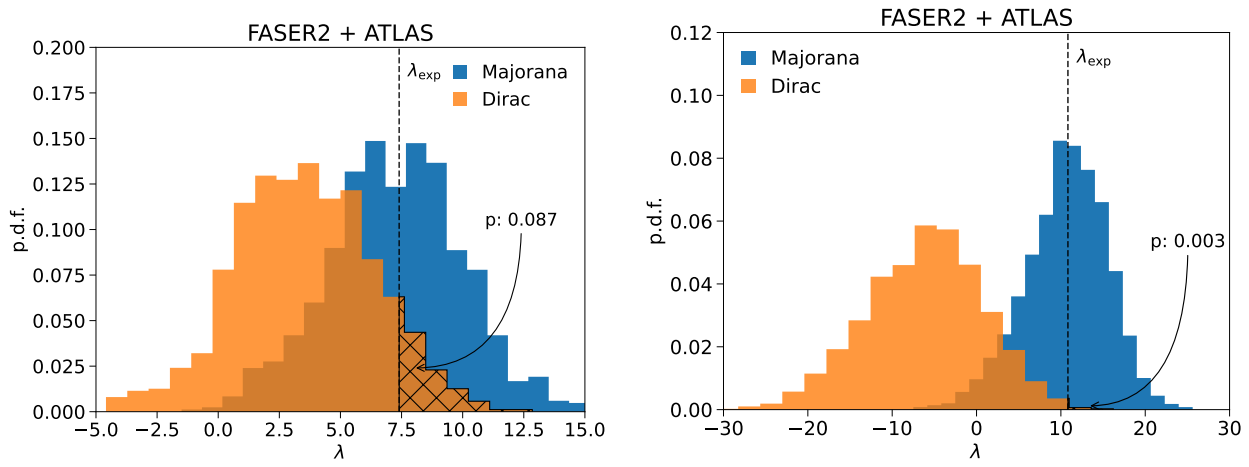


Figure 4.11: Probability distributions of the likelihood ratio between the Majorana and Dirac models, λ, using correlated FASER2 and ATLAS data, under the Majorana and Dirac hypotheses, respectively. The underlying data is assumed to be Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$. Left: Assuming HNL flux normalization and shape uncertainties of 60% and 10%, respectively, the Dirac hypothesis can be rejected at 91.3% CL (1.4σ). Right: Assuming a flux normalization uncertainty of $\sigma_\eta = 5\%$ and a negligible shape uncertainty σ_γ, the Dirac hypothesis can be rejected at 99.7% CL (2.7σ).

4.8 Conclusions

In this chapter, we have investigated the potential of FASER2 to measure the mass and coupling of an HNL and determine whether it is Majorana or Dirac in the event that HNL decays $N \to \mu\pi$ are detected. Additionally, we have explored the possibility of using a signal in FASER2 to trigger ATLAS, and using the correlated FASER2 and ATLAS data to improve model discrimination by correlating the charge of the muon produced in association with the HNL at ATLAS with the charge of the muon produced in HNL decay at FASER2.

When investigating the capabilities of FASER2 alone, we considered benchmark Model 1, a Majorana HNL with $m_N = 1.84$ GeV and $U_\mu = 0.0036$, which predicts approximately 8600 $N \to \mu\pi$ decays in FASER2 at the HL-LHC. We performed a 2D binned ML fit of the HNL model parameters to the expected dataset with events classified by the invariant mass and total energy of the $\mu\pi$ system. Assuming the currently large uncertainties in the forward hadron flux, with an overall flux uncertainty of $\sigma_\eta = 60\%$ and shape uncertainty of $\sigma_\gamma = 10\%$, we have found that the HNL's mass and couplings can be measured remarkably precisely to fractional uncertainties of approximately 0.1% and 3% at 95% CL, respectively, but FASER2 would only be able to favor the Majorana

model over the Dirac model at 87.8% CL (1.2σ). However, assuming that far-forward neutrino measurements are able to reduce forward hadron flux uncertainties to $\sigma_\eta = 5\%$ and a negligible σ_γ, the Dirac fermion hypothesis can be rejected at 99.8% CL (2.9σ).

In the regions of parameter space where the signal yield is insufficient to differentiate Majorana and Dirac HNLs with FASER2 data alone, we have also explored the possibility of using FASER2 to trigger ATLAS and using the correlated FASER2 and ATLAS events to improve model discrimination. In our work, we considered benchmark Model 2, a Majorana HNL with $m_N = 2.0$ GeV and $U_\mu = 0.002$ which predicts only 80 $N \to \mu\pi$ events at FASER2 at the HL-LHC. We determined the expected μ yield at ATLAS from HNL production, as well as the expected background from pileup interactions. Separating HNL events into those that appear LNC (with an OS muon observed in ATLAS) and those that appear LNV (with an SS muon observed in ATLAS), we found that FASER2 and ATLAS data combined were able to discriminate between Majorana and Dirac HNLs, even in the presence of misidentified background events (Figs. 4.9 and 4.10). We found that if the normalization uncertainty from forward hadron production was reduced to $\sigma_\eta = 5\%$ and the shape uncertainty σ_γ was negligible, the combined data from FASER2 and ATLAS are able to reject the Dirac hypothesis at 99.7% CL (2.7σ).

Further work could extend this analysis by consider more general HNL models as well as other HNL decays. For the sake of simplicity, we limited ourselves to the case of an HNL with only muon couplings. However, models with couplings to more than one flavor have been suggested as being more representative of realistic models that generate the observed neutrino masses and mixing angles [90] and thus pose an interesting avenue for further study. Additionally, we chose to focus on the decay channel $N \to \mu\pi$, where the decay products are fully visible and the N energy and momentum are therefore, in principle, fully reconstructible. However, another HNL decay mode, $\mu\rho$, although more complicated, also satisfies this criterion, and other decay modes, such as $\nu\mu\mu$, could be incorporated to further increase the overall signal yield and provide further characterization capability.

This work shows that the main LHC detectors and proposed auxiliary detectors can have a synergistic relation on an event-by-event basis. In Model 2, FASER2 would detect HNL events, but

ATLAS alone would not. However, by using FASER2 as a trigger for ATLAS, the combined data from FASER2 and ATLAS are able to detect the HNL over its entire lifetime from the cradle to the grave, that is, from its production to its decay. With this combined information, the combination of transverse and far-forward LHC detectors may be able to determine the HNL's mass and couplings and also establish that the HNL is a Majorana fermion. Such studies would have far-reaching implications for many topics, including neutrino masses, neutrinoless double beta decay experiments, baryogenesis, and our understanding of the underlying symmetries of the physical world.

Chapter 5

Conclusions and Outlook

In this dissertation, we considered three aspects of MeV-to-GeV HNL phenomenology relevant to the forward physics experiments at the LHC: modelling, discovery, and characterization. Chapter 3 introduced `HNLCalc`, a Python library that computes HNL production rates in pseudoscalar meson and tau decays, decay branching fractions, and total widths for general couplings to all three lepton flavors, and used it to extend the `FORESEE` forward-physics simulation framework to HNLs. Sensitivity reach contours were computed for FASER during Run 3, FASER throughout the HL-LHC era, and FASER2 at the HL-LHC, in benchmark coupling patterns [23]: electron-, muon-, and tau-flavor dominance, together with the mixed flavor scenarios. For the single-coupling benchmarks, the analysis reproduces and refines previous results [60, 88] and for the mixed-coupling 011 and 111 benchmarks, no comparable forward-LHC analysis existed before this work. Several extensions of this work are natural next steps. The same computational framework can be applied to the rest of the proposed FPF detector suite (FLArE, FORMOSA), enabling combined-experiment reach projections rather than the FASER-only contours computed here. Additionally, the framework could be used to generate projections for the FCC-hh, the proposed $\sim$ 100 TeV proton-proton successor to the LHC currently being studied at CERN [133].

Chapter 4 took the question one step further: assuming an HNL has been discovered, how well can FASER2 pin down its mass, its coupling, and its spinor nature? Three measurement techniques

were developed, each addressing a different aspect of the characterization problem. First, a two-dimensional binned likelihood fit to the invariant mass and total energy of the fully-reconstructible $N \to \mu\pi$ decay constrains the HNL mass and coupling at the percent level. Second, the Majorana-versus-Dirac question can be addressed by FASER2 alone through the impact of lepton-number-violating decays on the HNL kinematic spectrum, assuming that forward hadron flux uncertainty is reduced by FASERν and FPF measurements [109]. Third, when the FASER2-only signal is too small to lift this degeneracy, sensitivity is restored by using a FASER2 as a trigger for ATLAS to measure the charge of forward-muons produced in association with HNLs. When correlated with the charge of the muons detected in FASER2, this identifies each event as lepton-number-conserving or lepton-number-violating on a statistical basis and rejects the Dirac hypothesis with good confidence. Several extensions of this work are natural next steps. The same statistical framework can be applied to multi-flavor HNL models that more closely represent realistic see-saw completions [90], and to additional fully- or partially-reconstructible decay modes such as $N \to \mu\rho$ and $N \to \nu\mu^+\mu^-$ that could supplement the $N \to \mu\pi$ channel used here. More broadly, the proposal of using FASER2 as a trigger for ATLAS developed here is itself a detector capability that could extend beyond HNLs to other far-forward LLPs, such as dark photons or axion-like particles, whose characterization may similarly benefit from correlated central-detector information.

Reflections

Major discoveries in particle physics tend to unfold in two phases. An experiment first establishes an anomalous signal, and a longer program then extracts what it is. The Higgs boson is the recent precedent: ATLAS and CMS announced the observation of a new boson at 125 GeV in 2012 [134, 135], but the measurements of the couplings to fermions and gauge bosons, of the spin and parity, and of the production and decay rates took the better part of a decade and remain ongoing [136, 137]. The two technical chapters of this thesis are structured around this idea. Chapter 3 asks where to look for HNLs in the MeV–GeV window. Chapter 4 asks what we would do once we found one.

The forward-LHC LLP program is less developed on the second of these questions than on the

first. Sensitivity-reach calculations are well represented in the literature; quantitative studies of how to characterize a discovery, by measuring its mass, its couplings, its spinor nature, or its flavor structure, are comparatively few. The reason is straightforward. Small forward experiments are funded and scoped around discovery, and characterization is a problem one defers until there is a signal to characterize. This is reasonable at the level of any individual experiment, but it leaves the program collectively under-prepared for the moment of discovery. The contrast with the main LHC detectors is instructive: ATLAS and CMS were built with the capability to characterize a Higgs-like state, not only to discover one, with multiple decay channels, redundant subsystems, and precision tracking and calorimetry, because the theoretical case was strong enough to justify the investment. The forward-LLP program does not yet enjoy that level of theoretical confidence in any single target, but it does have a small set of well-motivated candidates, with HNLs prominent among them, for which the analogous case could be made.

Bibliography

[1] **KamLAND** Collaboration, S. Abe *et al.*, "Precision Measurement of Neutrino Oscillation Parameters with KamLAND," *Phys. Rev. Lett.* **100** (2008) 221803, arXiv:0801.4589 [hep-ex].

[2] E. Fernández-Martínez, M. González-López, J. Hernández-García, M. Hostert, and J. López-Pavón, "Effective portals to heavy neutral leptons," *JHEP* **09** (2023) 001, arXiv:2304.06772 [hep-ph].

[3] J. L. Feng *et al.*, "The Forward Physics Facility at the High-Luminosity LHC," *J. Phys. G* **50** no. 3, (2023) 030501, arXiv:2203.05090 [hep-ex].

[4] **FASER** Collaboration, H. Abreu *et al.*, "The FASER detector," *JINST* **19** no. 05, (2024) P05066, arXiv:2207.11427 [physics.ins-det].

[5] O. Salin, A. Barr, J. McFayden, J. Boyd, M. Vranjes Milosavljevic, N. Vranjes, M. D'Onofrio, C. Gwilliam, J. K. Anders, S. Jakobsen, H. Otono, M. Queitsch-Maitland, B. J. Wilson, K. Lohwasser, M. V. Diwan, M. Vicenzi, and W. Wu, "FASER2: Detector Design and Performance,". https://cds.cern.ch/record/2927003.

[6] J. L. Feng, A. Hewitt, F. Kling, and D. La Rocco, "Simulating heavy neutral leptons with general couplings at collider and fixed target experiments," *Phys. Rev. D* **110** no. 3, (2024) 035029, arXiv:2405.07330 [hep-ph].

[7] J. L. Feng, A. Hewitt, D. La Rocco, and D. Whiteson, "Characterizing heavy neutral leptons: Measuring parameters, discriminating Majorana versus Dirac, and using FASER2 as a trigger for ATLAS," *Phys. Rev. D* **113** no. 3, (2026) 035006, arXiv:2510.16107 [hep-ph].

[8] **Super-Kamiokande** Collaboration, Y. Fukuda *et al.*, "Evidence for oscillation of atmospheric neutrinos," *Phys. Rev. Lett.* **81** (1998) 1562–1567, arXiv:hep-ex/9807003.

[9] **SNO** Collaboration, Q. R. Ahmad *et al.*, "Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory," *Phys. Rev. Lett.* **89** (2002) 011301, arXiv:nucl-ex/0204008.

[10] **KamLAND** Collaboration, K. Eguchi *et al.*, "First results from KamLAND: Evidence for reactor anti-neutrino disappearance," *Phys. Rev. Lett.* **90** (2003) 021802, arXiv:hep-ex/0212021.

[11] **Daya Bay** Collaboration, F. P. An *et al.*, "Observation of electron-antineutrino disappearance at Daya Bay," *Phys. Rev. Lett.* **108** (2012) 171803, arXiv:1203.1669 [hep-ex].

[12] **K2K** Collaboration, M. H. Ahn *et al.*, "Measurement of Neutrino Oscillation by the K2K Experiment," *Phys. Rev. D* **74** (2006) 072003, arXiv:hep-ex/0606032.

[13] **MINOS** Collaboration, D. G. Michael *et al.*, "Observation of Muon Neutrino Disappearance with the MINOS Detectors and the NuMI Neutrino Beam," *Phys. Rev. Lett.* **97** (2006) 191801, arXiv:hep-ex/0607088.

[14] **T2K** Collaboration, K. Abe *et al.*, "Observation of Electron Neutrino Appearance in a Muon Neutrino Beam," *Phys. Rev. Lett.* **112** (2014) 061802, arXiv:1311.4750 [hep-ex].

[15] **NOvA** Collaboration, P. Adamson *et al.*, "First measurement of electron neutrino appearance in NOvA," *Phys. Rev. Lett.* **116** no. 15, (2016) 151806, arXiv:1601.05022 [hep-ex].

[16] **Particle Data Group** Collaboration, S. Navas *et al.*, "Review of particle physics," *Phys. Rev. D* **110** no. 3, (2024) 030001.

[17] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, "NuFit-6.0: updated global analysis of three-flavor neutrino oscillations," *JHEP* **12** (2024) 216, arXiv:2410.05380 [hep-ph].

[18] S. Weinberg, "Baryon and Lepton Nonconserving Processes," *Phys. Rev. Lett.* **43** (1979) 1566–1570.

[19] **KamLAND-Zen** Collaboration, S. Abe *et al.*, "Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen," *Phys. Rev. Lett.* **130** no. 5, (2023) 051801, arXiv:2203.02139 [hep-ex].

[20] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, "Dipole Portal to Heavy Neutral Leptons," *Phys. Rev. D* **98** no. 11, (2018) 115015, arXiv:1803.03262 [hep-ph].

[21] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, "Phenomenology of GeV-scale Heavy Neutral Leptons," *JHEP* **11** (2018) 032, arXiv:1805.08567 [hep-ph].

[22] A. M. Abdullahi *et al.*, "The present and future status of heavy neutral leptons," *J. Phys. G* **50** no. 2, (2023) 020501, arXiv:2203.08039 [hep-ph].

[23] J. Beacham *et al.*, "Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report," *J. Phys. G* **47** no. 1, (2020) 010501, arXiv:1901.09966 [hep-ex].

[24] S. Böser, C. Buck, C. Giunti, J. Lesgourgues, L. Ludhova, S. Mertens, A. Schukraft, and M. Wurm, "Status of Light Sterile Neutrino Searches," *Prog. Part. Nucl. Phys.* **111** (2020) 103736, arXiv:1906.01739 [hep-ex].

[25] S. Dodelson and L. M. Widrow, "Sterile-neutrinos as dark matter," *Phys. Rev. Lett.* **72** (1994) 17–20, arXiv:hep-ph/9303287.

[26] X.-D. Shi and G. M. Fuller, "A New dark matter candidate: Nonthermal sterile neutrinos," *Phys. Rev. Lett.* **82** (1999) 2832–2835, arXiv:astro-ph/9810076.

[27] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, "Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data," *Phys. Rev. D* **88** (2013) 043502, arXiv:1306.2314 [astro-ph.CO].

[28] K. Abazajian, G. M. Fuller, and M. Patel, "Sterile neutrino hot, warm, and cold dark matter," *Phys. Rev. D* **64** (2001) 023501, arXiv:astro-ph/0101524.

[29] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, "Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters," *Astrophys. J.* **789** (2014) 13, arXiv:1402.2301 [astro-ph.CO].

[30] M. Drewes *et al.*, "A White Paper on keV Sterile Neutrino Dark Matter," *JCAP* **01** (2017) 025, arXiv:1602.04816 [hep-ph].

[31] T. Asaka, S. Blanchet, and M. Shaposhnikov, "The nuMSM, dark matter and neutrino masses," *Phys. Lett. B* **631** (2005) 151–156, arXiv:hep-ph/0503065.

[32] T. Asaka and M. Shaposhnikov, "The νMSM, dark matter and baryon asymmetry of the universe," *Phys. Lett. B* **620** (2005) 17–26, arXiv:hep-ph/0505013.

[33] M. Fukugita and T. Yanagida, "Baryogenesis Without Grand Unification," *Phys. Lett. B* **174** (1986) 45–47.

[34] L. Calibbi and G. Signorelli, "Charged Lepton Flavour Violation: An Experimental and Theoretical Introduction," *Riv. Nuovo Cim.* **41** no. 2, (2018) 71–174, arXiv:1709.00294 [hep-ph].

[35] P. Minkowski, "$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?," *Phys. Lett. B* **67** (1977) 421–428.

[36] R. N. Mohapatra and G. Senjanovic, "Neutrino Mass and Spontaneous Parity Nonconservation," *Phys. Rev. Lett.* **44** (1980) 912.

[37] R. N. Mohapatra and J. W. F. Valle, "Neutrino Mass and Baryon Number Nonconservation in Superstring Models," *Phys. Rev. D* **34** (1986) 1642.

[38] M. Malinsky, J. C. Romao, and J. W. F. Valle, "Novel supersymmetric SO(10) seesaw mechanism," *Phys. Rev. Lett.* **95** (2005) 161801, arXiv:hep-ph/0506296.

[39] J. Kersten and A. Y. Smirnov, "Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation," *Phys. Rev. D* **76** (2007) 073005, arXiv:0705.3221 [hep-ph].

[40] R. E. Shrock, "A Test for the Existence of Effectively Stable Neutral Heavy Leptons," *Phys. Rev. Lett.* **40** (1978) 1688.

[41] R. E. Shrock, "New Tests For, and Bounds On, Neutrino Masses and Lepton Mixing," *Phys. Lett. B* **96** (1980) 159–164.

[42] R. E. Shrock, "General Theory of Weak Processes Involving Neutrinos. 2. Pure Leptonic Decays," *Phys. Rev. D* **24** (1981) 1275.

[43] A. Atre, T. Han, S. Pascoli, and B. Zhang, "The Search for Heavy Majorana Neutrinos," *JHEP* **05** (2009) 030, arXiv:0901.3589 [hep-ph].

[44] R. Abela, M. Daum, G. H. Eaton, R. Frosch, B. Jost, P. R. Kettle, and E. Steiner, "Search for an Admixture of Heavy Neutrino in Pion Decay," *Phys. Lett. B* **105** (1981) 263–266. [Erratum: Phys.Lett.B 106, 513 (1981)].

[45] Y. Asano *et al.*, "Search for a Heavy Neutrino Emitted in $K^+ \to \mu^+$ Neutrino Decay," *Phys. Lett. B* **104** (1981) 84–88.

[46] D. I. Britton *et al.*, "Improved search for massive neutrinos in $\pi^+ \to e^+\nu$ decay," *Phys. Rev. D* **46** (1992) R885–R887.

[47] **E949** Collaboration, A. V. Artamonov *et al.*, "Search for heavy neutrinos in $K^+ \to \mu^+\nu_H$ decays," *Phys. Rev. D* **91** no. 5, (2015) 052001, arXiv:1411.3963 [hep-ex]. [Erratum: Phys.Rev.D 91, 059903 (2015)].

[48] **PIENU** Collaboration, A. Aguilar-Arevalo *et al.*, "Improved search for heavy neutrinos in the decay $\pi \to e\nu$," *Phys. Rev. D* **97** no. 7, (2018) 072012, arXiv:1712.03275 [hep-ex].

[49] **NA62** Collaboration, E. Cortina Gil *et al.*, "Search for heavy neutral lepton production in K^+ decays," *Phys. Lett. B* **778** (2018) 137–145, arXiv:1712.00297 [hep-ex].

[50] G. Bernardi *et al.*, "Search for Neutrino Decay," *Phys. Lett. B* **166** (1986) 479–483.

[51] **CHARM** Collaboration, F. Bergsma *et al.*, "A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-GeV to 2.8-GeV," *Phys. Lett. B* **166** (1986) 473–478.

[52] **NuTeV, E815** Collaboration, A. Vaitaitis *et al.*, "Search for neutral heavy leptons in a high-energy neutrino beam," *Phys. Rev. Lett.* **83** (1999) 4943–4946, arXiv:hep-ex/9908011.

[53] **T2K** Collaboration, K. Abe *et al.*, "Search for heavy neutrinos with the T2K near detector ND280," *Phys. Rev. D* **100** no. 5, (2019) 052006, arXiv:1902.07598 [hep-ex].

[54] **MicroBooNE** Collaboration, P. Abratenko *et al.*, "Search for long-lived heavy neutral leptons and Higgs portal scalars decaying in the MicroBooNE detector," *Phys. Rev. D* **106** no. 9, (2022) 092006, arXiv:2207.03840 [hep-ex].

[55] S. Alekhin *et al.*, "A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case," *Rept. Prog. Phys.* **79** no. 12, (2016) 124201, arXiv:1504.04855 [hep-ph].

[56] **DELPHI** Collaboration, P. Abreu *et al.*, "Search for neutral heavy leptons produced in Z decays," *Z. Phys. C* **74** (1997) 57–71. [Erratum: Z.Phys.C 75, 580 (1997)].

[57] **ATLAS** Collaboration, G. Aad *et al.*, "Search for Heavy Neutral Leptons in Decays of W Bosons Using a Dilepton Displaced Vertex in s=13 TeV pp Collisions with the ATLAS Detector," *Phys. Rev. Lett.* **131** no. 6, (2023) 061803, arXiv:2204.11988 [hep-ex].

[58] **CMS** Collaboration, "Search for long-lived heavy neutral leptons with displaced vertices in pp collisions at $\sqrt{s} = 13\,\text{TeV}$ with the CMS detector,".

[59] **LHCb** Collaboration, R. Aaij *et al.*, "Search for Majorana neutrinos in $B^- \to \pi^+\mu^-\mu^-$ decays," *Phys. Rev. Lett.* **112** no. 13, (2014) 131802, arXiv:1401.5361 [hep-ex].

[60] **FASER** Collaboration, A. Ariga *et al.*, "FASER's physics reach for long-lived particles," *Phys. Rev. D* **99** no. 9, (2019) 095011, arXiv:1811.12522 [hep-ph].

[61] **FASER** Collaboration, A. Ariga *et al.*, "Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC," arXiv:1812.09139 [physics.ins-det].

[62] L. Buonocore, F. Kling, L. Rottoli, and J. Sominka, "Predictions for neutrinos and new physics from forward heavy hadron production at the LHC," *Eur. Phys. J. C* **84** no. 4, (2024) 363, arXiv:2309.12793 [hep-ph].

[63] L. A. Anchordoqui *et al.*, "The Forward Physics Facility: Sites, experiments, and physics potential," *Phys. Rept.* **968** (2022) 1–50, arXiv:2109.10905 [hep-ph].

[64] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, "ForwArd Search ExpeRiment at the LHC," *Phys. Rev. D* **97** no. 3, (2018) 035001, arXiv:1708.09389 [hep-ph].

[65] **FASER** Collaboration, H. Abreu *et al.*, "Detecting and Studying High-Energy Collider Neutrinos with FASER at the LHC," *Eur. Phys. J. C* **80** no. 1, (2020) 61, arXiv:1908.02310 [hep-ex].

[66] **FASER** Collaboration, R. Mammen Abraham *et al.*, "First Measurement of νe and $\nu\mu$ Interaction Cross Sections at the LHC with FASER's Emulsion Detector," *Phys. Rev. Lett.* **133** no. 2, (2024) 021802, arXiv:2403.12520 [hep-ex].

[67] **FASER** Collaboration, R. Mammen Abraham *et al.*, "Neutrino Rate Predictions for FASER," arXiv:2402.13318 [hep-ex].

[68] **SND@LHC** Collaboration, G. Acampora *et al.*, "SND@LHC: the scattering and neutrino detector at the LHC," *JINST* **19** no. 05, (2024) P05067, arXiv:2210.02784 [hep-ex].

[69] **FASER** Collaboration, H. Abreu *et al.*, "First Direct Observation of Collider Neutrinos with FASER at the LHC," *Phys. Rev. Lett.* **131** no. 3, (2023) 031801, arXiv:2303.14185 [hep-ex].

[70] F. Kling and S. Trojanowski, "Forward experiment sensitivity estimator for the LHC and future hadron colliders," *Phys. Rev. D* **104** no. 3, (2021) 035012, arXiv:2105.07077 [hep-ph].

[71] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner, "EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider," *Phys. Rev. C* **92** no. 3, (2015) 034906, arXiv:1306.0121 [hep-ph].

[72] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "Hadronic interaction model Sibyll 2.3d and extensive air showers," *Phys. Rev. D* **102** no. 6, (2020) 063002, arXiv:1912.03300 [hep-ph].

[73] S. Ostapchenko, "Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model," *Phys. Rev. D* **83** (2011) 014018, arXiv:1010.1869 [hep-ph].

[74] S. Roesler, R. Engel, and J. Ranft, "The Monte Carlo event generator DPMJET-III," in *International Conference on Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000)*. 12, 2000. arXiv:hep-ph/0012252.

[75] S. Alioli, P. Nason, C. Oleari, and E. Re, "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX," *JHEP* **06** (2010) 043, arXiv:1002.2581 [hep-ph].

[76] C. Bierlich *et al.*, "A comprehensive guide to the physics and usage of PYTHIA 8.3," *SciPost Phys. Codeb.* **2022** (2022) 8, arXiv:2203.11601 [hep-ph].

[77] J. D. Bjorken and C. H. Llewellyn Smith, "Spontaneously Broken Gauge Theories of Weak Interactions and Heavy Leptons," *Phys. Rev. D* **7** (1973) 887.

[78] R. Shrock, "Decay $L^0 \to \nu_l \gamma$ in gauge theories of weak and electromagnetic interactions," *Phys. Rev. D* **9** (1974) 743–748.

[79] D. A. Bryman and R. Shrock, "Constraints on Sterile Neutrinos in the MeV to GeV Mass Range," *Phys. Rev. D* **100** (2019) 073011, arXiv:1909.11198 [hep-ph].

[80] V. Gribanov, S. Kovalenko, and I. Schmidt, "Sterile neutrinos in tau lepton decays," *Nucl. Phys. B* **607** (2001) 355–368, arXiv:hep-ph/0102155.

[81] J. C. Helo, S. Kovalenko, and I. Schmidt, "Sterile neutrinos in lepton number and lepton flavor violating decays," *Nucl. Phys. B* **853** (2011) 80–104, arXiv:1005.1607 [hep-ph].

[82] M. Drewes and B. Garbrecht, "Combining experimental and cosmological constraints on heavy neutrinos," *Nucl. Phys. B* **921** (2017) 250–315, arXiv:1502.00477 [hep-ph].

[83] P. Ballett, T. Boschi, and S. Pascoli, "Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector," *JHEP* **03** (2020) 111, arXiv:1905.00284 [hep-ph].

[84] P. Coloma, E. Fernández-Martínez, M. González-López, J. Hernández-García, and Z. Pavlovic, "GeV-scale neutrinos: interactions with mesons and DUNE sensitivity," *Eur. Phys. J. C* **81** no. 1, (2021) 78, arXiv:2007.03701 [hep-ph].

[85] M. Ovchynnikov, J.-L. Tastet, O. Mikulenko, and K. Bondarenko, "Sensitivities to feebly interacting particles: Public and unified calculations," *Phys. Rev. D* **108** no. 7, (2023) 075028, arXiv:2305.13383 [hep-ph].

[86] R. 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," *Phys. Rev. D* **24** (1981) 1232.

[87] D. Gorbunov and M. Shaposhnikov, "How to find neutral leptons of the νMSM?," *JHEP* **10** (2007) 015, arXiv:0705.1729 [hep-ph]. [Erratum: JHEP 11, 101 (2013)].

[88] F. Kling and S. Trojanowski, "Heavy Neutral Leptons at FASER," *Phys. Rev. D* **97** no. 9, (2018) 095016, arXiv:1801.08947 [hep-ph].

[89] J. C. Helo, M. Hirsch, and Z. S. Wang, "Heavy neutral fermions at the high-luminosity LHC," *JHEP* **07** (2018) 056, arXiv:1803.02212 [hep-ph].

[90] M. Drewes, J. Klarić, and J. López-Pavón, "New benchmark models for heavy neutral lepton searches," *Eur. Phys. J. C* **82** no. 12, (2022) 1176, arXiv:2207.02742 [hep-ph].

[91] J.-L. Tastet and I. Timiryasov, "Dirac vs. Majorana HNLs (and their oscillations) at SHiP," *JHEP* **04** (2020) 005, arXiv:1912.05520 [hep-ph].

[92] C. Antel *et al.*, "Feebly-interacting particles: FIPs 2022 Workshop Report," *Eur. Phys. J. C* **83** no. 12, (2023) 1122, arXiv:2305.01715 [hep-ph].

[93] G. A. Vasquez and J. Zamora-Saa, "Unveiling the heavy neutrino nature at LHCb," *Phys. Rev. D* **108** no. 5, (2023) 053008, arXiv:2307.01871 [hep-ph].

[94] M. V. Diwan, L. Anchordoqui, A. Ariga, T. Ariga, A. Barr, *et al.*, "Forward Physics Facility,". https://www.osti.gov/biblio/1972463.

[95] **FASER** Collaboration, H. Abreu *et al.*, "Search for dark photons with the FASER detector at the LHC," *Phys. Lett. B* **848** (2024) 138378, arXiv:2308.05587 [hep-ex].

[96] **FASER** Collaboration, "Search for Axion-like Particles in Photonic Final States with the FASER Detector at the LHC,". https://cds.cern.ch/record/2892328.

[97] **FASER** Collaboration, H. Abreu *et al.*, "The FASER W-Si High Precision Preshower Technical Proposal," tech. rep., CERN, Geneva, Mar, 2022. https://cds.cern.ch/record/2803084.

[98] R. Barouki, G. Marocco, and S. Sarkar, "Blast from the past II: Constraints on heavy neutral leptons from the BEBC WA66 beam dump experiment," *SciPost Phys.* **13** (2022) 118, arXiv:2208.00416 [hep-ph].

[99] A. D. Dolgov, S. H. Hansen, G. Raffelt, and D. V. Semikoz, "Heavy sterile neutrinos: Bounds from big bang nucleosynthesis and SN1987A," *Nucl. Phys. B* **590** (2000) 562–574, arXiv:hep-ph/0008138.

[100] I. Boiarska, A. Boyarsky, O. Mikulenko, and M. Ovchynnikov, "Constraints from the CHARM experiment on heavy neutral leptons with tau mixing," *Phys. Rev. D* **104** no. 9, (2021) 095019, arXiv:2107.14685 [hep-ph].

[101] J. Adhikary *et al.*, "Scientific program for the Forward Physics Facility," *Eur. Phys. J. C* **85** no. 4, (2025) 430, arXiv:2411.04175 [hep-ex].

[102] **FPF Working Groups** Collaboration, L. A. Anchordoqui *et al.*, "The Forward Physics Facility at the Large Hadron Collider," arXiv:2503.19010 [hep-ex].

[103] G. Cottin, J. C. Helo, M. Hirsch, A. Titov, and Z. S. Wang, "Heavy neutral leptons in effective field theory and the high-luminosity LHC," *JHEP* **09** (2021) 039, arXiv:2105.13851 [hep-ph].

[104] Y.-n. Mao, K. Wang, and Z. S. Wang, "Can we discover lepton number violation with the LHC far detectors?," *Phys. Rev. D* **108** no. 9, (2023) 095025, arXiv:2305.03908 [hep-ph].

[105] H. K. Dreiner, D. Köhler, S. Nangia, M. Schürmann, and Z. S. Wang, "Recasting bounds on long-lived heavy neutral leptons in terms of a light supersymmetric R-parity violating neutralino," *JHEP* **08** (2023) 058, arXiv:2306.14700 [hep-ph].

[106] Z. S. Wang, Y. Zhang, and W. Liu, "Searching for heavy neutral leptons coupled to axion-like particles at the LHC far detectors and SHiP," *JHEP* **01** (2025) 070, arXiv:2409.18424 [hep-ph].

[107] N. Bernal, K. Deka, and M. Losada, "Heavy neutral leptons without prejudice," *Phys. Rev. D* **111** no. 7, (2025) 075031, arXiv:2412.12271 [hep-ph].

[108] S. K. A, S. Das, A. Das, and S. Mandal, "Right handed neutrino production from Z' interactions in forward search experiments," arXiv:2508.10734 [hep-ph].

[109] F. Kling, T. Mäkelä, and S. Trojanowski, "Investigating the fluxes and physics potential of LHC neutrino experiments," *Phys. Rev. D* **108** no. 9, (2023) 095020, arXiv:2309.10417 [hep-ph].

[110] O. Mikulenko and M. Marinichenko, "Measuring lepton number violation in heavy neutral lepton decays at the future muon collider," *JHEP* **01** (2024) 032, arXiv:2309.16837 [hep-ph].

[111] P. D. Bolton, F. F. Deppisch, M. Rai, and Z. Zhang, "Probing the nature of Heavy Neutral Leptons in direct searches and neutrinoless double beta decay," *Nucl. Phys. B* **1010** (2025) 116785, arXiv:2212.14690 [hep-ph].

[112] M. Drewes, "Distinguishing Dirac and Majorana Heavy Neutrinos at Lepton Colliders," *PoS* **ICHEP2022** (2022) 608, arXiv:2210.17110 [hep-ph].

[113] **SND@LHC** Collaboration, D. Abbaneo *et al.*, "Input from the SND@LHC collaboration to the 2026 Update to the European Strategy for Particle Physics," arXiv:2503.24233 [hep-ex].

[114] S. K. A., A. Datta, and D. Marfatia, "Distinguishing between Dirac and Majorana neutrinos at FASER," arXiv:2508.00170 [hep-ph].

[115] B. Pontecorvo, "Tagging Direct Neutrinos. A First Step to Neutrino Tagging," *Lett. Nuovo Cim.* **25** (1979) 257–259.

[116] **NA62** Collaboration, E. Cortina Gil *et al.*, "First detection of a tagged neutrino in the NA62 experiment," *Phys. Lett. B* **863** (2025) 139345, arXiv:2412.04033 [hep-ex].

[117] A. Barr and I. Vivarelli, private communication, 2025.

[118] J. L. Feng, A. Hewitt, F. Kling, and D. La Rocco, "HNLCalc: GitHub repository." https://github.com/laroccod/HNLCalc, 2024.

[119] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics," *Eur. Phys. J. C* **71** (2011) 1554, arXiv:1007.1727 [physics.data-an]. [Erratum: Eur.Phys.J.C 73, 2501 (2013)].

[120] **FASER** Collaboration, R. Mammen Abraham *et al.*, "First Measurement of the Muon Neutrino Interaction Cross Section and Flux as a Function of Energy at the LHC with FASER," *Phys. Rev. Lett.* **134** no. 21, (2025) 211801, arXiv:2412.03186 [hep-ex].

[121] G. Apollinari, O. Brüning, T. Nakamoto, and L. Rossi, "High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report," Tech. Rep. FERMILAB-DESIGN-2015-02, Fermi National Accelerator Laboratory, 2015. https://lss.fnal.gov/archive/design/fermilab-design-2015-02.pdf.

[122] I. Zurbano Fernandez *et al.*, "High-Luminosity Large Hadron Collider (HL-LHC): Technical design report,". https://e-publishing.cern.ch/index.php/CYRM/issue/view/127.

[123] **ATLAS** Collaboration, "ATLAS Phase-II Upgrade Scoping Document," tech. rep., CERN, Geneva, 2015. https://cds.cern.ch/record/2055248.

[124] **ATLAS** Collaboration, G. Aad *et al.*, "Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics," arXiv:0901.0512 [hep-ex].

[125] A. Henriques, "The ATLAS Tile Calorimeter," Tech. Rep. ATL-TILECAL-PROC-2015-002, CERN, 2015. https://cds.cern.ch/record/2004868/files/ATL-TILECAL-PROC-2015-002.pdf.

[126] **ATLAS** Collaboration, G. Aad *et al.*, "Operation and performance of the ATLAS tile calorimeter in LHC Run 2," *Eur. Phys. J. C* **84** no. 12, (2024) 1313, arXiv:2401.16034 [hep-ex].

[127] **ATLAS** Collaboration, "Performance of the ATLAS Hadronic Endcap Calorimeter in beam tests: Selected results," Tech. Rep. ATL-LARG-PUB-2022-001, CERN, 2022. https://cds.cern.ch/record/2811731/files/ATL-LARG-PUB-2022-001.pdf.

[128] D. M. Gingrich *et al.*, "Construction, assembly and testing of the ATLAS hadronic end-cap calorimeter," *JINST* **2** (2007) P05005.

[129] A. Henriques and the ATLAS Collaboration, "The ATLAS Tile Calorimeter," Tech. Rep. ATL-TILECAL-PROC-2015-002, CERN, 2015. https://cds.cern.ch/record/2004868/files/ATL-TILECAL-PROC-2015-002.pdf.

[130] F. Pastore and the ATLAS Collaboration, "ATLAS TDAQ upgrades for Phase-2," Tech. Rep. ATL-DAQ-PROC-2024-008, CERN, 2024. https://cds.cern.ch/record/2907275/files/ATL-DAQ-PROC-2024-008.pdf.

[131] **ATLAS** Collaboration, "Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System,". https://cds.cern.ch/record/2285584/files/ATLAS-TDR-029.pdf.

[132] **ATLAS** Collaboration, "Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer,". https://cds.cern.ch/record/2285580/files/ATLAS-TDR-026.pdf.

[133] **FCC** Collaboration, A. Abada *et al.*, "FCC-hh: The Hadron Collider: Future Circular Collider Conceptual Design Report Volume 3," *Eur. Phys. J. ST* **228** no. 4, (2019) 755–1107.

[134] **ATLAS** Collaboration, G. Aad *et al.*, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett. B* **716** (2012) 1–29, arXiv:1207.7214 [hep-ex].

[135] **CMS** Collaboration, S. Chatrchyan *et al.*, "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC," *Phys. Lett. B* **716** (2012) 30–61, arXiv:1207.7235 [hep-ex].

[136] **ATLAS** Collaboration, G. Aad *et al.*, "A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery," *Nature* **607** no. 7917, (2022) 52–59, arXiv:2207.00092 [hep-ex]. [Erratum: Nature 612, E24 (2022)].

[137] **CMS** Collaboration, A. Tumasyan *et al.*, "A portrait of the Higgs boson by the CMS experiment ten years after the discovery.," *Nature* **607** no. 7917, (2022) 60–68, arXiv:2207.00043 [hep-ex]. [Erratum: Nature 623, (2023)].

[138] Q. Chang, X.-N. Li, X.-Q. Li, and F. Su, "Decay constants of pseudoscalar and vector mesons with improved holographic wavefunction," *Chin. Phys. C* **42** no. 7, (2018) 073102, arXiv:1805.00718 [hep-ph].

[139] D. Ebert, R. N. Faustov, and V. O. Galkin, "Relativistic treatment of the decay constants of light and heavy mesons," *Phys. Lett. B* **635** (2006) 93–99, arXiv:hep-ph/0602110.

[140] T. Feldmann, "Quark structure of pseudoscalar mesons," *Int. J. Mod. Phys. A* **15** (2000) 159–207, arXiv:hep-ph/9907491.

[141] **CLEO** Collaboration, M. Artuso *et al.*, "Improved measurement of B (D+ —> mu+ nu) and the pseudoscalar decay constant f(D+)," *Phys. Rev. Lett.* **95** (2005) 251801, arXiv:hep-ex/0508057.

[142] S. Stone, "Measurement of $D_s^+ \to \ell^+\nu$ and the Decay Constant F_{D_s}," *Conf. Proc. C* **060726** (2006) 1032–1035, arXiv:hep-ex/0610026.

[143] **Particle Data Group** Collaboration, M. Tanabashi *et al.*, "Review of Particle Physics. Kinematics," *Phys. Rev. D* **98** no. 3, (2018) 030001. https://pdg.lbl.gov/2018/reviews/rpp2018-rev-kinematics.pdf.

[144] **SHiP** Collaboration, E. Graverini, E. Van Herwijnen, and T. Ruf, "Mass dependence of branching ratios into HNL for FairShip,". https://cds.cern.ch/record/2133817.

[145] N. Carrasco, P. Lami, V. Lubicz, L. Riggio, S. Simula, and C. Tarantino, "$K \to \pi$ semileptonic form factors with $N_f = 2 + 1 + 1$ twisted mass fermions," *Phys. Rev. D* **93** no. 11, (2016) 114512, arXiv:1602.04113 [hep-lat].

[146] **JLQCD** Collaboration, S. Aoki, G. Cossu, X. Feng, H. Fukaya, S. Hashimoto, T. Kaneko, J. Noaki, and T. Onogi, "Chiral behavior of $K \to \pi l \nu$ decay form factors in lattice QCD with exact chiral symmetry," *Phys. Rev. D* **96** no. 3, (2017) 034501, arXiv:1705.00884 [hep-lat].

[147] D. Melikhov and B. Stech, "Weak form-factors for heavy meson decays: An Update," *Phys. Rev. D* **62** (2000) 014006, arXiv:hep-ph/0001113.

[148] **ETM** Collaboration, N. Carrasco, P. Lami, V. Lubicz, E. Picca, L. Riggio, S. Simula, and C. Tarantino, "Scalar and vector form factors of $D \to \pi l \nu$ and $D \to K l \nu$ decays with $N_f = 2 + 1 + 1$ Twisted femions," *PoS* **LATTICE2015** (2016) 261, arXiv:1511.04877 [hep-lat].

[149] G. Duplancic and B. Melic, "Form factors of B, $B_s \to \eta^{(\prime)}$ and D, $D_s \to \eta^{(\prime)}$ transitions from QCD light-cone sum rules," *JHEP* **11** (2015) 138, arXiv:1508.05287 [hep-ph].

[150] **HPQCD** Collaboration, H. Na, C. M. Bouchard, G. P. Lepage, C. Monahan, and J. Shigemitsu, "$B \to D l \nu$ form factors at nonzero recoil and extraction of $|V_{cb}|$," *Phys. Rev. D* **92** no. 5, (2015) 054510, arXiv:1505.03925 [hep-lat]. [Erratum: Phys.Rev.D 93, 119906 (2016)].

[151] X. J. Chen, H. F. Fu, C. S. Kim, and G. L. Wang, "Estimating Form Factors of $B_s \to D_s^{(*)}$ and their Applications to Semi-leptonic and Non-leptonic Decays," *J. Phys. G* **39** (2012) 045002, arXiv:1106.3003 [hep-ph].

[152] M. A. Ivanov, J. G. Korner, and P. Santorelli, "The Semileptonic decays of the B_c meson," *Phys. Rev. D* **63** (2001) 074010, arXiv:hep-ph/0007169.

[153] **Particle Data Group** Collaboration, P. A. Zyla *et al.*, "Review of Particle Physics," *PTEP* **2020** no. 8, (2020) 083C01.

[154] R. N. Faustov and V. O. Galkin, "Weak decays of B_s mesons to D_s mesons in the relativistic quark model," *Phys. Rev. D* **87** no. 3, (2013) 034033, arXiv:1212.3167 [hep-ph].

[155] F. Herren and M. Steinhauser, "Version 3 of RunDec and CRunDec," *Comput. Phys. Commun.* **224** (2018) 333–345, arXiv:1703.03751 [hep-ph].

Appendix A

HNL Production and Decay Modeling

A.1 HNL Production Modes

In this appendix, we present a comprehensive compilation of expressions for the leading processes for HNL production. Specifically, in the following five subsections we present the branching fractions for the processes $P \to lN$, $P \to P'lN$, $P \to VlN$, $\tau \to PN$, and $\tau \to \nu lN, \bar{\nu} lN$, where P and P' denote pseudoscalar mesons, V denotes vector mesons, l is a charged lepton, and N is the HNL. The decays implemented are shown in Table 3.1.

A.1.1 2-Body Pseudoscalar Decays $P \to lN$

The 2-body leptonic decays of pseudoscalar mesons into HNLs have branching fractions [41, 87]

$$B(P \to l_\alpha N) = \tau_P \, |U_\alpha|^2 \, \frac{G_F^2 m_P m_N^2 |V_P|^2 f_P^2}{8\pi} \left[1 - \frac{m_N^2}{m_P^2} + 2\frac{m_l^2}{m_P^2} + \frac{m_l^2}{m_N^2}\left(1 - \frac{m_l^2}{m_P^2}\right)\right] \times \sqrt{\left(1 + \frac{m_N^2}{m_P^2} - \frac{m_l^2}{m_P^2}\right)^2 - 4\frac{m_N^2}{m_P^2}} \,, \tag{A.1}$$

where $\alpha = e, \mu, \tau$; τ_P is the meson lifetime; U_α parameterizes the mixing with the active neutrino; and G_F is the Fermi constant. The decay constants f_P are given in Table A.1, and V_P denotes the relevant CKM matrix element; for example, for $P = \pi^+$, $V_P = V_{ud}$.

A.1.2 3-Body Pseudoscalar Decays $P \to P'lN$

For pseudoscalar mesons decaying semi-leptonically to pseudoscalar mesons, the differential branching fraction is [87]

$$\begin{aligned}\frac{dB\left(P \to P' l_\alpha N\right)}{dE_N dq^2} = \tau_P \ |U_\alpha|^2 \ \frac{G_F^2 \left|V_{PP'}\right|^2}{64\pi^3 m_P^2} \ c_P \ \Big\{ & f_-^2\left(q^2\right) \cdot \left[q^2\left(m_N^2 + m_l^2\right) - \left(m_N^2 - m_l^2\right)^2\right] \\ & + 2 f_+\left(q^2\right) f_-\left(q^2\right) \left[m_N^2\left(2m_P^2 - 2m_{P'}^2 - 4E_N m_P - m_l^2 + m_N^2 + q^2\right)\right. \\ & \left. + m_l^2\left(4E_N m_P + m_l^2 - m_N^2 - q^2\right)\right] \\ & + f_+^2\left(q^2\right) \left[\left(4E_N m_P + m_l^2 - m_N^2 - q^2\right)\left(2m_P^2 - 2m_{P'}^2 - 4E_N m_P - m_l^2 + m_N^2 + q^2\right)\right. \\ & \left. - \left(2m_P^2 + 2m_{P'}^2 - q^2\right)\left(q^2 - m_N^2 - m_l^2\right)\right]\Big\} \ ,\end{aligned} \tag{A.2}$$

where $q^2 = (p_l + p_N)^2$, E_N is the HNL energy in the P COM frame, and $V_{PP'}$ is the appropriate CKM matrix element for the process (e.g., for $D^+ \to \overline{K}^0 l^+ N$, $V_{PP'} = V_{cs}$). The constant c_P and the form factors $f_-(q^2)$ and $f_+(q^2)$ are defined below. Eq. (A.2) corrects a minor typo in Eq. (B2) of Ref. [87], which was missing a plus sign. Additionally, we replace the masses m_K and m_π in Eq. (B2) of Ref. [87] with m_P and $m_{P'}$, respectively.

To determine the branching fractions we must integrate Eq. (A.2) over the region [143]

$$\begin{aligned}(m_l + m_N)^2 \leq \ q^2 \ \leq \ (m_H - m_{H'})^2 \ , \\ E_N(m_{NH'}^{2\,\min}, q^2) \leq \ E_N \leq \ E_N(m_{NH'}^{2\,\max}, q^2) \ ,\end{aligned} \tag{A.3}$$

P	f_P (MeV)	V	f_V (MeV)
π^0 [138]	130.3	ρ^0 [139]	220
π^+ [138]	130.3	ρ^+ [139]	220
K^+ [138]	156.4	ω [139]	195
η [140]	78.4	K^{*+} [138]	204
η' [140]	−95.7	ϕ [138]	229
D^+ [141]	222.6		
D_s^+ [142]	280.1		
B^+ [87]	190		
B_c^+ [87]	480		

Table A.1: The pseudoscalar and vector meson decay constants used in this study.

where

$$
\begin{aligned}
E_N(m^2_{NH'}, q^2) &= \frac{q^2 + m^2_{NH'} - m^2_{H'} - m^2_l}{2m_H} \ , \\
m^{2\,\mathrm{min}}_{NH'} &= (E^*_N + E^*_{H'})^2 - \left(\sqrt{E^{*2}_N - m^2_N} + \sqrt{E^{*2}_{H'} - m^2_{H'}}\right)^2 \ , \\
m^{2\,\mathrm{max}}_{NH'} &= (E^*_N + E^*_{H'})^2 - \left(\sqrt{E^{*2}_N - m^2_N} - \sqrt{E^{*2}_{H'} - m^2_{H'}}\right)^2 \ ,
\end{aligned}
\tag{A.4}
$$

where $E^*_N = (m^2_{lN} - m^2_l + m^2_N)/2m_{lN}$, and $E^*_{H'} = (m^2_H - m^2_{lN} - m^2_{H'})/2m_{lN}$ are the energies of N and H' in the COM frame of $l - N$ system. Eq. (A.3) corrects a typo found below Eq. (2) of Ref. [144], where the q^2 dependence of the integration bounds for E_N was neglected.

The hadronic form factors $f_-(q^2)$ and $f_+(q^2)$ appearing in Eq. (A.2) are defined by [144]

$$
f_+\left(q^2\right) = \frac{f_+(0)}{\left(1 - q^2/m^2_{V'}\right)} \ , \tag{A.5}
$$

$$
f_0\left(q^2\right) = \frac{f_0(0)}{\left(1 - q^2/m^2_S\right)} \ , \tag{A.6}
$$

$$
f_0\left(q^2\right) = f_+\left(q^2\right) + \frac{q^2}{m^2_P - m^2_{P'}} f_-\left(q^2\right) \ . \tag{A.7}
$$

Here $m_{V'}$ and m_S are the masses of the vector and scalar resonances, respectively. These are determined from the quark transition of the decay (e.g., for $D^+ \to \overline{K}^0 l^+ N$, which involves a $c \to s$ quark transition, $m_{V'} = m_{D^{*+}_s}$, and $m_S = m_{D^+_s}$). It is important to note that $f_+(0)$ and $f_-(0)$ are determined by Eqs. (A.5) to (A.7), and therefore, we only provide values for $f_0(0)$, which are listed in Table A.2.

The constants c_P are determined by the quark content of the initial and final state mesons. As an example, consider the decay $D^+ \to \eta l^+ N$, which depends on the matrix element $\langle \eta | \bar{d} \gamma^\mu P_L c | D^+ \rangle$. The quark content of η and η' is related to η_8 and η_1 through a rotation matrix:

$$
\begin{pmatrix} \eta \\ \eta' \end{pmatrix} = \begin{pmatrix} \cos\theta_P & -\sin\theta_P \\ \sin\theta_P & \cos\theta_P \end{pmatrix} \begin{pmatrix} \eta_8 \\ \eta_1 \end{pmatrix} \ . \tag{A.8}
$$

Inserting $\eta_8 = (u\bar{u} + d\bar{d} - 2s\bar{s})/\sqrt{6}$ and $\eta_1 = (u\bar{u} + d\bar{d} + s\bar{s})/\sqrt{3}$ we obtain,

$$\eta = \left(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}}\right) u\bar{u} + \left(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}}\right) d\bar{d} + \left(-\frac{2\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}}\right) s\bar{s} \ . \tag{A.9}$$

From this, we observe that the matrix element $\langle\eta|\bar{d}\gamma^\mu P_L c|D^+\rangle$ is proportional to $\left(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}}\right)\langle d|\bar{d}\gamma^\mu P_L c|c\rangle$, where $\langle d|\bar{d}\gamma^\mu P_L c|c\rangle$ is the same matrix element encountered in the decay $D^0 \to \pi^- l^+ N$ and thus shares the same form factors. This implies that the parameter c_P for the decay $D^+ \to \eta l^+ N$ is $\left(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}}\right)^2$. The values of c_P are listed in Table A.2.

A.1.3 3-Body Pseudoscalar Decays $P \to Vl N$

For pseudoscalar mesons decaying semi-leptonically to vector mesons, the differential branching fraction is [87]

$$\begin{aligned}
\frac{dB\left(P \to V l_\alpha N\right)}{dE_N dq^2} = \tau_P \ |U_\alpha|^2 \frac{G_F^2 \left|V_{PV}\right|^2}{32\pi^3 m_P} \ c_V \ \Bigg\{ & \frac{f_2^2}{2}\left(q^2 - m_N^2 - m_l^2 + \omega^2 \frac{\Omega^2 - \omega^2}{m_V^2}\right) \\
& + \frac{f_5^2}{2}\left(m_N^2 + m_l^2\right)\left(q^2 - m_N^2 + m_l^2\right)\left(\frac{\Omega^4}{4m_V^2} - q^2\right) \\
& + 2f_3^2 m_V^2 \left(\frac{\Omega^4}{4m_V^2} - q^2\right)\left(m_N^2 + m_l^2 - q^2 + \omega^2 \frac{\Omega^2 - \omega^2}{m_V^2}\right) \\
& + 2f_3 f_5 \left[m_N^2\omega^2 + \left(\Omega^2 - \omega^2\right) m_l^2\right]\left(\frac{\Omega^4}{4m_V^2} - q^2\right) + 2f_1 f_2 \left[q^2\left(2\omega^2 - \Omega^2\right) + \Omega^2\left(m_N^2 - m_l^2\right)\right] \\
& + f_1^2 \left[\Omega^4\left(q^2 - m_N^2 + m_l^2\right) - 2m_V^2\left[q^4 - \left(m_N^2 - m_l^2\right)^2\right] + 2\omega^2\Omega^2\left(m_N^2 - q^2 - m_l^2\right) + 2\omega^4 q^2\right]\Bigg\} \\
& + \frac{f_2 f_5}{2}\left[\omega^2 \frac{\Omega^2}{m_V^2}\left(m_N^2 - m_l^2\right) + \frac{\Omega^4}{m_V^2} m_l^2 + 2\left(m_N^2 - m_l^2\right)^2 - 2q^2\left(m_N^2 + m_l^2\right)\right] \\
& + f_2 f_3 \left[\Omega^2\omega^2 \frac{\Omega^2 - \omega^2}{m_V^2} + 2\omega^2\left(m_l^2 - m_N^2\right) + \Omega^2\left(m_N^2 - m_l^2 - q^2\right)\right] \ ,
\end{aligned} \tag{A.10}$$

where $\omega^2 = m_P^2 - m_V^2 + m_N^2 - m_l^2 - 2m_P E_N$, and $\Omega^2 = m_P^2 - m_V^2 - q^2$. The constants c_V are conceptually identical to the constants c_P discussed in Sec. A.1.2.

Decay Channel	$f_0(0)$	c_P
$K^+ \to \pi^0$ [145]	0.970	1/2
$K_S \to \pi^+$ [146]	0.9636	1/2
$K_L \to \pi^+$ [146]	0.9636	1/2
$\overline{D}^0 \to \pi^+$ [147]	0.69	1
$\overline{D}^0 \to K^+$ [148]	0.747	1
$D^+ \to \pi^0$ [147]	0.69	1/2
$D^+ \to \eta$ [147]	0.69	$(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}})^2$
$D^+ \to \eta'$ [147]	0.69	$(\frac{\sin\theta_P}{\sqrt{6}} + \frac{\cos\theta_P}{\sqrt{3}})^2$
$D^+ \to \overline{K}^0$ [148]	0.747	1
$D_s^+ \to \overline{K}^0$ [148]	0.747	1
$D_s^+ \to \eta$ [149]	0.495	1
$D_s^+ \to \eta'$ [149]	0.557	1
$B^+ \to \pi^0$ [147]	0.29	1/2
$B^+ \to \eta$ [147]	0.29	$(\frac{\cos\theta_P}{\sqrt{6}} - \frac{\sin\theta_P}{\sqrt{3}})^2$
$B^+ \to \eta'$ [147]	0.29	$(\frac{\sin\theta_P}{\sqrt{6}} + \frac{\cos\theta_P}{\sqrt{3}})^2$
$B^+ \to \overline{D}^0$ [150]	0.66	1
$B^0 \to \pi^+$ [147]	0.29	1
$B^0 \to D^-$ [150]	0.66	1
$B_s^0 \to K^-$ [147]	0.31	1
$B_s^0 \to D_s^-$ [151]	−0.65	1
$B_c^+ \to D^0$ [152]	0.69	1
$B_c^+ \to \eta_c$ [152]	0.76	1
$B_c^+ \to B^0$ [152]	−0.58	1
$B_c^+ \to B_s^0$ [152]	−0.61	1

Table A.2: The parameters $f_0(0)$ and c_P, which appear in the $P \to P'lN$ branching fraction expressions of Eqs. (A.2) and (A.5) to (A.7), where $\theta_P = -11.5°$ [153].

The form factors are

$$f_1 = \frac{V}{m_P + m_V} \ , \tag{A.11}$$

$$f_2 = \left(m_P + m_V\right) A_1 \ , \tag{A.12}$$

$$f_3 = -\frac{A_2}{m_P + m_V} \ , \tag{A.13}$$

$$f_4 = \left[m_V \left(2A_0 - A_1 - A_2\right) + m_P \left(A_2 - A_1\right)\right] \frac{1}{q^2} \ , \tag{A.14}$$

$$f_5 = f_3 + f_4 \ . \tag{A.15}$$

A_0 and V have the form

$$f\left(q^2\right) = \frac{f(0)}{\left(1 - q^2/m^2\right)\left(1 - \sigma_1 q^2/m^2 + \sigma_2 q^4/m^4\right)} \ , \tag{A.16}$$

where $m = m_S$ for A_0, and $m = m_{V'}$ for V. A_1 and A_2 have the form

$$f\left(q^2\right) = \frac{f(0)}{\left(1 - \sigma_1 q^2/m_{V'}^2 + \sigma_2 q^4/m_{V'}^4\right)} \ . \tag{A.17}$$

The parameters m_S and $m_{V'}$ are determined from the decay mode in a manner identical to $P \to P'lN$. The parameters σ_1 and σ_2 are unique for each of the form factors A_0, A_1, A_2, and V. A summary of the form factor parameters can be found in Table A.3.

The form factors for certain B_c decays are a special case and all of the form factors A_0, A_1, A_2, and V are chosen to have the form [152]

$$f\left(q^2\right) = \frac{f(0)}{1 - q^2/m_{\text{fit}}^2 - \delta(q^2/m_{\text{fit}}^2)^2} \ , \tag{A.18}$$

where δ and m_{fit} are fitting parameters unique to each of the form factors A_0, A_1, A_2, and V. These fitting parameters for the relevant decays can be found in Table A.4.

	A_0			A_1			A_2			V			
Decay Channel	$f_0(0)$	σ_1	σ_2	$f_0(0)$	σ_1	σ_2	$f_0(0)$	σ_1	σ_2	$f(0)$	σ_1	σ_2	c_V
$D^0 \to \rho^-$ [147]	0.66	0.36	0	0.59	0.50	0	0.49	0.89	0	0.90	0.46	0	1
$D^0 \to K^{*-}$ [147]	0.76	0.17	0	0.66	0.3	0	0.49	0.67	0	1.03	0.27	0	1
$D^+ \to \rho^0$ [147]	0.66	0.36	0	0.59	0.50	0	0.49	0.89	0	0.90	0.46	0	1/2
$D^+ \to \omega$ [147]	0.66	0.36	0	0.59	0.50	0	0.49	0.89	0	0.90	0.46	0	1/2
$D^+ \to \overline{K^{*0}}$ [147]	0.76	0.17	0	0.66	0.3	0	0.49	0.67	0	1.03	0.27	0	1
$D_s^+ \to K^{*0}$ [147]	0.67	0.2	0	0.57	0.29	0.42	0.42	0.58	0	1.04	0.24	0	1
$D_s^+ \to \phi$ [147]	0.73	0.10	0	0.64	0.29	0	0.47	0.63	0	1.10	0.26	0	1
$B^+ \to \rho^0$ [147]	0.30	0.54	0	0.26	0.73	0.1	0.29	1.4	0.5	0.31	0.59	0	1/2
$B^+ \to \omega$ [147]	0.30	0.54	0	0.26	0.54	0.1	0.24	1.40	0.50	0.31	0.59	0	1/2
$B^+ \to \overline{D^{*0}}$ [147]	0.69	0.58	0	0.66	0.78	0	0.62	1.04	0	0.76	0.57	0	1
$B^0 \to \rho^-$ [147]	0.30	0.54	0	0.26	0.54	0.1	0.24	1.40	0.50	0.31	0.59	0	1
$B^0 \to D^{*-}$ [147]	0.69	0.58	0	0.66	0.78	0	0.62	1.04	0	0.76	0.57	0	1
$B_s^0 \to K^{*-}$ [147]	0.37	0.60	0.16	0.29	0.86	0.6	0.26	1.32	0.54	0.38	0.66	0.30	1
$B_s^0 \to D_s^{*-}$ [154]	0.67	0.35	0	0.70	0.463	0	0.75	1.04	0	0.95	0.372	0	1
$B_c^+ \to D^{*0}$ [152]	0.56	0	0	0.64	0	0	−1.17	0	0	0.98	0	0	1

Table A.3: The parameters $f_0(0)$, σ_1, σ_2, and c_V which appear in the $P \to VlN$ branching fraction expressions Eqs. (A.10), (A.16) and (A.17).

	A_0			A_1			A_2			V			
Decay Channel	$f_0(0)$	δ	m_{fit}	$f_0(0)$	δ	m_{fit}	$f_0(0)$	δ	m_{fit}	$f(0)$	δ	m_{fit}	c_V
$B_c^+ \to J/\psi$ [152]	0.68	1.40	8.20	0.68	0.052	5.91	−0.004	−0.004	5.67	0.96	0.0013	5.65	1
$B_c^+ \to B^{*0}$ [152]	−0.27	0.13	1.86	0.6	−1.07	3.44	10.8	−0.09	1.73	3.27	−0.052	1.76	1
$B_c^+ \to B_s^{*0}$ [152]	−0.33	0.13	1.86	0.4	−1.07	3.44	10.4	−0.09	1.73	3.27	−0.052	1.76	1

Table A.4: Parameters $f_0(0)$, δ, m_{fit}, and c_V for certain $B_c \to VlN$ decays appearing in branching fraction expressions Eq. (A.10) and Eq. (A.18).

A.1.4 2-Body Tau Decays $\tau \to PN$

The branching fraction for the $\tau \to PN$ decays are [87]

$$B(\tau \to PN) = \tau_\tau \left|U_\tau\right|^2 \frac{G_F^2 m_\tau^3 \left|V_P\right|^2 f_P^2}{16\pi} \left[\left(1 - \frac{m_N^2}{m_\tau^2}\right)^2 - \frac{m_P^2}{m_\tau^2}\left(1 + \frac{m_N^2}{m_\tau^2}\right)\right] \times \sqrt{\left(1 - \frac{(m_P - m_N)^2}{m_\tau^2}\right)\left(1 - \frac{(m_P + m_N)^2}{m_\tau^2}\right)}\,, \tag{A.19}$$

where V_P is the CKM matrix element of P (e.g., if $P = \pi$, then $V_P = V_{ud}$). The τ lepton can also decay into a vector meson. In this case, the branching fraction is [87]

$$B(\tau \to VN) = \tau_\tau \left|U_\tau\right|^2 \frac{G_F^2 m_\tau^3 \left|V_V\right|^2 f_V^2}{8\pi} \left[\left(1 - \frac{m_N^2}{m_\tau^2}\right)^2 + \frac{m_V^2}{m_\tau^2}\left(1 + \frac{m_N^2 - 2m_V^2}{m_\tau^2}\right)\right] \times \sqrt{\left(1 - \frac{(m_V - m_N)^2}{m_\tau^2}\right)\left(1 - \frac{(m_V + m_N)^2}{m_\tau^2}\right)}\,, \tag{A.20}$$

where V_V is the CKM matrix element for the corresponding vector meson. Note that in the literature, this branching fraction is often written in terms of the coupling $g_V = m_V f_V$ [21].

A.1.5 3-Body Tau Decays $\tau \to \nu l N, \bar{\nu} l N$

Lastly, we consider 3-body τ decays. The differential branching fractions are [87]

$$\frac{dB\left(\tau \to \nu_\tau l_\alpha N\right)}{dE_N} = \tau_\tau \left|U_\alpha\right|^2 \frac{G_F^2 m_\tau^2}{2\pi^3} E_N \left(1 - \frac{m_l^2}{m_\tau^2 + m_N^2 - 2E_N m_\tau}\right)\sqrt{E_N^2 - m_N^2} \times \left(1 + \frac{m_N^2 - m_l^2}{m_\tau^2} - 2\frac{E_N}{m_\tau}\right)\,, \tag{A.21}$$

and

$$\frac{dB\left(\tau \to \bar{\nu}_\alpha l_\alpha N\right)}{dE_N} = \tau_\tau \left|U_\tau\right|^2 \frac{G_F^2 m_\tau^2}{4\pi^3} \left(1 - \frac{m_l^2}{m_\tau^2 + m_N^2 - 2E_N m_\tau}\right)^2 \sqrt{E_N^2 - m_N^2} \times \left[(m_\tau - E_N)\left(1 - \frac{m_N^2 + m_l^2}{m_\tau^2}\right) - \left(1 - \frac{m_l^2}{m_\tau^2 + m_N^2 - 2E_N m_\tau}\right)\left(\frac{(m_\tau - E_N)^2}{m_\tau} + \frac{E_N^2 - m_N^2}{3m_\tau}\right)\right]\,, \tag{A.22}$$

where $\alpha \neq \tau$. The bounds of integration are

$$m_N \leq \ E_N \leq \ \frac{m_\tau^2 + m_N^2 - m_l^2}{2m_\tau} \ . \tag{A.23}$$

A.2 HNL Decay Rates

HNL decays are induced by both CC and NC interactions. These interactions are a result of mixing with the SM neutrino gauge eigenstates, which can be seen in the following interacting Lagrangian:

$$-\mathcal{L}^{CC} = \frac{g}{\sqrt{2}} \sum_{\alpha=\ e,\mu,\tau} U_\alpha^* W_\mu^+ \overline{N^c} \gamma^\mu P_L l_\alpha + \text{h.c.} \ , \tag{A.24}$$

$$-\mathcal{L}^{NC} = \frac{g}{2\cos\theta_W} \sum_{\alpha=\ e,\mu,\tau} U_\alpha^* Z_\mu \overline{N^c} \gamma^\mu P_L \nu_\alpha + \text{h.c.} \ . \tag{A.25}$$

In our analysis of HNL decays, we adopt the convention to Ref. [83], where decay rates are summed over all final neutrino mass eigenstates, giving the relationship

$$\Gamma(N \to \nu X) = \sum_{i=\ 1,2,3} \Gamma(N \to \nu_i X) = \sum_{\alpha=\ e,\mu,\tau} \Gamma(N \to \nu_\alpha X) + \Gamma(N \to \overline{\nu}_\alpha X) \ . \tag{A.26}$$

The benefit of this choice is that it allows one to remain agnostic about whether light active neutrinos, ν_α are treated as Dirac or Majorana particles [83], a convention choice that leads to non-physical, factor-of-two discrepancies throughout the literature [21, 43, 80, 81, 87].

A.2.1 Leptonic Decays

HNLs decay into purely leptonic final states through the CC and NC processes shown in Fig. 3.3(a)-(c). The decay widths are [84]

$$\Gamma(N \to \nu\, l_\alpha^- l_\alpha^+) = \frac{G_F^2 m_N^5}{96\pi^3} \sum_{\beta=\ e,\mu,\tau} |U_\beta|^2 \left[(C_1^l + 2\sin^2\theta_W \delta_{\alpha\beta}) f_1(x_\alpha) + (C_2^l + \sin^2\theta_W \delta_{\alpha\beta}) f_2(x_\alpha) \right], \tag{A.27}$$

$$\Gamma(N \to \nu\, l_\alpha^- l_\beta^+) = \frac{G_F^2 m_N^5}{192\pi^3} \left[|U_\alpha|^2 I_1(0, x_\alpha^2, x_\beta^2) + |U_\beta|^2 I_1(0, x_\beta^2, x_\alpha^2) \right] \quad (\alpha \neq \beta)\ , \tag{A.28}$$

$$\Gamma(N \to \nu\overline{\nu}\nu) = \sum_{\alpha=\ e,\mu,\tau} \sum_{\beta=\ e,\mu,\tau} \Gamma_\pm(N \to \nu_\alpha \overline{\nu}_\beta \nu_\beta) = \frac{G_F^2 m_N^5}{96\pi^3} \left(\sum_{\alpha=\ e,\mu,\tau} |U_\alpha|^2 \right)\ , \tag{A.29}$$

where $x_i = \frac{m_i}{m_N}$, $C_1^l = (1 - 4\sin^2\theta_W + 8\sin^4\theta_W)/4$, and $C_2^l = (2\sin^4\theta_W - \sin^2\theta_W)/2$. The kinematic functions, $f_1(x)$, $f_2(x)$, and $I_1(x,y,z)$ are given by:

$$I_1(x,y,z) = 12 \int_{(\sqrt{x}+\sqrt{y})^2}^{(1-\sqrt{z})^2} \frac{ds}{s} (s - x - y)(1 + z - s)\, \lambda^{\frac{1}{2}}(1, x, y) \lambda^{\frac{1}{2}}(1, s, z), \tag{A.30}$$

$$f_1(x) = (1 - 14x^2 - 2x^4 - 12x^6)\sqrt{1 - 4x^2} + 12x^4(x^4 - 1)L(x), \tag{A.31}$$

$$f_1(x) = 4\left[x^2(2 + 10x^2 - 12x^4)\sqrt{1 - 4x^2} + 6x^4(1 - 2x^2 + 2x^4)L(x) \right], \tag{A.32}$$

where

$$L(x) = \ln\left(\frac{1 - 3x^2 - (1 - x^2)\sqrt{1 - 4x^2}}{x^2(1 + \sqrt{1 - 4x^2})} \right), \tag{A.33}$$

$$\lambda(x,y,z) = x^2 + y^2 + z^2 - 2xy - 2yx - 2xz\ . \tag{A.34}$$

For $m_N < m_\pi$, the dominant decay is the invisible decay $N \to \nu\overline{\nu}\nu$. However, for $m_N > m_\pi$, the 2-body semi-leptonic pion decay, $\nu\pi$, becomes kinematically accessible and begins to dominate the HNL decay width. This can be observed explicitly in Figs. 3.4 and 3.5.

A.2.2 Hadronic Decays

Single Meson Decay

HNL decays into hadronic final states occur through the NC and CC processes depicted in Fig. 3.3 (d)-(e). As mentioned in Sec. 3.4, for $m_N < 1.0$ GeV, the total hadronic width is calculated via decay into single meson final states. The formulae for decay of HNLs into single mesons are [84]

$$\Gamma(N \to \nu P^0) = \left(\sum_{\beta=e,\mu,\tau} |U_\beta|^2 \right) \frac{G_F^2 f_P^2 m_N^3}{16\pi} (1 - x_P^2)^2 \ , \tag{A.35}$$

$$\Gamma(N \to l_\alpha^\mp P^\pm) = |U_\alpha|^2 |V_{q\bar{q}}|^2 \frac{G_F^2 f_P^2 m_N^3}{16\pi} \lambda^{\frac{1}{2}}(1, x_P^2, x_\alpha^2) \left[1 - x_P^2 - x_\alpha^2 (2 + x_P^2 - x_\alpha^2) \right] \ , \tag{A.36}$$

$$\Gamma(N \to \nu V^0) = \left(\sum_{\beta=e,\mu,\tau} |U_\beta|^2 \right) \frac{G_F^2 f_V^2 \kappa_V^2 m_N^3}{16\pi} (1 + 2x_V^2)(1 - x_V^2)^2 \ , \tag{A.37}$$

$$\Gamma(N \to l_\alpha^\mp V^\pm) = |U_\alpha|^2 |V_{q\bar{q}}|^2 \frac{G_F^2 f_V^2 m_N^3}{16\pi} \lambda^{\frac{1}{2}}(1, x_V^2, x_\alpha^2) \left[(1 - x_V^2)(1 + 2x_V^2) + x_\alpha^2 (x_V^2 + x_\alpha^2 - 2) \right] \ , \tag{A.38}$$

where $P^0 = \pi^0, \eta, \eta\prime$, $P^\pm = \pi^\pm, K^\pm, D^\pm, D_s^\pm$, $V^0 = \rho, \omega, \phi$, $V^\pm = \rho^\pm, K^{*\pm}$, $\kappa_\rho = 1 - 2\sin^2\theta_W$, $\kappa_\omega = -2\sin^2\theta_W/3$, and $\kappa_\phi = -\sqrt{2}(1/2 - 2\sin^2\theta_W/3)$. The meson decay constants $f_{P,V}$ are given in Table A.1.

Multi-meson Decays

In the regime, $m_N > 1.0$ GeV, the total hadronic width is instead calculated via decays into quarks, since in this region decays into multi-meson states become relevant to the total width. The tree-level NC and CC decays into quarks are given by [21]

$$\Gamma(N \to \nu q\bar{q}) = \left(\sum_{\beta=e,\mu,\tau} |U_\beta|^2 \right) \frac{G_F^2 m_N^5}{32\pi^3} \left[C_1^q f_1(x_\alpha) + C_2^q f_2(x_\alpha) \right] \ , \tag{A.39}$$

$$\Gamma(N \to l_\alpha^- u\bar{d}) = \Gamma(N \to l_\alpha^+ \bar{u} d) = |U_\alpha|^2 |V_{ud}|^2 \frac{G_F^2 m_N^5}{64\pi^3} I_1(x_\alpha^2, x_u^2, x_d^2) \ , \tag{A.40}$$

where the coefficients C_1^q and C_2^q are given by

$$C_1^{u,c} = \frac{1}{4}\left(1 - \frac{8}{3}\sin^2\theta_W + \frac{32}{9}\sin^4\theta_W\right), \qquad C_2^{u,c} = \frac{1}{3}\sin^2\theta_W\left(\frac{4}{3}\sin^2\theta_W - 1\right), \tag{A.41}$$

$$C_1^{d,s,b} = \frac{1}{4}\left(1 - \frac{4}{3}\sin^2\theta_W + \frac{8}{9}\sin^4\theta_W\right), \qquad C_2^{d,s,b} = \frac{1}{6}\sin^2\theta_W\left(\frac{2}{3}\sin^2\theta_W - 1\right). \tag{A.42}$$

As is discussed in Refs. [21, 84], the total hadronic width can be estimated using this tree-level decay along with a QCD loop correction that accounts for the hadronization process. This correction is estimated from hadronic tau decay. The loop corrections can be defined as

$$1 + \Delta_{\text{QCD}} \equiv \frac{\Gamma(\tau \to \nu_\tau + \text{hadrons})}{\Gamma_{\text{tree}}(\tau \to \nu_\tau u\bar{d}) + \Gamma_{\text{tree}}(\tau \to \nu_\tau u\bar{s})}\ , \tag{A.43}$$

where the correction up to $\mathcal{O}(\alpha_s^3)$ has been calculated to be

$$\Delta_{\text{QCD}} = \frac{\alpha_s}{\pi} + 5.2\frac{\alpha_s^2}{\pi^2} + 26.4\frac{\alpha_s^3}{\pi^3}\ . \tag{A.44}$$

We apply this correction to the NC decays $\nu q\bar{q}$ with $q = u, d, s$, and the CC decays $lu\bar{d}$ and $lu\bar{s}$, where $\alpha_s = \alpha_s(m_N)$. The package `rundec` is used to compute the α_s running; see Ref. [155]. Additionally, for the decay νss, a phase space suppression factor, $\sqrt{1 - 4m_K^2/m_N^4}$, is included to not overestimate this mode's contribution for $m_N < 2m_K$. This same approach is also applied to the decays τud and τus, with kinematic thresholds at $m_\tau + 2m_\pi$ and $m_\tau + m_\pi + m_K$, respectively.

A.2.3 HNL Total Decay Width

The total decay width of the HNL is obtained by summing over all possible final state decays. Since we are considering Majorana HNLs, the CC decays $N \to l_\alpha^+ H^-(\overline{u}d)$, $\nu\, l_\alpha^+ l_\beta^-$ and their charge conjugates are both possible, and, since they have equal decay rates, a factor of 2 must be included

when calculating the total rate. In the region $m_N < 1.0$ GeV, the total decay width is

$$\Gamma_N = \Gamma(N \to \nu\overline{\nu}\nu) + \sum_{\alpha=e,\mu,\tau} \Gamma(N \to \nu\, l_\alpha^+ l_\alpha^-) + \sum_{P^0=\pi^0\eta,\eta'} \Gamma(N \to \nu\, P^0) + \sum_{V^0=\rho^0,\omega,\phi} \Gamma(N \to \nu\, V^0)$$
$$+ \sum_{\alpha\neq\beta=e,\mu,\tau} 2\Gamma(N \to \nu\, l_\alpha^- l_\beta^+) + \sum_{\substack{P=\pi,K,D,D_s \\ \alpha=e,\mu,\tau}} 2\Gamma(N \to l_\alpha^- P^+) + \sum_{\substack{V=\rho,K^* \\ \alpha=e,\mu,\tau}} 2\Gamma(N \to l_\alpha^- V^+)\ . \tag{A.45}$$

For $m_N > 1.0$ GeV, this becomes

$$\Gamma_N = \Gamma(N \to \nu\overline{\nu}\nu) + \sum_{\alpha=e,\mu,\tau} \Gamma(N \to \nu\, l_\alpha^+ l_\alpha^-) + \sum_{q=u,d,s,c,b} \Gamma(N \to \nu\, q\overline{q})$$
$$+ \sum_{\alpha\neq\beta=e,\mu,\tau} 2\Gamma(N \to \nu\, l_\alpha^- l_\beta^+) + \sum_{\substack{u=u,c \\ d=d,s,b}} 2\Gamma(N \to l_\alpha^- u\overline{d})\ . \tag{A.46}$$